Children in the Forest: Towards a Canonical Problem of Spatio-Temporal Collaboration

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

Canonical problems are simplified representations of a class of real world problems. They allow researchers to compare algorithms in a standard setting which captures the most important challenges of the real world problems being modeled. Such examples are the block world for planning, two-player games for algorithms which learn the behavior of the opponent agent, or the "split the pie" game for a large class of negotiation problems.

In this paper we focus on negotiating collaboration in space and time, a problem with many important real world applications. Although technically a multi-issue negotiation, we show that the problem can not be represented in a satisfactory manner by the split the pie model. We propose the "children in the rectangular forest" (CRF) model as a possible canonical problem for negotiating spatio-temporal collaboration. By exploring a centralized and a peer-to-peer negotiation based solution, we demonstrate that the problem captures the main challenges of the real world problems while allows us to simplify away some of the computationally demanding but semantically marginal features of real world problems.

1. INTRODUCTION

Canonical problems allow researchers to compare algorithms in a standard setting which captures the most important challenges of the real world problems being modeled. Such examples are the block world for planning problems or two-player games for algorithms which learn the behavior of the opponent agent [3]. Canonical problems are close relatives to the standardized test beds used in AI research, and frequently, the implementation of the test bed follows a canonical problem. The test bed based controlled experimentation approach had generated controversies [4], with arguments which are just as well applicable to the more theoretical canonical problems as well. Ultimately, the main danger is that the researchers are focusing on problems which are particular only to the testbed, with little relevance to the real world. While a valid criticism, this observation should only make us more careful on selecting our canonical problems, such that they are indeed representative of the real world challenges they represent. Several current initiatives such as the trading agent competition or the Robocup Rescue Simulation League are positioning themselves towards a more accurate modeling of real world problems.

One of the canonical problems for agent negotiation is the "split the pie" game [1, 6] where the participants are negotiating over the partitioning of a pie. The game can be extended in a straightforward game to cover more complex issues. Multi-issue negotiations can be handled by having to split multiple pies, the agents total utility being a function of the pie shares. For reasons related to the computational complexity, the utility function is commonly represented by a weighted sum over the pie shares received by each agent. The agents might or might not know the utility function of their negotiation partner, thus various complete and incomplete information scenarios can be represented. Negotiations with deadlines are represented by imposing a limit on the negotiation rounds. Another, frequently considered aspect is the discount factors, the cost of extended negotiation is represented by the pies shrinking after every negotiation round with a factor of δ [7].

The features which make the split the pie game a good canonical problem is that it is representative of a large class of real world applications. By its simplifying assumptions, it enables a formal analysis of the different components of the negotiation process: the negotiation procedure, the negotiation protocol, the strategies deployed by the negotiation partners, their preferences over the outcomes—usually represented by their utility function and so on. Furthermore, through reducing negotiation problems to the split the pie model the fundamental identity of some negotiation problems can be revealed (which might not be immediately obvious in their original formulation). In some cases, the problem is completely equivalent to the canonical problem; in other cases certain transformations, approximations and simplifying assumptions are needed.

For instance, the split multiple pies game is an immediate representative for the problem of pirates dividing the bounty. However, it can also represent the negotiation over the price of a car through the following transformation. Let us consider the manufacturer's suggested retail price P_{MSRP} of the car and the dealer's invoice price $P_{invoice}$. In effect, the "pie" will be represented by the amount of money $P_{MSRP} - P_{invoice}$, which is the amount of profit split by the dealer and the buyer when negotiating a deal between them. Naturally, extended negotiations reduce this profit through inflation (which corresponds more or less exactly to the shrinking pie model), and/or through the cost of storage to the dealer, cost of rental car for the buyer and so on. These latter phenomena do not map directly in the canonical problem, but they can be approximated reasonably well by it.

There are, however, cases when the splitting multiple pies model, or its natural extensions can not capture the essential challenges of

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

978-81-904262-7-5 (RPS) ©2007 IFAAMAS

a class of real world problems. In this paper, we are considering the negotiation problems where embodied agents are negotiating agreements regarding collaborative actions in the spatio-temporal domain. The issues under negotiation include actions such as meeting at certain locations at certain points in time, performing actions at certain locations before, at, or after specific timepoints, or traversing certain paths with certain speeds. Although these are technically multi-issue negotiations, the split multiple pies problem does not capture the essence of these problems. We propose an alternative canonical problem, the Children in the Rectangular Forest scenario, and we argue that (a) it represents many of the fundamental aspects of this class of problems and (b) it is simple enough to serve as the foundation of formal analysis. Through an example negotiation model we show how the CRF scenario can be used to analyse the components of a proposed negotiation approach, including the negotiation procedure, negotiation protocol, strategies and utility functions.

The remainder of this paper is organized as follows. Section 2 describes several real life problems involving negotiation in space and time, and we discuss the common properties of these problems which classify them as a group and, at the same time, prevent them of being modeled with the splitting pie or related models. Section 3 describes our proposed canonical problem, the Children in the Rectangular Forest (CRF) problem. In Section 4 we introduce an example negotiation model for the CRF, and we show how in order to solve the CRF the model needs to address the critical issues of negotiating collaborative actions in space and time. We conclude in Section 5.

