

Integrating Multiagent Dialogues, Planning and Plan Execution

Yuqing Tang

Dept. of Computer Science,
Graduate Center, City University of New York
365 Fifth Avenue, New York, NY 10016, USA
ytang@cs.gc.cuny.edu
Advised by: Simon Parsons

Abstract

Coming up with a plan for a team that operates in a non-deterministic environment is a complex process, and the problem is further complicated by the need for team members to communicate during the process of planning and the execution of plans. Dialogues can be used to coordinate both the planning process and the plan execution. In this work I develop a model for constructing joint plans for a team of agents that takes into account their communication needs while accommodating multiple agents' inconsistent specifications of transitions, initial states, and goals. The model builds on recent developments in symbolic non-deterministic planning, ideas that have not previously been applied to this problem.

Introduction

One of the fundamental problems in multiagent systems is how to get a team of agents to coordinate their behavior. While there are situations in which agents can do this without needing to communicate (Genesereth, Ginsberg, and Rosenschein 1988), in general coordination requires communication. Another important part of coordination is having the agents decide what to do. Since (Bratman, Israel, and Pollack 1988), the process of deciding what to do is considered to break down into two parts — deciding what goals to achieve, what (Bratman, Israel, and Pollack 1988) calls *deliberation*, and then deciding how those goals might best be achieved, which is usually described as *planning*. In this work I am interested in the planning part of the process while the problem of establishing joint intentions (Cohen and Levesque 1991) is tackled through non-monotonic reasoning on goals.

I am also greatly concerned with communication. Much recent work on agent communication uses argumentation-based dialogue (Parsons and McBurney 2003), and the long term goal of our work is to extend existing work on multiagent planning by developing models by which a team of agents can, in the course of an argumentation-based dialogue — by which we mean a process during which agents put forward suggested partial plans backed by reasons, as in (Tang and Parsons 2005) — develop a plan for the team. We want this to be done in a way that respects the non-deterministic

nature of the world, and which yields efficient implementation. This paper takes several steps towards this goal.

In particular, this work gives a mechanism, in both centralized and decentralized form, by which a multiagent team can construct plans that take into account the need to communicate to ensure that the plan is executed correctly (Parsons et al. 2008). By introducing new structures and variables for multiagent systems and dialogues, extending the non-deterministic planning, our approach can make use of new techniques from model-checking to provide efficient implementations. The extension incorporates the elements necessary to take multiple agents, and the necessary communication, into account. The use of a symbolic model makes it possible to turn the plan construction process into an argumentation-based dialogue in the future.

Building our approach on top of work in planning has advantages beyond ease and efficiency of implementation. By appropriating the underlying formal models, it is easy to acquire suitable formal guarantees for the planning model. It is straightforward, for example, to show that given an adequate description of the world, any plan that our planning process will construct is both a feasible and, in a specific sense an optimal, way to achieve the goals of the plan.

Multiagent Systems and Dialogues

This work uses a state-space model as a basis for the formalisation. This model is an adaptation of a model commonly used in non-deterministic planning (Cimatti et al. 2003). *States* are objects that capture some aspect of a system, and *actions* are transitions between states. States and actions together define a *state-space*. When action effects are non-deterministic then what one seeks for any state-space is a *policy*: i.e. a state-action table to specify which actions one should take in a given state. We define a non-deterministic state transition domain (NSTD) to be a tuple $\mathcal{M} = \langle \mathcal{P}, \mathcal{S}, \mathcal{A}, \mathcal{R} \rangle$ where:

- $\mathcal{P} = \mathcal{P}_S \cup \mathcal{P}_A$ is a finite set of propositions;
- $\mathcal{S} \subseteq 2^{\mathcal{P}_S}$ is the set of all possible states;
- $\mathcal{A} \subseteq 2^{\mathcal{P}_A}$ is the finite set of actions; and
- $\mathcal{R} \subseteq \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ is the state-transition relation.

We model how the agents can influence the external world as a *policy*, and we consider it to simply be a set of state-action

pairs,

$$\pi = \{\langle s_i, a_i \rangle\}$$

where $s_i \in \mathcal{S}$ and $a_i \in \mathcal{A}(s)$ with $\mathcal{A}(s) = \{a \mid \exists \langle s, a, s' \rangle \in \mathcal{R}\}$ that is the set of actions that are applicable in s . It is the state-action table used of (Cimatti et al. 2003). It is also related to what the literature on MDPs calls a policy (Boutilier, Dean, and Hanks 1999).

