

Planning for Decentralized Control of Multiple Robots Under Uncertainty

Christopher Amato¹, George D. Konidaris¹, Gabriel Cruz¹

Christopher A. Maynor², Jonathan P. How² and Leslie P. Kaelbling¹

¹CSAIL, ²LIDS
MIT
Cambridge, MA 02139

Abstract

We describe a probabilistic framework for synthesizing control policies for general multi-robot systems, given environment and sensor models and a cost function. Decentralized, partially observable Markov decision processes (Dec-POMDPs) are a general model of decision processes where a team of agents must cooperate to optimize some objective (specified by a shared reward or cost function) in the presence of uncertainty, but where communication limitations mean that the agents cannot share their state, so execution must proceed in a decentralized fashion. While Dec-POMDPs are typically intractable to solve for real-world problems, recent research on the use of macro-actions in Dec-POMDPs has significantly increased the size of problem that can be practically solved as a Dec-POMDP. We describe this general model, and show how, in contrast to most existing methods that are specialized to a particular problem class, it can synthesize control policies that use whatever opportunities for coordination are present in the problem, while balancing off uncertainty in outcomes, sensor information, and information about other agents. We use three variations on a warehouse task to show that a single planner of this type can generate cooperative behavior using task allocation, direct communication, and signaling, as appropriate.

Introduction

The decreasing cost and increasing sophistication of recently available robot hardware has the potential to create many new opportunities for applications where teams of relatively cheap robots can be deployed to solve real-world problems. Practical methods for coordinating such multi-robot teams are therefore becoming critical. A wide range of approaches have been developed for solving specific classes of multi-robot problems, such as task allocation [15], navigation in a formation [5], cooperative transport of an object [20], coordination with signaling [6] or communication under various limitations [33]. Broadly speaking, the current state of the art in multi-robot research is to hand-design special-purpose controllers that are explicitly designed to exploit some property of the environment or produce a specific desirable behavior. Just as in the single-robot case, it would be much more desirable to instead specify a world model and a cost metric, and then have a general-purpose planner automatically derive a controller that minimizes cost, while remaining robust to the uncertainty that is fundamental to real robot

systems [37].

The decentralized partially observable Markov decision process (Dec-POMDP) is a general framework for representing multiagent coordination problems. Dec-POMDPs have been studied in fields such as control [1, 23], operations research [8] and artificial intelligence [29]. Like the MDP [31] and POMDP [17] models that it extends, the Dec-POMDP model is very general, considering uncertainty in outcomes, sensors and information about the other agents, and aims to optimize policies against a general cost function. Dec-POMDP problems are often characterized by incomplete or partial information about the environment and the state of other agents due to limited, costly or unavailable communication. Any problem where multiple agents share a single overall reward or cost function can be formalized as a Dec-POMDP, which means a good Dec-POMDP solver would allow us to automatically generate control policies (including policies over when and what to communicate) for very rich decentralized control problems, in the presence of uncertainty. Unfortunately, this generality comes at a cost: Dec-POMDPs are typically infeasible to solve except for very small problems [3].

One reason for the intractability of solving large Dec-POMDPs is that current approaches model problems at a low level of granularity, where each agent’s actions are primitive operations lasting exactly one time step. Recent research has addressed the more realistic *MacDec-POMDP* case where each agent has *macro-actions*: temporally extended actions which may require different amounts of time to execute [3]. *MacDec-POMDPs* cannot be reduced to Dec-POMDPs due to the asynchronous nature of decision-making in this context — some agents may be choosing new macro-actions while others are still executing theirs. This enables systems to be modeled so that coordination decisions only occur at the level of deciding which macro-actions to execute. *MacDec-POMDPs* retain the ability to coordinate agents while allowing near-optimal solutions to be generated for significantly larger problems than would be possible using other Dec-POMDP-based methods.

Macro-actions are a natural model for the modular controllers often sequenced to obtain robot behavior. The macro-action approach leverages expert-designed or learned controllers for solving subproblems (e.g., navigating to a waypoint or grasping an object), bridging the gap between

traditional robotics research and work on Dec-POMDPs. This approach has the potential to produce high-quality general solutions for real-world heterogeneous multi-robot coordination problems by automatically generating control and communication policies, given a model.

The goal of this paper is to present this general framework for solving decentralized cooperative partially observable robotics problems and provide the first demonstration of such a method running on real robots. We begin by formally describing the Dec-POMDP model, its solution and relevant properties, and describe MacDec-POMDPs and a memory-bounded algorithm for solving them. We then describe a process for converting a robot domain into a MacDec-POMDP model, solving it, and using the solution to produce a SMACH [9] finite-state machine task controller. Finally, we use three variants of a warehouse task to show that a MacDec-POMDP planner allows coordination behaviors to emerge automatically by optimizing the available macro-actions (allocating tasks, using direct communication, and employing signaling, as appropriate). We believe the MacDec-POMDP represents a foundational algorithmic framework for generating solutions for a wide range of multi-robot systems.

Decentralized, Partially Observable Markov Decision Processes

Dec-POMDPs [8] generalize partially observable Markov decision processes to the multiagent, decentralized setting. Multiple agents operate under uncertainty based on (possibly different) partial views of the world, with execution unfolding over a bounded or unbounded sequence of steps. At each step, every agent chooses an action (in parallel) based purely on locally observable information, resulting in an immediate reward and an observation being obtained by each individual agent. The agents share a single reward or cost function, so they should cooperate to solve the task, but their local views mean that operation is decentralized during execution.

As depicted in Fig. 1, a Dec-POMDP [8] involves multiple agents that operate under uncertainty based on different streams of observations. We focus on solving sequential decision-making problems with discrete time steps and stochastic models with finite states, actions, and observations, though the model can be extended to continuous problems. A key assumption is that state transitions are *Markovian*, meaning that the state at time t depends only on the state and events at time $t - 1$. The reward is typically only used as a way to specify the objective of the problem and is not observed during execution.