2. NEGOTIATING COLLABORATION IN SPACE AND TIME

2.1 Problems of spatio-temporal collaboration

Convoy formation in disaster response applications

Efficient response in face of natural disasters such as Hurricane Katrina in New Orleans, the Asian tsunami or the earthquake in Pakistan requires participants to form teams and coordinate their actions. In the immediate aftermath of a disaster previously safe areas might turn into unsafe or unaccessible. The environment might contain new sources of danger in the form of natural obstacles (damaged buildings) or even hostile agents (such as looters or stray dogs).

The tasks facing the rescue teams appear unpredictably. The discovery of a wounded person at a dangerous location creates a new task with specific logistics, protection and medical facets.

The organization of the rescue teams can not be pre-planned, and more often than not, a centralized coordination is not possible. For instance, in case of Hurricane Katrina, the central dispatcher unit of the police was flooded; the police could use their radios only as short-distance walkie-talkies. Furthermore, although some of the disaster management teams are pre-established, trained together and have a clear pattern of command and control, many teams are assembled on an ad hoc basis, as a response to emerging tasks. Teams are composed from heterogeneous groups of entities: persons, vehicles, service animals, and so on. Team members might not report to the same chain of command, might have communication problems and their interests might not be completely aligned. For instance, the state police and guerilla groups might cooperate in a rescue operation but resume hostilities after the emergency.

Thus, the organization of disaster response activities requires negotiation between agents with different interests. [5, 2] explored the topic of negotiating convoy formation in disaster response applications. The assumption for this problem is that agents have tasks associated with geographic locations, and in order to achieve those, they need to traverse areas which are accessible only to convoys, but not to individual agents. The negotiation between agents is concerned about temporal commitments regarding specific locations. For instance, in order to successfully join a convoy at location L_{join} the agent will make a commitment to reach that location before time t_{join}, while the convoy will make the commitment that it will leave that location only *after* time t_{join}. To allow the agent to plan ahead towards its task, the convoy takes the commitment that it will reach the pre-agreed location L_{leave} before t_{leave}. As the convoy will carry a set of commitments towards all its members, these commitments need to be taken into consideration when new agents are joining the convoy. Naturally, not every commitment is feasible, and the feasibility of a set of commitments needs to be evaluated together.

Transportation for elderly and disabled persons[1]

Many local transportation companies in the United States are providing door-to-door transportation services for the elderly or disabled persons who can not use the fixed route bus service. For instance, in Orlando, the ACCESS LYNX program [8] is providing more than 3100 scheduled passenger trips per day, using a large number of shuttle type vehicles. The vehicles might be operated by external contractors. Naturally, these services can not follow the one-person/one-trip model followed by taxis, as that would be prohibitively expensive. Requests for transportation are submitted by phone by the passengers. These requests need to be satisfied using the shuttles currently in service. The shuttles need to organize their path and schedule dynamically, such that they provide the best possible service. An incoming request modifies the path of the shuttle, which needs to make a detour to pick up the passenger. The transfer of passengers from one shuttle to another needs to be scheduled dynamically, and the rendezvous of the shuttles of the transfer point agreed upon.

Let us now envision a negotiation-based solution to the problem of efficient scheduling of the passenger transportation. This assumes that the shuttles are competing with each other for business, using performance measures such as total passenger miles, total number of passengers served or total passenger miles / total miles. In addition, the goals of the dispatcher use different, global performance measures, such as average time before pickup or average time to destination.

One way to organize the negotiation process is to allow only pairwise negotiations between the dispatcher and the shuttle. To satisfy a new transportation request, the dispatcher might contact several shuttles, and negotiate a modification of the route in order to pick up a new passenger. If the transportation request can not be satisfied with a single shuttle, the dispatcher might negotiate a rendezvous of two shuttles in order to arrange for a transfer of the passenger. This is an example of a co-negotiation; the offers of the dispatcher in one negotiation are conditioned by the evolution of the other negotiation process. Finally, even previous agreements can be revisited based on the set of new requirements.

The issues under negotiation can be described with spatio-temporal constraints; for instance, an offer might look like this: "pick up a passenger at location L_1 after time t_1, drop him at location L_2 before time t_2 but do not leave location L_2 before time t_3".

[1]The authors want to express their thanks to Derrick Babb for this application example.

More complex negotiation patterns can also be deployed. For instance, the drivers might be able to negotiate directly among themselves, the passengers might get involved in the negotiation as well, and the negotiation might include incentives and disincentives as well.

These problems are only two examples from the much wider class of problems which involve negotiation about collaborative actions in space and time. For instance, the act of passing in soccer (human or robotic) requires the players to agree on the trajectory of the ball, and the future location of the receiver player at a specific time. The act of carrying a piano on the stairs requires the carriers to agree on specific forces to be applied at specific locations and moments.

We conclude that negotiation about collaborative actions in space and time is a large class of problems with important practical applications. In the next section, we argue that these negotiation problems can not be adequately modeled by the split the pie game, and we discuss some of the features which need to be mirrored by a canonical problem attempting to represent this class of applications.