Over the proposition variables, a propositional language \mathcal{L} with quantification extension can be defined by allowing standard connectives $\wedge, \vee, \rightarrow, \neg$ and quantifiers \exists, \forall . The resulting language is a logic of quantified boolean formulae (QBF) (Bryant 1992). A *symbol renaming operation*, which we use below, can be defined on \mathcal{L} . For a formula $\xi \in \mathcal{L}$, if \vec{x} and \vec{x}' are two vectors of propositional variables, then a variable renaming operation can be defined by $\xi[\vec{x}/\vec{x}']$ which means that all the appearances of variables \vec{x} are substituted by \vec{x}' .

Multiagent Systems

A multiagent system is composed of N agents, $AGS = \{T_1, \dots, T_N\}$. We model each agent T_i with a NSTD $\mathcal{M}_i = \langle \mathcal{P}_i, \mathcal{S}_i, \mathcal{A}_i, \mathcal{R}_i \rangle$ where $\mathcal{P}_i = \mathcal{P}_{i,S} \cup \mathcal{P}_{i,A}$, and associated with a local policy π_i . The multiagent system as a whole is modeled as a joint NSTD $\mathcal{M} = \langle \mathcal{P}, \mathcal{S}, \mathcal{A}, \mathcal{R} \rangle$ along with with a joint policy π . The set of proposition variables in the joint model is the union of those of individual agents: $\mathcal{P}_S = \bigcup_{i=1}^N \mathcal{P}_{i,S}$ and $\mathcal{P}_A = \bigcup_{i=1}^N \mathcal{P}_{i,A}$. The joint states, actions, and state transitions are the result of interpreting \mathcal{P}_S and \mathcal{P}_A as in the basic NSTD. The individual transition relation and the policy can be projected from the joint system: $\mathcal{R}_i = \exists_{-\mathcal{P}_i} \mathcal{R}$ and $\pi_i = \exists_{-\mathcal{P}_i} \pi$ where $-\mathcal{P}_i$ means the set of variables in the formula being quantified but not in \mathcal{P}_i . To distinguish from the dialogue model below, we call the individual models \mathcal{M}_i s and the joint model \mathcal{M} s, and the associated policies π_i s and π the *external* NSTDs and policies.

Dialogue Systems

In addition to the external transition system \mathcal{M}_i , each agent T_i is also associated with dialogue transition system $\mathcal{M}_{i,D} = \langle \mathcal{P}_{i,D}, \mathcal{S}_{i,D}, \mathcal{A}_{i,D}, \mathcal{R}_{i,D} \rangle$ and a dialogue policy $\pi_{i,D}$. The joint dialogue system is then defined as, $\mathcal{M}_D = \langle \mathcal{P}_D, \mathcal{S}_D, \mathcal{A}_D, \mathcal{R}_D \rangle$, along with a joint dialogue policy π_D . The set of proposition variables in the joint model is the union of those of individual agents: $\mathcal{P}_{S,D} = \bigcup_{i=1}^N \mathcal{P}_{i,S,D}$ and $\mathcal{P}_{A,D} = \bigcup_{i=1}^N \mathcal{P}_{i,A,D}$. The joint dialogue states, actions, and state transitions are the result of interpreting $\mathcal{P}_{S,D}$ and $\mathcal{P}_{A,D}$ as in the basic NSTD. The individual transition relation and the policy can be projected from the joint dialogue system: $\mathcal{R}_{i,D} = \exists_{-\mathcal{P}_{i,D}} \mathcal{R}_D$ and $\pi_{i,D} = \exists_{-\mathcal{P}_{i,D}} \pi_D$. Different from the external policy π , the requirement for a dialogue policy π_D is that the individual agent can execute the projected $\pi_{i,D}$ without knowing the other agents' dialogue states so that the dialogue can help the multiagent system coordinate the joint external policy execution.

Handling Inconsistent Specifications

In a multiagent system, different agents, acting on behalf of different users, may have inconsistent information about the

transitions, the initial and goal states. We need a mechanism to accommodate these inconsistent information, while the whole system can still reach an agreement on the joint model. Using the STRIPS-style specifications, for example STRIPS (Fikes and Nilsson 1971), PDDL (Ghallab et al. 1998), ADL (Pednault 1994) as the front-end specifications to our multiagent specifications, each agent T_i maintains a tuple of information

$$\langle TR_i, FR_i, MR_i, I_i, G_i \rangle$$

where

- $TR_i = \{TR_{i,k}\}$ is a list of local state transition specifications,
- $FR_i = \{FR_{i,k}\}$ is a list of frame specifications,
- $MR_i = \{MR_{i,k}\}$ is a list of multiagent interaction constraints,
- $I_i = \{I_{i,k}\}$ is a list of initial state specifications, and
- $G_i = \{G_{i,k}\}$ is a list of goal state specifications.

