

Exploiting Sparse Interactions for Optimizing Communication in Dec-MDPs

Francisco S. Melo
INESC-ID
Instituto Superior Técnico
Porto Salvo, Portugal
fmelo@inesc-id.pt

Matthijs T.J. Spaan
Delft University of Technology
Delft, The Netherlands
m.t.j.spaan@tudelft.nl

Stefan Witwicki
INESC-ID
Instituto Superior Técnico
Porto Salvo, Portugal
witwicki@inesc-id.pt

ABSTRACT

Decentralized partially observable Markov decision processes (Dec-POMDPs) provide powerful modeling tools for multiagent decision-making in the face of uncertainty, but solving these models comes at a very high computational cost. Two avenues for side-stepping the computational burden can be identified: structured interactions between agents and intra-agent communication. In this paper, we focus on the interplay between these concepts, namely how sparse interactions reflect in the communication needs. A key insight is that in domains with local interactions the amount of communication necessary for successful joint behavior can be heavily reduced, due to the limited influence between agents. We exploit this insight by deriving local POMDP models that optimize each agent’s communication behavior. Our experimental results show that our approach successfully exploits sparse interactions: we can effectively identify the situations in which it is beneficial to communicate, as well as trade off the cost of communication with overall task performance.

1. INTRODUCTION

Decentralized partially observable Markov decision processes (Dec-POMDPs) provide powerful modeling tools for multiagent decision-making with limited sensing capabilities in stochastic environments. However, the prohibitive computational cost required to compute an optimal decision rule renders them intractable except for the smallest of problems.¹ In the literature, two avenues for side-stepping the computational burden can be identified: localized interactions between agents [2, 16, 21, 23, 24] and intra-agent communication [7, 18, 20, 26]. In this paper, we focus on the *interplay* between these concepts, namely how sparse interactions reflect in the communication needs.

A key insight is that in domains with local interactions the amount of communication necessary for successful joint behavior can be heavily reduced, due to the limited influence between agents. Several previous works have implicitly relied on this observation, exploring sparse interactions by having agents share information locally [11, 13, 18, 21, 23]. In this work, we explicitly reason about the benefits of communication/information sharing in scenarios with sparse in-

¹Dec-MDPs are known to be NEXP-complete even in 2-agent scenarios [4].

teractions. Sparse interactions enable, to some extent, decoupling the decision-process of the different agents. We leverage such decoupling to derive local models that optimize each agent’s communication behavior, allowing it to overcome partial observability in those situations where decoupled decisions are not possible.

We provide a new way of optimizing communication by proposing a model in which agents need to plan about when to query other agents’ local states, which we call QUERY-POMDP. We observe that to execute optimal joint policies in fully observable scenarios—policies which can be computed efficiently—agents will generally need to reason about the state of other agents. Our agents construct a local POMDP model of the environment from the fully observable joint policy of all other agents. Solving this POMDP model allows the agent not only to determine how to solve the task at hand but also to determine when to query the local state of the environment. Our approach thus allows the agents to explicitly reason about communication, without incurring in the prohibitive computational cost of Dec-POMDP models that include communication [17]. Furthermore, in contrast to many methods in the literature [18, 26], QUERY-POMDP can properly handle noisy communication channels, and does not require strong independence assumptions [1]. Our empirical analysis on benchmark problems demonstrates the efficacy of QUERY-POMDP in balancing communication costs with coordination benefits.

2. BACKGROUND

We start by reviewing *decentralized partially observable Markov decision processes* (Dec-POMDPs) and related decision theoretic models. An N -agent Dec-POMDP \mathcal{M} is specified as a tuple $\mathcal{M} = (N, \mathcal{X}, (\mathcal{A}_k), (\mathcal{Z}_k), P, (O_k), r, \gamma)$, where:

- \mathcal{X} is the joint state-space;
- $\mathcal{A} = \times_{i=1}^N \mathcal{A}_i$ is the set of joint actions, with each \mathcal{A}_i the individual action set for agent $i, i = 1, \dots, N$;
- Each \mathcal{Z}_i represents the set of possible local observation for agent $i, i = 1, \dots, N$;
- $P(y | x, a)$ represents the transition probabilities from joint state x to joint state y when the joint action a is taken;
- Each $O_i(z_i | x, a)$ represents the probability of agent i making the local observation z_i when the joint state is x and the last joint action taken was a ;

- $r(x, a)$ represents the expected reward received by all agents for taking the joint action a in joint state x ;
- The scalar γ is a discount factor.

