

Integrating Belief-Desire-Intention Approaches with POMDPs: The Case of Team-Oriented Programs

Ranjit Nair and Milind Tambe and Stacy Marsella

Computer Science Department and Information Sciences Institute

University of Southern California

Los Angeles CA 90089

{nair,tambe}@usc.edu,marsella@isi.edu

Abstract

Integrating approaches based on belief-desire-intentions (BDI) logics with the more recent developments of distributed POMDPs is today a fundamental challenge in the multiagent systems arena. One common suggestion for such an integration is to use stochastic models (POMDPs) for generating agent behaviors, while using the BDI components for monitoring and creating explanations. We propose a completely inverse approach, where the BDI components are used to generate agent behaviors, and distributed POMDPs are used in an analysis mode. In particular, we focus on team-oriented programs for tasking multiagent teams, where the team-oriented programs specify hierarchies of team plans that the team and its subteams must adopt as their joint intentions. However, given a limited number of agents, finding a good way to allocate them to different teams and subteams to execute such a team-oriented program is a difficult challenge

We use distributed POMDPs to analyze different allocations of agents within a team-oriented program, and to suggest improvements to the program. The key innovation is to use the distributed POMDP analysis not as a black box, but as a glass box, offering insights into why particular allocations lead to good or bad outcomes. These insights help to prune the search space of different allocations, offering significant speedups in the search. We present preliminary experimental results to illustrate our methodology.

Introduction

Research in multiagent teamwork and cooperative multiagent systems in general has successfully used the belief-desire-intention (BDI) framework. Such BDI systems — explicitly or implicitly inspired by multi-modal logics based on beliefs, desires and intentions — have led to the development of several practical multiagent applications (Decker & Lesser 1993; Kitano *et al.* 1997). However, as we scale-up such multiagent teams, to 100s or 1000s of agents, robots or other entities, it becomes increasingly critical to provide analysis tools for such systems. For instance, in domains such as disaster rescue, such analysis will be important in order to specify how many agents and of what type to allocate to various roles in the team. These role allocations can have a drastic impact on the performance of the team and for

large teams, it will be difficult for human developers to even specify good allocation of roles within such teams.

Analysis of such role allocations is difficult for a variety of reasons. First, the challenge is not only to provide static allocations of agents to roles, but also to look ahead at what reallocations are necessary in the future as well. In particular, agents in a team may fail during their execution, and other agents may need to be reallocated to those roles. Role allocation must take into account the inevitable future costs of role reallocations. Second, there are significant uncertainties and costs associated with agents' execution of roles, and it is important to attempt to optimize such allocation (at least to some degree). Certainly, a simple check of finding the best match of capabilities to role requirements (Tidhar, Rao, & Sonenberg 1996) could lead to highly suboptimal team operations.

Fortunately, the recent emergence of distributed partially observable Markov decision processes (POMDPs) and MDPs have begun to provide tools for MAS analysis (Bernstein, Zilberstein, & Immerman 2000; Boutilier 1996; Peshkin *et al.* 2000; Pynadath & Tambe 2002; Xuan, Lesser, & Zilberstein 2001). These models can analyze complexity-optimality tradeoffs of multiagent coordination protocols. In particular, by encoding different multiagent coordination protocols as policies in distributed POMDPs, and comparing them against specific baseline policies, it is possible to investigate the potential for improvements in optimality and the computational costs that such improvements would engender. For instance, using as a baseline the globally optimal policy, we can identify if there are domain types where existing coordination protocols are (highly) suboptimal, when compared with the globally optimal policy. In addition, analysis of the computational complexity of the globally optimal policy informs us of the computational costs of such improvements.

Thus, the integration of BDI logic-based approaches for synthesis of multiagent behavior with POMDP-based analysis could lead to significant improvements in multiagent systems. Unfortunately, while the POMDP based analysis tools are powerful, they suffer from two key weaknesses. First, analysis in previous work focused on communication. In order to extend the analysis to other types of coordination and in particular to role allocation and reallocation in teams, we define *Role-based Markov Team Decision*

Problem(RMTDP). As important, we demonstrate an overall methodology by which such generalization gets done. A second, more critical problem in these analyses is that the globally optimal policy is often impossible to generate practically and thus any potential gains identified using this policy would be hard to obtain in practice. In other words, solving the distributed POMDPs to obtain an optimal joint policy is impractical. Furthermore, operationalizing of such policies can also be extremely difficult within a BDI system, as it may not allow all the flexibility that a globally optimal policy requires.

Our approach to addressing the second problem is to decompose the problem of role allocation and reallocation into two separate components, and attack each component separately. Such a divide-and-conquer strategy is certainly computationally attractive. Furthermore, it often matches the logical separation into two components seen in multiagent systems: First, there is an organizational structure and second, there are algorithms by which the organization coordinates its behavior. In this view, analysis can make suggestions both about improvements in organizational structure as well as more detailed suggestions about how the organization coordinates. It is important to stress that this decomposition into two components can be either explicit or implicit in the design of a MAS. Explicit representation has sometimes taken the form of a Team Oriented Program (TOP) (Tambe, Pynadath, & Chauvat 2000; Tidhar 1993; Tavares & Demazeau 2002) that specifies the organizational structure of team and the conditions which lead to organizational change. Second, there is a TOP “interpreter”, specifically coordination algorithms that interpret the organizational structure to ensure that the agents coherently coordinate in pursuit of goals. However, even in systems where there is not an explicit TOP, there is still an implicit organization in terms of the allocation of limited number of agents to different subteams that perform the tasks. Although further discussion will be made using the notion of an explicit TOP, the methodology and conclusions apply to other team planning approaches as well.

