

Sample Bounded Distributed Reinforcement Learning for Decentralized POMDPs

Bikramjit Banerjee¹, Jeremy Lyle², Landon Kraemer¹, Rajesh Yellamraju¹

¹School of Computing, The University of Southern Mississippi, Hattiesburg, MS 39402

²Department of Mathematics, The University of Southern Mississippi, Hattiesburg, MS 3940

Abstract

Decentralized partially observable Markov decision processes (Dec-POMDPs) offer a powerful modeling technique for realistic multi-agent coordination problems under uncertainty. Prevalent solution techniques are *centralized* and assume prior knowledge of the *model*. We propose a distributed reinforcement learning approach, where agents take turns to learn best responses to each other's policies. This promotes decentralization of the policy computation problem, and relaxes reliance on the full knowledge of the problem parameters. We derive the relation between the sample complexity of best response learning and error tolerance. Our key contribution is to show that sample complexity could grow exponentially with the problem horizon. We show empirically that even if the sample requirement is set lower than what theory demands, our learning approach can produce (near) optimal policies in some benchmark Dec-POMDP problems.

Introduction

Decentralized partially observable Markov decision processes (Dec-POMDPs) offer a powerful modeling technique for realistic multi-agent coordination and decision making problems under uncertainty. Because solving Dec-POMDPs is NEXP-complete (Bernstein et al. 2002), exact solution techniques for finite horizon problems require significant time and memory resources (Szer and Charpillet 2006; Oliehoek et al. 2010; Spaan, Oliehoek, and Amato 2011). However, solution techniques for Dec-POMDPs (exact or approximate) suffer from less acknowledged limitations as well: that most of them are *centralized* and assume prior knowledge of the *model*. That is, a single program computes the optimal joint policy, with the full knowledge of the problem parameters. While these techniques have had success in benchmark problems with comprehensively defined domain parameters, such strict definitions may be difficult and tedious in many real-world problems. In these cases, the problem parameters may need to be first estimated from experience and then exact/approximate solvers may be applied. However, this problem of model estimation can be complex in POMDPs because states are unobservable (Chrisman 1992), and additionally in Dec-POMDPs agents are incapable of observing each other's actions and

observations. Not surprisingly, model estimation is largely ignored in the Dec-POMDP literature. Furthermore, the task of computing optimal policies in Dec-POMDPs has seldom been decentralized in any meaningful way (e.g., see (Emery-Montemerlo et al. 2004).)

In this paper we propose a distributed reinforcement learning approach to solving finite horizon Dec-POMDPs. When agents *learn* their *own* policies, not only is the task of policy computation distributed, but also the problem parameters do not need to be known a priori. In lieu of the knowledge of problem parameters, access to a simulator, or simply the ability to draw samples from unknown distributions would be sufficient. Effectively, estimation of the problem parameters is built into the learning algorithm. Policy learning in *finite* horizon tasks is justified due to the same reasons as finite horizon reinforcement learning, viz., that agents can learn policies in offline simulations before applying them in the real domain. Furthermore, unlike many exact and approximate solution approaches for Dec-POMDPs, the memory usage of a learning approach is not much larger than the size of a *single* policy per agent at any time, which makes it relatively more memory efficient. Thus we posit distributed reinforcement learning as a more practical alternative to the traditional Dec-POMDP solvers.

Our initial experiments with concurrent *independent* reinforcement learning (Claus and Boutilier 1998) in benchmark Dec-POMDP problems have yielded unsatisfactory results, some of which are included in the experiments section. In this paper we propose a non-concurrent independent learning approach, where agents take turn in learning best response policies to each other via a semi-model based Monte Carlo algorithm, but no agent explicitly attempts to model the other agent. We show theoretically that Monte Carlo reinforcement learning for best response has a complexity of $O(T^3|A|^T|\Omega|^{3T-1})$, where T is the horizon length, $|A|$ and $|\Omega|$ are respectively the number of actions and observations available to an agent. While it is intuitive that this expression would be exponential in T , we show that even the time spent on each trajectory could be exponential in T . Our analysis shows that this depends on the number of distinct scenarios that the other (non-learning and invisible) agent may encounter. We also show empirically that a "few" alternations of best response learning produce (near) optimal policies in some benchmark problems, although in general

alternate best response learning can converge to local optima.

Related Work

Reinforcement learning (RL) has been applied in infinite horizon POMDPs in both model based (Chrisman 1992; McCallum 1995; Shani, Brafman, and Shimony 2005) and model free (Meuleau et al. 1999) ways. Model based methods learn a model (e.g., hidden Markov models or utility suffix memories for POMDPs) of the environment first and then compute a policy based on the learned model, while model free methods learn a policy directly. Model learning can be more complex in Dec-POMDPs because the actions and the observations of the other agents are unobservable. We use a semi-model based approach, where we do not attempt to estimate the Dec-POMDP parameters directly, owing to their hidden parts, but instead learn intermediate functions that capture the visible parts of the dynamics (see equations 3, 4 given later) via Monte Carlo estimation, and compute a policy based on these functions.