An N -agent *Decentralized Markov decision process* (Dec-MDP) is a particular class of Dec-POMDP in which the state is *jointly fully observable*, however, optimal solving of the model is in the same complexity class as the Dec-POMDP model. Formally this can be translated into the following condition: for every joint observation $z \in \mathcal{Z}$, with $\mathcal{Z} = \times_{i=1}^N \mathcal{Z}_i$, there is a state $x \in \mathcal{X}$ such that

$$\mathbb{P}[X(t) = x \mid Z(t) = z] = 1,$$

where $X(t)$ is the joint state of the process at time t and $Z(t)$ the corresponding joint observation. A *partially observable Markov decision process* (POMDP) is a 1-agent Dec-POMDP and a *Markov decision process* (MDP) is a 1-agent Dec-MDP. Finally, an N -agent *multiagent MDP* (MMDP) is an N -agent Dec-MDP that is *fully observable*, *i.e.*, for every individual observation $z_i \in \mathcal{Z}_i$ there is a state $x \in \mathcal{X}$ such that

$$\mathbb{P}[X(t) = x \mid Z_i(t) = z_i] = 1.$$

In this partially observable multiagent setting, an individual (non-Markov) policy for agent i is a mapping

$$\pi_i : \mathcal{H}_i \rightarrow \Delta(\mathcal{A}_i),$$

where $\Delta(\mathcal{A}_i)$ is the space of probability distributions over \mathcal{A}_i , and \mathcal{H}_i is the set of all possible finite histories for agent i . The purpose of all agents is to determine a joint policy π that maximizes the total sum of discounted rewards. In other words, considering a distinguished initial state $x^0 \in \mathcal{X}$ that is assumed common knowledge among all agents, the goal of the agents is to maximize

$$V^\pi = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r(X(t), A(t)) \mid X(0) = x^0 \right]. \quad (1)$$

For a more detailed introduction to Dec-POMDPs and related models see, e.g., [19].

3. A MOTIVATING EXAMPLE

We motivate our ideas in a simple navigation scenario, corresponding to the environment depicted in Fig. 1. In this scenario, two robots (Robot 1 and Robot 2) must navigate to their corresponding goal states (marked as Goals 1 and 2). At the same time, they must avoid colliding in the narrow doorway (the central state), since it leads to a large penalty. Each robot has 4 possible actions (namely “Move North”, “Move South”, “Move East” and “Move West”) that move the robot in the corresponding direction. The motion of one robot does not depend on the position or action of the other robot except in the doorway: if the robots collide in the doorway, then their actions have an increasing failure probability. Complicating matters, initially each robot starts uniformly at random in one of the 10 locations on its side of the doorway.

In a fully observable situation, the agents will move toward their respective goals. When reaching the doorway, if the other robot is also close to the doorway one of the two will




Figure 1: H-Environment, where two robots need to interact only around the narrow doorway to reach their corresponding goals. The shaded arrows correspond to a possible policy for Robot 2 in the absence of Robot 1.

stop so that the other can safely traverse.² It will then resume its trajectory to its goal.

In order for the agents to actually execute the policy just described, they only need to reason about the state of the other agent when reaching the darker area in their starting side of the environment. And then, once one robot is in the doorway, it can just proceed toward its goal, independently of the state of the other robot. Moreover, even if the robots are generally unable to observe the position of the other robot, but they are able to *query it*, they can reasonably assume that the other robot will behave more or less as in the fully observable scenario. This observation is the departing point for the model and approach proposed in this paper and described in the continuation.

4. A MODEL FOR STATE QUERYING

We depart from an N -agent Dec-MDP model, and address the problem of when communication can be beneficial to improve the performance in such a model. For the purposes of our study, we momentarily focus on the decision processes of all except one agent, which we refer to as agent k . Unlike other communication-based approaches to Dec-POMDPs (e.g., [18, 27]), we adopt a relatively general communication model, in which the messages received by an agent are taken as part of its local (noisy) observation. Also, messages received by agent k depend on explicit information-querying actions executed by k .

Throughout this section, we represent the (finite) state-space of the Dec-MDP as a set \mathcal{X} and assume that it can be factorized as $\mathcal{X} = \mathcal{X}_k \times \mathcal{X}_{-k}$, where the elements $x_k \in \mathcal{X}_k$ correspond to agent k 's local state. The state at time t , $X(t)$, is thus a pair $\langle X_k(t), X_{-k}(t) \rangle$. We also assume that the observations of each agent do not depend on the actions of the remaining agents, *i.e.*,

$$\mathbb{P}[Z_i(t) = z_i \mid X(t), A(t)] = \mathbb{P}[Z_i(t) = z_i \mid X(t), A_i(t)],$$

for all $i = 1, \dots, N$. Therefore, we can simply write the observation probabilities as $\mathbf{O}_i(z_i \mid x, a_i), i = 1, \dots, N$.

²Which one stops is determined by the joint policy they adopt.

