ECAI 2010
H. Coelho et al. (Eds.)
IOS Press, 2010

doi:10.3233/978-1-60750-606-5-1001

On the use of Argumentation in Multi-Agent Planning

Óscar Sapena and **Eva Onaindia** and **Alejandro Torreño** [1]

1 Introduction

Multi-Agent Planning (MAP) emphasizes the problem of distributing planning among several agents, finding a plan for each agent that achieves its private goals, and merging the resulting local plans to come up with a plan that meets the global goals as well [3]. Mostly, the emphasis in MAP is on how to manage the interdependencies between the agents' plans and coordinate the local plans.

In contrast to these approaches, we propose an argumentation-based model for Cooperative Distributed Planning [4], i.e. building a global plan amongst a set of agents who will contribute differently to the joint task based on their abilities, knowledge and private interests. In our model, planning is achieved by cooperative agents that are distributed functionally or spatially, and can have private goals as well. Argument-based interactions like persuading an agent to adopt a course of action, or negotiating on the use of resources [5] are used in a dialectical process to attain a collective behavior when devising a joint plan.

Our proposal to address a CDP problem is to make use of argumentation schemes and associated critical questions [6], following the computational representation of practical argumentation presented in [2, 1]. The argumentation-based model allows planning agents to propose, discuss and refine partial solutions according to their local knowledge.

2 Multi-agent planning model

A MAP problem is described as follows: given an initial state, a set of global goals, a set of (at least) two agents, and for each agent a set of its abilities and (probably) its private goals, find a plan that achieves both global and private goals. Therefore, a MAP problem can be regarded as a CDP task in which agents have their own planning tasks.

Definition 1 *A CDP task is a tuple $\mathcal{T} = \langle \mathcal{AG}, \Theta, \mathcal{P}, \mathcal{I}, \mathcal{G}, \mathcal{F} \rangle$ where $\mathcal{AG} = \{1 \ldots n\}$ is a finite, non-empty set of planning agents, Θ is the set of actions of the agents, $\mathcal{P}$ is a finite set of propositional state variables, $\mathcal{I} \subseteq \mathcal{P}$ is the initial state, $\mathcal{G} \subseteq \mathcal{P}$ denotes the problem goals and $\mathcal{F}$ is a utility function to choose between plans.*

Each agent is actually solving a different planning task because each has its own local knowledge and a partial view of the problem's overall state. In our framework, the planning model of an agent also comprises some information on the abilities of the other agents in order to promote a coherent coordination towards a joint plan.

A CDP task $\mathcal{T}$ can thus be seen as solving as many planning tasks as agents in $\mathcal{AG}$. Each agent $i \in \mathcal{AG}$ has its own planning task $\mathcal{T}_i = \langle \Theta_i, \mathcal{I}_i, \mathcal{G}_i, \mathcal{F} \rangle$ such that solving $\mathcal{T}$ implies solving $\bigcup_{\forall i \in \mathcal{AG}} \mathcal{T}_i$:

- $\Theta_i \subseteq \Theta$ is the set of actions in the planning model of agent i. We define $\Theta_i = \Gamma_i \cup \Delta_i$, where Γ_i denotes the actions executable by agent i, and Δ_i denotes the rest of actions known by agent i.
- $\mathcal{I}_i \subseteq \mathcal{I}$ denotes the local knowledge of agent i, its partial perspective of the environment. Formally, we can define $\mathcal{I} = \bigcup_{\forall i \in \mathcal{AG}} \mathcal{I}_i$.
- $\mathcal{G}_i = \mathcal{G} \cup \mathcal{PG}_i$, where $\mathcal{G}$ are the goals of the CDP task, and $\mathcal{PG}_i \subseteq \mathcal{P}$ are agent i's private goals. $\mathcal{PG}_i$ is an optional parameter.
- $\mathcal{F}$ is the global utility function, which will be used by the participating agents to argue over the partial solutions.

In our MAP model, agents' contributions to the plan for the CDP task will follow the Partial Order Planning (POP) paradigm.

3 Argumentation framework

In this section we propose an adaptation of the argumentation model presented in [2, 1] to solve a CDP task. Our framework presents a partial plan as a presumptive justification of the course of action given by a particular agent as a contribution to the task. The goal is to decide whether the contribution plan of an agent to a base plan is a good alternative or not. Given a base plan (initially, the empty plan Π_0), a particular agent i suggests a refinement which, according to its model, represents a good step towards the goals (and also for its private interests). Thus, agent i attempts to persuade the others to adopt the resulting plan. Since the agent's contribution is a presumptive argument based on its local knowledge, the rest of agents will likely argue against its argument. An argument is defined as follows:

Definition 2 *Given a CDP task $\mathcal{T}$, an argument is a pair $A = \langle \Pi_b, \Pi_i \rangle$ where Π_b is the current base plan for $\mathcal{T}$ (the premise), and Π_i is the presumptive justification of agent i in favour of the conclusion that $\Pi_b \circ \Pi_i$ is a valid refinement of Π_b.*

As in [2], we define an argumentation scheme (AS) that responds to our definition of argument: in the current circumstances given by the base plan Π_b, we should proceed with plan Π_i, resulting in a new refinement plan Π_r, which will realise some subgoals G and promote some values V. AS is interpreted as follows:

- The base plan Π_b expresses the real current situation agreed by all the agents, so it is never a source of disagreement. This means the premise of an argument can never be attacked.
- Π_i is the presumptive justification argued by an agent i to proceed towards the resolution of the CDP task $\mathcal{T}$.
- $\Pi_r = \Pi_b \circ \Pi_i$ is the composition of the base plan and agent i's contribution plan. If Π_r is finally accepted by all the participants, it will become the new base plan.
- $G = opengoals(\Pi_b) \setminus opengoals(\Pi_r)$ contains the open goals of Π_b solved by Π_i. Some of the goals in G can also belong to $\mathcal{G}$.
- V is a set of two function values: progress ($prog$) and cost ($cost$). $prog(\Pi_r)$ measures to what extent further plan compositions can be applied over Π_r. $cost(\Pi_r)$ is the cost of the best solution plan reachable from Π_r, in terms of $\mathcal{F}$. Higher values of $prog$ and lower values of $cost$ are preferable. Both are estimated values.

Attacks will arise in situations perceived by the agents as conflicting in the elaboration of the global plan. The following section details the challenges that agents will pose to a given argument.

[1] Universidad Politécnica de Valencia, Spain

3.1 Argument evaluation

Given a CDP task $\mathcal{T}$, and an argument $A = \langle \Pi_b, \Pi_i \rangle$ proposed by agent i, agents will discuss on the convenience of proceeding through $\Pi_r = \Pi_b \circ \Pi_i$ to solve $\mathcal{T}$. In our model, actions are deterministic and agents' beliefs are always consistent, so they can only be altered through the refinements of the joint plan. We can thus affirm there are neither discrepancies on the current situation nor disputes on the truth of the agents' knowledge. Hence, we only consider attacks concerned with the choice of action, particularly the ones classified as "side-effects of the action" and "interferences with other actions" in [2].

Attacks discuss the feasibility of Π_r. Since agents have a partial view of the world, an agent's contribution might prevent the others from refining it, and hence, Π_r should be rejected. More precisely, agent j will attack the proposal of agent i $\Pi_r = \Pi_b \circ \Pi_i$ in one of the following situations:

AT1: Inability of an agent j to solve an open goal p in Π_r
- **Attack:** $\exists p \in opengoals(\Pi_r)/p \in opengoals(\Pi_r \circ \Pi_j), \forall \Pi_j$.
- **Justification:** The inability of j to solve p demotes $prog(\Pi_r)$ unless some other agent can do it; otherwise, Π_r is unfeasible.
- **Counter-attack:** Another agent k, $k \neq j$, can achieve p: $\exists \Pi_k/p \notin opengoals(\Pi_r \circ \Pi_k)$.

AT2: Π_r prevents agent j from promoting $prog$
- **Attack:** $\exists p \in opengoals(\Pi_r) \rightarrow \exists g \in \mathcal{G}_j/g \in opengoals(\Pi_r \circ \Pi_j), \forall \Pi_j$
- **Justification:** It exists an open goal p in Π_r that will prevent agent j from from making a refinement that achieves a (global or private) goal g in $\mathcal{G}_j$, and hence it cannot promote $prog(\Pi_r)$.
- **Counter-attack:** if g has been notified as a global goal, another agent k, $k \neq j$, can make a counter-attack if it can solve g despite of the presence of p in Π_r: $\exists \Pi_k/g \notin opengoals(\Pi_r \circ \Pi_k)$.

AT3: Non-affordable cost of Π_r for agent j
- **Attack:** $\sum_{\forall a \in A/ag(a)=j} cost(a) > cost\_limit_j$.
- **Justification:** Agent j argues that the cost of its actions in Π_r is excessive. AT3 demotes $prog(\Pi_r)$ as well as $cost(\Pi_r)$ since any further refinement of Π_r proposed by j will have equal or higher cost. AT3 cannot be counter-attacked, so Π_r must be rejected.

A plan is considered unfeasible when it receives an attack that it is not invalidated by any counter-attack. The next section presents a basic protocol, based on this process, for solving a CDP task.