Peshkin et al. (2000) introduced a gradient ascent based policy search algorithm for multi-agent reinforcement learning in Dec-POMDP-like settings, using finite state controllers as the policy language. In this and other RL based techniques, the agents are assumed to be able to observe the global rewards, while Dec-POMDP solvers assume that rewards are unobservable when agents execute the optimal policy. We reconcile these positions by noting that it is reasonable for agents to observe global rewards in offline simulations for the purpose of *learning only*, as long as the ultimate execution of the learned policies does not rely on reward observations. This is the position underlying our approach as well.

Among work using centralized sampling-based techniques for solving Dec-POMDPs, Wu et al. (2010) is one of the most recent. They have applied a centralized Monte Carlo approach to learning finite state controllers for finite horizon problems, while our approach is more decentralized. More recently, Zhang and Lesser (2011) have applied distributed reinforcement learning to a variant of the finite horizon Dec-POMDP problem, where agents are organized in a network, and agents’ influence on each other are limited to cliques. This factored structure as well as the known communication structure of the domain is exploited to solve such problems more scalably than regular Dec-POMDPs. Zhang and Lesser also exploited the known communication structure to coordinate the sub-teams via distributed constraint optimization, and produced a more efficient learning-based alternative to the regular solvers. While our goal is similar and we also consider finite horizon problems, we focus on less structured and unfactored Dec-POMDPs that are inherently less scalable.

Our work is most closely related to JESP (Nair et al. 2003) which is based on a redefinition of a belief state as a cross product of the hidden state with the other agents’ policies. JESP uses dynamic programming to compute the best response policy of one agent holding the others’ policies fixed, and alternate the computation of best responses in a hill-climbing fashion. Our approach can be considered

the “learning counterpart” of JESP with sample complexity analysis, although our approach outperforms JESP due to specialized initialization (Kraemer and Banerjee 2012) as opposed to random initialization of JESP. We omit these experiments here due to space constraint.

Decentralized POMDP

The Decentralized POMDP (Dec-POMDP) formalism is defined as a tuple $\langle n, S, A, P, R, \Omega, O \rangle$, where:

- n is the number of agents playing the game.
- S is a finite set of (unobservable) environment states.
- $A = \times_i A_i$ is a set of joint actions, where A_i is the set of individual actions that agent i can perform.
- $P(s'|s, \vec{a})$ gives the probability of transitioning to state $s' \in S$ when joint action $\vec{a} \in A$ is taken in state $s \in S$.
- $R : S \times A \rightarrow \mathfrak{R}$, where $R(s, \vec{a})$ gives the immediate reward the agents receive upon executing action $\vec{a} \in A$ in state $s \in S$.
- $\Omega = \times_i \Omega_i$ is the set of joint observations, where Ω_i is the finite set of individual observations that agent i can receive from the environment.
- $O(\vec{\omega}|s', \vec{a})$ gives the probability of the agents jointly observing $\vec{\omega} \in \Omega$ if the current state is $s' \in S$ and the previous joint action was $\vec{a} \in A$.

The reward function R , transition model P , and observation model O are defined over joint actions and/or observations, which forces the agents to coordinate. Additionally, for finite horizon problems, horizon T is also specified. The goal of the Dec-POMDP problem is to find a policy for each agent (*joint policy*) that maximizes the total expected reward over T steps of interaction, given that the agents cannot communicate their observations and actions to each other. A joint policy Π is a set of individual policies, π_i , which maps the histories of action-observation pairs of agent i to actions in A_i . The problem of finding an optimal joint policy has been proven to be NEXP-complete (Bernstein et al. 2002).

Reinforcement Learning

Reinforcement learning problems are modeled as *Markov Decision Processes* or MDPs (Sutton and Barto 1998). Conceptually, an MDP is a POMDP where an observation uniquely identifies the hidden state. It is given by the tuple $\langle S, A, R, P \rangle$, where S is the set of visible environmental states that an agent can be in at any given time, A is the set of actions it can choose from at any state, $R : S \times A \mapsto \mathfrak{R}$ is the reward function, i.e., $R(s, a)$ specifies the reward from the environment that the agent gets for executing action $a \in A$ in state $s \in S$; $P : S \times A \times S \mapsto [0, 1]$ is the state transition probability function specifying the probability of the next state in the Markov chain following the agent’s selection of an action in a state. In finite horizon problems, the agent’s goal is to learn a *non-stationary* policy $\pi : S \times t \mapsto A$ that maximizes the sum of current and future rewards from any state s , given by,

$$V^\pi(s^0, t) = E_P[R(s^0, \pi(s^0, t)) + R(s^1, \pi(s^1, t+1)) + \dots + R(s^{T-t}, \pi(s^{T-t}, T))]$$

where s^0, s^1, \dots, s^{T-t} are successive samplings from the distribution P following the Markov chain with policy π .