More formally, a Dec-POMDP is described by a tuple $\langle I, S, \{A_i\}, T, R, \{\Omega_i\}, O, h \rangle$, where

- I is a finite set of agents.
- S is a finite set of states with designated initial state distribution b_0 .
- A_i is a finite set of actions for each agent i with $A = \times_i A_i$ the set of joint actions, where \times is the Cartesian product operator.




Figure 1: Representation of n agents in a Dec-POMDP setting with actions a_i and observations o_i for each agent i along with a single reward r .

- T is a state transition probability function, $T : S \times A \times S \rightarrow [0, 1]$, that specifies the probability of transitioning from state $s \in S$ to $s' \in S$ when the actions $\vec{a} \in A$ are taken by the agents. Hence, $T(s, \vec{a}, s') = \Pr(s' | \vec{a}, s)$.
- R is a reward function: $R : S \times A \rightarrow \mathbb{R}$, the immediate reward for being in state $s \in S$ and taking the actions $\vec{a} \in A$.
- Ω_i is a finite set of observations for each agent, i , with $\Omega = \times_i \Omega_i$ the set of joint observations.
- O is an observation probability function: $O : \Omega \times A \times S \rightarrow [0, 1]$, the probability of seeing observations $\vec{o} \in \Omega$ given actions $\vec{a} \in A$ were taken which results in state $s' \in S$. Hence $O(\vec{o}, \vec{a}, s') = \Pr(\vec{o} | \vec{a}, s')$.
- h is the number of steps until the problem terminates, called the horizon.

Note that while the actions and observation are factored, the state need not be. This flat state representation allows more general state spaces with arbitrary state information outside of an agent (such as target information or environmental conditions). Because the full state is not directly observed, it may be beneficial for each agent to remember a history of its observations. Specifically, we can consider an action-observation history for agent i as

$$H_i^A = (s_i^0, a_i^1, \dots, s_i^{l-1}, a_i^l).$$

Unlike in POMDPs, it is not typically possible to calculate a centralized estimate of the system state from the observation history of a single agent, because the system state depends on the behavior of all of the agents.

Solutions

A solution to a Dec-POMDP is a *joint policy*—a set of policies, one for each agent in the problem. Since each policy is a function of history, rather than of a directly observed state, it is typically represented as either a policy tree, where the vertices indicate actions to execute and the edges indicate transitions conditioned on an observation, or as a finite state controller which executes in a similar manner. An example of each is given in Figure 2.

As in the POMDP case, the goal is to maximize the total cumulative reward, beginning at some initial distribution




Figure 2: A single agent’s policy represented as (a) a policy tree and (b) a finite-state controller with initial state shown with a double circle.

over states b_0 . In general, the agents do not observe the actions or observations of the other agents, but the rewards, transitions, and observations depend on the decisions of all agents. The work discussed in this paper (and the vast majority of work in the Dec-POMDP community) considers the case where the model is assumed to be known to all agents.

The value of a joint policy, π , from state s is

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{h-1} \gamma^t R(\vec{a}^t, s^t) | s, \pi \right],$$

which represents the expected value of the immediate reward for the set of agents summed for each step of the problem given the action prescribed by the policy until the horizon is reached. In the finite-horizon case, the discount factor, γ , is typically set to 1. In the infinite-horizon case, as the number of steps is infinite, the discount factor $\gamma \in [0, 1)$ is included to maintain a finite sum and $h = \infty$. An *optimal policy* beginning at state s is $\pi^*(s) = \arg \max_\pi V^\pi(s)$.

Unfortunately, large problem instances remain intractable: some advances have been made in optimal algorithms [1, 2, 4, 10, 12, 27], but optimally solving a Dec-POMDP is NEXP-complete, so most approaches that scale well make very strong assumptions about the domain (e.g., assuming a large amount of independence between agents) [13, 24, 26] and/or have no guarantees about solution quality [28, 34, 38].

Macro-Actions for Dec-POMDPs

Dec-POMDPs typically require synchronous decision-making: every agent repeatedly determines which action to execute, and then executes it within a single time step. This restriction is problematic for robot domains for two reasons. First, robot systems are typically endowed with a set of controllers, and planning consists of sequencing the execution of those controllers. However, due to both environmental and controller complexity, the controllers will almost always execute for an extended period, and take differing amounts of time to run. Synchronous decision-making would thus require us to wait until all robots have completed their controller execution before we perform the next action selection, which is suboptimal and may not even always be possible (since the robots do not know the system state and staying in place may be difficult in some domains). Second, the

planning complexity of a Dec-POMDP is doubly exponential in the horizon. A planner that must try to reason about all of the robots’ possible policies at every time step will only ever be able to make very short plans.

Recent research has extended the Dec-POMDP model to plan using *options*, or temporally extended actions [3]. This MacDec-POMDP formulation models a group of robots that must plan by sequencing an existing set of controllers, enabling planning at the appropriate level to compute near-optimal solutions for problems with significantly longer horizons and larger state-spaces.

We can gain additional benefits by exploiting known structure in the multi-robot problem. For instance, most controllers only depend on locally observable information and do not require coordination. For example, consider a controller that navigates a robot to a waypoint. Only local information is required for navigation—the robot may detect other robots but their presence does not change its objective, and it simply moves around them—but choosing the target waypoint likely requires the planner to consider the locations and actions of all robots. Macro-actions with independent execution allow coordination decisions to be made only when necessary (i.e., when choosing macro-actions) rather than at every time step. Because we build on top of Dec-POMDPs, macro-action choice may depend on history, but during execution macro-actions may depend only on a single observation, depend on any number of steps of history, or even represent the actions of a set of robots. That is, macro-actions are very general and can be defined in such a way to take advantage of the knowledge available to the robots during execution.

