

# Quantification of Transfer in Reinforcement Learning via Regret Bounds for Learning Agents

**Adrienne Tuynman**

Université Paris-Saclay, ENS Paris-Saclay  
France

adrienne.tuynman@ens-paris-saclay.fr

**Ronald Ortner**

Lehrstuhl für Informationstechnologie  
Montanuniversität Leoben  
Austria

rortner@unileoben.ac.at

## Abstract

We present an approach for the quantification of the usefulness of transfer in reinforcement learning via regret bounds for a multi-agent setting. Considering a number of  $N$  agents operating in the same Markov decision process, however possibly with different reward functions, we consider the regret each agent suffers with respect to an optimal policy maximizing her average reward. We show that when the agents share their observations the total regret of all agents is smaller by a factor of  $\sqrt{N}$  compared to the case when each agent has to rely on the information collected by herself. This result demonstrates how considering the regret in multi-agent settings can provide theoretical bounds on the benefit of sharing observations in transfer learning.

**Keywords:** transfer, reinforcement learning, regret

## 1. Introduction

Transfer learning is a wide and important area of machine learning with a lot of empirical studies but few theoretical guarantees available. This holds true even more in the specific area of transfer in reinforcement learning (RL). Accordingly, the two survey papers of [Taylor and Stone \(2009\)](#) and [Lazaric \(2012\)](#) are mainly dedicated to establishing a taxonomy of various settings as well as used performance criteria investigated in the literature. One particular problem is the quantification of the benefit of transfer in RL, for which there are hardly any general results. While [Taylor and Stone \(2009\)](#) discuss different criteria for the evaluation of transfer learning methods (including *total reward* and *transfer ratio* that will be closest to our alternative suggestion of considering the *regret*), these are rather discussed with respect to application to empirical studies.

This paper aims to take a step towards closing this gap by first suggesting a multi-agent setting with simultaneously learning agents in a Markov decision process (MDP) with shared information but different tasks. For the quantification of the usefulness of the shared information we propose to compare worst case regret bounds with respect to the optimal policies of each agent to the standard (single-agent) RL setting. We present an optimistic learning algorithm that is an adaptation of a single-agent RL algorithm to the multi-agent case and for which we can show regret bounds that are able to quantify the benefit of transfer in the considered setting. Implicitly our multi-agent algorithm also contributes to an issue pointed out by [Lazaric \(2012\)](#) who stated that “*the problem of how the exploration on one task should be adapted on the basis of the knowledge of previous related tasks is a problem which received little attention so far.*”

While [Taylor and Stone \(2009\)](#) mention multi-agent settings as an interesting application area for transfer (cf. also the related survey of [Stone and Veloso, 2000](#)), there is also a particular survey of [da Silva and Costa \(2019\)](#), which discusses transfer for multi-agent RL. In general, whereas we will concentrate on MDPs in which all agents share common transition functions that only depend on the chosen action by each individual agent, the survey considers more complex scenarios where transitions depend on the joint actions chosen by the agents. This also gives rise to different goals reaching from Nash equilibria and cooperative learning to adversarial settings in which one is interested in algorithms for a single self-interested agent.

Following the classification of multi-agent settings in (da Silva and Costa, 2019), we assume that all agents share the same state and action space and that the transferred knowledge is the experience of each single agent (also called *instance transfer*) that is shared with all other agents.

As already mentioned we are interested in regret bounds that measure for any time step  $T$  the difference of the accumulated rewards of an agent to the total reward an optimal policy (for the respective agent) would have achieved. We will see that while it is difficult to say how much benefit the experience of a single agent contributes to learning another task by a different agent, when considering the total regret over all agents the worst case regret bounds are smaller by a factor of  $\sqrt{\aleph}$  for a total of  $\aleph$  agents. Apart from the application to transfer, these are the first regret bounds for multi-agent settings in MDPs we are aware of.

The paper is organized as follows: The precise setting is introduced in the following Section 2. Then we provide our algorithm and present respective performance bounds in the main Section 3 before we conclude with a discussion in Section 4.

## 2. A Multi-Agent Learning Approach

In this section we will introduce a multi-agent learning setting. We first introduce the respective RL setting for a single agent however.

### 2.1 Preliminaries on Single-Agent-RL in MDPs

In the standard setting of RL, the learning agent acts in a Markov decision process (MDP).