Reinforcement learning algorithms often evaluate an action-quality value function Q given by

$$Q(s, a, t) = R(s, a) + \max_{\pi} \gamma \sum_{s'} P(s, a, s') V^{\pi}(s', t + 1) \quad (1)$$

This quality value stands for the sum of rewards obtained when the agent starts from state s at step t , executes action a , and follows the optimal policy thereafter. Action quality functions are preferred over value functions, since the optimal policy can be calculated more easily from the former. Learning algorithms can be model based or model free. Model based methods explicitly estimate $R(s, a)$ and $P(s, a, s')$ functions, and hence estimate $Q(s, a, t)$. Model free methods directly learn $Q(s, a, t)$, often by online dynamic programming, e.g., *Q-learning*. In this paper, we use (semi-) model based learning for Dec-POMDPs with the associated analysis of sample complexity, thus establishing a baseline for RL in Dec-POMDPs. Model based reinforcement learning algorithms have been analyzed in many domains before, but to the best of our knowledge such analyses have not been performed for decentralized POMDPs, where partial observability of the learner's environment is compounded by the unobservability of the other agents' observations and actions.

RL for Dec-POMDPs

Solution techniques for Dec-POMDPs have been mostly centralized (Szer and Charpillet 2006; Oliehoek et al. 2010; Spaan, Oliehoek, and Amato 2011), in that a single program computes the optimal joint policy, with the full knowledge of the problem parameters, viz., P, R, O . While these techniques have had success in benchmark problems with comprehensively defined P, R, O , such strict definitions may be difficult and tedious in real-world problems. In this paper we address this issue by applying reinforcement learning to the policy computation problem. The main distinguishing characteristics of our approach are

- Instead of a single program computing the optimal joint policy, each agent learns its own policy. In this paper agents learn distributedly, but not concurrently. That is, they share the task of policy learning, by only learning their own policies, but do not update policies concurrently. Concurrent learning will effectively parallelize Dec-POMDP solution, but it is also challenging due to potential oscillation. Our experiments show that concurrent learning is not as efficient as the proposed distributed learning approach. We leave the improvement of concurrent learning in Dec-POMDPs as a future avenue.
- Instead of using knowledge of P, R, O , agents learn on the basis of sampling these unknown functions. This allows our approach to be readily applicable in tasks where these parameters are unknown, or hard to compute. However, for evaluation purposes, we still consider well-defined benchmark problems in this paper.

- Most Dec-POMDP solvers maintain many policies in memory at any time, partly or wholly. Even memory bounded techniques (Seuken 2007) maintain multiple policies, although of a bounded total size. Instead, a learner only needs to maintain sufficient information in memory to construct *one* policy. However, for finite horizon problems, this policy has a size exponential in T .

Although the agents are unaware of P, R, O , we assume that the agents know the size of the problem, i.e., $|A|, |S|, |\Omega|$, the maximum magnitude over all rewards, R_{\max} , and that they are capable of signalling to each other so that no two agents are learning at the same time.

Since states are not visible, a reinforcement learning agent can use the policy representation of finite horizon Dec-POMDPs, and learn a mapping from histories of its past actions and observations to actions (Zhang and Lesser 2011). For simplicity of notation, we assume two agents only, and identical action and observation sets for both agents. Given the policy of the other agent, π , the quality of a learner's action a at a given level- t history h_t is given by

$$Q_t^*(h_t, a|\pi) = R_t^*(h_t, a|\pi) + \sum_{\omega} H_t^*(h_t, a, h_{t+1}|\pi) \cdot \max_b Q_{t+1}^*(h_{t+1}, b|\pi) \quad (2)$$

where h_{t+1} is a level- $t + 1$ history produced by the concatenation of h_t and (a, ω) , i.e., $h_{t+1} = (h_t, a, \omega)$. The best response policy of the learner, π_{ℓ} , to the other agent's policy π is given by

$$\pi_{\ell}(h_t) = \arg \max_a Q_t^*(h_t, a|\pi).$$

The functions R_t^* and H_t^* represent level- t reward and history transition functions for the learner, given by

$$R_t^*(h_t, a|\pi) = \sum_{s, h_-} P(s|h_t, h_-) P(h_-|h_t, \pi) R(s, \vec{a}) \quad (3)$$

$$H_t^*(h_t, a, h_{t+1}|\pi) = \sum_{s, s', h_-} P(s|h_t, h_-) P(h_-|h_t, \pi) \cdot P(s'|s, \vec{a}) \sum_{\omega_-} O(\vec{\omega}|s', \vec{a})(4)$$

where h_- is the history of action-observations encountered by the other agent, $\vec{\omega} = \langle \omega, \omega_- \rangle$ and $\vec{a} = \langle a, \pi(h_-) \rangle$ are the joint observation and action respectively. A learning agent is unaware of every factor on the right hand sides of equations 3, 4, and must estimate R^* and H^* solely from its own experience of executing actions and receiving observations and rewards.