2.2 Defining characteristics of negotiating collaboration in space and time

Let us investigate the main reasons why the split the pie game can not serve as a valid canonical problem for negotiations concerning spatio-temporal collaborations. Although, there are many immediate differences in the formulations of the problems, not all of these are fundamental. For instance, our problem domain involves collaboration, while the split the pie model apparently involves a radical conflict of interests. This difference however, is only superficial. While the single issue "split the pie" is a zero-sum game, the multiple pie games are not[2]. With an appropriate utility function, the split the multiple pies game can be used to model the negotiation of collaborative activities. There are however, more fundamental differences, from which we highlight the following five:

(1) Heterogeneous types of issues.
(2) Non-monotonic valuation of issues.
(3) Evolving world (vs. discount factors)
(4) Offers need to be verified for feasibility
(5) Interaction between the negotiation time and physical time

Let us discuss these characteristics in more detail, by contrasting them to the set of negotiation problems modeled by the split the pie game.

(1) Heterogeneous types of issues and **(2) Non-monotonic valuation of the issues**

For the multi-issue split the pie game, all the issues are represented by a numerical value in the $[0, 1]$ interval. There is an assumption that all the different pies have an intrinsic, positive value; the ultimate goal of the negotiation partners being to acquire 100% of all the pies. Of course, the different agents might have different valuations for the different pies, and in a stretch, the utility function might be a non-linear function of the shares[3]. The issues in a split the pie game can be therefore characterized as *worth* values. It makes perfect sense to define the partial derivative of the utility function of agent a with respect to every component of the offer vector. All these partial derivatives will be non-negative, as the game assumes that the utility of a pie can not be negative.

$$\frac{\partial U^a([x_1 \dots x_n])}{\partial x_i} \geq 0 \quad \forall i \in \{1 \dots n\} \tag{1}$$

That is, the utility function of the agent in the split the pie game is monotonic in all the components of an offer. In fact, when the utility is a linear combination, the partial derivative will be a constant, and exactly the corresponding weight in the linear combination of utilities:

$$\frac{\partial U^a([x_1 \dots x_n])}{\partial x_i} = K_i^a \tag{2}$$

However, the situation is different for the case of negotiating spatio-temporal collaboration. Here, the issues under negotiation (that is, the components of an offer) can represent either (a) worth, (b) time values, and (c) points in the 2-dimensional or 3-dimensional space. For the worth-type values, the monotonicity considerations still apply. Things are somewhat more complicated for time values. If the time value represents, for instance, the arrival time to a destination, and we state that it is the goal of the agent to arrive as early as possible, the time value can be immediately mapped into a worth-type issue. However, if the time represents the time of a rendezvous (for instance, catching an airplane), the contribution of the issue to the utility corresponds to a step function: any value smaller than the target has the same value, while every value later than the target is worth 0.

The situation is even more complex for the spatial values. Although there are instances in which a location can be mapped to a worth value (for instance, considering the distance to the final destination), this worth value can not represent the point in the negotiation. Two agents can not agree to rendez-vous at "200 miles from New-York", they need to decide on a specific location. There is no objective, positive or negative value in a certain rendez-vous point, its value becomes evident only in the context of the remainder of the offer (the time of the rendez-vous, the path of the convoy after the rendez-vous and so on).

(3) Evolving world (vs. discount factors)

Many negotiation models consider that the utility of a certain offer depends on the moment in the negotiation process when it was presented. Most studies of the split the pie game consider that the value of the issues under negotiation decreases in time, this feature being modeled with discount factors δ, which shrink the pie after every negotiation round. This is a good model of many (but not all[4]) practical situations, and it has the analytical advantage that it models the incentive of the participants to conclude a negotiation. In the case of negotiating spatio-temporal collaboration, we have to consider that (a) the agents may be moving during the negotiation and (b) the time passes. That is, the value of the offer depends on the current location of the agent, as well as the current time. Note that this can not be modeled with discount factors; the value of the offer does not necessarily decrease in time. For instance, the value of a rendezvous point increases if the agent moves on a path which takes it closer to the proposed point, and it starts to decrease once the agent passes the closest location and the distance increases.

(4) Feasibility of the offers

In the split the pie model every correctly formed offer is feasible. However, for spatio temporal collaborations, there are offers which, altough potentially of high value, are not feasible because of the physical world limitations. For instance, one of the participants might propose a rendez-vous at a location and time which

[2] Because different agents can have different valuations of the different pies, and thus they can reach deals which are advantageous to both of them.

[3] Although most research studies consider the utility to be an additive, linear combination of the values.

[4] Even in purely worth oriented domains, it is possible that the value of the "pie" increases during the negotiation, consider for instance negotiations concerning real-estate deals.

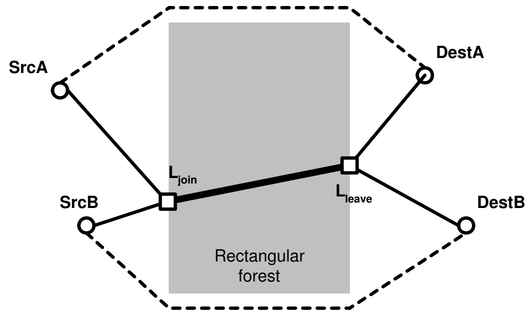

Figure 1. The Children in the Rectangular Forest problem. The trajectories associated with the conflict deal are shown with an interrupted line, while the trajectories corresponding to a possible agreement are shown with a continuous line.

can not be made by the other participant. The feasibility of an offer can not be evaluated in advance, as it is dependent on the current state of the world. In incomplete information settings, the feasibility needs to be verified by all participants separately. Verifying whether an offer is feasible or not can be computationally expensive, as it might involve path planning and estimation of the future state of the world. This has a significant impact in the offer generation step, as the offeror needs to evaluate and verify for feasibility the offers before making them.