4.1 Query Actions and Observations

For the purpose of allowing our agent to reason about communication, we assume that each agent has the ability to *query* the other agents for their local state information. In order to make this explicit, we differentiate between *communication actions* and the remaining actions—henceforth referred as *primitive actions*, and write the set of individual actions for agent k as the cartesian product of the set of communication actions, \mathcal{A}_k^C , and the set of primitive actions, \mathcal{A}_k^P , *i.e.*, $\mathcal{A}_k = \mathcal{A}_k^C \times \mathcal{A}_k^P$. We also assume that transition probabilities are independent of the communication actions,

$$\mathbb{P}(y | x, \langle a_{-k}, (a_k^C, a_k^P) \rangle) = \mathbb{P}(y | x, \langle a_{-k}, (b_k^C, a_k^P) \rangle)$$

for any $x, y \in \mathcal{X}$, $a_{-k} \in \mathcal{A}_{-k}$, $a_k^P \in \mathcal{A}_k^P$ and $a_k^C, b_k^C \in \mathcal{A}_k^C$.

We also differentiate between *communication observations* (*i.e.*, observations that result from communication actions) and *primitive observations*, that do not depend on the communication actions. Formally, we write the set of individual observations for agent k as the cartesian product of the set of communication observations, \mathcal{Z}_k^C , and primitive observations, \mathcal{Z}_k^P , *i.e.*, $\mathcal{Z}_k = \mathcal{Z}_k^C \times \mathcal{Z}_k^P$. We consider that communication observations do not depend on primitive actions, and that primitive observations do not depend on communication actions. This means that we can decouple the observation probabilities as

$$\mathbb{O}_k((z_k^C, z_k^P) | x, (a_k^C, a_k^P)) = \mathbb{O}_k^C(z_k^C | x, a_k^C) \mathbb{O}_k^P(z_k^P | x, a_k^P),$$

where

$$\begin{aligned} \mathbb{O}_k^C(z_k^C | x, a_k^C) &= \mathbb{P}[Z_k^C(t) = z_k^C | X(t) = x, A_k^C(t) = a_k^C] \\ \mathbb{O}_k^P(z_k^P | x, a_k^P) &= \mathbb{P}[Z_k^P(t) = z_k^P | X(t) = x, A_k^P(t) = a_k^P]. \end{aligned}$$

Finally, we assume that the reward function can also be decomposed as the sum of two components. The first component, denoted r^C , concerns the *cost of communication* and is independent on the primitive actions of agent k and on the actions of the other agents. The second component, denoted as r^P corresponds to the “regular” (or domain-level) reward defining the overall goal of the agents. It is assumed independent of the communication actions of agent k . Formally, if $a = \langle a_{-k}, a_k \rangle$ and $a_k = (a_k^C, a_k^P)$, this means that the reward r can be written as

$$r(x, a) = r^P(x, \langle a_{-k}, a_k^P \rangle) + r^C(x, a_k^C). \quad (2)$$

Figure 2 depicts a dynamic Bayesian network that summarizes all above considerations.

Following the discussion in Section 3, and for the purpose of its planning process, agent k will treat all remaining agents as if they follow a Markov policy, π_{-k} , that corresponds to the optimal policy for the underlying MMDP. This policy, being Markovian, depends only on the state of the system at time t , $X(t)$, *i.e.*,

$$\begin{aligned} \mathbb{P}[A_{-k}(t) = a_{-k} | H(t)] \\ = \mathbb{P}[A_{-k}(t) = a_{-k} | X(t) = x] = \pi_{-k}(x, a_{-k}), \end{aligned} \quad (3)$$

where $A_{-k}(t)$ denotes the action taken by all agents other than k at time t , $H(t)$ denotes the whole history of the process up to time t and $a_{-k} \in \mathcal{A}_{-k}$. From this perspective, the decision process for agent k can be modeled as a (single-agent) POMDP that we describe in the next section.




Figure 2: The factored decision model, from agent k 's perspective.

4.2 POMDP Model for a Single Agent

Let $\mathcal{M} = (N, \mathcal{X}, (\mathcal{A}_k), (\mathcal{Z}_k), \mathbb{P}, (\mathbb{O}_k), r, \gamma)$ be a Dec-MDP that meets the above assumptions. Let π_{-k} denote the (state-dependent) joint MMDP policy for all agents other than k . The single-agent POMDP model for agent k is a tuple $\mathcal{M}_k = (\mathcal{X}, \mathcal{A}_k, \mathcal{Z}_k, \mathbb{P}_k, \mathbb{O}_k, r_k, \gamma)$, where:

- \mathcal{X} corresponds to the original Dec-MDP state-space.
- \mathcal{A}_k is the individual action-space for agent k .
- \mathcal{Z}_k is the individual observation-space for agent k .
- \mathbb{P}_k are the transition probabilities obtained from the original transition probabilities. In particular, given an action $a_k = (a_k^C, a_k^P)$, we have

$$\mathbb{P}_k(y | x, a_k) = \sum_{a_{-k} \in \mathcal{A}_{-k}} \pi_{-k}(x, a_{-k}) \mathbb{P}(y | x, \langle a_{-k}, a_k^P \rangle).$$

- \mathbb{O}_k are the observation probabilities for agent k , that match the original Dec-MDP observation probabilities. In particular, given an action $a_k = (a_k^C, a_k^P)$, we have

$$\mathbb{O}_k(z_k | x, a_k) = \mathbb{O}_k^C(z_k^C | x, a_k^C) \mathbb{O}_k^P(z_k^P | x, a_k^P), \quad (4)$$

where $z_k = (z_k^C, z_k^P)$.