For brevity, we call the following expression β .

$$\beta = |A| \left(\frac{(|A||\Omega|)^T - 1}{|A||\Omega| - 1} \right). \quad (5)$$

β gives the maximum number of (h, a) pairs of all lengths that a learner may encounter. We now give the definition of a key parameter that appears in the complexity expression of our algorithm.

Definition 1. *The minimum reachability over all feasible states at level t for a fixed policy of the other agent, π , is given by*

$$\rho_{t, \pi} = \min_{s, h_t, h_-|\pi} P(s|h_t, h_-) P(h_-|h_t, \pi)$$

Feasibility excludes unreachable states, and therefore ensures that always $\rho_{t,\pi} > 0$.

When a learner takes action a at history h_t (with the other agent executing $\pi(h_-)$) and the resulting joint observation is $\vec{w} = \langle w, w_- \rangle$, then reachability can be propagated as

$$\begin{aligned} & P(s|h_t, h_-)P(h_-|h_t, \pi)P(s'|s, \vec{a})O(\vec{w}|s', \vec{a}) \\ = & P(s'|h_{t+1}, h'_-)P(h'_-|h_{t+1}, \pi) \end{aligned}$$

where $h_{t+1} = (h_t, a, \omega)$ and $h'_- = (h_-, \pi(h_-), \omega_-)$. Clearly, the minimum reachability at level $t + 1$ is

$$\rho_{t+1,\pi} \leq \rho_{t,\pi},$$

forming a monotonically decreasing sequence with increasing t . Therefore, we refer to the minimum reachability over all steps, $\rho_{T-1,\pi}$, simply as ρ dropping both subscripts when π is clear from the context.

The Algorithm: MCQ-alt

In this paper we present an approach where agents take turn to learn best response to each others policies, using an R-Max (Brafman and Tennenholtz 2002) like approach to learn the best response Q-values. The learner records immediate rewards and history transitions at every history encountered, providing samples of R^* and H^* given in equations 3, 4 respectively. These samples are incorporated into running averages to maintain estimates \hat{R} and \hat{H} respectively. For histories of length $T - 1$ (i.e., full length), h_{T-1} , if the pair (h_{T-1}, a) has been encountered

$$\begin{aligned} N = & \max(|S|^2|\Omega|^{T-1}, 4/\rho) \frac{(4R_{\max}T|S|)^2|\Omega|^{T+1}}{\alpha^2} \\ & \ln(16|S|^2|\Omega|^T\beta/\delta) \end{aligned} \quad (6)$$

times (N is derived in the next section), then the learner sets $Q_{T-1}(h_{T-1}, a)$ to the average of the immediate rewards received (i.e., $\hat{R}_{T-1}(h_{T-1}, a)$). It then marks (h_{T-1}, a) as “Known”. If (h_{T-1}, a) is “Known” for every a , then h_{T-1} is marked “Known”. For an intermediate length history, h_t , if every history, h_{t+1} , produced by concatenating h_t with (a, ω) for all combinations of action-observations encountered is “Known”, then h_t is marked “Known”. For a “Known” intermediate history h_t , $Q_t(h_t, a)$ is updated for every a as

$$Q_t(h_t, a) = \hat{R}_t(h_t, a) + \sum_{h'} \hat{H}_t(h_t, a, h') \max_b Q_{t+1}(h', b)$$

The learner’s exploration strategy is as follows. For a “Known” history, action selection is greedy, i.e., the Q -maximizing action is selected. For a history that is not yet marked “Known”, the least frequently taken action (ties broken randomly) is executed. The learner freezes its current policy when the empty history is marked “Known”, and signals to the other agent to start learning, while it executes its current policy without exploration. The above algorithm learns best response directly without modeling the other agents in the environment, and only modeling the visible parts of the environment’s dynamics. Due to alternation

of best response learning, and the use of Monte Carlo estimation, we call this algorithm “Monte Carlo Q alternating”, or “MCQ-alt” in short.