**Definition 1** A Markov decision process  $M = (\mathcal{S}, \mathcal{A}, p, r, s_1)$  consists of a set of states  $\mathcal{S}$  and a set of actions  $\mathcal{A}$  available in each of the states.<sup>1</sup> When choosing an action  $a \in \mathcal{A}$  in state  $s \in \mathcal{S}$ , the agent obtains a random reward with mean  $r(s, a)$  and moves to a random new state  $s'$  that is determined by the transition probabilities  $p(s' | s, a)$ . The learning agent starts in an initial state  $s_1$ .

A *policy* in an MDP defines which actions are taken when acting in the MDP. In the following, we will only consider *stationary* policies  $\pi : \mathcal{S} \rightarrow \mathcal{A}$  that pick in each state  $s$  a fixed action  $\pi(s)$ . For the average reward criterion we will consider, it is sufficient to consider stationary policies (cf. Section 2.1.2 below).

#### 2.1.1 ASSUMPTIONS AND THE DIAMETER

In order to allow the agent to learn, we make further assumptions about the underlying MDP. First, with unbounded rewards it is difficult to learn (as it is easy to miss out a single arbitrarily large reward at some step), so we assume bounded rewards. Renormalizing, in the following we make the assumption that all rewards are contained in the unit interval  $[0, 1]$ .

Further, we assume that it is always possible to recover when taking a wrong action. This is possible if the MDP is *communicating*, that is, any state is reachable from any other state with positive probability when selecting a suitable policy. More precisely this means that given a communicating MDP and two states  $s, s'$  there is always a policy  $\pi$  such that the expected number of steps  $T_\pi(s, s')$  that an agent executing policy  $\pi$  needs to reach state  $s'$  when starting in state  $s$  is finite. A related parameter we will use in the following is the *diameter* of the MDP that measures the maximal distance between any two states in a communicating MDP.

**Definition 2** The diameter  $D$  in a communicating MDP  $M$  is defined as

$$D(M) := \max_{s, s' \in \mathcal{S}} \min_{\pi} T_\pi(s, s').$$

---

1. It is straightforward but notationally a bit awkward to generalize our setting as well as the results to the case where different actions are available in each of the states.

### 2.1.2 THE LEARNING GOAL AND THE NOTION OF REGRET

The goal of each learning agent is to maximize her accumulated reward after any  $T$  steps. For the sake of simplicity, in the following we will however consider the *average reward*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T r(s_t, a_t),$$

where  $s_t$  and  $a_t$  are state and action at step  $t$ , respectively. Unlike the accumulated  $T$ -step reward, the average reward is known to be maximized by a stationary policy  $\pi^*$  (Puterman, 1994). The optimal average reward  $\rho^*$  is a good proxy for the (expected) optimal  $T$ -step reward, as  $T\rho^*$  differs from the latter at most by a term that is of order  $D$ , as noted by Jaksch et al. (2010). Accordingly, we measure the performance of the learning agent by considering its *regret* after any  $T$  times steps defined as

$$R_T := T\rho^* - \sum_{t=1}^T r_t,$$

where  $r_t$  is the reward obtained by the agent at step  $t$ .

## 2.2 The Multi-Agent Setting

We are now finally ready to introduce the multi-agent setting. There are  $\aleph$  learning agents  $\alpha = 1, 2, \dots, \aleph$  all of which act in the same Markov decision process  $M$ , only that each agent  $\alpha$  has her individual reward function  $r^\alpha$ . We assume however that all agents share their experience. That is, each agent has immediately access to all the observations made by each other agent  $\alpha$ , in particular her history  $(s_t^\alpha, a_t^\alpha, r_t^\alpha)_{t \geq 0}$  specifying the state  $s_t^\alpha$  visited at step  $t$ , the action  $a_t^\alpha$  chosen and the reward  $r_t^\alpha$  obtained. (As each learning agent aims at maximizing her own rewards, knowledge of the other agents' rewards is only of use in the special case when all agents have the same reward function, cf. Theorem 3 below.) For each agent  $\alpha$  we consider her individual regret

$$R_T^\alpha := T\rho^{*,\alpha} - \sum_{t=1}^T r_t^\alpha,$$

where  $\rho^{*,\alpha}$  is the optimal average reward under the reward function  $r^\alpha$ .