- r_k is the reward function obtained from the original Dec-MDP reward function after averaging over the other agents' policy, π_{-k} , *i.e.*,

$$r_k(x, a_k) = \sum_{a_{-k} \in \mathcal{A}_{-k}} \pi_{-k}(x, a_{-k}) r(x, \langle a_{-k}, a_k \rangle).$$

Given this POMDP model, we can use standard POMDP solution techniques to explore the trade-off between the costs and benefits of communication for agent k .

4.3 Results for the H-environment Example

Continuing the example of Section 3, the application of our model allows us to better understand under which circumstances the benefits of using communication compensate for its costs. For this purpose, we fix the policy of Agent 2 as shown in Fig. 1, which corresponds one possible joint MMDP policy for this environment. As explained above, given such a policy we can construct a POMDP from the point of view of Agent 1, in which it can query Agent 2's states at any time step, at a particular communication cost. For illustration purposes, the initial state of Agent 2 is selected randomly on the right half of the environment. We




Figure 3: Results for the H-environment. (a)-(d) Query frequency in each state for Agent 1, varying in deterministic (Det. P) or noisy (Noisy P) transitions and communication cost.

test several experimental conditions that include the presence or absence of transition noise and different costs for the communication actions.

We examine in which states Agent 1 queries Agent 2’s state. When communication is free (Figs. 3(a) and (b)), Agent 1 queries in all the states it passes through.³ With a communication cost of 0.3 (Fig. 3(d)), however, it only queries when near to and left of the doorway. In these states it is crucial to know Agent 2’s location to avoid potential collisions, an intuition that is exploited automatically by our model.

The use of a POMDP model in this context ensures that the agent explicitly reasons about information gathering actions which, in our setting, translates in weighting the benefits of communication in terms of the overall task against the costs associated with it. It is also worth noting that the optimal communication policy depends on the communication costs, on the task at hand, the transition model, and the query model (*i.e.*, the observation probabilities associated with the query actions). The planning framework proposed takes into consideration all these aspects, and the right trade-off between communication costs and benefits comes out as a natural outcome of the policy computed by the agent.

4.4 Computing Policies for Multiple Agents

In the previous section we proposed using a POMDP model to compute the policy for one agent k , treating all other agents as if they were following the optimal joint policy

³We note that, due to the transition noise, an agent can remain in the same state more than one consecutive time-step, and hence the values larger than 1.

for the underlying MMDP. Given this POMDP model for agent k we can compute the corresponding optimal policy using any preferred POMDP solution technique. We use this approach to better understand the communication needs of one agent in a simple multiagent navigation scenario, and to determine in which situations the cost of communication outweighs its value.

We now want to extend these ideas and actually compute the policy for *all* agents in the Dec-MDP. The idea of using POMDP models to plan in multiagent scenarios has been previously explored in the Dec-POMDP literature [6, 14]. The general difficulty with these approaches arises from the fact that each agent has only a local observation of the joint state of the world. This implies that, when planning for agent k , the POMDP model necessary to properly capture the behavior of all agents other than k can either be prohibitively large, require agent k to reason about how the other agents reason about agent k ’s state, leading to infinitely nested beliefs, or both [6, 14].

In our approach, we rely on the intuition discussed in Section 3, according to which the use of *active communication* allied with *sparse interactions* may actually alleviate the difficulties associated with planning in multiagent systems with partial observability. We plan for each agent k while treating all other agents as if following the *optimal joint MMDP policy*. In scenarios where interactions are sparse, the general behavior of the agents is expected to roughly follow the MMDP policy, as discussed in Section 3 and in those situations where coordination is necessary, agents can resort to communication, but weighting the benefits of such communication with the associated costs.

Several previous works have already studied the benefits of exploiting communication and structured interactions separately (see, for example, [7, 21, 24]).⁴ The novelty in our approach lies precisely on the fact that we can explicitly exploit the interplay between these two aspects (communication and sparse interactions) to attain efficient planning in multiagent problems. Section 5 describes the application of our approach in several navigation scenarios of different dimensions. Our results empirically show that our approach is indeed able to make effective use of communication and attain a performance that indeed approaches that observed in fully observable settings.