We first demonstrate improvements to the coordination component. Here, we focus on an automated approach to creating “locally optimal” algorithms that estimate the globally optimal policy. Such policies have been shown to be effective estimates that provide significant performance improvements with less significant computational costs than the globally optimal policy. In particular, Pynadath and Tambe (Pynadath & Tambe 2002) describe an algorithm for finding a locally optimal policy that provides an exponential speedup over the globally optimal policy. However, this previous work manually defined and generated a local optimal. Here we not only demonstrate an *automated* approach to generating alternative locally optimal but incorporate it into a methodology for doing team analysis. Our approach is based on examination of critical coordination events, “trigger” events for teamwork, and the perturbation of the team’s behavior at these coordination events.

Next we focus on making improvements to the team-oriented program itself. We address this problem by presenting an approach to analyzing and improving the TOP,

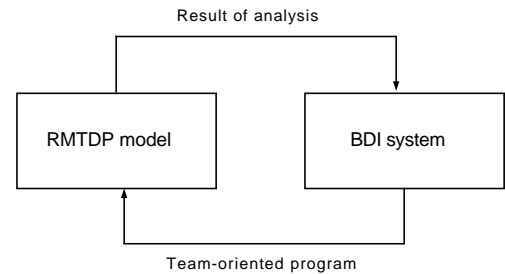


Figure 1: Integration of BDI and POMDP

using our analysis techniques. Analysis of alternative TOPs, of course, makes the intractability issue even more pronounced. Thus we take an even more radical approach here to addressing the tractability problem. Again, our approach is based on a perturbation-like process, but here we perturb the TOP itself and then seek to evaluate it in the RMTDP. Thus, our approach involves the following basic steps:

1. Start with developer specified team oriented program
2. Represent team oriented program as an RMTDP policy
3. Search space of team oriented programs using RMTDP for evaluation

A key contribution of our work is that we can significantly improve upon the process of searching this space of team-oriented programs. In particular, rather than treating the RMTDP as a black-box, we treat it as a glass-box, diagnosing the results of the evaluation to understand why specific results were obtained. We use these results to prune the search for an improvement in team-oriented programs. In particular, the diagnosis aids in exploiting the structure of the team oriented program to come up with component-wise upper bounds on performance. These heuristic upper bounds are obtained using the RMTDP evaluations and can be used to prune the team oriented programming space. We provide comparison of various approaches for searching the space of team oriented programs theoretically and empirically in a concrete domain involving teams of helicopter pilots. We illustrate that both techniques introduced in this paper lead to potential improvements in the search.

Our approach in this paper involves integrating the BDI-logic based Team-oriented Programming approach with a decentralized POMDP model, RMTDP, in a unique way. Rather than using the RMTDP for generating agent behavior and the BDI components for monitoring and creating explaining, we invert the roles as shown in Figure 1. Thus, the BDI system is used to come up with a Team-oriented Program, which is then evaluated and analyzed using the RMTDP model. The results of this analyses are then fed back to the BDI system.

Team Oriented Program

A Team Oriented Program specifies three key aspects of a team: (i) a team organization hierarchy; (ii) a team (reactive) plan hierarchy; and (iii) assignments of agents to execute plans.

We consider a group of helicopters involved in a mission of transporting cargo through enemy terrain as a running example throughout this paper to explain its different parts. We start with a fixed number of helicopters, say 6. These helicopters can be used as either transport or scouting helicopters and have a joint goal to get from a point X to point Y. There are 3 paths of different lengths and different risks to crashes that the helicopters can use to get from X to Y. When a scouting helicopter moves along a path the traversed portion becomes safe for other helicopters to travel on. A helicopter may either move from its current position to the next point on its path with a single move action or may remain where it is. When a scout fails (e.g., it crashes) it can be replaced by a transport however transports cannot be replaced by scouts. More details of this domain are presented in the next section.

The team organization hierarchy consists of roles for individuals and for groups of agents. For example, Figure 2 illustrates the organization hierarchy of the roles of helicopters involved in a mission of transporting cargo through enemy terrain. Each leaf node corresponds to a role for an individual agent, while the internal nodes correspond to (sub)teams of these roles. *Task Force* is thus the highest level team in this organization and *SctAI* is an individual role.

The second aspect of a team-oriented program involves specifying a hierarchy of reactive team plans. While these reactive team plans are much like reactive plans for individual agents, the key difference is that they explicitly express joint activities of the relevant team. The reactive team plans require that the developer specify the: (i) initiation conditions under which the plan is to be proposed; (ii) termination conditions under which the plan is to be ended, specifically, conditions which cause the reactive team plan to be achieved, irrelevant or unachievable; and (iii) team-level actions to be executed as part of the plan. Figure 3 shows an example from the evacuation scenario (please ignore the bracketed names [] for now). Here, high-level reactive team plans, such as **Execute Mission**, typically decompose into other team plans, such as **DoScouting**. **DoScouting** is itself achieved via other sub-plans such as **UseRoute1**. There may be additional relationships between sub-plans. An AND relationship is indicated with a solid arc while an OR relationship is indicated with a dotted arc. Thus, **DoScouting**, **DoTransport** and **RemainingScouts** must all three be done while any of **UseRoute1**, **UseRoute2** or **UseRoute3** need be performed.

The software developer must also specify any domain-specific coordination constraints in the execution of team plans. In the example program in Figure 3, the plan **Execute Mission** has three subplans: **DoScouting** which involves trying to make one path from X to Y safe for the transports, **DoTransport** to move the transports along a scouted path, and **RemainingScouts** for the scouts which haven't reached the destination yet to get there. The developer must represent the domain-specific constraint that a subteam assigned to perform **DoTransport** cannot do so until the other subteam assigned **DoScouting** has reached its masking locations and begun observing. However, **DoTransport** and **RemainingScouts** can be done in parallel.