### 2.2.1 MOTIVATION: INTERPRETATION OF REGRET BOUNDS IN TERMS OF TRANSFER LEARNING

As all agents share their information on the one hand and each of the agents acts to maximize her own rewards, the difference in regret between the single agent and the multi-agent setting specifies the usefulness of the information that comes from the other learning agents. In the simplest case with just two agents with different learning tasks (i.e., reward functions) in the same transition structure, the goal would be to quantify how useful the information collected by one agent for learning her task is for learning the other agent's task. While this seems to be difficult when considering each agent in isolation, the regret bounds we are going to derive will show the usefulness of sharing information for all agents simultaneously.

## 3. Regret Bounds for Multi-Agent Learning

In this section, we propose an algorithm for the introduced multi-agent setting and present upper bounds on the total regret of all learning agents.

### 3.1 UCRL2 with Shared Information

For our approach we propose a variant of the reinforcement learning algorithm UCRL2 (Jaksch et al., 2010) that is able to use shared information. That is, each agent will select her actions according to UCRL2 (briefly introduced below) however using not only information collected by herself but also by all other agents.

---

**Algorithm 1** Multi-agent-UCRL with shared information
 

---

- 1: **Input:** State space  $\mathcal{S}$ , action space  $\mathcal{A}$ , number of agents  $\aleph$ , confidence parameter  $\delta$ .
- 2: **Initialization:** Observe the initial state  $s_1$ .  
In the following, let  $t$  be the current time step.
- 3: **for** episodes  $k = 1, 2, \dots$  **do**
- 4: Compute estimates  $\hat{r}_k^\alpha(s, a)$  for rewards of each agent  $\alpha$  and estimates  $\hat{p}_k(\cdot | s, a)$  for transition probabilities (common to all agents).

**Compute policy  $\tilde{\pi}_k^\alpha$  for each agent  $\alpha$ :**

- 5: Let  $\mathcal{M}_k^\alpha$  be the set of plausible MDPs  $\tilde{M}$  with rewards  $\tilde{r}(s, a)$  and transition probabilities  $\tilde{p}(\cdot | s, a)$  close to the computed estimates

$$\begin{aligned} |\tilde{r}(s, a) - \hat{r}_k^\alpha(s, a)| &\leq \text{conf}_r^\alpha(s, a, S, A, \aleph, \delta, t), \\ \|\tilde{p}(\cdot | s, a) - \hat{p}_k(\cdot | s, a)\|_1 &\leq \text{conf}_p(s, a, S, A, \aleph, \delta, t). \end{aligned}$$

- 6: For each agent  $\alpha$  compute an MDP  $\tilde{M}_k^\alpha$  in  $\mathcal{M}_k^\alpha$  with

$$\rho^*(\tilde{M}_k^\alpha) = \max\{\rho^*(M) \mid M \in \mathcal{M}_k^\alpha\}.$$

Further, let  $\tilde{\pi}_k^\alpha$  be the respective optimal policy in  $\tilde{M}_k^\alpha$ .<sup>2</sup>

**Let each agent  $\alpha$  execute policy  $\tilde{\pi}_k^\alpha$ :**

- 7: Let each agent  $\alpha$ 
    - ▷ choose action  $a_t^\alpha = \tilde{\pi}_k^\alpha(s_t^\alpha)$ ,
    - ▷ obtain reward  $r_t^\alpha$ , and
    - ▷ observe state  $s_{t+1}^\alpha$ .
  - 8: If the number of visits of some agent in some state-action pair has been doubled since the start of the episode, terminate the current episode and start a new one.
  - 9: **end for**
- 