(5) The interaction between the negotiation time and physical time

The shrinking pie model abstracts away the physical time, and replaces it with the discrete time model of the negotiation turns. This is a very powerful feature of the model, and a major help in analysis. However, in the class of problems considered by us, we can not make this simplifying step. As we have shown in point (3) above, the agents are acting in an evolving physical world concommitently with the negotiation process. The time taken for the negotiation, including the overhead of offer exchanging and the computational time to generate and evaluate the offers have a direct impact on the outcome of the negotiation. For instance, in fast, real time applications, such as Robocup soccer there is simply no time for exchanging and evaluating multiple offers. In fact, the real-world soccer is increasingly moving towards pre-trained "set pieces", showing that at the speed of human path planning there is no time for evaluating even a single offer - the only offers which can be made are the ones whose values are pre-calculated through previous agreements in the training sessions!

Even if there is time to evaluate several offers, the outcome of negotiation will be different for agents with slow or fast computational facilities (software and/or hardware), and naturally, the outcome is different if one of the agents has more powerful computational facilities than the other.

We conclude that the class of problems representing negotiation about spatio-temporal collaboration have a series of features which are not correctly represented by the "split the pie" model. We are looking for a canonical problem which reflects these features, and at the same time is as simple as possible.

3. THE PROPOSED CANONICAL PROBLEM: CHILDREN IN THE RECTANGULAR FOREST

Let us assume that two children *A* and *B* are going from their departure locations *SrcA* and *SrcB* at one side of a rectangular forest of size $h \times w$, and they are going to their destinations *DestA* and *DestB* on the other side of the forest. The children were told not to go alone in the forest, but they can potentially traverse the forest together. The walking speed of the children can be different, when together, they will walk with the velocity of the slower child. The problem of the children is to use negotiation to agree on whether they will join up to traverse the forest together, and if yes, the join point, the join time, the point where they separate and the time they separate. If the negotiation fails, the conflict deal is a path which goes around the forest (see Figure 1).

We assume a rational behavior from the two children, that is, they will prefer the choice which takes them faster to their destinations.

Let us consider some properties of the problem.

Property 1. *The optimal trajectories of the conflict deal and the collaboration deal are a sequence of straight segments.*

The proof of this property is relying on elementary geometrical properties and due to lack of space it is left to the reader. What remains to be discussed is whether a rational agent would choose a curvilinear trajectory *during the negotiations* (note that the Property 1 only talks about the trajectories associated with the *deals*). The surprising answer is, yes. Let us consider an agent which might estimate the probability of a deal while waiting for an answer from a negotiation partner. An agent which is almost sure of a deal might move towards the predicted rendez-vous point, while an agent which is almost sure of the conflict would move in the direction of the conflict deal trajectory. Between these two extremes, the agents might plan for an optimal trajectory, which strikes a balance between these choices. As the agents are moving during the negotiation time, the probability and utility of the deal changes in time. An optimal path therefore will be curvilinear, with edge points corresponding to events in the negotiation, such as the receiving of a new offer or the finishing of a utility calculation.

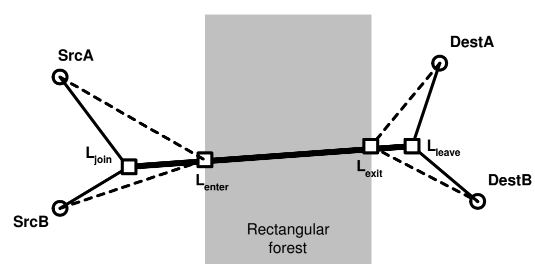

Figure 2. For any possible join and leave location, there is a join or leave location on the edge of the forest of equal or larger utility for both agents.

Property 2. *Deals where the join location is not on the edge of the forest are not Pareto optimal, or there is a deal where the join location is on the edge of the forest which provides the same utility to the agents.*

Proof: The proof of this property is a simple application of the triangle inequality. Let us make the assumption that there is deal

$(L_{join}, t_{join}, L_{leave}, t_{leave})$, where the join location L_{join} is not on the edge of the forest (see Figure 2). Let us consider the point of entering the forest L_{enter} and the time to enter the forest t_{enter}. Then, from the triangle inequality:

$$dist(SrcA, L_{join}) + dist(L_{join}, L_{enter}) \geq dist(SrcA, L_{enter})$$
$$dist(SrcB, L_{join}) + dist(L_{join}, L_{enter}) \geq dist(SrcB, L_{enter})$$

That is we can build an offer $(L_{enter}, t_{enter}, L_{leave}, t_{leave})$, which is at least as good as the previous offer. In fact, if the strong inequality holds, we can build an offer which contains time values which are lower, that is, the offer as a whole has a higher utility.

□.

Property 3. *Deals where the leave location is not on the edge of the forest are not Pareto optimal, or there is a deal where the leave location is on edge of the forest which provides the same utility to the agents.*

The proof of this property is analogous to the previous one (see Figure 2).

In the following, we evaluate the proposed problem in the light of the five characteristics of the spatio-temporal collaboration problems we have highlighted in the previous section.

(1) Heterogeneous issues

The CRF problem, as stated above, is a 4-issue negotiation, with two issues being points in the 2-dimensional space and two issues being time values. Depending of the assumptions of the problem, this can be farther simplified. For instance, if the velocities of both agents are known, the leave time is completely determined by the join and leave locations and the join time, effectively reducing the problem to a 3-issue negotiation. By applying Properties 2 and 3, we can reduce the negotiation of the locations to a negotiation only on the y axis, another important simplifying factor.