4 MAP protocol

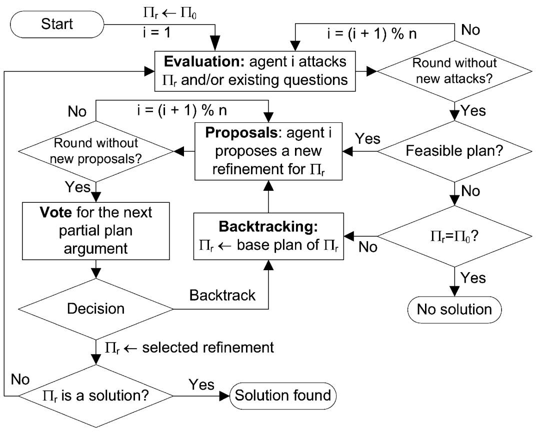

Figure 1. Multi-agent planning protocol overview

In this section, we present an argumentation-based protocol for solving a MAP problem (figure 1), which is divided into two stages:

1. **Plan evaluation**. During their turn, agents can attack the current argument, $\Pi_r = \Pi_b \circ \Pi_i$, and/or the associated critical questions. Initially, Π_r will be the empty plan Π_0. When an agent finishes its attacks to Π_r, the turn is passed. The process ends after a round without new attacks. If Π_r is proved to be unfeasible, the process backtracks to Π_b. If Π_0 is proved to be unfeasible, the process finishes without a solution.
2. **Proposals of refinement plans**. If Π_r is evaluated successfully, agents will propose refinements to Π_r. Each refinement (except for the first one) must be justified by explaining why it improves the previous one. After a round without new proposals, agents will vote to choose the next base plan, or to backtrack. The most voted refinement will become the next argument to evaluate, unless it is a solution plan, in which case the process ends successfully.

We identify various different justifications to promote a refinement Π_{r2} over another one Π_{r1}. Let $G1 = opengoals(\Pi_r) \setminus opengoals(\Pi_r \circ \Pi_{r1})$ be the goals solved by Π_{r1}, and $G2 = opengoals(\Pi_r) \setminus opengoals(\Pi_r \circ \Pi_{r2})$ be the goals solved by Π_{r2}:

RJ1: Alternative way of promoting V
- **Premise:** $G1 \neq G2$.
- **Justification:** Π_{r2} promotes progress and/or cost, i.e. $prog(\Pi_r \circ \Pi_{r2}) > prog(\Pi_r \circ \Pi_{r1})$ and/or $cost(\Pi_r \circ \Pi_{r2}) < cost(\Pi_r \circ \Pi_{r1})$.

RJ2: Alternative way of achieving $G1$
- **Premise:** $G1 = G2$.
- **Justification:** Π_{r2} promotes cost and/or progress.

These justifications are speculative because an agent cannot guarantee that extending a plan to completion will reach a good-quality solution. Voting is necessary because agents can have different opinions on the refinement according to their beliefs.

5 Conclusions

This paper defines a MAP problem and presents an argumentation model to address it. Our proposal aims at solving planning problems cooperatively, while considering self-interested agents.

Agents argue plan refinements and agree on the presumptively best joint plan. Our framework, designed in terms of argument schemes and attacks, presents some advantages to similar models; (1) the argumentation scheme is instanced to a set of elements rather than to a single one; (2) the choice of goal is seen as a means to forward the discussion to, presumptively, better argumentation lines, and (3) the evaluation of attacking situations takes account of future plans that are precluded. These contributions are aimed to promote cooperation for a collective resolution of a planning problem.

REFERENCES

[1] K. Atkinson and T. Bench-Capon, 'Practical reasoning as presumptive argumentation using action based alternating transition systems', *Artificial Intelligence*, **171**, 855–874, (2007).

[2] K. Atkinson, T. Bench-Capon, and P. McBurney, 'Computational representation of practical argument', *Synthese*, **152**(2), 157–206, (2006).

[3] M. de Weerdt, A. ter Mors, and C. Witteveen, 'Multi-agent planning. an introduction to planning and coordination', in *Handouts of the European Agent Systems Summer School (EASSS-05)*, pp. 1–32, (2005).

[4] M.E. desJardins, E.H. Durfee, C.L. Ortiz, and M.J. Wolverton, 'A survey of research in distributed continual planning', *AI Magazine*, **20**(4), 13–22, (1999).

[5] I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg, 'Argumentation-based negotiation', *The Knowledge Engineering Review*, **18**(4), 343–375, (2004).

[6] D.N. Walton, *Argumentation Schemes for Presumptive Reasoning*, Lawrence Erlbaum Associates, Mahwah, NJ, 1996.

[2] **Acknowledgments**. This work has been supported by the Spanish MICINN under projects TIN2008-06701-C03-03 and Consolider Ingenio 2010 CSD2007-00022, the Valencian Prometeo project 2008/051.