Figure 2: TOP: Organization hierarchy with roles

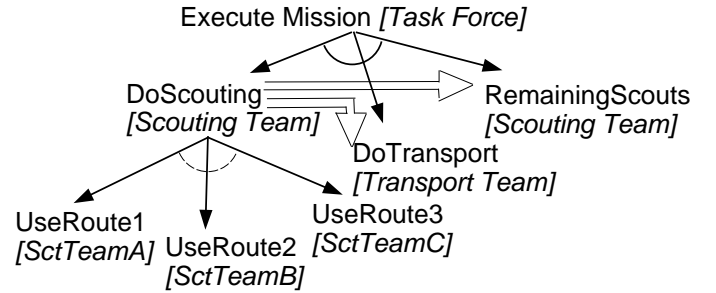


Figure 3: TOP: Partial reactive team plan hierarchy

The final aspect of team-oriented programming is assignments of agents to plans. This is done by first assigning the roles in the organization hierarchy to plans and then assigning agents to roles. Assigning only abstract roles rather than actual agents to plans provides a useful level of abstraction: new agents can be more quickly (re)assigned when needed. Figure 3 shows the assignment of roles to the reactive plan hierarchy for the helicopter domain (in brackets [] adjacent to the plans). For instance, *Task Force* team is assigned to jointly perform **Execute Mission**. Not all roles need be fulfilled however all roles within AND relationships and at least one within OR relationships must be fulfilled.

Markov Team Decision Problem

For quantitative analysis of role allocation and reallocation, we extend the Markov Team Decision Problem (MTDP) (Pynadath & Tambe 2002). While our extension focuses on role (re)allocation, it also illustrates a general methodology for analysis of other aspects of team coordination. Note that, while we use MTDP, other possible distributed POMDP models could potentially also serve as a basis (Bernstein, Zilberstein, & Immerman 2000; Xuan, Lesser, & Zilberstein 2001).

Given a team of agents α , an MTDP (Pynadath & Tambe 2002) is defined as a tuple: $\langle S, A, P, \Omega, O, R \rangle$. It consists of a finite set of states $S = \Xi_1 \times \dots \times \Xi_n$. Each agent i can perform an action from its set of actions A_i . $P(s, \langle a_1, \dots, a_{|\alpha|} \rangle, s')$ gives the probability of transitioning from state s to state s' given that the agents perform the actions $\langle a_1, \dots, a_{|\alpha|} \rangle$ jointly. Each agent i receives an observation $\omega_i \in \Omega_i$ based on the function $O(s, \langle a_1, \dots, a_{|\alpha|} \rangle, \omega_1, \dots, \omega_{|\alpha|})$, which gives

the probability that the agents receive the observations, $\omega_1, \dots, \omega_{|\alpha|}$ given that the world state is s and they perform $\langle a_1, \dots, a_{|\alpha|} \rangle$ jointly. The agents receive a single joint reward $R(s, a_1, \dots, a_{|\alpha|})$.

The state of the world, s need not be observable to the agent. Thus, each agent i chooses its actions based on its local *policy*, π_i , which is a mapping of its observation history to actions. Thus, at time t , agent i will perform action $\pi_i(\omega_i^0, \dots, \omega_i^t)$. $\pi = \langle \pi_1, \dots, \pi_{|\alpha|} \rangle$ refers to the joint policy of the team of agents.

Extension for explicit coordination: \mathcal{RL}

Beginning with MTDP, the next step in our methodology is to make an explicit separation between domain-level actions and the coordination actions of interest. Earlier work introduced the COM-MTDP model (Pynadath & Tambe 2002) where the coordination action was fixed to be the communication action. However, other coordination actions could also be separated from domain-level actions in order to investigate their impact. Thus, to investigate role allocation and reallocations, actions for allocating agents to roles and to reallocate such roles are separated out. To that end, we define RMTDP (Role-based Markov Team Decision Problem) as a tuple, $\langle S, A, P, \Omega, O, R, \mathcal{RL} \rangle$ with a new component, \mathcal{RL} . In particular, $\mathcal{RL} = \{r_1, \dots, r_s\}$ is a set of all roles that the agents can undertake. Each instance of role r_j may be assigned some agent i to fulfill it. Agents' actions are now distinguishable into two types:

Role-Taking actions: $\Upsilon = \prod_{i \in \alpha} \Upsilon_i$ is a set of combined role-taking actions, where $\Upsilon_i = \{v_{ir_j}\}$ contains the role-taking actions for agent i . $v_{ir_j} \in \Upsilon_i$ means that agent i takes on the role $r_j \in \mathcal{RL}$.

Role-Execution Actions: $\Phi = \prod_{i \in \alpha} \Phi_i$ is a set of combined execution actions, where $\Phi_i = \bigcup_{r_j \in \mathcal{RL}} \Phi_{ir_j}$ contains the execution actions for agent i . Φ_{ir_j} is the set of agent i 's actions for executing role $r_j \in \mathcal{RL}$.

Thus, in RMTDP, successive epochs alternate between role-taking (Υ) and role-execution actions (Φ). If the time index is divisible by 2, agents are in the role-taking epoch, executing role-taking actions, and otherwise they are in the role-execution epoch. Although this sequencing of role-taking and role-execution epochs restricts different agents from running role-taking and role-execution actions in the same epoch, it is conceptually simple and synchronization is automatically enforced. More importantly, the distinction between role-taking and -execution actions is critical to enable a separation in their costs, so as to more easily analyze the costs of role-taking actions. To this end, in RMTDP, reward is role-taking reward, $R_\Upsilon(s, a_1, \dots, a_{|\alpha|})$, for even time indices and role-execution reward, $R_\Phi(s, a_1, \dots, a_{|\alpha|})$, otherwise. We view the role-taking reward as the cost (negative reward) for taking up different roles in different teams. For instance, in our example domain, when transports change roles to be scouts, there is cost for dumping its cargo and loading scout equipment. However, such change of roles may potentially provide significant future rewards.