UCRL2 (sketched below) is a model-based algorithm that implements the idea of optimism in the face of uncertainty. Basically, UCRL2 uses the collected rewards as well as the transition counts for each state-action pair  $(s, a)$  in order to compute estimates  $\hat{r}$ ,  $\hat{p}$  for rewards and transition probabilities, respectively. While the original algorithm only considers a single agent, in the shared information setting the estimates are obviously computed from the collected observations of all agents. Based on the estimates and respective confidence intervals a set  $\mathcal{M}^\alpha$  of plausible MDPs is computed for each agent  $\alpha$ . Different choices of confidence intervals can be made, we are going to use

$$\text{conf}_r^\alpha(s, a, S, A, \aleph, \delta, t) := \sqrt{\frac{7 \log\left(\frac{2S\aleph t}{\delta}\right)}{2 \max\{1, N_t^\alpha(s, a)\}}}, \quad (1)$$

$$\text{conf}_p(s, a, S, A, \aleph, \delta, t) := \sqrt{\frac{14S \log\left(\frac{2At}{\delta}\right)}{\max\{1, N_t(s, a)\}}} \quad (2)$$

for the confidence intervals used by agent  $\alpha$  at step  $t$ , where  $N_t^\alpha(s, a)$  is the number of samples taken by agent  $\alpha$  in  $(s, a)$ , and  $N_t(s, a) := \sum_{\alpha=1}^{\aleph} N_t^\alpha(s, a)$  is the total number of samples taken by agents in  $(s, a)$ . As all agents have the same transition structure, but the reward function is individual, the confidence intervals for rewards depend only on the number of visits in  $(s, a)$  by the agent  $\alpha$ , while the confidence intervals for the transition probabilities depend on the number of samples shared by all agents.

Each agent  $\alpha$  then optimistically picks an MDP  $\tilde{M}^\alpha \in \mathcal{M}^\alpha$  and a policy  $\tilde{\pi}^\alpha$  that maximize the average reward. This policy is played until the number of visits in some state-action pair has been doubled for some agent, when a new policy is computed. The phases between computation and execution of a new policy are called *episodes*.

---

2. The computation of  $\tilde{\pi}_k^\alpha$  and  $\tilde{M}_k^\alpha$  can be done using extended value iteration as introduced in (Jaksch et al., 2010). For the proof of the regret bound we will assume that extended value iteration is employed with accuracy  $1/\sqrt{t}$ .

While the algorithm is described as a central algorithm controlling all  $\aleph$  agents, it could also be rewritten as a local algorithm for each single agent. As all the agents share their observations, this is straightforward.

### 3.2 Performance Guarantees

For the combined regret of all agents when using Multi-Agent-UCRL the following upper bound can be shown to hold.

**Theorem 1** *With probability  $1 - \delta$ , after any  $T$  steps the regret of all agents controlled by Multi-agent-UCRL with confidence intervals (1) and (2) is upper bounded by*

$$\sum_{\alpha=1}^{\aleph} R_T^\alpha \leq 15 \left( D\sqrt{S} + \sqrt{\aleph} \right) \sqrt{S \aleph T \log \left( \frac{8 \aleph T}{\delta} \right)}.$$

**Corollary 2** *If  $\aleph < D\sqrt{S}$ , the average regret per agent after  $T$  steps is  $\tilde{O} \left( \frac{DS\sqrt{AT}}{\sqrt{\aleph}} \right)$ .*

Corollary 2 shows that the performance guarantee improves over a simple sum over the regret bounds of the individual agents, as the regret of a single agent in the standard reinforcement learning setting for the same type of algorithm is  $\tilde{O}(DS\sqrt{AT})$ , see (Jaksch et al., 2010). Appendix A provides a sketch of the proof of Theorem 1, which is based on the analysis of Jaksch et al. (2010). The main reason for the improvement being possible is that learning the transition function in general is harder in the sense that the respective error of the corresponding estimates has a larger influence on the regret. As several agents are able to learn the common transition function faster, this results in improved bounds.

Further improvement is possible if the agents also have a joint reward function. Then all agents share the same MDP so that it makes sense using common confidence intervals also for the rewards that do not depend on the individual visits  $N^\alpha(s, a)$  of an agent  $\alpha$  but on the number of total visits  $N(s, a)$  of all agents. That is, confidence intervals

$$\text{conf}_r^\alpha(s, a, S, A, \aleph, \delta, t) := \sqrt{\frac{7 \log \left( \frac{2SAt}{\delta} \right)}{2 \max \{1, N_t(s, a)\}}} \quad (3)$$

can be used for the rewards. Accordingly, also the episode termination criterion should be changed to depend not on the visits of an individual agent but on the total number of visits. That is, an episode will be terminated if the total number of visits of all agents has been doubled since the start of the episode. With these modifications, one can obtain regret bounds of order  $\tilde{O}(DS\sqrt{ANT})$ .

