Towards Flexible Negotiation in Teamwork

Zhun Qiu Milind Tambe Hyuckchul Jung
Information Sciences Institute & Computer Science Department
University of Southern California
4676 Admiralty Way, Marina del Rey, CA 90292, USA
zqiu@isi.edu tambe@isi.edu jungh@isi.edu

ABSTRACT

In a complex, dynamic multi-agent setting, coherent team actions are often jeopardized by agents' conflicting beliefs about different aspects of their environment, about resource availability, and about their own or teammates' capabilities and performance. Team members thus need to communicate and negotiate to restore team coherence. This paper focuses on the problem of negotiations in teamwork to resolve such conflicts. The basis of such negotiations is inter-agent argumentation based on Toulmin's argumentation pattern. There are several novel aspects in our approach. First, our approach to argumentation exploits recently developed general, explicit teamwork models, which make it possible to provide a generalized and reusable argumentation facility based on teamwork constraints. Second, an emphasis on collaboration in argumentation leads to novel argumentation strategies geared towards benefiting the team rather than the individual. Third, our goal, to realize argumentation in practice in an agent team, has led to decision theoretic and pruning techniques to reduce argumentation overhead. Our approach is implemented in a system called CONSA.

Keywords

Argumentation, negotiation, multi-agent, collaboration, teamwork

1. INTRODUCTION

Collaboration or teamwork in multi-agents is a critical capability in a large number of applications. Unfortunately, in a complex, dynamic multi-agent setting, coherent team action is often jeopardized by conflicts in agents' beliefs, plans, and actions. While recent research on teamwork has made progress in flexible agent coordination and communication, collaborative negotiations in a team to resolve such conflicts remains an open and difficult challenge. This paper focuses on such collaborative negotiations. Our approach is based on *argumentation*, as it appears more relevant in addressing concerns of teamwork. The system we have designed and implemented is called CONSA (**CO**llaborative **N**egotiation **S**ystem based on **A**rgumentation). While CONSA builds on the past work of argumentation-based approach to negotiation in the multi-agents arena, it advances the state of the art in several ways. First, it uses recently developed teamwork models as a source of reusable argumentation knowledge. Second, it introduces novel argumentation strategies more suitable for collaborative settings. Third, it introduces novel techniques, including decision-theoretic techniques, to improve argumentation efficiency.

2. CONCRETE NEED FOR NEGOTIATION

For the past few years, we have been developing agent teams for several applications, including teams of helicopter pilot agents for real-world battlefield simulation [2], player teams for the RoboCup soccer simulation and software agent teams for (simulated) unmanned aircraft. These teams are based on a state-of-the-art, implemented teamwork model called STEAM [2] that provides significant teamwork flexibility. Yet STEAM or other teamwork models do not address the problem of agent negotiations to resolve conflicts. Thus inter-agent conflicts still unexpectedly arise in some cases. For example, a helicopter pilot team in combat simulations must plan firing positions, i.e., positions for individual helicopters to hide and attack enemy. Typically, the commander pilot agent plans all these positions, one per each team member, and sends them to team members. For instance, a team member named *cheetah102* may obtain a position with coordinates to hide behind a small hill. Once all of the positions are communicated, the commander asks its team to proceed. Unfortunately, in one run, one team member (*cheetah102*) never received its position due to radio interference. Thus, the commander thought the message was sent, but *cheetah102* never received it, leading to a conflict in beliefs.

3. THE FRAMEWORK OF CONSA

CONSA' s framework is based on studies of argumentation in philosophy and AI. In CONSA, the argumentation process involves one agent (sender) making a proposal to another agent (receiver) possibly with an attached justification. The receiver evaluates the proposal, and either accepts or refutes it. If refuting the proposal, the receiver again attaches a justification. In addition, the receiver may send back a counter-proposal to the sender, continuing this cycle of proposals and counter-proposals.

CONSA's generation of proposal or its refutation is accomplished via actions called argumentation *moves*. Two moves that have been used in the literature are *rebutting* and *undercutting*. *Rebutting* refutes the teammate' s claim (proposal) directly, with some justifications. In contrast, *undercutting* attacks the attached justification of the proposal, rather than the proposal itself.

CONSA' s representation of arguments, for generating and evaluating proposals, is based on the philosopher Toulmin' s [3] argumentation pattern (TAP). In a TAP, an argument consists of *claim* (conclusion), *data* (foundation facts), *warrant* (the authority for taking the step from the data to the claim) and *qualification*

(degree of force which is conferred on the claim based on the data and warrant).