Within this model, we can represent the specialized behaviors associated with each role, e.g. a transport vs. a scout

role. While filling a particular role, r_j , agent i can only perform role-execution actions, $\phi \in \Phi_{ir_j}$, which may be different from the role-execution actions Φ_{ir_l} for role r_l . These different roles can produce varied effects on the world state (modeled via transition probabilities, P) and the team's utility. Thus, the policies must ensure that agents for each role have the capabilities that benefit the team the most.

Complexity results with RMTDP

While previous sections qualitatively emphasized the difficulty of role (re)allocation, RMTDP helps in understanding the complexity more precisely. In particular, we can define a role-taking policy, $\pi_{i\Upsilon}$ for each agent's role-taking action, a role-execution policy, $\pi_{i\Phi}$ for each agent's role-execution action. The goal in RMTDP is then to come up with joint policies π_Υ and π_Φ that will maximize the total reward over a finite horizon T . Such an optimal role taking policy not only provides for role allocation, but it also takes into account optimal future role reallocations. The following theorem illustrates the complexity of finding such optimal joint policies.

Theorem 1 *The decision problem of determining if there exist policies, π_Υ and π_Φ , for an RMTDP, that yield a reward at least K over some finite horizon T is NEXP-complete.*

Proof: Proof follows from the reduction of MTDP (Pynadath & Tambe 2002) to/from RMTDP. To reduce MTDP to RMTDP, we set RMTDP's role taking actions, Υ' to null. To reduce RMTDP to MTDP, we generate a new MTDP whose state space contains an additional feature to indicate if the current state corresponds to a role-taking or -executing stage of the RMTDP. The transition function, P' , augments the original function P : $P'(\langle \xi_{1b}, \dots, \xi_{nb}, \text{taking} \rangle, v_1, \dots, v_{|\alpha|}, \langle \xi_{1e}, \dots, \xi_{ne}, \text{executing} \rangle) = P(\langle \xi_{1b}, \dots, \xi_{nb} \rangle, v_1, \dots, v_{|\alpha|}, \langle \xi_{1e}, \dots, \xi_{ne} \rangle)$ where $v_1, \dots, v_{|\alpha|}$ is a role-taking action in the RMTDP (similarly from executing to taking). Finding the required policy in MTDP is NEXP-complete (Pynadath & Tambe 2002). \square

While the previous theorem focused on the complexity of combined role-taking and role execution actions, we can focus on the complexity of just determining the role taking actions, given fixed role-execution actions. Unfortunately, as the following theorem illustrates, determining an optimal policy for even this restricted problem has the same complexity as the original RMTDP.

Theorem 2 *The decision problem of determining if there exists a role-taking policy, π_Υ , for an RMTDP, that yields a reward at least K together with a fixed role-execution policy π_Φ , over some finite horizon T is NEXP-complete.*

Proof sketch: We begin with an MTDP and reduce it to an RMTDP with a fixed role-execution policy (in the simplest such fixed role-execution policy, agents execute NO-OPs). \square

Note that Theorem 2 refers to a completely general *globally optimal* role-taking policy, where any number of agents can change roles at any point in time. Given the above result, in general the globally optimal role-taking policy will be of

doubly exponential complexity, and so we may be left no choice but to run a brute-force policy search, i.e. to enumerate all the role-taking policies and then evaluate them, which together determines the run-time of finding the globally optimal policy. The number of policies is $\left(|\Upsilon|^{\frac{|\Omega|^T-1}{|\Omega|-1}}\right)^{|\alpha|}$, i.e.

doubly exponential in the finite horizon and the number of agents. This clearly illustrates the point made in Section , that the search for a globally optimal policy is intractable.

Note that, in the worst case, cost of evaluating a single policy can be given by $O\left((|S| \cdot |\Omega|)^T\right)$ (Pynadath & Tambe 2002). We will in general assume a fixed procedure for policy evaluation and primarily focus on the number of policies being evaluated. Improvement in policy evaluation could be an orthogonal dimension of investigation.

Constructing an RMTDP

Constructing an RMTDP for evaluating a TOP is a key step in our approach. To that end, we must define each of the elements of the RMTDP tuple, specifically, $\langle S, A, P, \Omega, O, R, \mathcal{RL} \rangle$, by a process that relies on both the TOP plans as well as the underlying domain. While this step has not been automated, we briefly describe mapping techniques based on the work on our two domains.

First, we need to define the set of states S . To this end, it is critical to model the variables tested in the preconditions and termination conditions of the TOP plans. For complex domains, it is useful to consider abstract descriptions of the state modeling only the significant variables. Agents’ role-taking and -execution actions in RMTDP are defined as follows. For each role in the TOP organization hierarchy, we define a role-taking action in each state s . The role-execution actions are those allowed for that role in the TOP plan hierarchy given the variable values in state s .

To illustrate these steps, consider the plans in Figure 3. The preconditions of plans such as **UseRoute1** and others test the start location of the helicopters to be at start location X, while the termination conditions test that scouts are at end location Y. Thus, the locations of all the helicopters are critical variables modeled in our set of states S . For role-taking, each helicopter can perform one of four actions, i.e. being a member of one of the three scouting teams or of the transport team. Role-execution actions are the TOP actions for the plan that the agent’s role is assigned in the TOP. In our case, the role execution policy for the scout role is to always go forward until it reaches Y, while for the transport role the policy is to wait at X until it obtains observation of a signal that one scouting subteam has reached Y.