The problem can be immediately extended to include a worth type issue, for instance by one of the agents offering some compensation to the other agent in exchange for a more favorable leave location.

(2) Non-monotonic valuation of issues

The issues under negotiation do not contribute linearly and monotonically to the utility of the agents. For instance, the join location and time has only an indirect impact on the time of arrival to destination, by its impact over what leave locations and times are feasible.

(3) Evolving world

The agents are moving during the negotiation, which makes the value of an offer dependent on the time at which it is evaluated and the state of the world. For instance, if an agent decided that an agreement is very likely, it moves towards the predicted join location, through this action increasing the value of the predicted deal. Alternatively, if an agent assumes that a deal is highly unlikely, it will move on the conflict deal trajectory, making the value of the offer lower and lower as it moves farther and farther away from the proposed join location.

(4) Offers need to be verified for feasibility

Not every offer is feasible in the CRF world, due to the limited velocities of the agents'. The feasibility conditions of an offer $(L_{enter}, t_{enter}, L_{leave}, t_{leave})$ when the locations of the agents are L_A and L_B and the current time is $t_{current}$ are described by the following inequalities:

$$\frac{dist(L_A, L_{join})}{v_A} \leq t_{join} - t_{current}$$
$$\frac{dist(L_B, L_{join})}{v_B} \leq t_{join} - t_{current}$$
$$\frac{dist(L_{join}, L_{leave})}{v_A} \leq t_{leave} - t_{join}$$
$$\frac{dist(L_{join}, L_{leave})}{v_B} \leq t_{leave} - t_{join} \quad (3)$$

Notice that in a scenario where the agents do not know the other agents location and velocity, the first and the third condition can be evaluated only by agent A, while the second and fourth condition only by agent B. We also note, however, that while the feasibility of an offer is described by multiple conditions and exhibits interesting variations depending on the amount of information disclosure, the calculations themselves are simple and computationally inexpensive.

(5) Interaction between negotiation time and physical time

The negotiation between two agents happens in the physical time of the movement. If a negotiation round i takes t_i time, the agents will move $v_A t_i$ and $v_B t_i$ respectively on their planned trajectories. How much each negotiation round takes depends on the algorithms deployed by the agents, the computational power of the agent whose turn it is, and the messaging overhead. Various scenarios can be modelled with relative ease (negotiation between a fast and a slow agent, negotiations on different physical scales etc).

While there is a rich set of modelling possibilities, the calculations are sufficiently simple to make both simulation based and (at least as long as straight segment based trajectories are considered) analytical approaches possible.

We conclude that the Children in the Rectangular Forest problem exhibits all the five characteristics of the class of problems of negotiating for spatio-temporal collaboration. Through immediate and natural extensions many additional interesting and behaviors can be modelled. At the same time, significant simplifications can be applied to the calculations of utility and feasibility, and the model is sufficiently simple to make analytical study possible.

The CRF problem can be seen as a simplified version of the disaster response convoy formation problem. However, with some transformations, it can also serve as the model for the other problems mentioned before. By considering one of the agents to be the shuttle and the other agent the passenger it can model the problem of the transportation of elderly persons. One additional modification would be that in this case that the two agents will move with the velocity of the *faster* agent after joining. The case of soccer pass can be modeled by considering one of the agents to be the ball while the other agent being the player waiting for a pass. The common feature of these problems is that they are all dealing with negotiation about rendez-vous at certain point and time - one possible name for these problems being *spatio-temporal coincidence problems*.

4. AGENT STRATEGIES FOR THE CRF PROBLEM

4.1 Generalities

In the following we introduce several ways to approach the CRF problem. It is not our intention to give a definite solution to the problem. In fact, our proposition is that the CRF problem, depending on the assumptions, is a rich source of problems, whose solutions can then be applied to the whole collection of real world applications dealing with spatio-temporal coalitions. Thus, the remainder of this section is an attempt for a beginning of a consistent analysis of the problem.

We start with the assumption that the agent's goal is to reach as soon as possible their destination locations *DestA* and *DestB*.

Let us first consider a complete information setting where all the agents have a complete knowledge of all the relevant information about the current state of CRF world. That includes the current location of the agents L_A and L_B, their destinations *DestA* and *DestB* and there velocities v_A and v_B. A coalition offer under these conditions can be expressed through the pair of locations (L_{join}, L_{leave}), where L_{join} will be on the entrance edge of the forest while L_{leave} on the exit edge. There is no need for time values in the offer because it is natural that the agents will move with their greatest available velocities when alone and they will move with the minimum of their shared velocities while together. Thus the time values can be calculated from the location values.