5. EXPERIMENTS

In this section we illustrate the application of our method to several navigation scenarios from the POMDP and Dec-POMDP literature. We use robot navigation scenarios (see Fig. 4), since our model is particularly suited for modeling multi-robot problems. Furthermore, results can be easily visualized and interpreted in this class of problems. We consider only two-agent scenarios for ease of interpretation although, as argued, our approach can be used with any number of agents. Moreover, for our purposes (*i.e.*, investigating how the proposed use of communication can alleviate the computational burden of solving Dec-MDPs), two-agent scenarios are already illustrative, as they are NEXP-complete [4].

⁴We refer to Section 6 for a detailed discussion of related approaches.




Figure 4: Environments used in the experiments.

Experimental Setup

In each of the test scenarios, two robots must each reach one specific state, marked either by a boxed number or a cross, \times . Each robot has 4 actions that move the robot in one of the four possible direction with probability 0.8 and fail with probability 0.2, plus a fifth “NoOp” action. As for the communication model, the agent’s query actions will provide the local state of the other agent with probability 0.8 in the shaded areas. With a probability 0.2, communication will fail and the query action will provide no observation. Outside the shaded areas, the query actions always yield no communication. Since the communication model is available to the agents, planning takes into consideration communication limitations.⁵

The darker cells correspond to states where the agents receive a penalty of -20 when standing there simultaneously, in which case the rate of action failure is also increased to 0.4 for both agents.

All agents have full local state observability. When an agent queries another agent, it incurs a cost of -0.1 and, in the shaded areas, successfully observes the local state of the queried agent with a probability of 0.8. With a probability of 0.2 it receives no observation about the state of the other. In the white cells, an agent is never able to perceive the state of the other, but still incurs a penalty of -0.1 if it attempts to communicate. When an agent reaches its goal position, it receives a reward of 10 and moves to a rewardless absorbing state. Throughout the experiments, we used $\gamma = 0.95$.

For each of the test scenarios, following the approach in Section 4, we compute the optimal MMDP joint policy that we use to determine a POMDP model describing the decision

⁵We note that, although for the experimental results we consider a relatively simple communication model, it already includes state-dependent noise in the resulting observations. Generally, the framework accommodates for arbitrary communication models.

Table 1: Main differences between the groups of agents used.

Agents	Comm.	Succ.	Failed
	Freq.	Comm.	Comm.
QUERYPOMDP	Variable	POMDP	POMDP
NEVER COMM	Never	–	Indiv. MDP
ALWAYS COMM	1 step	MMDP	Indiv. MDP
COMM $k = 2$	2 steps	MMDP	Indiv. MDP
COMM $k = 3$	3 steps	MMDP	Indiv. MDP
COMM $k = 4$	4 steps	MMDP	Indiv. MDP

process for each individual agent. This POMDP is then solved using the PERSEUS approximate solver [22]. We test our QUERYPOMDP policy for 100 independent trials of 250 steps each and measure the obtained performance in terms of total discounted reward. We also test the performance of other sets of agents that communicate at different (but fixed) frequencies (see Table 1):

- “NEVER COMM” agents *never communicate*. These agents observe only their local state, and each follows the optimal policy for the underlying single-agent MDP obtained by disregarding the other agent in the environment;
- “ALWAYS COMM” agents *communicate at every time-step*, incurring the corresponding penalty. As QUERYPOMDP agents, they are subject to communication errors/limitations and, as such, are not always able to perceive the state of the other agent. When communication fails, the agent observes only its local state and adopts the individual MDP policy. When communication succeeds, it adopts the underlying MMDP policy.

Table 2: Total discounted reward for each set of agents in each of the test-scenarios. Entries in *italic* in the same column are not statistically different.

Environment	MAP 1	MAP 2	MAP 3	CIT	ISR	MIT	PENT.	SUNY
MMDP	5.787	5.253	6.608	<i>5.305</i>	<i>6.817</i>	3.182	7.606	<i>5.297</i>
QUERYPOMDP	5.132	3.598	6.156	<i>5.260</i>	<i>6.755</i>	<i>2.964</i>	6.444	<i>5.328</i>
NEVER COMM	-1.834	0.900	1.917	<i>5.306</i>	6.663	<i>2.959</i>	5.641	<i>5.283</i>
ALWAYS COMM	1.961	2.248	3.276	3.286	4.779	1.116	5.038	3.297
COMM $k = 2$	-0.069	1.097	3.001	4.306	5.839	2.141	5.578	4.294
COMM $k = 3$	-0.127	1.707	1.564	4.666	6.114	2.426	5.246	4.646
COMM $k = 4$	-0.785	1.289	3.295	4.324	5.760	2.106	5.448	4.317

- “COMM $k = 2, 3, 4$ ” agents query the state of the other agent every k steps. Except for the different communication frequency, they are otherwise similar to “ALWAYS COMM” agents.