Further, the types of observations for each agent must be defined. We define the set of observations to be the variables tested in the preconditions and termination conditions of the TOP plans and individual agent plans. For instance, the transport helos may observe the status of scout helos (normal or destroyed), as well as a signal that a path is safe. Finally, we must define the transition, observation and reward functions. Determining these functions requires some combination of human domain expertise and empirical data on the domain behavior. However, as shown later in Section , even

an approximate dynamic and observational model, is sufficient to deliver significant benefits. Defining the reward and transition function may sometimes require additional state variables to be modeled. In our helicopter domain, the time at which each the scouting and transport mission was completed determined the amount of reward and hence time was included as a state variable.

Analysis using Model

RMTDP can be shown to have NEXP-complete complexity, from reductions similar to COM-MTDP paper. In order to reduce the complexity, it is useful to define events, such as failure of an agent, which will act like a coordination trigger. This approach is seen often in implemented systems, for example, restrict the problem of “Team formation and reformation” to “Role Replacement”. For example, STEAM(Tambe 1997) assumes an initial team formation performed by a human, and focuses on reformation via role replacement, where a failed agent must be replaced by another. Similarly, the SharedPlans theory focuses on unreconciled actions(Grosz & Kraus 1996), where an agent or a subteam considers substituting for an unfilled (or failed) role.

In STEAM, an agent in role R will replace a failed agent in role F only if the following inequality holds:

$$\begin{aligned} \text{Criticality}(F) - \text{Criticality}(R) &> 0 \\ \text{Criticality}(x) &= 1 \text{ if } x \text{ is critical; } = 0 \text{ otherwise} \end{aligned} \quad (1)$$

In other words, replacement occurs if role F is considered critical and role R is not critical. Thus STEAM’s classification of coordination triggers is role-failure-critical or role-failure-not-critical. Our classification is more fine-grained, i.e., in a scenario, it may involve first-role-failure vs second-role-failure vs third-failure, etc.

Earlier methodology, as specified in COM-MTDP(Pynadath & Tambe 2002) dictated that we first derive by hand, an algorithm for a “locally optimal” policy. However, there are two problems in this derivation: (i) Deriving such a complex expression by hand is cumbersome and hinders analysis; (ii) The focus here remains on a single decision, and no guidance is provided on multiple decisions. It is possible to automatically generate various locally optimal policies by perturbing the response of a particular coordination trigger. For example, we can perturb STEAM’s response to the first failure that occurs and replace it by an optimal policy given that the response to all other triggers remains the same as STEAM’s. Various such locally optimal policies can be automatically generated by perturbing the response to one or more coordination trigger.

Apart from coming up with various locally optimal policies automatically, RMTDP is also useful in the analysis of the complexity and optimality of various approaches to the “Role Replacement” problem. For example in STEAM, criticality is determined in $O(|\alpha|)$ by processing role-dependency conditions supplied by a human. This is clearly very tractable especially when compared to the “globally optimal” policy, allows any number of agents to perform *any* role taking action at any point in time (even before the actual failure). The time complexity for finding the

globally optimal joint policy by searching this space is thus:

$$O \left(\left(|\Upsilon_\alpha|^{\frac{|\Omega_\alpha|^T - 1}{|\Omega_\alpha| - 1}} \right)^{|\alpha|} \cdot (|S| \cdot |\Omega_\alpha|^T) \right), \text{ i.e. doubly exponential in the finite horizon and the number of agents. The complexity of a locally optimal policy depends on the how "fine-grained" its response to the trigger is. For example, the complexity of a locally optimal policy that varies its response depending on what time the trigger occurred has complexity } O(2^T (|S| \cdot |\Omega_\alpha|^T)) \text{ while a locally optimal policy that varies its response depending on both time of the trigger and which trigger has complexity } O(2^{|\text{triggers}| * T} (|S| \cdot |\Omega_\alpha|^T)).$$

To further demonstrate the utility of RMTDPs, we now consider the example domain involving helicopter agents. These agents must decide whether to do a role replacement when a failure occurs. We compare the performance of various policies, across a space of distinct domains obtained by varying the cost of replacement and the probability of a failure.

In this experiment, we start with 6 helicopters and various starting configurations. For example, 3 scouts all assigned to path 2 (**UseRoute2**). When a scout fails (e.g., it crashes) it can be replaced by a transport by incurring a *Role replacement cost* for expending additional man-power. Once a transport becomes a scout it cannot transform back. (We assume that there is an ordering that determines which transport will perform a role replacement.) Further, we assume that a helicopter can fail at any unscouted point x between X and Y based on some known (uniform) probability distribution. To ensure a focus on role replacement, we assume that the policies for role execution and communication are the same for all approaches. A helicopter’s role execution policy while assigned to a scout role is that it will always go forward until it reaches Y , while the transport role execution policy is to wait at X until any one scout reaches Y . The reward is higher if more helicopters reach the destination safely and if they reach early rather than late.