Let us now calculate the time of reaching the derivation for the two agents considering the offer (L_{join}, L_{leave}). The time is composed of three parts: (a) the time to the rendez-vous, which will be the time it takes for the *last* agent to reach the rendez-vous point, (b) the time of the common travel of the agents from L_{join} to L_{leave}, which will happen with the velocity of the *slowest* agent and (c) the travel of the agent from L_{leave} to its respective destination, which will happen with the agents highest available velocity.

$$t_A = \max\left(\frac{dist(L_A, L_{join})}{v_A}, \frac{dist(L_B - L_{join})}{v_B}\right) + \frac{dist(L_{join}, L_{leave})}{\min(v_A, v_B)} + \frac{dist(L_{leave}, DestA)}{v_A}$$

$$t_B = \max\left(\frac{dist(L_A - L_{join})}{v_A}, \frac{dist(L_B - L_{jo\,in})}{v_B}\right) + \frac{dist(L_{join} - L_{leave})}{\min(v_A, v_B)} + \frac{dist(L_{leave} - DestB)}{v_B}$$

where $dist(L_1, L_2) = \sqrt{(X_1 - X_2)^2 + (Y_1 - Y_2)^2}$ denotes the Euclidean distance between locations L_1 and L_2.

We note that in the context of the CRF problem the only non-constant components are the y coordinates of the join and leave locations, thus we have $t_A = t_A(y_{join}, y_{leave})$ and $t_B = t_B(y_{join}, y_{\mathrm{leave}})$. This opens a door to a whole series of mathematical calculations; for instance agent A can calculate the join and leave points which would be optimal from its point of view by finding the values of y_{join} and y_{leave} which minimize the value of $t_A(y_{join}, y_{leave})$. Unfortunately, due to the several min and max functions in its expression, the function t_A, although continuous in every point, it is not differentiable in every point, making analythical solutions difficult.

4.2 Peer-to-Peer Negotiation with Incomplete Information for the CRF Problem

Let us now outline a negotiation-based solution for the CRF problem. As before, our goal is not to find a definitive solution, but to start an investigation into the nature of protocols, strategies and analysis methods which can be deployed in the setting on this problem.

We start with the assumption of self-interested agents which initially do not have any information about the opponent agent. That is, an agent A knows its current location L_A, its destination *DestA* and its velocity v_A, but does not know these values for the negotiation partner B. During negotiation, however, the A can acquire some information at least about the preferences of B, through the *offers* made by B and through the *evaluations* of the offers made by A. With our assumptions, we will describe an offer made by agent A at negotiation turn i as the quadruplet of the join location, join time, leave location and leave time: $O_i^A = \{L_{join}, t_{join}, L_{leave}, t_{leave}\}$. The evaluation of the offer made by agent A by agent B follows the same format as the original offer, but with the time values potentially changed: $E_i^A = \{L_{join}, t'_{join}, L_{leave}, t'_{leave}\}$. The evaluation of an offer, intuitively, is the "best attempt" of the agent to meet the requirements of an offer.

The negotiation protocol we are considering requires the agents to exchange messages, which contain the pair of (a) a new offer generated by the agent and (b) the evaluation of the offer made by the other agent in the previous round. Any of these components can be missing.

Let us now consider the flow of the negotiation between the agents A and B, by describing the choices of agent A at a certain moment in the negotiation process. We will separate our discussion in two cases: (a) the last message the agent received was of the form (O_{i-1}^B, E_{i-2}^A) or $(O_{i-1}^B, \emptyset)$ and (b) the last message received was of the form $(\emptyset, E_{i-2}^A)$.

(a) The last message received by A **was** (O_{i-1}^B, E_{i-2}^A) **or** $(O_{i-1}^B, \emptyset)$

If agent B sent an evaluation of the previous offer, that involves that the new offer is feasible for agent B and it represents an improvement over previous offers. This information might be useful for A for it's strategy to generate future offers. However, the fact that a new offer is present in B's message means that it does not accept the offer.

Agent A proceeds to evaluate the offer $(O_{i-1}^B$. Depending on the outcome of this evaluation, A might choose to answer with different messages.

- B's offer is worse than the conflict deal for A. A generates a new offer and sends a message of the type $(O_i^A, \emptyset)$.
- B's offer is better than the conflict deal, but it does not satisfy A. A generates a new offer and sends it back together with the evaluation of B's offer (O_i^A, E_{i-1}^B).
- A decides to accept B's offer. The pre-requisite for this is that the offer is *feasible* for A and better than the conflict deal. In this case A does not generate a new offer and only sends B its evaluation $(\emptyset, E_{i-1}^B)$.
- A decides to stop the negotiation, by sending a message $(\emptyset, \emptyset)$.

(b) The last message received by A **was** $(\emptyset, E_{i-2}^A)$

By sending a message with the evaluation but without a counteroffer, B had indicated its willingness to accept the join and leave locations proposed by A. However the evaluation might contain different time points for joining and leaving, showing the best effort of the agent B. At this point, agent A needs to decide whether the evaluated offer is still desirable for it. If true, an agreement message is sent, and the two agents will rendez-vous at the join point.

If the evaluated offer is not attractive for agent A, it can continue the negotiation by sending a message of type $(O_i^A, \emptyset)$, or it can stop the negotiation and take the conflict deal by sending $(\emptyset, \emptyset)$.

4.3 Examples of Negotiation Strategies

In the previous section we described a possible negotiation protocol for the CRF problem. The protocol, however, describes only the negotiation flow. For a complete description of the negotiation, we also need to know the negotiation strategies of the agents. The negotiation strategy is responsible for deciding which offers are acceptable for the agent, and how much it can concede in the counteroffer and what kind of new offers will the agent create. The negotiation strategy is, naturally, dependent on the negotiation protocol - for instance, a protocol which allows multiple alternative offers to be sent in a single message would require a different strategy than a protocol which allows only one. But for any negotiation protocol a large number of different strategies are possible. The negotiation partners necessarily need to conform to the same protocol, but they may have different strategies.