Model

We first consider macro-actions that only depend on a single robot’s information. This is an extension the *options framework* [36] to multi-agent domains while dealing with the lack of synchronization between agents. The options framework is a formal model of a macro-actions [36] that has been very successful in aiding representation and solutions in single robot domains [19]. A MacDec-POMDP with local options is defined as a Dec-POMDP where we also assume M_i represents a finite set of options for each agent, i , with $M = \times_i M_i$ the set of joint options [3]. A *local option* is defined by the tuple:

$$M_i = (\beta_{m_i}, \mathcal{I}_{m_i}, \pi_{m_i}),$$

consisting of stochastic termination condition $\beta_{m_i} : H_i^A \rightarrow [0, 1]$, initiation set $\mathcal{I}_{m_i} \subset H_i^A$ and option policy $\pi_{m_i} : H_i^A \times A_i \rightarrow [0, 1]$. Note that this representation uses action-observation histories of an agent in the terminal and initiation conditions as well as the option policy. Simpler cases can consider reactive policies that map single observations to actions as well as termination and initiation sets that depend only on single observations. This is especially appropriate when the agent has knowledge about aspects of the state necessary for option execution (e.g., its own location when navigating to a waypoint causing observations to be location estimates). As we later discuss, initiation and termi-

nal conditions can depend on global states (e.g., also ending execution based on unobserved events).

Because it may be beneficial for agents to remember their histories when choosing which option to execute, we consider policies that remember option histories (as opposed to action-observation histories). We define an *option history* as

$$H_i^M = (h_i^0, m_i^1, \dots, h_i^{l-1}, m_i^l),$$

which includes both the action-observation histories where an option was chosen and the selected options themselves. The option history also provides an intuitive representation for using histories within options. It is more natural for option policies and termination conditions to depend on histories that begin when the option is first executed (action-observation histories) while the initiation conditions would depend on the histories of options already taken and their results (option histories). While a history over primitive actions also provides the number of steps that have been executed in the problem (because it includes actions and observations at each step), an option history may require many more steps to execute than the number of options listed. We can also define a (stochastic) local policy, $\mu_i : H_i^M \times M_i \rightarrow [0, 1]$ that depends on option histories. We then define a joint policy for all agents as μ .

Because option policies are built out of primitive actions, we can evaluate policies in a similar way to other Dec-POMDP-based approaches. Given a joint policy, the primitive action at each step is determined by the high level policy which chooses the option and the option policy which chooses the action. The joint policy and option policies can then be evaluated as:

$$V^\mu(s) = \mathbb{E} \left[\sum_{t=0}^{h-1} \gamma^t R(\bar{a}^t, s^t) | s, \pi, \mu \right].$$

For evaluation in the case where we define a set of options which use observations (rather than histories) for initiation, termination and option policies (while still using option histories to choose options) see Amato, Konidaris and Kaelbling [3].

Algorithms

Because Dec-POMDP algorithms produce policies mapping agent histories to actions, they can be extended to consider options instead of primitive actions. Two such algorithms have been extended [3], but other extensions are possible.

In these approaches, deterministic policies are generated which are represented as policy trees (as shown in Figure 2). A policy tree for each agent defines a policy that can be executed based on local information. The root node defines the option to choose in the known initial state, and another option is assigned to each of the legal terminal states of that option; this continues for the depth of the tree. Such a tree can be evaluated up to a desired (low-level) horizon using the policy evaluation given above, which may not reach some nodes of the tree due to the differing execution times of some options.

A simple exhaustive search method can be used to generate hierarchically optimal deterministic policies. This algorithm is similar in concept to the dynamic programming




Figure 3: Policies for a single agent after (a) one step and (b) two steps of dynamic programming using options m_1 and m_2 and (deterministic) terminal states as β^s .

algorithm used in Dec-POMDPs [16], but full evaluation and pruning (removing dominated policies) are not used. Instead the structure of options is exploited to reduce the space of policies considered. That is, to generate deterministic policies, trees are built up as in Figure 3. Trees of increasing depth are constructed until all of the policies are guaranteed to terminate before the desired horizon. When all policies are sufficiently long, all combinations of these policies can be evaluated as above (by flattening out the policies into primitive action Dec-POMDP policies, starting from some initial state and proceeding until h). The combination with the highest value at the initial belief state, b_0 , is a hierarchically optimal policy. Note that the benefit of this approach is that only legal policies are generated using the initiation and terminal conditions for options.

Memory-bounded dynamic programming (MBDP) [34] has also been extended to use options as shown in Algorithm 1. This approach bounds the number of trees that are generated by the method above as only a finite number of policy trees are retained (given by parameter *MaxTrees*) at each tree depth. To increase the tree depth to $t + 1$, all possible trees are considered that choose some option and then have the trees retained from depth t as children. Trees are chosen by evaluating them at states that are reachable using a heuristic policy that is executed for the first $h - t - 1$ steps of the problem. A set of *MaxTrees* states is generated and the highest-valued trees for each state are kept. This process continues, using shorter heuristic policies until all combinations of the retained trees reach horizon h . Again, the set of trees with the highest value at the initial belief state is returned.

The MBDP-based approach is potentially suboptimal because a fixed number of trees are retained, and trees optimized at the states provided by the heuristic policy may be suboptimal (because the heuristic policy may be suboptimal and the algorithm assumes the states generated by the heuristic policy are known initial states for the remaining policy tree). Nevertheless, since the number of policies at each step is bounded by *MaxTrees*, MBDP has time and space complexity linear in the horizon. As a result, this approach has been shown to work well in many large MacDec-POMDPs [3].