We compared the performance of the various policies for different initial configurations of transport and scout helicopters. In the *STEAM* policy, the transports use the inequality 1 to determine whether to replace a failed scout. In *STEAM*, failure of the last remaining scout would be seen as *critical* and all other roles as non-critical. In *Seow* (Seow & How 2002), we consider a fixed utility function based approach to determine if a transport should replace a scout based on the configuration and the rewards for successful completion of scouts and transports. This is very similar to the role replacement policy used by FC Portugal in its championship RoboCup soccer team. *PertSteam1*, *PertSteam2* and *PertSteam3* are locally optimal perturbations of the *STEAM* policy where the response to the first, second or third failure, respectively, is replaced by the locally optimal response. Here the locally optimal response depends on the time of the failure and is obtained by calculating the expected future reward. The computational cost of calculating these policies is $O(2^T (|S| \cdot |\Omega_\alpha|^T))$. We compared these policies to a locally optimal policy where the

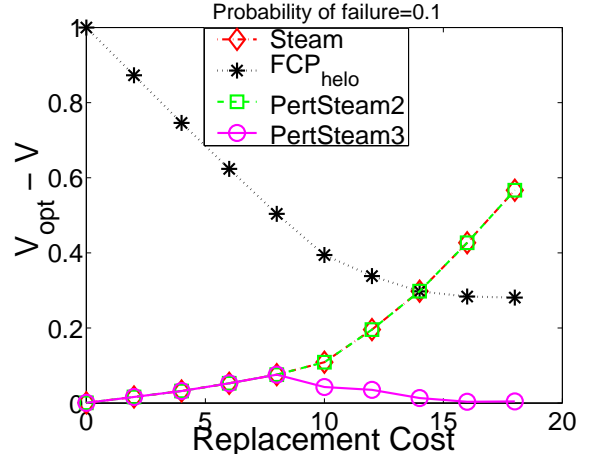


Figure 4: Sub-optimality of replacement policies with probability of failure fixed at 0.1

response to failure depends on which failure it was and also on the time of failure and obtained by calculating the expected future response. This policy is very expensive to compute $O(2^{|\text{failures}|T} (|S| \cdot |\Omega_\alpha|^T))$. The number of failures, $|\text{failures}|$ is $|\alpha|$, because in the worst case all the agents would fail. Note that this is cheaper than the globally optimal policy. We use this policy as the benchmark for the comparison of the various policies.

In figures 4 and 5 we compare the sub-optimality of various replacement policies when the initial configuration was 4 scouts and 2 transports. All 4 scouts were assigned to route 2. In both figures we plot the sub-optimality with respect to the benchmark policy on the Y-axis. In figure 4 we varied the replacement cost, keeping the probability of failure fixed at 0.1. Here we found that at low replacement costs *STEAM* and the locally optimal perturbations of *STEAM* were very close to benchmark while the *Seow* policy did quite poorly initially. However, with increasing replacement cost we find that *Seow* starts getting very close to the benchmark, while *STEAM* becomes more and more suboptimal. In figure 5 we varied the probability of failure keeping the replacement cost fixed at 10. As seen in this figure, all the policies do quite well at low failure rates. However when the probability of failure increases, the *Seow* policy does worse than *STEAM* and other locally optimal perturbations of *STEAM*.

Searching Team Oriented Program Space

In the previous section, we explored the derivation of locally optimal policies for reorganizing a team based only replacing failed roles. Whereas these dealt with finding a policy for team reorganization issue, it did not address the issue of finding a good initial allocation of agents to roles within a team plan. We now will look at this problem of finding the best initial configuration assuming that we have a fixed role replacement strategy. We will describe this search with respect to Team Oriented Program (TOP).

The TOP is used in several ways. First it reduces the num-

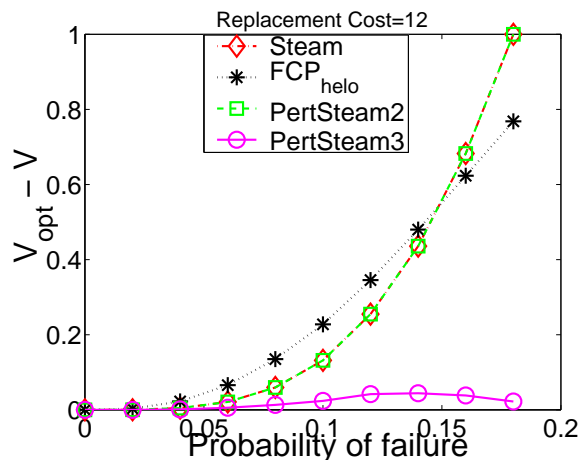


Figure 5: Sub-optimality of replacement policies with role replacement cost fixed

ber of policies that are explored in the RMTDP by using the TOP and perturbations to the TOP to suggest “local” policies to derive. Second, the TOP structure suggests alternative heuristics to be employed in the search for an optimal policy. Third, the TOP structure is used to simplify the search for an optimal policy by using the structure to decompose the RMTDP search space.

The TOP structure and perturbation of that structure can help guide the search for the optimal policy because the TOP by itself is in essence an abstract policy which imposes constraints on any specific policy consistent with it. To appreciate this fact, let’s consider again the TOP for our helicopter example. See Figure 3. Note the TOP imposes a hierarchical team organization that assigns roles to various domain agents as well as the number of agents in those roles. In addition, coordination constraints are imposed which specify which subteam tasks need to be performed, which subteam tasks are substitutes for each other as well as the order in which these tasks need to be completed. Not all teams and team policies will be consistent with this organizational structure and more specifically not all policies will explore the state space in a fashion consistent with the ordering constraints. Therefore the TOP constrains the space of policies and further any policy consistent with it can be used to evaluate the TOP.

We restrict the space of team oriented programs by varying only certain dimensions like initial team structure, initial role allocation and role relationships. This helps improve tractability. We assume a reasonable developer specified starting team plan. In order to demonstrate this methodology we restrict the dimensions of the team oriented program space to initial team structure and role allocation.