The entries in TR_i and FR_i depends only on the variables in \mathcal{P}_i , and entries in MR_i can have any variables in \mathcal{P} in general. Here, we take $FR_{i,k} = p_{i,k} \leftrightarrow p'_{i,k}$ for each state variable in $\mathcal{P}_{i,S}$ to specify that $p_{i,k}$ doesn't change by default.

Each entry in TR_i, FR_i, MR_i, I_i and G_i is associated with a label variable. The association is in the form of a one-to-one mapping

$$\langle LABEL_{i,k}, INPUT_{i,k} \rangle$$

stands for $\langle lt_{i,k}, TR_{i,k} \rangle, \langle lf_{i,k}, FR_{i,k} \rangle, \langle lm_{i,k}, MR_{i,k} \rangle, \langle li_{i,k}, I_{i,k} \rangle$, and $\langle lg_{i,k}, G_{i,k} \rangle$. Correspondingly, these variables are grouped into \mathcal{P}_{i,L^*} where $*$ stands any of TR, FR, MR, I, G . With these labeling variables, we can encode the selection of a set of specifications using QBFs. Let $\sigma \subseteq \bigcup_{i=1}^N (TR_i \cup FR_i \cup MR_i)$, we define

$$SEL(\sigma) =$$

$$\bigwedge_{TR_{i,k} \notin \sigma} (\neg lt_{i,k}) \wedge \bigwedge_{FR_{i,k} \notin \sigma} (\neg lf_{i,k}) \wedge \bigwedge_{MR_{i,k} \notin \sigma} (\neg lm_{i,k}) \\ \wedge \bigwedge_{TR_{i,k} \in \sigma} (lt_{i,k}) \wedge \bigwedge_{FR_{i,k} \in \sigma} (lf_{i,k}) \wedge \bigwedge_{MR_{i,k} \in \sigma} (lm_{i,k})$$

$$SEP(\sigma) = SEL(\sigma) \wedge$$

$$\bigwedge_{TR_{i,k} \in \sigma} TR_{i,k} \bigwedge_{FR_{i,k} \in \sigma} FR_{i,k} \bigwedge_{MR_{i,k} \in \sigma} MR_{i,k}$$

Let

$$C\{INPUT\}_i = \bigwedge_k (\{LABEL\}_{i,k} \rightarrow \{INPUT\}_{i,k})$$

where $\{INPUT\}$ stands any of TR, FR, MR, I, G , and $\{LABEL\}$ stands for the corresponding label variables to encode the local state transition specification combinations (CTR), frame specifications (CFR), inter-agent specifications respectively (CMR), initial states (CI) and goal states (CG). Now we can encode all the consistent combinations with linear number of BDD operations:

Algorithm 1 Weak Planning

```
1: function weakPlan( $R, I, G$ ) { (1)  $R$ : The transition relation;
  (2)  $I$ : The set of initial states; (3)  $G$ : The set of goals states }
2:  $FT \leftarrow G, SA \leftarrow \emptyset, CoveredStates \leftarrow FT$ 
3: repeat
4:    $newSA \leftarrow computePreImage(R, FT')$ 
5:    $FT \leftarrow \exists_{\mathcal{P}_A} newSA$ 
6:    $\Delta SA \leftarrow prune(SA, newSA)$ 
7:    $SA \leftarrow SA \vee \Delta SA$ 
8:    $CoveredStates \leftarrow CoveredStates \vee FT$ 
9: until  $\Delta SA == \emptyset$  OR ( $I \subseteq CoveredStates$ )
10: return  $SA$ 
11: end function
```

Proposition 1. Let $CONSR = \bigwedge_{i=1}^N (CTR_i \wedge CFR_i \wedge CMR_i)$, we have

$$CONSR = \bigvee_{\sigma \subseteq \bigcup_{i=1}^N (TR_i \cup FR_i \cup MR_i)} SEP(\sigma)$$

Proof. By $a \rightarrow b \equiv a \rightarrow (a \wedge b)$ and disjunctive expansion. \square

As $SEP(\sigma) = FALSE$ for any inconsistent combination, inconsistent combinations are automatically excluded from $CONSR$. However, $CONSR$ contains the empty combinations. We can overcome this by introducing a specification combination criteria SCC :

$$CONSR^+ = CONSR \wedge SCC$$

There can be many choices for SCC . One of them is $SCC = \neg ZTR_1 \wedge \neg ZTR_2 \wedge \dots \wedge \neg ZTR_N$ where $ZTR_i = \bigwedge_{TR_{i,k} \in TR_i} (\neg lt_{i,k})$. It is to encode that each agent must contribute at least one specification of their local transitions.

