

# Exploration Bonus for Regret Minimization in Undiscounted Discrete and Continuous Markov Decision Processes

**Jian Qian**

*Sequel Team - Inria Lille*

JIAN.QIAN@ENS.FR

**Ronan Fruit**

*Sequel Team - Inria Lille*

RONAN.FRUIT@INRIA.FR

**Matteo Pirotta\***

*Facebook AI Research*

PIROTTA@FB.COM

**Alessandro Lazaric**

*Facebook AI Research*

LAZARIC@FB.COM

**Editor:**

## Abstract

We introduce and analyse two algorithms for exploration-exploitation in discrete and continuous Markov Decision Processes (MDPs) based on *exploration bonuses*. SCAL<sup>+</sup> is a variant of SCAL (Fruit et al., 2018b) that performs efficient exploration-exploitation in any *unknown weakly-communicating* MDP for which an upper bound  $c$  on the span of the *optimal bias function* is known. For an MDP with  $S$  states,  $A$  actions and  $\Gamma \leq S$  possible next states, we prove that SCAL<sup>+</sup> achieves the same theoretical guarantees as SCAL (i.e., a high probability regret bound of  $\tilde{O}(c\sqrt{\Gamma SAT})$ ), with a much smaller computational complexity. Similarly, C-SCAL<sup>+</sup> exploits an exploration bonus to achieve sublinear regret in any undiscounted MDP with continuous state space. We show that C-SCAL<sup>+</sup> achieves the same regret bound as UCCRL (Ortner and Ryabko, 2012) while being the first implementable algorithm with regret guarantees in this setting. While optimistic algorithms such as UCRL, SCAL or UCCRL maintain a high-confidence set of plausible MDPs around the true unknown MDP, SCAL<sup>+</sup> and C-SCAL<sup>+</sup> leverage on an exploration bonus to directly plan on the empirically estimated MDP, thus being more computationally efficient.

## 1. Introduction

While learning in an unknown environment, a reinforcement learning (RL) agent must trade off the *exploration* needed to collect information about the dynamics and reward, and the *exploitation* of the experience gathered so far to gain as much reward as possible. The performance of an online learning agent is usually measured in terms of cumulative regret which compares the rewards accumulated by the agent with the rewards accumulated by an optimal agent. A popular strategy to deal with the exploration-exploitation dilemma (i.e., minimize regret) is to follow the *optimism in the face of uncertainty* (OFU) principle.

Optimistic approaches have been widely studied in the context of stochastic multi-armed bandit (MAB) problems. In this setting, OFU-based algorithms maintain optimistic estimates of the expected reward of each action  $a$  (i.e., arm), and play the action with highest

---

\*. Work done while at INRIA Lille (Sequel Team).

optimistic estimate (see e.g., Bubeck and Cesa-Bianchi, 2012; Lattimore and Szepesvári, 2018). These optimistic estimates are usually obtained by adding a high probability *confidence bound*  $b(a)$  to the empirical average reward  $\hat{r}(a)$  i.e.,  $\hat{r}(a) + b(a)$ . The confidence bound plays the role of an *exploration bonus*: the higher  $b(a)$ , the more likely  $a$  will be explored. As an example, based on Hoeffding’s inequality, the Upper-Confidence Bound (UCB) algorithm uses  $b(a) = \tilde{\Theta}(r_{\max}/\sqrt{N(a)})$  where  $N(a)$  is the total number of times action  $a$  has been played before and all rewards are assumed to lie between 0 and  $r_{\max}$  with probability 1. UCB can be shown to achieve nearly-optimal regret guarantees.