In order to find the best TOP we could check the evaluation of each and every initial configuration (initial team structure and initial role allocation) using RMTDP and choose the configuration that returns the highest value. This involves a brute force search through the entire search space of team oriented programs. The number of possible ini-

tial configurations is exponential in the number of agents and hence this method is not computationally viable when there are many agents and roles. Other methods like using combinatorial auctions to determine the initial configuration (Hunsberger & Grosz 2000) are faced with the same problems.

Pruning Search Space using evaluations

The process of searching this space of team oriented programs can be further improved using heuristic approaches that exploit the structure of the team oriented program to come up with component-wise upper bounds on performance. These heuristic upper bounds are obtained using the RMTDP evaluations and can be used to prune the team oriented programming space. The estimates are then summed up to get a overestimate of the expected reward for that internal node which we refer to as the max value. This is first done at the level just above the leaf nodes and can then be propagated up for all internal nodes. The process of arriving at component-wise estimates including the different heuristics that we applied are described in the following subsection.

Once this estimate is available for all internal nodes we begin evaluation of the leaf nodes. If a leaf node has a value higher than that of the previous best leaf we check if any pruning of the TOP space is possible. For this, we compare each internal node’s max value with the evaluation of the new best leaf node. If lesser, then we can eliminate the leaf node and all its descendants from the TOP space thus resulting in fewer leaf nodes to evaluate. Clearly pruning nodes higher in the hierarchy is more useful as this will likely result in greater pruning.

1. Parents \leftarrow list of parent nodes
2. for each parent in Parents do:
3. for each component in Team-Oriented Program
4. Find estimate of maximum expected reward
5. max[parent] \leftarrow Sum component-wise estimates
6. bestVal $\leftarrow -\infty$
7. for each parent in Parents do:
8. if done[parent] = true or pruned[parent] = true
9. continue
10. child \leftarrow parent \rightarrow nextChild()
11. if child = null
12. done[parent] \leftarrow true
13. continue
14. childVal \leftarrow Evaluate(child)
15. if childVal > bestVal
16. bestVal \leftarrow childVal
17. best \leftarrow child
18. for each parent1 in Parents do:
19. if max[parent1] < bestVal
20. pruned[parent1] \leftarrow true
21. return best

Figure 6: Algorithm for Searching TOP space.

Heuristics for Pruning

The process of obtaining over-estimates for each component relies on treating each component independently of the oth-

ers. In the case of components that are performed in sequence, the end states of the first component are the start states of the resulting components. Thus in order to accurately obtain an overestimate of the second component we need to be able to determine its start states. The key here is to avoid doing a full-scaled reachability analysis. The key idea to identify the start states for a component is to arrive at test to determine if a given state is valid start state. This test can be easily devised based on the start conditions of the component and the end-conditions for the successful completion of the previous component. The set of start states can be further reduced because not all the state variables of the state will be affected by the component. By examining the reward function, the model dynamics and the observation function we can identify which state variables are not affected by the component. We can put in dummy values for these state variables thus combining all states which have the same values for the relevant state variables. This will result in a much smaller set of state variables that can be obtained by making just one pass through the set of states.

We assume that components of the TOP that are performed in parallel cannot affect each other.

There are two main heuristics that we have applied for pruning. These are:

1. Component-wise maximum expected reward (MAXEXP heuristic)
2. Component-wise expected reward assuming no failures (NOFAIL heuristic)

In order to calculate the MAXEXP heuristic, we first identify the start states of each component. Then we evaluate each component separately from each of its start states and use the maximum evaluation as the MAXEXP heuristic. In order to calculate the NOFAIL heuristic, we evaluate each component separately from each of its start states but we assume that the probability of any action failing is 0. Thus all actions are considered to be deterministic. This will result in much less branching and hence evaluations will proceed much quicker. The NOFAIL heuristic only works if the evaluation without failures is always higher than with failures which should normally be the case.

The evaluation of the MAXEXP and NOFAIL heuristics for space of TOPs which are perturbations of the TOP described in figures 2 and 3 is shown in figure 7.

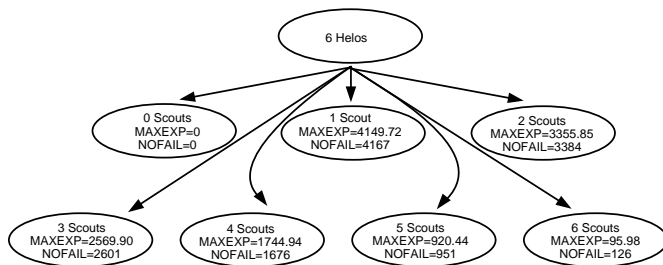


Figure 7: MAXEXP and NOFAIL heuristic values for different TOPs

| | MAXEXP | NOFAIL | NOPRUNE |
|--------------------------|--------|--------|---------|
| 6 Agents, 2 Comp. | 16 | 16 | 54 |
| 6 Agents, 3 Comp. | 16 | 16 | 54 |
| 7 Agents, 2 Comp. | 17 | 17 | 63 |

Table 1: Number of nodes of the TOP space evaluated

Experimental Results

In our experiments we used the same domain described in section 2 and 3. However here, we are trying to determine what the best initial assignment of agents to roles will be assuming that the replacement strategy and domain policy is fixed. For the purposes of this experiment we fixed the replacement strategy to be the same as STEAM, i.e. a transport will replace a failed scout only if the failure is critical. As can be seen in fig 3, there are 3 main components- **DoScouting, DoTransport and RemainingScouts**. Scouting is done first until one of the scouts reaches the destination. This is followed by Transport and RemainingScouts done in parallel. We tried two settings of the problem, with or without the RemainingScouts component. Also we tried starting with either 6 or 7 scouts.