Solving Multi-Robot Problems with MacDec-POMDPs

The MacDec-POMDPs framework is a natural way to represent and generate behavior for general multi-robot systems. A high-level description of this process is given in

Algorithm 1 Option-based memory bounded dynamic programming

```
1: function OPTIONMBDP( $MaxTrees, h, H_{pol}$ )
2:    $t \leftarrow 0$ 
3:    $someTooShort \leftarrow true$ 
4:    $\mu_t \leftarrow \emptyset$ 
5:   repeat
6:      $\mu_{t+1} \leftarrow \text{GenerateNextStepTrees}(\mu_t)$ 
7:     Compute  $V^{\mu_{t+1}}$ 
8:      $\hat{\mu}_{t+1} \leftarrow \emptyset$ 
9:     for all  $k \in MaxTrees$  do
10:       $s_k \leftarrow \text{GenerateState}(H_{pol}, h - t - 1)$ 
11:       $\hat{\mu}_{t+1} \leftarrow \hat{\mu}_{t+1} \cup \arg \max_{\mu_{t+1}} V^{\mu_{t+1}}(s_k)$ 
12:    end for
13:     $t \leftarrow t + 1$ 
14:     $\mu_t \leftarrow \hat{\mu}_{t+1}$ 
15:     $someTooShort \leftarrow \text{testLength}(\mu_t)$ 
16:  until  $someTooShort = false$ 
17:  return  $\mu_t$ 
18: end function
```

Figure 4. We assume an abstract model of the system is given in the form of macro-action representations, which include the associated policies as well as initiation and terminal conditions. These macro-actions are controllers operating in (possibly) continuous time with continuous actions and feedback, but their operation is discretized for use with the planner. This discretization represents an underlying discrete Dec-POMDP which consists of the primitive actions, states of the system and the associated rewards. The Dec-POMDP methods discussed above typically assume a full model is given, but in this work, we make the more realistic assumption that we can simulate the macro-actions in an environment that is similar to the real-world domain. As a result, we do not need a full representation of the underlying Dec-POMDP and use the simulator to test macro-action completion and evaluate policies. In the future, we plan to remove this underlying Dec-POMDP modeling and instead represent the macro-action initiation, termination and policies using features directly in the continuous robot state-space. In practice, models of each macro-action’s behavior can be generated by executing the corresponding controller from a variety of initial conditions (which is how our model and simulator was constructed in the experiment section). Given the macro-actions and simulator, the planner then automatically generates a solution which optimizes the value function with respect to the uncertainty over outcomes, sensor information and other agents. This solution comes in the form of SMACH controllers [9] which are hierarchical state machines for use in a ROS [32] environment. Each node in the SMACH controller represents a macro-action which is executed on the robot and each edge corresponds to a terminal condition. In this way, the trees in Figure 3 can be directly translated into SMACH controllers, one for each robot. Our system is thus able to automatically generate SMACH controllers, which are typically designed by hand, for complex, general multi-robot systems.

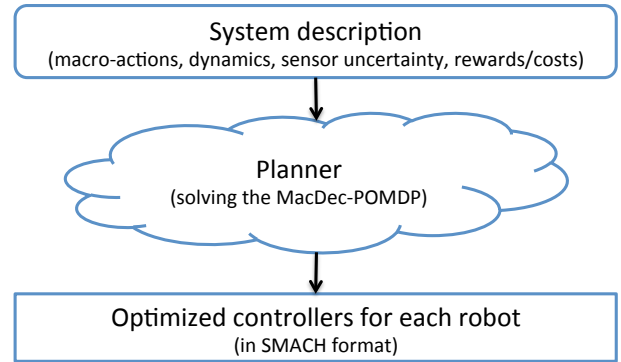


Figure 4: A high level system diagram.

It is also worth noting that our approach can incorporate existing solutions for more restricted scenarios as macro-actions. For example, our approach can build on the large amount of research in single and multi-robot systems that has gone into solving difficult problems such as navigation in a formation [5] or cooperative transport of an object [20]. The solutions to these problems could be represented as macro-actions in our framework, building on existing research to solve even more complex multi-robot problems.

Planning using MacDec-POMDPs in the Warehouse Domain

We test our methods in a warehousing scenario using a set of iRobot Creates (Figure 5), and demonstrate how the same general model and solution methods can be applied in versions of this domain with different communication capabilities. This is the first time that Dec-POMDP-based methods have been used to solve large multi-robot domains. We do not compare with other methods because other Dec-POMDP cannot solve problems of this size and current multi-robot methods cannot automatically derive solutions for these multifaceted problems. The results demonstrate that our methods can automatically generate the appropriate motion and communication behavior while considering uncertainty over outcomes, sensor information and other robots.

The Warehouse Domain

We consider three robots in a warehouse that are tasked with finding and retrieving boxes of two different sizes: large and small. Robots can navigate to known depot locations (rooms) to retrieve boxes and bring them back to a designated drop-off area. The larger boxes can only be moved effectively by two robots (if a robot tries to pick up the large box by itself, it will move to the box, but fail to pick it up). While the locations of the depots are known, the contents (the number and type of boxes) are unknown. Our planner generates a SMACH controller for each of the robots offline which are then executed online in a decentralized manner.

In each scenario, we assumed that each robot could observe its own location, see other robots if they were within (approximately) one meter, observe the nearest box when




Figure 5: The warehouse domain with three robots.

in a depot and observe the size of the box if it is holding one. These observations were implemented within a Vicon setup to allow for system flexibility, but the solutions would work in any setting in which these observations are generated. In the simulator that is used by the planner to generate and evaluate solutions, the resulting state space of the problem includes the location of each robot (discretized into nine possible locations) and the location of each of the boxes (in a particular depot, with a particular robot or at the goal). The primitive actions are to move in four different directions as well as pickup, drop and communication actions. Note that this primitive state and action representation is used for evaluation purposes and not actually implemented on the robots (which just utilize the SMACH controllers). Higher fidelity simulators could also be used. The three-robot version of this scenario has 1,259,712,000 states, which is several orders of magnitude larger than problems typically solvable by Dec-POMDP solvers. These problems are solved using the option-based MBDP algorithm initialized with a hand coded heuristic policy. Experiments were run on a single core of a 2.5 GHz machine with 8GB of memory. Exact solution times were not calculated, but average solution times for the policies presented below were approximately one hour.