The infimum reachability, ρ used in equation 6, may not be known in many problems. Since it decreases geometrically with increasing T , it may even be hard to determine whether $|S|^2|\Omega|^{T-1}$ dominates $4/\rho$. In such cases, it may be possible to ignore it, but this may be inadequate for higher T . For the domains and horizons used for experiments in this paper, ρ does not appear to be a dominating factor, so our analysis focuses on the dependence on T instead.

An important feature of MCQ-alt is its well-defined stopping criterion, viz., when the empty history becomes “Known”, which is controlled by a single parameter, N . In contrast, Q-learning is controlled by multiple parameters, and its stopping criterion can be affected by oscillation or non-convergence.

Note that in learning the best response, a learner attempts to cover every (h_t, a) encountered equally well, to guarantee arbitrarily small errors in the value function. However, there are at least two reasons why the value function may not need to be accurate to an arbitrary degree: (1) policies usually converge long before value functions, which we verify in this paper experimentally, and (2) some less likely paths may have little impact on the value function and N could be lowered for these paths; we leave this for future work.

Analysis

We focus on sample complexity analysis of the best response learning process. As in (Nair et al. 2003), the fixed policy of the other agent effectively reduces the Dec-POMDP to a POMDP, in which the state space is the cross product of the hidden states and the other’s policy. Since the other agent’s actions and observations are unobservable to the learner, the number of scenarios encountered by the other agent (referred to as K) becomes a key parameter in the sample complexity analysis, as we demonstrate in this section.

First we note that the number of episodes needed for the empty history to be “Known” is

$$\geq N|\Omega|^{T-1}|A|^T$$

since the number of distinct (h_{T-1}, a) tuples is $|\Omega|^{T-1}|A|^T$. Also, given the backup process of histories becoming “Known” in our algorithm, when all (h_{T-1}, a) become “Known” the empty history must also become “Known”, and this takes N visitations of each tuple. The actual number of episodes needed is, however, likely to be greater than $N|\Omega|^{T-1}|A|^T$, because only part of the exploration process is under the learner’s control, where it can select a but not ω . Thus it can be led to revisit paths that are already “Known”.

While it is fairly intuitive that the episode complexity given above should be exponential in T , we will now show that so could N . In order to demonstrate this, we first present a generic sampling process, and use the resulting sample complexity expression to derive N .

The Basic Sampling Process

Consider the following sampling process, with K classes of random variables, $\{X_{jl}\}_{j=1}^K$, such that $Y =$

$\frac{\sum_{i=1}^{N_1} X_{1i} + \sum_{i=1}^{N_2} X_{2i} + \dots + \sum_{i=1}^{N_K} X_{Ki}}{N}$, where $N = \sum_j^K N_j$. The process generates a sample of some X_{jl} at each of the N iterations, where the probability that the sample belongs to class $j \in [1, K]$ is p_j . Therefore, all X_{jl} as well as all N_j are random variables. Suppose that $E[X_{jl}] = M_j$, $\forall l$, and that the maximum magnitude of any X_{jl} is $X_{\max} > 0$. We wish Y to estimate the unknown value $\sum_j p_j M_j$, therefore we call $|Y - \sum_j p_j M_j|$ the *estimation error*. We claim the following sufficient condition for bounding the estimation error, but omit the proof due to lack of space.

Theorem 1. *If the total number of samples is set $N \geq \max(K\eta, 4\eta/\min_j p_j)$ where*

$$\eta = \frac{4X_{\max}^2 K}{\epsilon^2} \ln(8K/\delta),$$

then the estimation error is bounded, i.e.,

$$P(|Y - \sum_j p_j M_j| > \epsilon) < \delta$$

Derivation of N

In our analysis we shall use the max norm function, i.e., $\|f - g\|$ represents $\max_x |f(x) - g(x)|$. We first establish the dependence of the error in the Q functions on the errors in our estimates \hat{H} and \hat{R} .

Lemma 2. *If $\|\hat{H}_\tau - H_\tau^*\| \leq \epsilon_1$ and $\|\hat{R}_\tau - R_\tau^*\| \leq \epsilon_2$ for all τ , then at any step t*

$$\|Q_t - Q_t^*\| \leq (T-t)\epsilon_2 + (T-t-1)|\Omega|R_{\max}\epsilon_1$$

Proof: By induction. For basis, $t = T-1$. Since T is the last step in an episode, $\|Q_{T-1} - Q_{T-1}^*\| = \|\hat{R}_{T-1} - R_{T-1}^*\| \leq \epsilon_2$, hence true.