**Theorem 3** *In a setting where all agents share the same reward function, using Multi-agent-UCRL with confidence intervals (3) and a modified episode termination criterion, with probability  $1 - \delta$  the regret of all agents after any  $T$  steps is upper bounded by*

$$\sum_{\alpha=1}^{\aleph} R_T^\alpha \leq 34DS\sqrt{ANT \log \left( \frac{8 \aleph T}{\delta} \right)}.$$

We highlight the necessary modifications in the proof of Theorem 1 to obtain Theorem 3 in Appendix B. By Theorem 3, Corollary 2 holds without any condition on the number of agents in the setting of shared rewards.

## 4. Discussion

The results presented in the previous section provide a mean for the quantification of the benefit of transfer. In general it is difficult to quantify how much worth are samples collected for one task when learning a new one. This is because it is possible that there are agents  $\alpha, \alpha'$  such that the samples collected by  $\alpha$  are of little use for  $\alpha'$ . That way it is in

general difficult to quantify the benefit of the samples collected by some agent  $\alpha$  for another agent  $\alpha'$  when learning her task. However when considering the total regret over all agents there is a collective benefit as our bounds on the regret show. This on the one hand confirms the difficulty of quantifying “local” transfer, but on the other hand shows that a general “global” transfer is possible.

The assumption of agents that learn simultaneously allows a straightforward adaptation of the analysis known for optimistic algorithms like UCRL. However, in principle the setting could be easily adapted to the case where the agents learn successively or with time delays. Such delays could also be considered to come from the communication between the agents, which in practice might not be immediate and could also happen in batches. While such modified settings of course demand adaptations in the analysis, these can be done without any major changes with respect to the results we have presented.

In general, our approach can serve as a blueprint for further investigations of multi-agent settings as a means for progress in transfer for RL in more general settings.

## References

- Felipe Leno da Silva and Anna Helena Reali Costa. A survey on transfer learning for multiagent reinforcement learning systems. *J. Artif. Intell. Res.*, 64:645–703, 2019.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *J. Mach. Learn. Res.*, 11:1563–1600, 2010.
- Alessandro Lazaric. Transfer in reinforcement learning: A framework and a survey. In Marco A. Wiering and Martijn van Otterlo, editors, *Reinforcement Learning*, volume 12 of *Adaptation, Learning, and Optimization*, pages 143–173. Springer, 2012.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1994.
- Peter Stone and Manuela M. Veloso. Multiagent systems: A survey from a machine learning perspective. *Auton. Robots*, 8(3):345–383, 2000.
- Matthew E. Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *J. Mach. Learn. Res.*, 10:1633–1685, 2009.

## Appendix A. Proof of Theorem 1

First, note that the difference between the observed rewards  $r_t^\alpha$  when agent  $\alpha$  chooses action  $a$  in state  $s$  at step  $t$  and the respective mean rewards  $r^\alpha(s, a)$  is a martingale difference sequence and hence can be bounded by Azuma-Hoeffding (cf. e.g. Lemma 10 of Jaksch et al., 2010), so that

$$\sum_{\alpha} \sum_t (r_t^\alpha - r^\alpha(s_t^\alpha, a_t^\alpha)) \leq \sqrt{\frac{5}{2} T \aleph \log \frac{8 T \aleph}{\delta}}$$

with probability at least  $1 - \left(\frac{\delta}{8 \aleph T}\right)^{5/4} \geq 1 - \frac{\delta}{12(\aleph T)^{5/4}}$ , similarly to eq. (7) of Jaksch et al. (2010). Accordingly, writing  $v_k^\alpha(s, a)$  for the number of visits of agent  $\alpha$  in  $(s, a)$  during episode  $k$  and defining  $\Delta_k^\alpha := v_k^\alpha(s, a)(\rho^{*,\alpha} - r^\alpha(s, a))$  to be the regret of agent  $\alpha$  in episode  $k$ , we can bound

$$\sum_{\alpha} R_T^\alpha \leq \sum_{\alpha} \sum_k \Delta_k^\alpha + \sqrt{\frac{5}{2} T \aleph \log \frac{8 \aleph T}{\delta}} \tag{4}$$

with probability  $1 - \frac{\delta}{12(\aleph T)^{5/4}}$ .