In our Dec-POMDP model, navigation has a small amount of noise in the amount of time required to move to locations (reflecting the real-world dynamics): this noise increases when the robots are pushing the large box (reflecting the need for slower movements and turns in this case). We defined macro-actions that depend only on the observations above, but option selection depends on the history of options executed and observations seen as a result (the option history).

Scenario 1: No Communication

In the first scenario, we consider the case where robots could not communicate with each other. Therefore, all cooperation is based on the controllers that are generated by the planner (which knows the controllers generated for all robots when planning offline) and observations of the other robots (when executing online). The macro-actions were as follows:

- Go to depot 1.
- Go to depot 2.

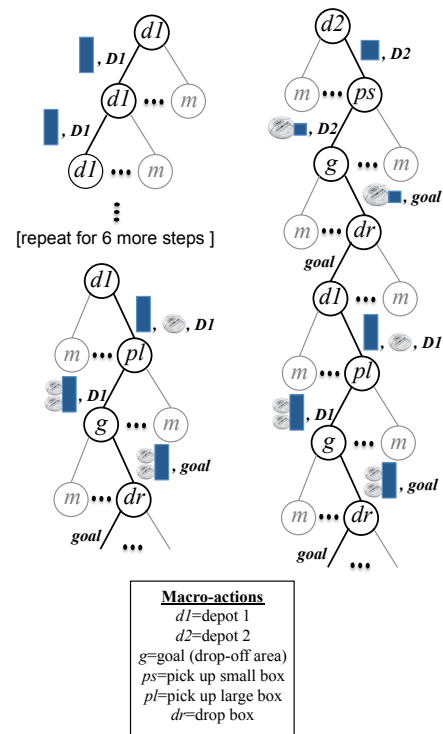


Figure 7: Path executed in policy trees. Only macro-actions executed (nodes) and observations seen (edges, with boxes and robots given pictorially) are shown.

- Go to the drop-off area.
- Pick up the small box.
- Pick up the large box.
- Drop off a box.

The depot macro-actions are applicable anywhere and terminate when the robot is within the walls of the appropriate depot. The drop-off and drop macro-actions are only applicable if the robot is holding a box, and the pickup macro-actions are only applicable when the robot observes a box of the particular type. Navigation is stochastic in the amount of time that will be required to succeed (as mentioned above). Picking up the small box was assumed to succeed deterministically, but this easily be changed if the pickup mechanism is less robust. These macro-actions correspond to natural choices for robot controllers.

This case¹ (seen in Figure 6 along with a depiction of the executed policy in Figure 7) uses only two robots to more clearly show the optimized behavior in the absence of communication. The policy generated by the planner assigns one robot to go to each of the depots (Figure 6(a)). The robots then observe the contents of the depots they are in (Figure 6(b)). If there are two robots in the same room as a large box, they will push it back to the goal. If there is only one robot in a depot and there is a small box to push, the robot

¹Videos for all scenarios can be seen at http://youtu.be/istb8TIp_jw

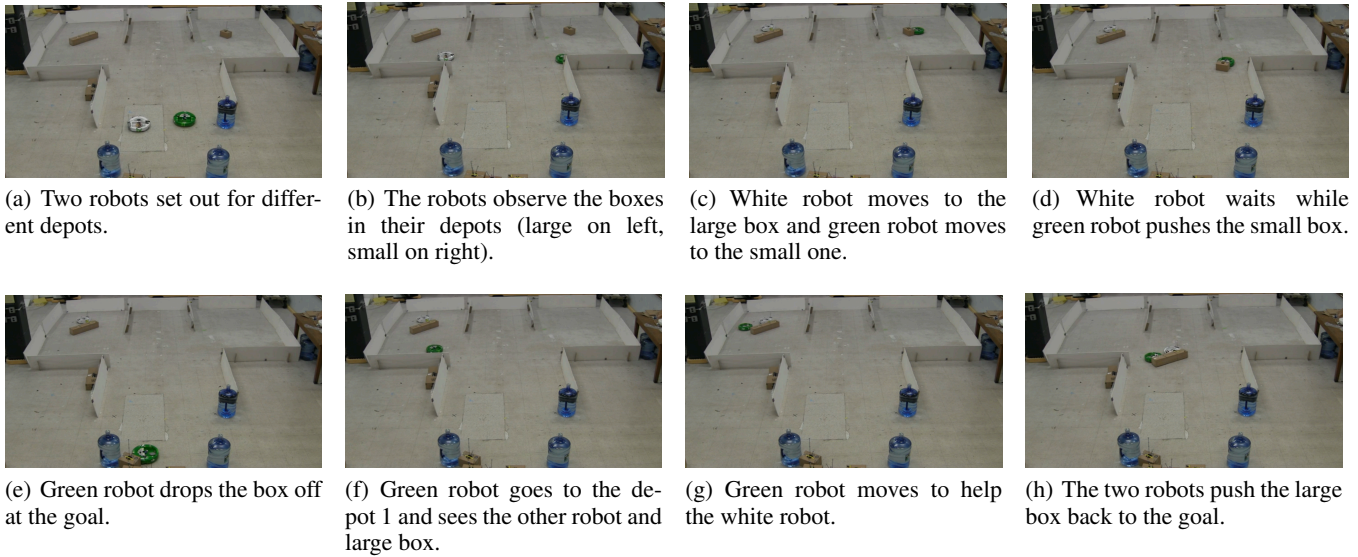


Figure 6: Video captures from the no communication version of the warehouse problem.

will push the small box (Figure 6(c)). If the robot is in a depot with a large box and no other robots, it will stay in the depot, waiting for another robot to come and help push the box (Figure 6(d)). In this case, once the other robot is finished pushing the small box (Figure 6(e)), it goes back to the depots to check for other boxes or robots that need help (Figure 6(f)). When it sees another robot and the large box in the depot on the left (depot 1), it attempts to help push the large box (Figure 6(g)) and the two robots are successful pushing the large box to the goal (Figure 6(h)). In this case, the planner has generated a policy in a similar fashion to task allocation—two robots go to each room, and then search for help needed after pushing any available boxes. However, in our case this behavior was generated by an optimization process that considered the different costs of actions, the uncertainty involved and the results of those actions into the future.