Comparisons between these different agents will allow us to analyze (i) the impact that communication costs can have on performance, if communication is not optimized; and (ii) the impact that communication can have in mitigating partial observability. Comparing directly against other methods such as the one of Roth et al. [18] would not be very informative, as they do not attempt to trade off communication costs with task performance.

Results and Discussion

The performance of the 6 agent groups in terms of total discounted reward is summarized in Table 2. As a reference against which to assess the quality of our computed policy we also provide the results for the MMDP optimal policy in the different environments, providing a performance upper bound [15]. The QUERYPOMDP approach performs very favorably, outperforming all other policies and coming close to the MMDP upper bound in several of the tested scenarios.

The results in Table 2 prompt several interesting observations. First, comparing the performance of the MMDP policy against that of the group that never communicates provides an important indication of how critical coordination is in a given scenario. NEVERCOMM agents act individually, disregarding the existence of other agents in the environment. In environments where coordination is critical, NEVERCOMM agents will perform poorly. MMDP agents, on the other hand, always act in a perfectly coordinated manner, in which coordination does not come at a cost. In an environment where little coordination is needed, the difference between these two groups is going to be small. In contrast, scenarios that require significant coordination will cause the performance of the two groups to significantly differ.

From Table 2, we can see that coordination is critical in the smaller environments (Maps 1-3). In the larger environments, such as CIT, MIT and SUNY, coordination is less critical. The results in the smaller environments illustrate the impact of effective communication in mitigating the effects of partial observability. Our method is actually able to attain a performance very close to that of the MMDP agents, even paying for communication. Additionally, our approach uses communication efficiently, since the performance of all other communicating agents is significantly inferior. In contrast, in CIT, MIT and SUNY, non-communicating agents actually attain optimal performance. The difference in performance to the communicating groups can be explained by the com-

munication penalty. Again, in these scenarios, our approach is able to effectively manage communication needs, as it performs similarly to non-communicating agents.

A second observation is that the MMDP performance is an upper bound on the optimal Dec-MDP performance. This means that in those scenarios where our approach performs close to or as well as the MMDP group, we can immediately conclude that it is also performing close to or as well as the optimal Dec-MDP policy. A general comparison of the performance of our method against that of the MMDP group indicates that our method, if not optimal, must be very close to optimal in most scenarios tested. This, in turn, indicates that approximating the behavior of our agents with that of MMDP agents does provide a solid basis for planning.

Summarizing, our results show that, in scenarios with sparse interactions like the ones analyzed, our agents behave approximately as MMDP agents, effectively using communication to mitigate the effects of partial observability.

6. RELATED WORK

In the Dec-POMDP literature, early approaches introduced the idea of transition and reward independence as forms of simplified interactions [3]. Further examples of models with sparse interactions include *interaction-driven Markov games* [11, 21], *distributed POMDPs with coordination locales* [23], *transition-decoupled POMDPs* [24], *factored Dec-POMDPs* [16], and models relying on event-driven interactions [2].

Our representation is closest to interaction-driven Markov games (IDMGs) [11, 21]. This model leverages the independence between the different agents in a Dec-POMDP to decouple the decision process in significant portions of the joint state-space. In those situations in which the agents interact, IDMGs rely on communication to bring down the the computational complexity of the joint decision process. The use of communication to overcome partial observability differentiates this approach from other approaches that also exploit local interactions among the agents. However, communication is assumed to *always* take place and to be error-free [21]. In our case, we add explicit query actions to each agent’s action repertoire, enabling it to query another agent’s state, subject to certain constraints. For instance, two robots may only be able to share information when they are physically close. Furthermore, we assume that communication is subject to errors and comes at a cost that must be considered.

Explicit communication in multiagent planning was already addressed in [17], where the proposed Com-MTDP model allows agents to explicitly reason about communica-

tion in Dec-POMDP scenarios. However, being a generalization of Dec-POMDPs, it shares the discouraging computational complexity of the latter model. The actual process of communication has been investigated in [8]. Roth et al. [18] propose to exploit a factored Dec-MDP model and policy representation, in which agents query other agents’ local states when this knowledge is required for choosing their local actions. Although this work already seeks to optimize communication, this optimization is conducted parallel with the underlying decision process. Therefore, the cost of communication does not directly translate in the agent’s task performance, as in our proposed approach, rendering the tradeoff between communication costs and benefits unclear. Another closely related work is that of Wu et al. [26] where communication is used as a means to decrease the planning complexity in Dec-POMDP models. Like in our proposed approach, this work considers that communication may not always be available. However, unlike our approach, this work does not consider explicitly optimization of communication. Finally, Mostafa and Lesser [13] do optimize communication, while considering the presence of communication limitations. However, this optimization is also conducted parallel with the underlying decision process, without directly impacting in the agent’s task performance. Also, none of the aforementioned methods considers noisy communication channels.