Analogously to App. B.1 of (Jaksch et al., 2010) using our slightly adapted confidence intervals (3) for the rewards, it can be shown that at each step  $t$  the probability that there is an agent  $\alpha$  whose true MDP  $M^\alpha$  is not contained in her set of plausible MDPs  $\mathcal{M}^\alpha(t)$  is bounded by  $\frac{\delta}{15t^6}$ . As shown in Sec. 4.2 of (Jaksch et al., 2010), this implies that the respective regret caused by failing confidence intervals is bounded by

$$\sum_{\alpha} \sum_k \Delta_k^\alpha \mathbb{1}\{M^\alpha \notin \mathcal{M}_k^\alpha\} \leq \aleph \sqrt{T} \quad (5)$$

with probability  $1 - \frac{\delta}{12(\aleph T)^{5/4}}$ .

For the episodes  $k$  in which  $M^\alpha \in \mathcal{M}_k^\alpha$ , one has analogously to eqs. (14)–(18) of (Jaksch et al., 2010) that

$$\begin{aligned} \sum_{\alpha} \sum_k \Delta_k^\alpha \mathbb{1}\{M^\alpha \notin \mathcal{M}_k^\alpha\} &\leq \\ &\left( \sqrt{14 \log \left( \frac{2S\aleph T}{\delta} \right)} + 2 \right) \sum_{\alpha} \sum_k \sum_{s,a} \frac{v_k^\alpha(s,a)}{\sqrt{\max\{1, N_{t_k}^\alpha(s,a)\}}} \\ &+ D \sqrt{14S \log \left( \frac{2\aleph T}{\delta} \right)} \sum_{\alpha} \sum_k \sum_{s,a} \frac{v_k^\alpha(s,a)}{\sqrt{\max\{1, N_{t_k}^\alpha(s,a)\}}} \\ &+ \sum_{\alpha} \sum_k \sum_{t=t_k}^{t_{k+1}-1} (p(\cdot | s_t^\alpha, a_t^\alpha) - \mathbf{e}_{s_t^\alpha}) \mathbf{w}_k^\alpha, \end{aligned} \quad (6)$$

where  $t_k$  denotes the initial time step of episode  $k$ ,  $\mathbf{e}_i$  is the unit vector with  $i$ -th coordinate 1 and all other  $S-1$  coordinates 0, and  $\mathbf{w}_k^\alpha$  is a modified value vector (computed by extended value iteration for episode  $k$ , cf. footnote 2) for agent  $\alpha$  with  $\|\mathbf{w}_k^\alpha\|_\infty \leq D/2$ .

The last term of (7) is a martingale difference sequence and another application of Azuma-Hoeffding gives (similar to eq. 19 of (Jaksch et al., 2010) but now for a sequence of length  $\aleph T$ )

$$\begin{aligned} \sum_{\alpha} \sum_k \sum_{t=t_k}^{t_{k+1}-1} (p(\cdot | s_t^\alpha, a_t^\alpha) - \mathbf{e}_{s_t^\alpha}) \mathbf{w}_k^\alpha &\leq \\ &D \sqrt{\frac{5}{2} \aleph T \log \left( \frac{8\aleph T}{\delta} \right)} + DS\aleph \log_2 \left( \frac{8T}{SA} \right) \end{aligned} \quad (8)$$

with probability at least  $1 - \frac{\delta}{12T^{5/4}}$ , where the last term comes from a bound over the number of episodes, similar to Appendix B.2 of (Jaksch et al., 2010) (but now counted extra for each agent).

Finally, for the two triple sums in (7), we have as shown in (Jaksch et al., 2010) that

$$\sum_k \sum_{s,a} \frac{v_k^\alpha(s,a)}{\sqrt{\max\{1, N_{t_k}^\alpha(s,a)\}}} \leq (1 + \sqrt{2}) \sum_{(s,a)} \sqrt{N_T^\alpha(s,a)}$$

and similarly

$$\sum_{\alpha} \sum_k \sum_{s,a} \frac{v_k^\alpha(s,a)}{\sqrt{\max\{1, N_{t_k}^\alpha(s,a)\}}} \leq (1 + \sqrt{2}) \sum_{(s,a)} \sqrt{N_T(s,a)}.$$