Scenario 2: Local Communication

In scenario 2, robots can communicate when they are within one meter of each other. The macro-actions are the same as above, but we added ones to communicate and wait for communication. The resulting macro-action set is:

- Go to depot 1.
- Go to depot 2.
- Go to the drop-off area.
- Pick up the small box.
- Pick up the large box.
- Drop off a box.
- Go to an area between the depots (the “waiting room”).
- Wait in the waiting room for another robot.
- Send signal #1.

- Send signal #2.

Here, we allow the robots to choose to go to a “waiting room” which is between the two depots. This permits the robots to possibly communicate or receive communications before committing to one of the depots. The waiting-room macro-action is applicable in any situation and terminates when the robot is between the waiting room walls. The depot macro-actions are now only applicable in the waiting room, while the drop-off, pick up and drop macro-actions remain the same. The wait macro-action is applicable in the waiting room and terminates when the robot observes another robot in the waiting room. The signaling macro-actions are applicable in the waiting room and are observable by other robots that are within approximately a meter of the signaling robot. Note that *we do not specify what sending each communication signal means*.

The results for this domain are shown in Figure 8. We see that the robots go to the waiting room (Figure 8(a)) (because we required the robots to be in the waiting room before choosing to move to a depot) and then two of the robots go to depot 2 (the one on the right) and one robot goes to depot 1 (the one on the left) (Figure 8(b)). Note that because there are three robots, the choice for the third robot is random while one robot will always be assigned to each of the depots. Because there is only a large box to push in depot 1, the robot in this depot goes back to the waiting room to try to find another robot to help it push the box (Figure 8(c)). The robots in depot 2 see two small boxes and they choose to push these back to the goal (Figure 8(d)). Once the small boxes are dropped off (Figure 8(e)), one of the robots returns to the waiting room (Figure 8(f)) and then is recruited by the other robot to push the large box back to the goal (Figure 8(g)). The robots then successfully push the large box back to the goal (Figure 8(h)). Note that in this case the planning process *determines how the signals should be used to*

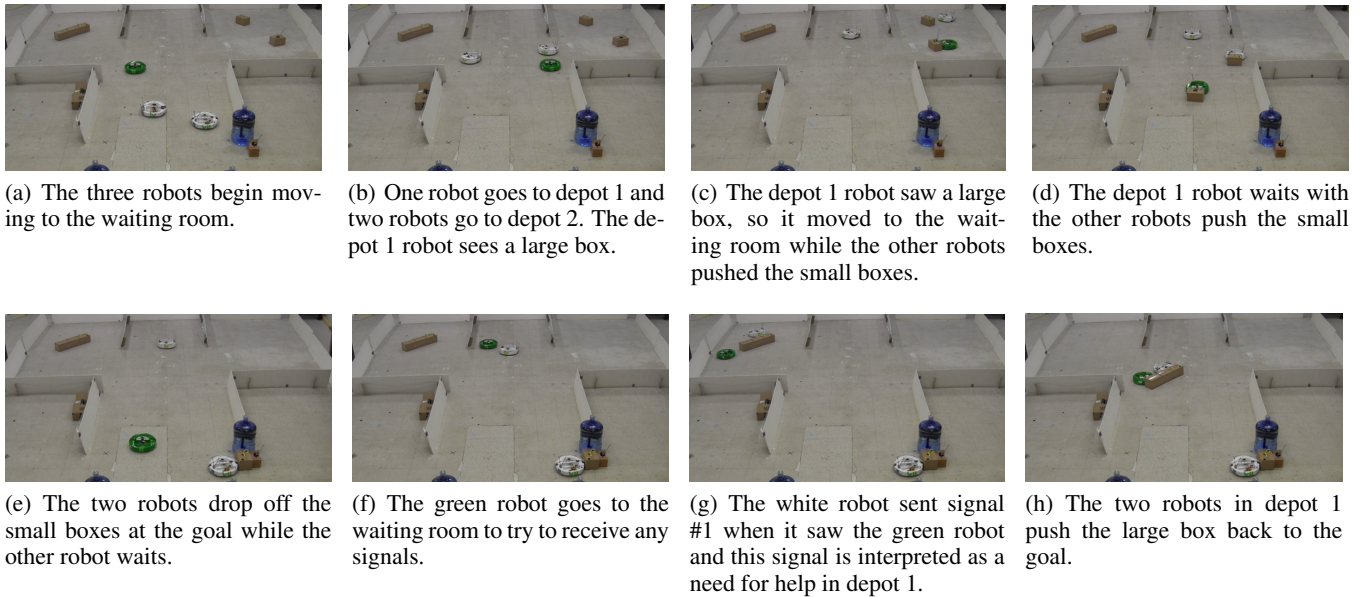


Figure 8: Video captures from the limited communication version of the warehouse problem.

perform communication.

Scenario 3: Global Communication

In the last scenario, the robots can use signaling (rather than direct communication). In this case, there is a switch in each of the depots that can turn on a blue or red light. This light can be seen in the waiting room and there is another light switch in the waiting room that can turn off the light. (The light and switch were simulated in software and not incorporated in the physical domain.) As a result, the macro-actions in this scenario were as follows:

- Go to depot 1.
- Go to depot 2.
- Go to the drop-off area.
- Pick up the small box.
- Pick up the large box.
- Drop off a box.
- Go to an area between the depots (the “waiting room”).
- Turn on a blue light.
- Turn on a red light.
- Turn off the light.