A key point in our approach is that, although we use the MMDP policy in our planning, its computation is significantly more efficient than computing a centralized policy for the actual partially observable decision problem. The fact that we plan individually for each agent is somewhat related to several works that use round-robin policy optimization to individually optimize the policy of different agents in Dec-POMDP settings. One of the early examples is the JESP algorithm [14], which also models agents individually as POMDPs, but does not use communication. Round-robin policy optimization has been used to learn communication primitives in Dec-POMDPs whose base models are transition and observation independent [20], but which are coupled through the communication actions agents can choose to execute. In that case, however, agents have to learn when sending a particular message will be beneficial for team performance, which is far from trivial given that the policy of the receiving agent does not exploit the information provided by incoming messages. In our case, however, agents can opt to query other agents’ states, and it is much easier to determine when doing so improves performance. Secondly, we consider a much richer model where agents also “physically” influence each other, instead of only through communication.

7. CONCLUSIONS

In this paper, we analyzed the interplay between sparse interactions and communication in multiagent planning. We observed that, in scenarios where interactions among agents are sparse (*i.e.*, intra-agent action coordination is only infrequently necessary), the distributed execution of an MMDP policy seldom requires full-state information. As such, if each agent is (individually) allowed to query other agents for their local state information when necessary, it may be possible to partly mitigate partial state observability and leverage more efficient planning approaches.

Relying on this insight, we proposed the use of a POMDP model to analyze the communication needs of an agent in a

Dec-MDP scenario where the interaction between the agents is sparse. Our model accommodates communication costs and failures—the agent must explicitly reason about these factors in its decision process. QUERYPOMDP allows agents to optimize communication, explicitly trading-off its costs with its benefits in terms of the underlying task.

We used our approach to optimize communication in the simple scenario of Fig. 1, where our approach was successfully able to capture the intuition that the fundamental states for coordination are those around the doorway. We further explored the usefulness of this approach in computing policies for larger and more general Dec-MDPs. We built POMDP models for each agent by considering the other agents to behave as if in an MMDP, and use the obtained POMDP optimal policies. Our results show that our agents are able to effectively using communication to mitigate the effects of partial observability, behaving approximately as MMDP agents.

This work raises several interesting research questions. First of all, it would be interesting to generalize these techniques beyond Dec-MDPs, accommodating scenarios in which agents can query other agents’ *observations* instead of states. Another important issue is that, currently, we assume that the query observation model is known. While this assumption may be reasonable in a planning setting, the performance of the QUERYPOMDP agents is critically dependent on its communication capabilities. Therefore, if the agent designer is given the possibility of enabling communication in critical situations, how could these be computed? Early works [5, 9, 10] proposed methods that allow agents to automatically learn where information about others may be useful. A more recent method proposed the use of constrained optimization to explicitly determine when communication can be used to construct informative local beliefs [12]. An interesting avenue for future work is the extension of such approaches to the QUERYPOMDP setting.

Finally, one other important aspect that will be explored in future work is the extent to which sparsity of interaction affects the optimality of our planning approach. In fact, while our results indicate that our approach is able to effectively leverage the particular structure that may exist in multiagent settings, it is not clear how the performance of our approach depends on the sparsity of interactions. In fact, our method leverages such existing structure only implicitly. Characterizations such as the one in [25] may provide insights into which situations can be successfully addressed using our approach.

8. ACKNOWLEDGMENTS

The authors acknowledge the helpful suggestions by the anonymous reviewers. This work was funded in part by Fundação para a Ciência e a Tecnologia (INESC-ID multi-annual funding) through the PIDDAC Program funds and the project CMU-PT/SIA/0023/2009 under the Carnegie Mellon-Portugal Program. M.S. is funded by the FP7 Marie Curie Actions Individual Fellowship #275217 (FP7-PEOPLE-2010-IEF).

9. REFERENCES

- [1] R. Becker, A. Carlin, V. Lesser, and S. Zilberstein. Analyzing myopic approaches for multi-agent communications. *Computational Intelligence*, 25(1):31–50, 2009.