For the inductive case, we see that $|Q_t(h_t, a) - Q_t^*(h_t, a)|$

$$\begin{aligned} &= |\hat{R}_t(h_t, a) + \sum_{h'} \hat{H}_t(h_t, a, h') \max_b Q_{t+1}(h', b) - R_t^*(h_t, a) - \sum_{h'} H_t^*(h_t, a, h') \max_b Q_{t+1}^*(h', b)| \\ &\leq \|\hat{R}_t - R_t^*\| + |\sum(\hat{H}_t \max Q_{t+1} - H_t^* \max Q_{t+1}^*)| \\ &= \|\hat{R}_t - R_t^*\| + |\sum(\hat{H}_t \max Q_{t+1} - \hat{H}_t \max Q_{t+1}^* + \hat{H}_t \max Q_{t+1}^* - H_t^* \max Q_{t+1}^*)| \\ &\leq \epsilon_2 + |\sum(\hat{H}_t \max Q_{t+1} - \hat{H}_t \max Q_{t+1}^*)| + |\sum(\hat{H}_t \max Q_{t+1}^* - H_t^* \max Q_{t+1}^*)| \\ &\leq \epsilon_2 + \max |Q_{t+1} - Q_{t+1}^*| + |\sum(\hat{H}_t - H_t^*) \max Q_{t+1}^*| \end{aligned}$$

In the last expression, the second term is upper bounded by $(T-t-1)\epsilon_2 + (T-t-2)|\Omega|R_{\max}\epsilon_1$, by the induction hypothesis. In the third term, $\max Q_{t+1}^* \leq R_{\max}$, and the sum is taken over all observations. Therefore the third term is upper bounded by $|\Omega|R_{\max}\epsilon_1$. Adding the bounds of the three terms we get the result. \square

Lemma 2 implies that the error bound increases for smaller histories, and therefore is maximum at the empty history. This is why the learner must continue until Q_0 is sufficiently accurate, i.e., the empty history becomes ‘‘Known’’. In the following analysis, we characterize ‘‘sufficiently accurate’’, to derive a bound on N used by the algorithm (equation 6).

Theorem 3. *To ensure that $\|Q_0 - Q_0^*\| \leq \alpha$ w.p. $\geq 1 - \delta$, it is sufficient to set*

$$N \geq \max(|S|^2|\Omega|^{T-1}, 4/\rho) \frac{(4R_{\max}T|S|)^2|\Omega|^{T+1}}{\alpha^2} \ln(16|S|^2|\Omega|^T\beta/\delta)$$

in our algorithm, where β is given in equation 5 and ρ results from Definition 1.

Proof: By Lemma 2,

$$\|Q_0 - Q_0^*\| \leq T\epsilon_2 + (T-1)|\Omega|R_{\max}\epsilon_1.$$

To achieve the α bound on the error, it is sufficient to set $\epsilon_1 \leq \alpha/2(T-1)|\Omega|R_{\max}$, and $\epsilon_2 \leq \alpha/2T$.

Now the number of \hat{H} and \hat{R} entries that need to be learned are $|\Omega|\beta$ and β respectively, where β is given in equation 5. Therefore, it is sufficient to require both of the following for any t :

$$P(\|\hat{R}_t - R_t^*\| > \alpha/2T) < \delta/2\beta \quad (7)$$

$$P(\|\hat{H}_t - H_t^*\| > \alpha/2(T-1)|\Omega|R_{\max}) < \delta/2\beta|\Omega|$$

First consider $\|\hat{R}_t - R_t^*\|$ and equation 3. The estimation of any \hat{R}_t in our algorithm matches the description of the sampling process, with $E[X_{jl}] = R(s, \vec{a})$ and $p_j = P(s|h_t, h_-)P(h_-|h_t, \delta)$; the last quantity being the reachability, of which the infimum is ρ (Definition 1). Note that p_j cannot be set to $\sum_{h_-} P(s|h_t, h_-)P(h_-|h_t, \delta)$, since each sample X_{jl} received corresponds to a specific history h_- encountered by the other agent. Therefore in this case, the number of variable classes in the sampling process is $K = |S||\Omega|^t \leq |S||\Omega|^{T-1}$. This is the maximum possible number of terms in the summation of equation 3 which corresponds to full length histories that can be encountered by the other agent, given its (fixed) policy. Making the appropriate substitutions in Theorem 1, we see that to ensure equation 7, it is sufficient to set

$$N \geq \max(|S||\Omega|^{T-1}, 4/\rho) \frac{(4R_{\max}T)^2|S||\Omega|^{T-1}}{\alpha^2} \ln(16|S||\Omega|^{T-1}\beta/\delta)$$