The first seven macro-actions are the same as for the communication case except we relaxed the assumption that the robots had to go to the waiting room before going to the depots (making both the depot and waiting room macro-actions applicable anywhere). The macro-actions for turning the lights on are applicable in the depots and the macro-actions for turning the lights off are applicable in the waiting room. While the lights were intended to signal requests for help in each of the depots, we did not assign a particular


color to a particular depot. In fact, we did not assign them any specific meaning, allowing the planner to set them in any way that improves performance.

The results are shown in Figure 9. Because one robot started ahead of the others, it was able to go to depot 1 to sense the size of the boxes while the other robots go to the waiting room (Figure 9(a)). The robot in depot 1 turned on the light (red in this case, but not shown in the images) to signify that there is a large box and assistance is needed (Figure 9(b)). The green robot (the first other robot to the waiting room) sees this light, interprets it as a need for help in depot 1, and turns off the light (Figure 9(c)). The other robot arrives in the waiting room, does not observe a light on and moves to depot 2 (also Figure 9(c)). The robot in depot 2 chooses to push a small box back to the goal and the green robot moves to depot 1 to help the other robot (Figure 9(d)). One robot then pushes the small box back to the goal while the two robots in depot 1 begin pushing the large box (Figure 9(e)). Finally, the two robots in depot 1 push the large box back to the goal (Figure 9(f)). This behavior is optimized based on the information given to the planner. *The semantics of all these signals as well as the movement and signaling decisions were decided on by the planning algorithm to maximize value.*


Related Work

There are several frameworks that have been developed for multi-robot decision making in complex domains. For instance, behavioral methods have been studied for performing task allocation over time in loosely-coupled [30] or tightly-coupled [35] tasks. These are heuristic in nature and make strong assumptions about the type of tasks that will be completed.


One important related class of methods is based on using



(a) One robot starts first and goes to depot 1 while the other robots go to the waiting room. (b) The robot in depot 1 sees a large box, so it turns on the red light (the light is not shown).



(c) The green robot sees the light first, so it turns it off and goes to depot 1 while the white robot goes to depot 2. (d) The robots in depot 1 move to the large box, while the robot in depot 2 begins pushing the small box.



(e) The robots in depot 1 begin pushing the large box and the robot in depot 2 pushes a small box to the goal. (f) The robots from depot 1 successfully push the large box to the goal.

Figure 9: Video captures from the signaling version of the warehouse problem.

linear temporal logic (LTL) [7, 21] to specify behavior for a robot; from this specification, reactive controllers that are guaranteed to satisfy the specification can be derived. These methods are appropriate when the world dynamics can be effectively described non-probabilistically and when there is a useful discrete characterization of the robot’s desired behavior in terms of a set of discrete constraints. When applied to multiple robots, it is necessary to give each robot its own behavior specification. Other logic-based representations for multi-robot systems have similar drawbacks and typically assume centralized planning and control [22].

Market-based approaches use traded value to establish an optimization framework for task allocation [11, 15]. These approaches have been used to solve real multi-robot problems [18], but are largely aimed to tightly-coupled tasks, where the robots can communicate through a bidding mechanism.

Emery-Montemerlo et al. [14] introduced a (cooperative) game-theoretic formalization of multi-robot systems which resulted in solving a Dec-POMDP. An approximate forward search algorithm was used to generate solutions, but scalability was limited because a (relatively) low-level Dec-POMDP was used. Also, Messias et al. [25] introduce

an MDP-based model where a set of robots with controllers that can execute for varying amount of time must cooperate to solve a problem. However, decision-making in their system is centralized.

Conclusion

We have demonstrated—for the first time—that complex multi-robot domains can be solved with Dec-POMDP-based methods. The MacDec-POMDP model is expressive enough to capture multi-robot systems of interest, but also simple enough to be feasible to solve in practice. Our results show that a general purpose MacDec-POMDP planner can generate cooperative behavior for complex multi-robot domains with task allocation, direct communication, and signaling behavior emerging automatically as properties of the solution for the given problem model. Because all cooperative multi-robot problems can be modeled as Dec-POMDPs, MacDec-POMDPs represent a powerful tool for automatically trading-off various costs, such as time, resource usage and communication while considering uncertainty in the dynamics, sensors and other robot information. These approaches have great potential to lead to automated solution methods for general multi-robot coordination problems with large numbers of heterogeneous robots in complex, uncertain domains.

In the future, we plan to explore incorporating these state-of-the-art macro-actions into our MacDec-POMDP framework as well as examine other types of structure that can be exploited. Other topics we plan to explore include increasing scalability by making solution complexity depend on the number of agent interactions rather than the domain size, and having robots learn models of their sensors, dynamics and other robots. These approaches have great potential to lead to automated solution methods for general multi-robot coordination problems with large numbers of heterogeneous robots in complex, uncertain domains.

References

- [1] Amato, C.; Chowdhary, G.; Geramifard, A.; Ure, N. K.; and Kochenderfer, M. J. 2013. Decentralized control of partially observable Markov decision processes. In *Proceedings of the Fifty-Second IEEE Conference on Decision and Control*, 2398–2405.
- [2] Amato, C.; Dibangoye, J. S.; and Zilberstein, S. 2009. Incremental policy generation for finite-horizon DEC-POMDPs. In *Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling*, 2–9.
- [3] Amato, C.; Konidaris, G.; and Kaelbling, L. P. 2014. Planning with macro-actions in decentralized POMDPs. In *Proceedings of the Thirteenth International Conference on Autonomous Agents and Multiagent Systems*.
- [4] Aras, R.; Dutech, A.; and Charpillet, F. 2007. Mixed integer linear programming for exact finite-horizon planning in decentralized POMDPs. In *Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling*, 18–25.