- [2] R. Becker, V. Lesser, and S. Zilberstein. Decentralized Markov decision processes with event-driven interactions. In *Proc. Int. Conf. Auton. Agents and Multiagent Systems*, pages 302–309, 2004.
- [3] R. Becker, S. Zilberstein, V. Lesser, and C. Goldman. Solving transition independent decentralized Markov decision processes. *J. Artificial Intelligence Research*, 22:423–455, 2004.
- [4] D. Bernstein, R. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, 27(4):819–840, 2002.
- [5] Y. De Hauwere, P. Vrancx, and A. Nowé. Learning multiagent state space representations. In *Proc. 9th Int. Conf. Autonomous Agents and Multiagent Systems*, pages 715–722, 2010.
- [6] P. Doshi and P. Gmytrasiewicz. On the difficulty of achieving equilibrium in interactive POMDPs. In *Proc. 21st AAAI Conf. Artificial Intelligence*, pages 1131–1136, 2006.
- [7] C. Goldman and S. Zilberstein. Optimizing information exchange in cooperative multiagent systems. In *Proc. 2nd Int. Conf. Autonomous Agents and Multiagent Systems*, pages 137–144, 2003.
- [8] C. Goldmann, M. Allen, and S. Zilberstein. Learning to communicate in a decentralized environment. *J. Auton. Agents and Multiagent Systems*, 15(1):47–90, 2007.
- [9] J. Kok, P. Hoen, B. Bakker, and N. Vlassis. Utile coordination: Learning interdependencies among cooperative agents. In *IEEE Symp. Computational Intelligence and Games*, pages 61–68, 2005.
- [10] F. Melo and M. Veloso. Learning of coordination: Exploiting sparse interactions in multiagent systems. In *Proc. 8th Int. Conf. Autonomous Agents and Multiagent Systems*, pages 773–780, 2009.
- [11] F. Melo and M. Veloso. Decentralized MDPs with sparse interactions. *Artificial Intelligence*, 175(11):1757–1789, 2011.
- [12] J. Messias, M. Spaan, and P. Lima. Efficient offline communication policies for factored multiagent POMDPs. In *Adv. Neural Information Proc. Systems*, volume 24, pages 1917–1925, 2011.
- [13] H. Mostafa and V. Lesser. Offline planning for communication by exploiting structured interactions in decentralized MDPs. Technical Report TR 2009-020, Comp. Science Dep. Univ. Massachusetts, 2009.
- [14] R. Nair, M. Tambe, M. Yokoo, D. Pynadath, and S. Marsella. Taming decentralized POMDPs: Towards efficient policy computation for multiagent settings. In *Proc. 18th Int. Joint Conf. Artificial Intelligence*, pages 705–711, 2003.
- [15] F. Oliehoek, M. Spaan, and N. Vlassis. Optimal and approximate Q -value functions for decentralized POMDPs. *J. Artificial Intelligence Research*, 32:289–353, 2008.
- [16] F. Oliehoek, M. Spaan, S. Whiteson, and N. Vlassis. Exploiting locality of interaction in factored Dec-POMDPs. In *Proc. Int. Conf. Autonomous Agents and Multiagent Systems*, pages 517–524, 2008.
- [17] D. Pynadath and M. Tambe. The communicative multiagent team decision problem: Analyzing teamwork theories and models. *J. Artificial Intelligence Research*, 16:389–423, 2002.
- [18] M. Roth, R. Simmons, and M. Veloso. Exploiting factored representations for decentralized execution in multiagent teams. In *Proc. Int. Conf. Auton. Agents and Multiagent Systems*, pages 469–475, 2007.
- [19] S. Seuken and S. Zilberstein. Formal models and algorithms for decentralized decision making under uncertainty. *Autonomous Agents and Multiagent Systems*, Feb. 2008.
- [20] M. Spaan, G. Gordon, and N. Vlassis. Decentralized planning under uncertainty for teams of communicating agents. In *Proc. Int. Conf. Autonomous Agents and Multiagent Systems*, pages 249–256, 2006.
- [21] M. Spaan and F. Melo. Interaction-driven Markov games for decentralized multiagent planning under uncertainty. In *Proc. Int. Conf. Auton. Agents and Multiagent Systems*, pages 525–532, 2008.
- [22] M. Spaan and N. Vlassis. Perseus: Randomized point-based value iteration for POMDPs. *J. Artificial Intelligence Res.*, 24:195–220, 2005.
- [23] P. Varakantham, J. Kwak, M. Taylor, J. Marecki, P. Scerri, and M. Tambe. Exploiting coordination locales in distributed POMDPs via social model shaping. In *Proc. 19th Int. Conf. Automated Planning and Scheduling*, pages 313–320, 2009.
- [24] S. Witwicki and E. Durfee. Influence-based policy abstraction for weakly-coupled Dec-POMDPs. In *Proc. Int. Conf. Automated Planning and Scheduling*, 2010.
- [25] S. Witwicki and E. Durfee. Towards a unifying characterization for quantifying weak coupling in Dec-POMDPs. In *Proc. 10th Int. Joint Conf. Autonomous Agents and Multiagent Systems*, pages 29–36, 2011.
- [26] F. Wu, S. Zilberstein, and X. Chen. Multi-agent online planning with communication. In *Proc. Int. Conf. Automated Planning and Scheduling*, pages 321–329, 2009.
- [27] P. Xuan and V. Lesser. Multi-agent policies: From centralized ones to decentralized ones. In *Proc. 1st Int. Joint Conf. Autonomous Agents and Multiagent Systems*, pages 1098–1105, 2002.