Observing that  $\sum_{s,a} N_T^\alpha(s,a) = T$  for each  $\alpha$  and  $\sum_{s,a} N_T(s,a) = \aleph T$  it follows by Jensen's inequality that

$$\sum_{\alpha} \sum_k \sum_{s,a} \frac{v_k^\alpha(s,a)}{\sqrt{\max\{1, N_{t_k}^\alpha(s,a)\}}} \leq (1 + \sqrt{2}) \aleph \sqrt{SAT} \quad (9)$$

and

$$\begin{aligned} \sum_{\alpha} \sum_k \sum_{s,a} \frac{v_k^\alpha(s,a)}{\sqrt{\max\{1, N_{t_k}^\alpha(s,a)\}}} &= \sum_k \sum_{s,a} \frac{\sum_{\alpha} v_k^\alpha(s,a)}{\sqrt{\max\{1, N_{t_k}^\alpha(s,a)\}}} \\ &\leq (1 + \sqrt{2}) \sqrt{S\aleph T}. \end{aligned} \quad (10)$$

In summary, we obtain from eqs. (4)–(10) that the total regret is bounded by

$$\begin{aligned} \sum_{\alpha} R_T^{\alpha} &\leq (D+1)\sqrt{\frac{5}{2}\aleph T \log \frac{8\aleph T}{\delta}} + \aleph\sqrt{T} + DSA\aleph \log_2\left(\frac{8T}{SA}\right) \\ &\quad + (1+\sqrt{2})\left(\sqrt{14 \log\left(\frac{2\aleph T}{\delta}\right)} + 2\right)\aleph\sqrt{SAT} \\ &\quad + (1+\sqrt{2})D\sqrt{14S \log\left(\frac{2\aleph T}{\delta}\right)}\sqrt{S\aleph T} \end{aligned}$$

with probability  $1 - \frac{\delta}{4T^{5/4}}$ . Summing over all  $T \geq 2$  shows that this bound holds simultaneously for all  $T \geq 2$  with probability at least  $1 - \delta$ .

It remains to simplify the bound. Summarizing terms, noting that  $DSA\aleph \log_2\left(\frac{8T}{SA}\right) \leq \frac{2}{34}DS\sqrt{\aleph AT \log \frac{\aleph T}{\delta}}$  if  $T \leq 34^2\aleph \log \frac{\aleph T}{\delta}$  (otherwise the theorem holds trivially, cf. App. B of [Jaksch et al., 2010](#)) we obtain

$$\sum_{\alpha} R_T^{\alpha} \leq 14\aleph\sqrt{SAT \log\left(\frac{2\aleph T}{\delta}\right)} + 15DS\sqrt{S\aleph T \log\left(\frac{8\aleph T}{\delta}\right)},$$

which completes the proof of Theorem 1.

## Appendix B. Proof of Theorem 3

The proof of Theorem 3 follows the same line as the proof of Theorem 1 with a few minor adaptations due to the modified confidence intervals (3) for the rewards, which however only causes slight changes in the constants. The main difference is that due to the modified confidence intervals, all terms linear in  $\aleph$  can be avoided:

First, all agents now share the same true MDP  $M$  as well as the set of plausible MDPs  $\mathcal{M}_k$ , so that the probability that  $M$  is not in  $\mathcal{M}_k^{\alpha}$  neither depends on the single agents nor on their number  $\aleph$ . The respective regret over all agents then is bounded by  $\sqrt{\aleph T}$  instead of  $\aleph\sqrt{T}$  as in (5), however with the same error probability.

Further, using the modified confidence intervals instead of eq. (9) one obtains

$$\begin{aligned} \sum_{\alpha} \sum_k \sum_{s,a} \frac{v_k^{\alpha}(s,a)}{\sqrt{\max\{1, N_{t_k}(s,a)\}}} &= \sum_k \sum_{s,a} \frac{v_k(s,a)}{\sqrt{\max\{1, N_{t_k}(s,a)\}}} \\ &\leq (1+\sqrt{2})\sqrt{S\aleph T}, \end{aligned}$$

which also replaces the linear dependence on  $\aleph$  by  $\sqrt{\aleph}$  and allows to summarize terms similar to the proof of the original regret bounds of [Jaksch et al. \(2010\)](#). Details of the derivation are straightforward and are therefore skipped.