Similarly, for $\|\hat{H}_t - H_t^*\|$, the sampling process is characterized by $E[X_{jl}] = 1$ and $p_j = P(s|h_t, h_-)P(h_-|h_t, \delta)P(s'|s, \vec{a})P(\vec{\omega}|s', \vec{a})$. The last quantity is the propagated reachability, of which the infimum is also ρ . Since $t \leq T-2$, this yields $K = |S|^2|\Omega|^{t+1} \leq |S|^2|\Omega|^{T-1}$, and we have

$$N \geq \max(|S|^2|\Omega|^{T-1}, 4/\rho) \frac{(4R_{\max}(T-1)|S|)^2|\Omega|^{T+1}}{\alpha^2} \ln(16|S|^2|\Omega|^T\beta/\delta)$$

Combining the two, we get the result. \square

It is interesting to note that N is polynomial in most problem parameters, except that it is logarithmic in $|A|$, and exponential in T . Although Theorem 3 suggests that $N = O(T^3|\Omega|^{2T})$, our experiments suggest that it does not require to grow in some domains due to simpler structure. Even in the domains where it does need to grow, the rate of growth could be lower.





Figure 1: Plots of average relative errors against N in DEC-TIGER (left) and RECYCLING-ROBOTS (right).

Initial Policy

If agents alternate in learning best responses, the agent that does not learn initially must play some previously specified fixed policy. Our experiments show that if this policy is random then the final outcome is unpredictably poor. Instead, we simply let the two agents perform concurrent reinforcement learning to learn initial policies, on a slightly simpler Dec-POMDP. This Dec-POMDP reduces the set of observations to one dummy observation. In other words, the agents simply ignore the observations, and learn a mapping from their own past action histories to actions ($\pi(a_1, a_2, \dots, a_t) = a$). This policy can be readily translated to the regular policy language, by setting $\pi(a_1, \omega_1, a_2, \omega_2, \dots, a_t, \omega_t) = a$ for all possible chains of observations $(\omega_1, \omega_2, \dots, \omega_t)$. This is the policy used by the initially non-learning agent. See (Kraemer and Banerjee 2012) for more details.

Experimental Results

We present experimental results from two benchmark domains: DEC-TIGER (Nair et al. 2003) and RECYCLING-ROBOTS (Amato, Bernstein, and Zilberstein 2007). We used the initial policy learned by concurrent reinforcement learning (as described above) over 200000 episodes to perform alternating Monte-Carlo Q learning as described in this paper, for values of N ranging from 10 to 1000. There were 2 alternations, i.e., each agent learned best response to the other’s policy exactly once. In each experiment, the resulting joint policy after 2 alternations was evaluated to yield $|v_{jpol} - v_{opt}|/|v_{opt}|$, i.e., the relative error based on known optimal values for horizons 3, 4 and 5. The plots in Figure 1 show these relative errors averaged over 50 runs. For comparison, we also show the result from concurrent Q-learning (referred to as “Q-conc”), with $\alpha = 0.01$, $\epsilon = 0.005$, which were found to produce best results in the selected settings. Table 1 also shows the average relative error rates with the initial policy (derived by concurrent learning), to verify that MCQ-alt does indeed improve these policies.

Each setting of N and T makes MCQ-alt finish in a certain number of episodes, say $e_{N,T}$, i.e., until the empty history becomes “Known” in each alternation. The average of these numbers of episodes over all agents and all runs is used to determine the length of the Q-conc runs. The average relative error of Q-conc for a given N and T is reported at the

end of (average) $e_{N,T}$ episodes.

In Figure 1 (left) for DEC-TIGER, first we note that horizons 3 and 4 are solved accurately with $N \geq 200$ and 1000 respectively, by MCQ-alt. Q-conc solves horizon 3 accurately with a number of episodes corresponding to $N = 1000$, but is unable to solve horizon 4. Neither achieves 0 error for horizon 5. MCQ-alt is also clearly more efficient than Q-conc. More importantly, we see that for a given N , the relative error increases with increasing horizon. This is clear with MCQ-alt, but not so clear with Q-conc with even a hint of non-convergence (error increases for $T = 4$). For MCQ-alt, this implies that N needs to increase to produce the same error on increasing horizons. This is direct evidence for the claim made earlier in this paper, although the rate at which N needs to increase falls short of the $O(T^3|\Omega|^{2T})$ rate established in this paper. This is explained by the fact that $O(T^3|\Omega|^{2T})$ is a *sufficient* rate for *value convergence*; it is not *necessary*, and policies can converge sooner.

	DEC-TIGER			RECYCLING-ROBOTS		
	T=3	T=4	T=5	T=3	T=4	T=5
Initial policy relative error	2.16	2.67	2.42	0.37	0.39	0.33
% unreachable histories	0	0	0	27.9	39.8	47.8

Table 1: Relative errors of initial policies, and the proportion of unreachable histories.