In the following, we succinctly describe two simple negotiation strategies, both of them using the protocol outlined in the previous section. Again, these strategies are presented as illustrations of the space of approaches which can be modelled with the CRF problem.

First, we will define a normalized utility function. The arrival time to destination can not serve at this purpose, because it takes different values for the two agents. We propose the utility function of an agent A which receives an offer O to be:

$$\mathrm{U^A(O_B)} = \frac{t_{conflict} - t_O}{t_{\mathrm{conflict}} - t_{\mathrm{opt}}} \tag{4}$$

where t_O is the time to destination for the agent if the offer is accepted, $t_{conflict}$ is the time to destination in case of the conflict deal and t_{opt} is the optimal time to destination. t_{opt} is the time to destination of an agent assuming that the forest is removed; in practical terms this can be achieved if the negotiation partner agrees on leave and join points on the straight line connecting the agent to its destination and its velocity is larger or equal to the velocity of the agent. The maximum value of the utility is 1, but it can have negative values, which represent offers which are worse than the conflict deal.

4.4 Strategy 1: Supervised negotiation

One of the problems faced in every negotiation process is once the negotiation reaches the conciliation stage (that is, the exchanged offers are better then the conflict deal for both agents), there is a need for some sort of external pressure to force the selection of one of the possible deals. The first strategy we introduce uses an external mediator to select the offer on which the agreement will be based. While the two agents are judging offers based on their own utility functions U^A and U^B, the mediator agent will select offers based on its own, "supervisors' view" utility function, which in our case will be defined as the sum of the two utilities:

$$U^S(O) = U^A(O) + U^B(O) \tag{5}$$

An equivalent way of saying is that the supervisor tries to minimize the sum of the time to destination of the two agents (for a shape of this surface in function of the join and leave points, see Figure 3).

To guarantee that the supervisor chooses a deal which is rational for both agents (that is, it has positive utility), the mediator will enter the negotiation process only after the conciliation phase is entered. At this point the two agents will concede a fixed amount *concedingPace* at every negotiation turn, until the offers meet. At every turn, the agents are sending their offer evaluations to the mediator. At the end of the negotiation, the mediator selects the offer which maximizes the U^S utility function, and this will be the deal on which the agents will agree upon. The algorithm of this negotiation strategy is described by the following pseudocode.

agent A received a message from agent B;
if the message is $(O^B_{i-1}, \emptyset)$ **then**
    $E^B_{i-1} \leftarrow evaluation(O^B_{i-1})$ and $U_{EB} \leftarrow Utility(E^B_{i-1})$;
    if $U_{EB} \geq 0$ **then**
        propose$((O^A_i, E^B_{i-1}))$ where O^A_i concedes from O^A_{i-1} towards E^B_{i-1};
    else
        propose$((O^A_i, \emptyset))$ where O^A_i concedes from O^A_{i-1} towards E^B_{i-1};
    end if
    i=i+1;
end if
if the message is (O^B_{i-1}, E^A_{i-1}) **then**
    $E^B_{i-1} \leftarrow evaluation(O^B_{i-1})$ and $U_{EB} \leftarrow Utility(E^B_{i-1})$;
    if $U_{EB} \geq 0$ **then**
        Enter the conciliation phase;
        if $O^B_{i-1} = E^A_{i-1}$ in location **then**
            contact mediator, which finds the point that $\max(U_{EA} + U_{EB})$;
            form an agreement;
        else
            concede the offer, propose$((O^A_i, E^B_{i-1}))$ where O^A_i concedes from O^A_{i-1} towards E^B_{i-1};
        end if
    else
        insist on offer $O^A_i = O^A_{i-1}$ until agent B's offer has a positive utility;
    end if
    i=i+1;
end if

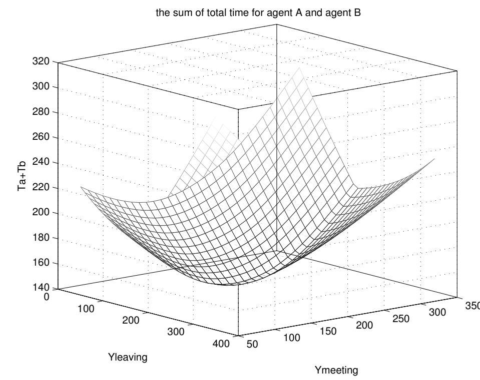

Figure 3. The sum of the time to destination of the two agents in function of the location of the join and leave points on the two edges of the forest. The supervisor tries to optimize this function subject to the constraint that the two times should be individually better than the conflict deal for both agents.

This method is time-consuming because it amounts to an exhaustive exploration of the solution space. The time complexity of the method is $O(|O^A_1 - O^B_1|/concedingPace)$, where $|O^A_1 - O^B_1|$ is the distance between the first offer of the two agent.

4.5 Strategy 2: Internal urgency criteria

In the previous approach we used the authority of an external mediator agent to enforce a certain agreement (which needs to be rational for both agents). Let us now consider a strategy which is relying on an internal urgency criteria of the agent.

Each agent separately decides on the maximum number of negotiation rounds in which it will participate, *roundLimit*. Note that this is not the same as an externally enforced limited round negotiation - the opponent agents do not know the number of rounds the opponent agent is considering (in fact, they don't know anything about the nature of the other agents' strategy). Based on this number, the agent calculates a value OPT, which is defined as remaining rounds / total number of rounds. This value starts at 1 and decreases to 0 over the negotiation rounds.