In CONSA, claims are agents' individual or mutual beliefs. An agent' s proposal usually involves a claim, with a justification consisting of the supporting warrant and data from the TAP. In general, the data may itself be another claim (belief), so that a recursive tree of TAP structure emerges in support of a claim. With respect to qualifications on claims, in CONSA, claims have high, medium and low strength and may need to be evaluated recursively if there is a recursive TAP structure.

4. COLLABORATIVE NEGOTIATIONS

CONSA exploits recent advances in multi-agent teamwork to address argumentation about teamwork (§4.1). Such argumentation about teamwork is expected to be collaborative (§4.2), if agents are continuing to participate in their team.

4.1 Teamwork Model for CONSA

STEAM [2] is the basis of our agent team implementations and consists of two components, both currently realized in the Soar [1] architecture. The first is the enhanced agent architecture with explicit representation of team' s joint intentions, mutual beliefs, and team goals. The second is the domain-independent teamwork knowledge to enable individual agents' flexible teamwork [2]. This knowledge takes the form of two types of actions. *Coherence preserving* actions require agents to coherently activate and terminate joint intentions, while *maintenance and repair* actions lead to team reorganization when necessary. These actions require an explicit specification of the dependency relationship of the joint intention on individual team members' activities, based on the notion of a *role*. Three primitive role-relationships (AND-combination, OR-combination, role-dependency) can currently be specified in STEAM and can be combined.

STEAM has significant implications for CONSA in two areas. First, STEAM provides reusable argumentation knowledge for negotiations in teamwork. In particular, it provides warrants for constructing TAPs. For instance, the AND-combination role relationship can imply some warrants that an agent can utilize in argumentation.

- ω_2: Joint-intention(τ) $\wedge$ AND-combination(τ) $\wedge$ achieved(τ)

  $\rightarrow$ All-roles-fulfilled(τ)

ω_2 infers the fulfillment of all roles when another agent declares that a joint intention τ with an AND-combination is achieved.

The second implication of STEAM for CONSA is categorizing the types of conflicts that arise in teamwork. In particular conflicts can be categorized into the following three types: (1) Team members may have conflicting beliefs about jointly terminating a team operator. (2) Team members may have conflicting beliefs about initiating a team operator. (3) Agents executing individual operators may unintentionally interfere with each other' s role performance i.e., the conflict is in execution, but not in initiation or termination.

4.2 Collaborative Argumentation Strategies

That the argumentation is collaborative has a significant impact on CONSA' s design, from its detection of conflicts to generation of counter-proposals. With respect to conflict detection, agents must attempt to detect only those conflicts where argumentation will improve team utility. In particular, team members often send proposals to coherently initiate and terminate joint activities, based on the information that is often unavailable to others. While such information may contradict an agent's own beliefs, if a teammate' s proposal is based on more credible information, that should be accepted. Thus, agents should engage in argumentation when they detect true conflicts. Thus, representing strengths of beliefs, as discussed in §3, is critical for CONSA's collaborative argumentation, as it enables CONSA to efficiently detect true conflicts.

Having detected a conflict, CONSA's next step is to decide upon and execute an argumentation move. Here, the collaboration has an impact because, in addition to the earlier rebut and undercut moves, collaboration also frequently requires a third strategy, referred to here as "improve support". In particular, an agent receiving a proposal from its team member, may very well agree with and accept the proposal, but may have a support or justification for the proposal that is better than the one offered by the sender.

Of course, in reality, more than one of the argumentation moves suggested so far may be applicable to an argument. Furthermore, for any type of move, more than one instantiation (number of moves) may be possible. The key issue here is ordering. For collaborative argumentation, the agent must attempt to make progress in the negotiation as efficiently as possible. Therefore, the ordering heuristic used is to present the strongest argument first, where strength is derived from the strength of the claim in the associated TAPs.

Finally, once an agent determines an argumentation move, it should not automatically execute it. CONSA relies on a decision theoretic approach to weigh three choices once it detects a conflict. First, an agent may avoid negotiation but not accept the teammates' belief either. Second, it can again avoid negotiation, but accept the teammates' belief. Third, an agent can negotiate in detail with its teammate and reach the right conclusion. Under different circumstances, the cost and utility of each of these choices will vary.

5. CONSA IMPLEMENTATION

CONSA is currently implemented and fully integrated with the STEAM teamwork implementation. It consists of 63 Soar rules at present time. Our implementation has enabled agents to begin negotiations in all the test cases we used in helicopter pilot and RoboCup soccer domain.

6. REFERENCES

[1] Newell, A. 1990. *Unified Theories of Cognition*. Harvard Univ. Press.

[2] Tambe, M. 1997. *Towards flexible teamwork*. Journal of Artificial Intelligence Research (JAIR), 7:83-124

[3] Toulmin, S. 1958. *The uses of argument*. Cambridge: Cambridge University Press.