In Figure 1 (right) for RECYCLING-ROBOTS, we see something more interesting. Although MCQ-alt is still more efficient than Q-conc, the relative errors now *decrease* with increasing horizon, for a given N . This is true with both MCQ-alt and Q-conc. This is partly due to the fact that there are many unreachable histories in this domain. Table 1 shows the increasing proportion of unreachable histories in RECYCLING-ROBOTS with increasing T , which suggests that K (and hence the error) grows much slower than Theorem 3 assumes. As it turns out, N does not need to grow with increasing horizons for small errors in this domain. Instead, the increasing values of $e_{N,T}$ even for a fixed N are sufficient to actually *reduce* the average errors with increasing T , as seen in Figure 1 (right). However, it must be noted that alternating best response learning also almost always converges to local optima in this domain, so 0 aver-

age error was never observed.

Conclusion

We have presented a distributed Monte Carlo based reinforcement learning algorithm for solving decentralized POMDPs. We have derived the sample complexity that guarantees arbitrarily accurate best response policies, and shown empirically that 2 alternations of best response learning can produce (near) optimal joint policies in some benchmark problems. In the future, more judicious use of samples with a variable N will be explored, and a more elaborate investigation into the convergence behavior of concurrent learning will be undertaken.

Acknowledgments: We thank the anonymous reviewers for helpful comments and pointers. This work was supported in part by the U.S. Army under grant #W911NF-11-1-0124.

References

- Amato, C.; Bernstein, D.; and Zilberstein, S. 2007. Optimizing memory-bounded controllers for decentralized POMDPs. In *Proc. UAI*.
- Bernstein, D. S.; Givan, R.; Immerman, N.; and Zilberstein, S. 2002. The complexity of decentralized control of markov decision processes. *Mathematics of Operations Research* 27:819–840.
- Brafman, R. I., and Tennenholtz, M. 2002. R-max - A general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research* 3:213–231.
- Chrisman, L. 1992. Reinforcement learning with perceptual aliasing: The perceptual distinctions approach. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, 183–188. San Jose, CA: AAAI Press.
- Claus, C., and Boutilier, C. 1998. The dynamics of reinforcement learning in cooperative multiagent systems. In *Proceedings of the 15th National Conference on Artificial Intelligence*, 746–752. Menlo Park, CA: AAAI Press/MIT Press.
- Emery-Montemerlo, R.; Gordon, G.; Schneider, J.; and Thrun, S. 2004. Approximate solutions for partially observable stochastic games with common payoffs. *Autonomous Agents and Multiagent Systems, International Joint Conference on* 1:136–143.
- Kraemer, L., and Banerjee, B. 2012. Informed initial policies for learning in dec-pomdps. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence Student Abstract and Poster Program*. To appear.
- McCallum, A. K. 1995. *Reinforcement Learning with Selective Perception and Hidden State*. Ph.D. Dissertation, Department of Computer Science, University of Rochester.
- Meuleau, N.; Peshkin, L.; Kim, K.; and Kaelbling, L. 1999. Learning finite-state controllers for partially observable environments. In *Proc. UAI*, 427–436.
- Nair, R.; Tambe, M.; Yokoo, M.; Pynadath, D.; and Marsella, S. 2003. Taming decentralized pomdps: Towards efficient policy computation for multiagent settings. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, 705–711.
- Oliehoek, F. A.; Spaan, M. T. J.; Dibangoye, J. S.; and Amato, C. 2010. Heuristic search for identical payoff bayesian games. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-10)*, 1115–1122.
- Peshkin, L.; Kim, K.; Meuleau, N.; and Kaelbling, L. 2000. Learning to cooperate via policy search. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI '00)*.
- Seuken, S. 2007. Memory-bounded dynamic programming for dec-pomdps. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, 2009–2015.
- Shani, G.; Brafman, R.; and Shimony, S. 2005. Model-based online learning of POMDPs. In *Proceedings of the European Conference on Machine Learning (ECML)*, volume Lecture Notes in Computer Science 3720, 353–364. Springer.
- Spaan, M. T. J.; Oliehoek, F. A.; and Amato, C. 2011. Scaling up optimal heuristic search in Dec-POMDPs via incremental expansion. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-11)*, 2027–2032.
- Sutton, R., and Barto, A. G. 1998. *Reinforcement Learning: An Introduction*. MIT Press.
- Szer, D., and Charpillet, F. 2006. Point-based dynamic programming for dec-pomdps. In *Proceedings of the 21st National Conference on Artificial Intelligence*, 1233–1238.
- Wu, F.; Zilberstein, S.; and Chen, X. 2010. Rollout sampling policy iteration for decentralized POMDPs. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI-10)*, 666–673.
- Zhang, C., and Lesser, V. 2011. Coordinated multi-agent reinforcement learning in networked distributed POMDPs. In *Proc. AAAI-11*.