We then compared the performance of the MAXEXP and NOFAIL heuristics for pruning with doing no pruning at all (NOPRUNE). The performance is compared in terms of number of nodes in the TOP space evaluated and these results are shown in table . All 3 techniques were able to find the best allocation. As can be seen, a lot of pruning was obtained by using these heuristic techniques thus demonstrating their usefulness.

Related Work

While the research in this paper focused on team-oriented programming(Tambe, Pynadath, & Chauvat 2000; Tidhar 1993) it is relevant to other techniques of modeling and tasking collectives of agents. For instance, Lesser et al's TAEMS approach for modeling tasks and coordination relationships is analogous to a team-oriented program and the analysis proposed here may provide benefits in terms of allocating or reallocating roles to agents.

Another key area of related work is team formation, which has relied on search using matching capabilities (Tidhar, Rao, & Sonenberg 1996) or combinatorial auctions(Hunsberger & Grosz 2000) to form teams. There are several key differences of this search process from the one discussed in our work. First, the search in the TOP space described in this paper uses stochastic models (RMTDPs) to compute costs. One key advantage is that RMTDPs enable the computation of not only the immediate benefits of team formation, but also the costs of reformation upon failure. Second, a major innovation in our current work is to use RMTDPs as a *glass box*, to extract key heuristics to prune the search space. The use of such models or their use in pruning has been absent in prior work.

Research on adaptive agent organizations is relevant as well. For instance, Horling et al illustrate heuristic techniques for modifying an agent organization. Barber and

MacMahon illustrate the difficult challenge of such adaptations in an agent organization (Barber & Martin 2001). Fortunately, RMTDP begins to provide the missing analytical tools that can help search the space of agent organizations more efficiently.

Finally, in terms of research on markov decision processes, previous sections have already discussed the relationship of RMTDP to COM-MTDP. We can also discuss RMTDP's relationship to other such distributed models. Given that the MTDP model is identical to the POIPSG model (Peshkin *et al.* 2000) and DEC-POMDP (Bernstein, Zilberstein, & Immerman 2000), RMTDP could be seen to enhance this model to enable explicit consideration of role allocation and reallocation.

Conclusion

Integrating approaches based on belief-desire-intentions (BDI) logics with the more recent developments of distributed POMDPs is today a fundamental challenge in the multiagent systems arena. We address this issue by using distributed POMDPs to analyze and direct the process of arriving at a good BDI based team plan.

In particular, we have presented an approach to analyzing and improving teamwork and empirically demonstrated the effectiveness of the approach. Our approach employs POMDP-based analysis but makes two key improvements to prior work in this area. First, we presented a formal framework that allows analysis of any aspect of team coordination. In contrast, prior work was restricted to the analysis of communication. Second, we addressed the central issue impacting the practicality of POMDP analysis, the cost of finding the globally optimal policy. We addressed this issue by decomposing the problem of finding a policy into a coordination component and a team oriented program component. Effective perturbation-based techniques for attacking each of these components of the overall problem were presented. This work is extended and described in more detail in (Nair, Tambe, & Marsella 2003).

Moving forward, we envision that this decomposition based approach could be realized within a larger iterative analysis process. Such a process would improve the team oriented program, in turn improve the coordination and then repeat the process.

Acknowledgment

This research was supported by grant #0208580 from the National Science Foundation. We thank Jim Blythe and David Pynadath for discussions related to the paper.

References

- Barber, S., and Martin, C. 2001. Dynamic reorganization of decision-making groups. In *Agents*.
- Bernstein, D. S.; Zilberstein, S.; and Immerman, N. 2000. The complexity of decentralized control of MDPs. In *UAI*.
- Boutilier, C. 1996. Planning, learning & coordination in multiagent decision processes. In *TARK*.
- Decker, K., and Lesser, V. 1993. Quantitative modeling of complex computational task environments. In *AAAI*.

- Grosz, B., and Kraus, S. 1996. Collaborative plans for complex group action. *Artificial Intelligence* 86(2):269–357.
- Hunsberger, L., and Grosz, B. 2000. A combinatorial auction for collaborative planning. In *ICMAS*.
- Kitano, H.; Asada, M. Kuniyoshi, Y.; Noda, I.; and Osawa, E. 1997. Robocup: The robot world cup initiative. In *IJCAI*.
- Nair, R.; Tambe, M.; and Marsella, S. 2003. Role allocation and reallocation in multiagent teams: Towards a practical analysis. In *AAMAS*.
- Peshkin, L.; Meuleau, N.; Kim, K.-E.; and Kaelbling, L. 2000. Learning to cooperate via policy search. In *UAI*.
- Pynadath, D., and Tambe, M. 2002. Multiagent teamwork: Analyzing the optimality complexity of key theories and models. In *AAMAS*.
- Seow, K., and How, K. 2002. Collaborative assignment: a multi-agent negotiation approach using bdi concepts. In *AAMAS*.
- Tambe, M.; Pynadath, D.; and Chauvat, N. 2000. Building dynamic agent organizations in cyberspace. *IEEE Internet Computing* 4(2).
- Tambe, M. 1997. Towards flexible teamwork. *JAIR* 7:83–124.
- Tavares, J., and Demazeau, Y. 2002. Vowels co-ordination model. In *AAMAS*.
- Tidhar, G.; Rao, A.; and Sonenberg, E. 1996. Guided team selection. In *ICMAS*.
- Tidhar, G. 1993. Team-oriented programming: Social structures. Technical Report 47, Australian A.I. Institute.
- Xuan, P.; Lesser, V.; and Zilberstein, S. 2001. Communication decisions in multiagent cooperation. In *Agents*.