$CONSR^+$ contains the union of all consistent combinations which over-branches multiagent system into uncertainties. We can filter $CONSR^+$ by only keeping the maximal consistent sets of specifications

$$\mathcal{R}^{MR} = Max(CONSR^+, \xi(\subseteq)[\mathcal{P}_{LMR}])$$

$$\mathcal{R} = Max(\mathcal{R}^{MR}, \xi(\subseteq)[\mathcal{P}_{LTR}]) \quad (1)$$

The set maximal function Max can be computed using QBF/BDDs as showed in (Tang, Norman, and Parsons 2010). The joint maximal or minimal consistent initial states I and the goal states G can be computed in a similar manner which we omitted here due to the length limit. .

Multiagent Planning

Centralized Planning

We can then feed the I , G , and \mathcal{R} computed above into a basic weak planning process, as Algorithm 1 adapted from (Cimatti et al. 2003), to obtain a joint multiagent policy and then utilize the coordination dialogues proposed in (Tang, Norman, and Parsons 2009a) to coordinate the necessary inter-agent behaviors.

Algorithm 2 Agent T_i 's Pre-Image Computation

```
1: function computePreImage( $CTR_i, FT'$ ) { (1)  $CTR_i$ :
  Agent  $T_i$ 's local combination on transition specifications; (2)
   $FT'$ : The frontier states }
2: Compute  $CTR_i \wedge FT'$  and  $\neg \widehat{CTR}_i \wedge FT'$ 
3: for  $j = 1 \dots N$ , and  $j \neq i$  do
4:   Send  $CTR_i \wedge FT'$  to agent  $T_j$ 
5:   Send  $\neg \widehat{CTR}_i \wedge FT'$  to agent  $T_j$ 
6: end for
7: Wait for  $CTR_k \wedge FT'$  and  $\neg \widehat{CTR}_k \wedge FT'$  from all other agent
   $T_k$  ( $k = 1, \dots, N$  and  $k \neq i$ )
8:  $SA \leftarrow \exists_{\mathcal{P}'_S, \mathcal{P}_{LTR}, \mathcal{P}_{LMR}, \mathcal{P}_{LFR}} (R \wedge FT')$ 
9: return  $SA$ 
10: end function
```

Decentralized Planning

By decomposing the expression of \mathcal{R} (Equation 1), we have

$$R \wedge FT' = SHR_{TR} \wedge \bigwedge_{i=1}^N (CTR_i \wedge FT') \wedge \forall_{\mathcal{P}_{LTR}} \left(\bigvee_{i=1}^N (\neg \widehat{CTR}_i \wedge FT') \vee (\neg \hat{\xi}_{LTR} \vee \hat{\xi}_{LTR}) \wedge FT' \right)$$

where SHR_{TR} is an expression for the information shared by all the agents, $\hat{\xi}_{LTR}$ and $\hat{\xi}_{LTR}$ are QBFs to encode the subset (\subseteq) and superset (\supseteq) relations on the selection of $TR_{i,k}$ s. Assuming that each agent T_i knows all the labeling symbols of \mathcal{P}_{LTR} , then

- $CTR_i \wedge FT'$ can be computed by agent T_i locally
- $\neg \widehat{CTR}_i \wedge FT'$ can be computed by agent T_k locally

With this decomposition, we can then have a distributed version of *computePreImage* as in Algorithm 2.

Proposition 2. If every agent follows the algorithm 2 and algorithm 1, the every agent computes a same joint policy.

Proof. Every agent obtains the same joint state frontier FT' at each search step. The agents following algorithm 2 computes the same pre-image. Therefore, when the algorithm 1 end, all the agents will compute a same joint policy. \square

Note that the concern of the decentralized planning is not to save computations but to have all the agents reach consensus on a joint policy.

Integrating Dialogues

In (Tang, Norman, and Parsons 2009a) I integrate the planning of external policies with the planning of dialogue policies. The approach incorporates the information of external states, state transitions and policy into the dialogue states. Dialogue actions are then used to communicate the truth values of the variables carrying this information between agents. Dialogue state transitions are automatically generated regarding the ontology of the external states and actions, and the external policy. A dialogue policy is planned by setting the goal dialogue states in which the external joint states and actions are

- fully communicated with each other, or
- partially communicated to a level in which every agent can compute a unique local action agreeing with their shared joint policy

Implementations

I have been working on a C++ implementation of the system (Tang, Norman, and Parsons 2009b). The focus is on providing utilities to specify the complicated ontology over the propositional variables needed to implement multiagent planning and dialogues. In the implementation, I first specify the ontology for planning and dialogues, and then the system will automatically generate the BDD variables and BDDs for the QBFs. Currently the underlying BDD computation is implemented using my own naive implementation for the convenience of tuning the ontology manipulations. After the model becomes stable, I plan to bridge our implementation to a proven BDD package for efficiency and stability.

