Satisficing Coalition Formation among Agents

Leen-Kiat Soh
Computer Science and Engineering Department
University of Nebraska
115 Ferguson Hall, Lincoln, NE 68588-0115
Tel: 402-472-6738
lksoh@cse.unl.edu

Costas Tsatsoulis
Department of Electrical Engineering and Computer Science
Information & Telecommunication Technology Center
University of Kansas
2335 Irving Hill Road, Lawrence, KS 66045
Tel: 785-864-7749
tsatsoul@ittc.ku.edu

ABSTRACT
In a multiagent system where each agent has only an incomplete view of the world, optimal coalition formation is difficult. Coupling that with real-time and resource constraints often makes the rationalization process infeasible or costly. We propose a coalition formation approach that identifies and builds sub-optimal yet satisficing coalitions among agents to solve a problem detected in the environment. All agents are peers and autonomous. Each is motivated to conserve its own resources while cooperating with other agents to achieve a global task or resource allocation goal. The (initiating) agent—that detects a problem—*hastily* forms an initial coalition by selecting neighboring agents that it considers to have high potential utilities, based on the capability of each neighbor and its respective inter-agent relationships. The initiating agent next finalizes the coalition via multiple concurrent 1-to-1 negotiations with only neighbors of high potential utility, during which constraints and commitments are exchanged in an argumentation setting.

Categories and Subject Descriptors
I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence – *coherence and coordination, multiagent systems.*

General Terms
Algorithms, Experimentation.

Keywords
Coalition Formation, Negotiation, Case-Based Reasoning, Utility Theory, Multiagent System.

1. INTRODUCTION
In this paper, we present a coalition formation strategy that facilitates sub-optimal yet satisficing collaborations among agents. The need for such collaborations arises due to the *reactiveness* requirement of agents in a time-critical and resource-constrained environment. In a non-hierarchical multiagent system, agents are peers and autonomous. Each controls a set of resources, monitors its part of the world, and reacts to events that it detects. For events that require collaborations, an (initiating) agent forms a coalition out of its neighbor agents. However, each agent only has an incomplete (and sometimes outdated and noisy) view of the world and its neighbors—making rational optimality in its coalition formation process impossible. So, in a situation where an optimal solution does not exist, where an optimal solution cannot be computed, or where an optimal solution cannot be derived on time, we have to look to sub-optimal solutions that are satisficing in that they meet minimum requirements rather than achieving maximum performance.

Our proposed coalition formation approach breaks the process into three stages: initial coalition formation, coalition finalization, and coalition acknowledgment. The objective of the initial coalition formation is to *hastily* identify potential candidates and rank them accordingly to their potential utilities. However, since the agent maintains only a partial view of the world, it needs to determine with certainty whether any of the candidates is willing to help. Instead of approaching all neighbors to compute an optimal coalition, the agent only selectively negotiates with top-ranked candidates to finalize the coalition so as to conserve both computational resources and communication usage. Some candidates may refuse to participate, some may agree. When a coalition is formed, all remaining negotiations are terminated. When a coalition no longer can be formed (due to some negotiation failures), all remaining negotiations are terminated as well. This strategy is thus opportunistic and high-risk since the formation of a coalition cannot be guaranteed. To help alleviate this weakness, our agents are equipped with learning mechanisms (case-based learning and reinforcement learning) to learn to form better coalitions faster, and a constant monitoring module for risk mitigation. Note that with this modular approach, we refine the validity of our coalition gradually as time permits. After the initial coalition formation step, an agent already has a potentially working coalition. As the agent gains more information about its neighbors (through negotiation), it is able to eliminate or confirm candidates.

AAMAS '02, July 15-19, 2002, Bologna, Italy.

2. COALITION FORMATION
In our approach, an initiator starts the coalition formation process first by selecting members in the agent's *neighborhood* that are qualified to be part of an initial coalition. Second, it evaluates these members to rank them in terms of their respective potential utility values to the coalition. Third, it initiates negotiation requests to the top-ranked candidates,

trying to convince them to join the coalition. In the end, the coalition may fail to form because of the candidates' refusal to cooperate, or may form successfully when enough members reach a deal with the initiating agent. Finally, the agent sends an acknowledgement message to the coalition members to announce the success or failure of the proposed coalition. If it is a success, then all coalition members that have agreed to join will carry out their respective tasks at planned time steps.

This approach is opportunistic as the goal is to obtain a satisficing coalition and the success of the formation is not guaranteed. This is the risk that our agent is willing to take: the utility of responding timely to a problem is dominating the utility gained from the quality of the solution, since the domain is time-critical and dynamic.