At every round, the agent calculates the utilities of E^B_{i-1} (denoted with U_{EB}) and E^A_{i-1} (denoted U_{EA}). By comparing the U_{EB}, U_{EA} and OPT, the agent A decides whether it insists on its former offer O^A_{i-1}, concedes from its former offer towards O^B_{i-1} or agrees to the other agents' offer O^B_{i-1}. In general, the agent will agree on any offer which is as good as or better than its own offer *or* OPT. As $OPT \geq 0$, the rationality of the agent is assured. The OPT value represents the urgency of the agent to conclude a negotiation; towards the end of the negotiation, the agents will agree on offers with lower utility than at the beginning. In the last negotiation step, the agent will agree on any deal, provided that it is rational (it has a positive utility). The algorithm for this strategy is described by the following pseudocode.

agent A received a message from agent B;
if the message is $(O^B_{i-1}, \emptyset)$ **then**
    $E^B_{i-1} \leftarrow evaluation(O^B_{i-1})$ and $U_{EB} \leftarrow Utility(E^B_{i-1})$;
    if $U_{EB} \geq 0$ **then**
        propose$((O^A_i, E^B_{i-1}))$ where O^A_i concedes from O^A_{i-1} towards E^B_{i-1};
    else
        $propose((O^A_i, \emptyset))$ where O^A_i concedes from O^A_{i-1} towards E^B_{i-1};
    end if
    i=i+1;
end if
if the message is (O^B_{i-1}, E^A_{i-1}) **then**
    $E^B_{i-1} \leftarrow evaluation(O^B_{i-1})$;
    $U_{EB} \leftarrow Utility(E^B_{i-1})$ and $U_{EA} \leftarrow Utility(E^A_{i-1})$;
    UPDATE OPT=(roundLimit-i)/roundLimit;
    if $U_{EB} \geq 0$ **then**
        if $U_{EB} \geq OPT$ **then**
            $propose((\emptyset, E^B_{i-1}))$ to request an agreement;
        else
            if $U_{EA} \geq OPT and U_{EB} \leq OPT$ **then**
                insist current offer.
            end if
        else
            if $U_{EA} \leq OPT and U_{EB} \leq OPT$ **then**
                concede current offer
            end if
        end if
    else
        insist offer $O^A_i = O^A_{i-1}$ until agent B's offer has a positive utility;
    end if
    i=i+1;
end if
if the message is $(\emptyset, E^A_{i-1})$ **then**
    evaluate to to check if the utility of E^A_{i-1} is positive.
    form an agreement.
end if

5. CONCLUSIONS

In this paper we considered a class of problems which use negotiation to establish collaboration in space and time. We have shown that the "split the pie" game, frequently used as a canonical problem in studies concerning multi-issue negotiation, can not adequately model this class of problems and we highlighted five characteristics of the class of problems which need to be modelled by a representative canonical problem. We proposed the "Children in the Rectangular Forest" (CRF) model as a canonical problem for the class of applications we are considering, and we have shown that it exhibits the identified characteristics, while keeping the intervening formulas simple and of a low computational complexity. Through an example negotiation protocol and two example negotiation strategies we illustrated the type of negotiation models which can be applied to the proposed canonical problem.

The value of a canonical problem is ultimately dependent on its aesthetic and utilitarian characteristics. We find the CRF problem aesthetically appealing in its simplicity, ability to represent complex collaboration patterns and the fact that it is mathematically analysable. Its utility, however, needs to be reflected in the number of real world problems whose solutions can be reduced to CRF and, conversely, the number of algorithms developed in the context of CRF which can be applied to real world problems. This is our future work.

Acknowledgments

The authors want to thank the reviewers of this paper. Their insightful comments made not only this paper better, but also suggested additional points of view and representation approaches. Many reviews are just a yes/no decision; in this case reading them was the equivalent of a brainstorming session about future research directions.

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 heron.

References

[1] K. Binmore, M. Osborne, and A. Rubinstein. Noncooperative models of bargaining. In R. J. Aumann and S. Hart, editors, *Handbook of Game theory with Economic Applications*, volume 1, pages 179–225. Elsevier, 1992.

[2] L. Bölöni, M. Khan, and D. Turgut. Agent-based coalition formation in disaster response applications. In *Proceedings of the IEEE Workshop on Distributed Intelligent Systems*, pages 259–264, June 2006.

[3] D. Carmel and S. Markovich. Learning models of intelligent agents. In H. Shrobe and T. Senator, editors, *Proceedings of the Thirteenth National Conference on Artificial Intelligence and the Eighth Innovative Applications of Artificial Intelligence Conference, Vol. 2*, pages 62–67, Menlo Park, California, 1996. AAAI Press.

[4] S. Hanks, M. E. Pollack, and P. R. Cohen. Benchmarks, test beds, controlled experimentation, and the design of agent architectures. *AI Magazine*, 14(4):17–42, 1993.

[5] M. A. Khan and L. Bölöni. Negotiation-based coalitions in the physical world. In P. Stone and G. Weiss, editors, *Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2006)*, pages 411–413, May 2006.

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

[7] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

[8] Lynx access paratransit service. URL `http://www.golynx.com/?pid=1155583`.