- [5] Balch, T., and Arkin, R. C. 1998. Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation* 14(6):926–939.
- [6] Becker, R.; Holland, O.; and Deneubourg, J.-L. 1994. From local actions to global tasks: Stigmergy and collective robotics. In *Artificial life IV*, volume 181, 189.
- [7] Belta, C.; Bicchi, A.; Egerstedt, M.; Frazzoli, E.; Klavins, E.; and Pappas, G. J. 2007. Symbolic planning and control of robot motion [grand challenges of robotics]. *Robotics & Automation Magazine, IEEE* 14(1):61–70.
- [8] Bernstein, D. S.; Givan, R.; Immerman, N.; and Zilberstein, S. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research* 27(4):819–840.
- [9] Bohren, J. 2010. SMACH. <http://wiki.ros.org/smach/>.
- [10] Boularias, A., and Chaib-draa, B. 2008. Exact dynamic programming for decentralized POMDPs with lossless policy compression. In *Proceedings of the Eighteenth International Conference on Automated Planning and Scheduling*.
- [11] Dias, M. B., and Stentz, A. T. 2003. A comparative study between centralized, market-based, and behavioral multirobot coordination approaches. In *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '03)*, volume 3, 2279 – 2284.
- [12] Dibangoye, J. S.; Amato, C.; Buffet, O.; and Charpillet, F. 2013a. Optimally solving Dec-POMDPs as continuous-state MDPs. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- [13] Dibangoye, J. S.; Amato, C.; Doniec, A.; and Charpillet, F. 2013b. Producing efficient error-bounded solutions for transition independent decentralized MDPs. In *Proceedings of the Twelfth International Conference on Autonomous Agents and Multiagent Systems*.
- [14] Emery-Montemerlo, R.; Gordon, G.; Schneider, J.; and Thrun, S. 2005. Game theoretic control for robot teams. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, 1163–1169.
- [15] Gerkey, B., and Matarić, M. 2004. A formal analysis and taxonomy of task allocation in multi-robot systems. *International Journal of Robotics Research* 23(9):939–954.
- [16] Hansen, E. A.; Bernstein, D. S.; and Zilberstein, S. 2004. Dynamic programming for partially observable stochastic games. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, 709–715.
- [17] Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101:1–45.
- [18] Kalra, N.; Ferguson, D.; and Stentz, A. T. 2005. Hoplites: A market-based framework for planned tight coordination in multirobot teams. In *Proceedings of the International Conference on Robotics and Automation*, 1170 – 1177.
- [19] Kober, J.; Bagnell, J. A.; and Peters, J. 2013. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research* 32(11):1238 – 1274.
- [20] Kube, C., and Bonabeau, E. 2000. Cooperative transport by ants and robots. *Robotics and Autonomous Systems* 30(1-2):85–101.
- [21] Loizou, S. G., and Kyriakopoulos, K. J. 2004. Automatic synthesis of multi-agent motion tasks based on ltl specifications. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 1, 153–158. IEEE.
- [22] Lundh, R.; Karlsson, L.; and Saffiotti, A. 2008. Autonomous functional configuration of a network robot system. *Robotics and Autonomous Systems* 56(10):819–830.
- [23] Mahajan, A. 2013. Optimal decentralized control of coupled subsystems with control sharing. *IEEE Transactions on Automatic Control* 58:2377–2382.
- [24] Melo, F. S., and Veloso, M. 2011. Decentralized MDPs with sparse interactions. *Artificial Intelligence*.
- [25] Messias, J. V.; Spaan, M. T.; and Lima, P. U. 2013. GSMDPs for multi-robot sequential decision-making. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*.
- [26] Nair, R.; Varakantham, P.; Tambe, M.; and Yokoo, M. 2005. Networked distributed POMDPs: a synthesis of distributed constraint optimization and POMDPs. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*.
- [27] Oliehoek, F. A.; Spaan, M. T. J.; Amato, C.; and Whiteson, S. 2013. Incremental clustering and expansion for faster optimal planning in Dec-POMDPs. *Journal of Artificial Intelligence Research* 46:449–509.
- [28] Oliehoek, F. A.; Whiteson, S.; and Spaan, M. T. J. 2013. Approximate solutions for factored Dec-POMDPs with many agents. In *Proceedings of the Twelfth International Conference on Autonomous Agents and Multiagent Systems*.
- [29] Oliehoek, F. A. 2012. Decentralized POMDPs. In Wiering, M., and van Otterlo, M., eds., *Reinforcement Learning: State of the Art*, volume 12 of *Adaptation, Learning, and Optimization*. Springer. 471–503.
- [30] Parker, L. E. 1998. ALLIANCE: An architecture for fault tolerant multirobot cooperation. *IEEE Transactions on Robotics and Automation* 14(2):220–240.
- [31] Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience.
- [32] Quigley, M.; Conley, K.; Gerkey, B.; Faust, J.; Foote, T.; Leibs, J.; Wheeler, R.; and Ng, A. Y. 2009. ROS: an

- open-source robot operating system. In *ICRA workshop on open source software*, volume 3.
- [33] Rekleitis, I.; Lee-Shue, V.; New, A. P.; and Choset, H. 2004. Limited communication, multi-robot team based coverage. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 4, 3462–3468. IEEE.
- [34] Seuken, S., and Zilberstein, S. 2007. Memory-bounded dynamic programming for DEC-POMDPs. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, 2009–2015.
- [35] Stroupe, A. W.; Ravichandran, R.; and Balch, T. 2004. Value-based action selection for exploration and dynamic target observation with robot teams. In *Proceedings of the International Conference on Robotics and Automation*, volume 4, 4190–4197. IEEE.
- [36] Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence* 112(1):181–211.
- [37] Thrun, S.; Burgard, W.; and Fox, D. 2005. *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press.
- [38] Velagapudi, P.; Varakantham, P. R.; Sycara, K.; and Scerri, P. 2011. Distributed model shaping for scaling to decentralized POMDPs with hundreds of agents. In *Proceedings of the Tenth International Conference on Autonomous Agents and Multiagent Systems*.