2.1 Initial Coalition Formation

The goal of the initial coalition formation process is to find a group of neighbors that can be of help to the initiating agent. In addition, these *candidates* are ranked according to their potential utility values to accomplish the task at hand. Thus, the ranking criteria include the capability of a candidate to provide useful resources and the past and current negotiation relationships between the initiating agent and that candidate. As a result, even if a candidate is the most capable among all candidates, its potential utility is reduced if it does not have a good relationship with the initiating agent.

First, upon detecting a resource or task allocation problem, the initiator describes the problem using its monitors. The parametric description helps guide the identification of the coalition candidates. Second, to establish who can provide useful resources or perform certain tasks, the initiator identifies them from its knowledge profile of its neighborhood. This knowledge profile consists of the following for each neighbor: (a) the name and communication id, (b) the resources it has, and (c) the tasks that it can perform. By matching the neighbor profile with the profile of the problem, the initiator selects useful neighbors and forms the initial coalition.

Since computational resources are limited, and negotiating consumes CPU and bandwidth, the initiator does not start negotiation with all members of the coalition, but first ranks them and then initiates negotiation selectively with the top-ranked ones. Ranking of the coalition members is done using a multi-criterion utility-theoretic evaluation technique.

The potential utility of a candidate is a weighted sum of (1) its ability to help towards the problem at hand, (2) its past relationship with the initiator, and (3) its current relationship with the initiator. Each of these items is based on a set of observable environmental parameters.

Each agent keeps a profile of its neighborhood, and current and past relationships with its neighbors, and the selection of the potential members of an initial coalition is based on this profile. The current relationship is based on the negotiation strains and leverage between two agents at the time when the coalition is about to be formed. The past relationship is collected over time and enables an agent to adapt to form coalitions more effectively.

2.2 Coalition Finalization

After obtaining the initial coalition and the coalition candidates ranked according to their respective potential utility values, the initiator invokes the coalition finalization step. This step consists of negotiations. To conserve computational resources and communication resources, the initiator only approaches the top-ranked candidates. When a candidate agrees to negotiate, the initiator proceeds with a one-to-one negotiation guided by a negotiation strategy. However, since the initiator conducts multiple, concurrent 1-to-1 negotiations, each negotiation process has access to the coalition status. This awareness allows the agent to modify its instructions for each negotiation process.

Our agents use a variation of the *argumentative negotiation model* in which it is not necessary for them to exchange their inference model with their negotiation partners. Note that after the initial coalition formation, the initiator knows who can help. The goal of negotiations is to find out who is willing to help. To do so, first the initiator contacts a coalition candidate to start a negotiating session. When the candidate or responder agrees to negotiate, it computes a *persuasion threshold* that indicates the degree to which it needs to be convinced in order to free or share a resource or perform a task (alternatively, one can view the persuasion threshold as the degree to which an agent tries to hold on to a resource). Subsequently, the initiator attempts to convince the responder by sharing parts of its local information. The responder, in turn, uses a set of domain-specific rules to establish whether the information provided by the initiator pushes it above a resource's persuasion threshold, in which case it frees the resource. If the responder is not convinced by the evidential support provided by the initiator, it requests more information that is then provided by the initiator. The negotiation continues based on the established strategy and eventually either the agents reach an agreement, in which case a resource or a percentage of a resource is freed, or the negotiation fails. Note that, motivated to cooperate, the responder also counter-offers when it realizes that the initiator has exhausted its arguments or when time is running out for the particular negotiation. How to negotiate successfully is dictated by a negotiation strategy, which each agent derives using case-based reasoning (CBR). CBR greatly limits the time needed to decide on a negotiation strategy, which is necessary in our real-time domain since the agent does not have to compute its negotiation strategy from scratch.

3. CONCLUSIONS

We have described a coalition formation strategy that aims at obtaining satisficing solution for time-critical, noisy, and incomplete resource or task allocation problem. Initially, coalition candidates are selected hastily from an agent's neighborhood and subsequently ranked according to their respective potential utilities. Next, during the finalization phase, the coalition is refined and verified through negotiations, where information is exchanged between two agents to clarify commitments and constraints. The agent is able to coordinate directly and indirectly through a coalition awareness link with its negotiation threads.

Finally, we have built a multiagent system complete with end-to-end agent behavior. Our preliminary results are promising in that an initiator was able to form satisficing coalitions quickly given its constraints.