Conclusions

This paper has presented a model to integrate multiagent dialogues, planning, and plan execution. The model is symbolic, and capable of handling non-deterministic actions. In addition to the model, we have provided procedures for creating joint plans in a centralized or decentralized manner. In combination with our multiagent dialogue models, the plans can include the communication necessary for plan execution — that is the detection and communication of information relevant to the execution of the plan. I believe this is the first time that this kind of planning model, drawn from the literature of non-deterministic planning, has been combined with a communication model and then applied to multiagent teams.

I am working to extend the non-deterministic transition model to Markov Decision Processes (MDPs) (Tang and Parsons 2009) to utilize the statistical information. Another area of future work is incorporating prior work on cognitive models of non-monotonic reasoning and dialogues, the argumentation-based approaches, on planning (Tang and Parsons 2005), which assumes a simple, deterministic model of actions, with the work I have described here. Another area of future work, which addresses the main area in which this model falls short of a model of teamwork, is to consider the formation of joint intentions. Here there is a rich vein of work to draw on, for instance (Cohen and Levesque 1990; Grosz and Kraus 1999), and I will seek to incorporate this into the model.

References

- Boutilier, C.; Dean, T.; and Hanks, S. 1999. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research* 11:1–94.
- Bratman, M. E.; Israel, D. J.; and Pollack, M. E. 1988. Plans and resource-bounded practical reasoning. *Computational Intelligence* 4.
- Bryant, R. E. 1992. Symbolic boolean manipulation with ordered binary-decision diagrams. *ACM Computing Surveys* 24(3):293–318.
- Cimatti, A.; Pistore, M.; Roveri, M.; and Traverso, P. 2003. Weak, strong, and strong cyclic planning via symbolic model checking. *Artificial Intelligence* 147(1-2):35–84.
- Cohen, P. R., and Levesque, H. J. 1990. Intention is choice with commitment. *Artificial Intelligence* 42(2–3):213–261.
- Cohen, P. R., and Levesque, H. J. 1991. Teamwork. *Nous* 25(4):487–512.
- Fikes, R., and Nilsson, N. J. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. In *Proceedings of the 2nd International Joint Conference on Artificial Intelligence*, 608–620.
- Genesereth, M. R.; Ginsberg, M. L.; and Rosenschein, J. S. 1988. Cooperation without communication. In Bond, A. H., and Gasser, L., eds., *Distributed Artificial Intelligence*. San Francisco, CA: Morgan Kaufmann. 220–226.
- Ghallab, M.; Isi, C. K.; Penberthy, S.; Smith, D. E.; Sun, Y.; and Weld, D. 1998. PDDL - the planning domain definition language. Technical report, CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control.
- Grosz, B., and Kraus, S. 1999. The evolution of Shared-Plans. In Rao, A., and Woolridge, M., eds., *Foundations and Theories of Rational Agency*, 227–262.
- Parsons, S., and McBurney, P. 2003. Argumentation-based dialogues for agent co-ordination. *Group Decision and Negotiation* 12(5):415 – 439.
- Parsons, S.; Poltrock, S.; Bowyer, H.; and Tang, Y. 2008. Analysis of a recorded team coordination dialogue. In *Proceedings of the Second Annual Conference of the ITA*.
- Pednault, E. P. D. 1994. Adl and the state-transition model of action. *J. Log. Comput.* 4(5):467–512.
- Tang, Y., and Parsons, S. 2005. Argumentation-based dialogues for deliberation. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, 552–559. New York, NY, USA: ACM Press.
- Tang, Y., and Parsons, S. 2009. An MDP model for planning team actions with communication. Technical report, International Technology Alliance in Network and Information Science.
- Tang, Y.; Norman, T. J.; and Parsons, S. 2009a. A model for integrating dialogue and the execution of joint plans. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems*.
- Tang, Y.; Norman, T. J.; and Parsons, S. 2009b. Towards the implementation of a system for planning team activities. In *Proceedings of the Second Annual Conference of the ITA*.
- Tang, Y.; Norman, T. J.; and Parsons, S. 2010. Computing argumentation in polynomial number of bdd operations: A preliminary report. In *Seventh International Workshop on Argumentation in Multiagent Systems*.