Strehl and Littman (2008) later generalized the idea of enforcing exploration by using a bonus on the reward to the RL framework. They analysed the *infinite-horizon  $\gamma$ -discounted setting* and introduced the Model Based Interval Estimation with Exploration Bonus (MBIE-EB) algorithm. MBIE-EB plays the optimal policy of the empirically estimated MDP where for each state-action pair  $(s, a)$ , a bonus  $b(s, a)$  is added to the empirical average reward  $\hat{r}(s, a)$  i.e., the immediate reward associated to  $(s, a)$  is  $\hat{r}(s, a) + b(s, a)$ . Unlike in MAB where the optimal arm is the one with maximal immediate reward, the goal of RL is to find a policy maximizing the cumulative reward i.e., the  $Q$ -function. Therefore, the bonus needs to account for the uncertainty in both the rewards and transition probabilities and so  $b(s, a) = \tilde{\Theta}\left(\frac{r_{\max}}{1-\gamma}\sqrt{\frac{1}{N(s, a)}}\right)$  where  $\frac{r_{\max}}{1-\gamma}$  is the range of the  $Q$ -function. Strehl and Littman (2008) also derived PAC guarantees on the sample complexity of MBIE-EB. More recently, *count-based methods* (e.g., Bellemare et al., 2016; Tang et al., 2017; Ostrovski et al., 2017; Martin et al., 2017) tried to combine the idea of MBIE-EB with Deep RL (DRL) techniques to achieve a good exploration-exploitation trade off in high dimensional problems. The exploration bonus usually used has a similar form  $\tilde{\Theta}\left(\frac{\beta}{\sqrt{N}}\right)$  where  $\beta$  is now an hyper-parameter tuned for the specific task at hand, and the visit count  $N$  is approximated using discretization (e.g., hashing) or density estimation methods.

Exploration bonuses have also been successfully applied to *finite-horizon problems* (Azar et al., 2017; Kakade et al., 2018; Jin et al., 2018). In this setting, the planning horizon  $H$  is known to the learning agent and the range of the  $Q$ -function is  $r_{\max}H$ . A natural choice for the bonus is then  $b(s, a) = \tilde{\Theta}(r_{\max}H/\sqrt{N(s, a)})$ . UCBV1.1 introduced by Azar et al. (2017) uses such a bonus and achieves near-optimal regret guarantees  $\tilde{O}(H\sqrt{SAT})$ . Extensions of UCBV1.1 exploiting the variance instead of the range of the  $Q$ -function achieve a better regret bound  $\tilde{O}(\sqrt{HSAT})$  (Azar et al., 2017; Kakade et al., 2018; Jin et al., 2018).

Both the finite horizon setting and infinite horizon discounted setting assume that there exists an *intrinsic horizon* (respectively  $H$  and  $\frac{1}{1-\gamma}$ ) known to the learning agent. Unfortunately, in many common RL problems it is not clear how to define  $H$  or  $\frac{1}{1-\gamma}$  and it is often desirable to set them as big as possible (e.g., in episodic problem, the time to the goal is not known in advance and random in general). As  $H$  tends to infinity the regret (of UCBV1.1, etc.) will become linear while as  $\gamma$  tends to 1 the sample complexity (of MBIE-EB, etc.) tends to infinity (not to mention the numerical instabilities that may arise). In this paper we focus on the much more natural infinite horizon undiscounted setting (Puterman, 1994, Chap. 8) which generalizes the two previous settings to the case where  $H \rightarrow +\infty$  and  $\gamma \rightarrow 1$  respectively. Several algorithms implementing the OFU principle in the infinite horizon undiscounted case have been proposed in the literature (e.g., Jaksch et al., 2010; Ortner and Ryabko, 2012; Fruit et al., 2017, 2018b; Talebi and Maillard, 2018), but none

of these approaches exploits the idea of an exploration bonus. Instead, they all construct an *extended* MDP<sup>1</sup> with continuous action space, which can be interpreted as the concatenation of all possible MDPs compatible with some high probability confidence bounds on the transition model, among which is the true MDP. The policy executed by the algorithm is the optimal policy of the extended MDP. UCRL (Jaksch et al., 2010) achieves a regret of order<sup>2</sup>  $\tilde{O}(r_{\max}D\sqrt{\Gamma SAT})$  after  $T$  time steps where  $D$ ,  $\Gamma$ ,  $S$  and  $A$  are respectively the diameter of the true MDP, the maximum number of reachable next states from any state, the number of states and the number of actions. (Fruit et al., 2018b) showed an improved bound for SCAL  $\tilde{O}(\min\{r_{\max}D, c\}\sqrt{\Gamma SAT})$  when a known upper bound on the optimal bias span  $c \geq sp\{h^*\}$  is known to the learning agent. Although such algorithms can be efficiently implemented in the tabular case, it is difficult to extend them to more scalable approaches like DRL. In contrast, as already mentioned, the exploration bonus approach is simpler to adapt to large scale problems and inspired count based methods in DRL.

In this paper we introduce and analyse SCAL<sup>+</sup>, the first algorithm that relies on an exploration bonus to efficiently balance exploration and exploitation in the infinite-horizon undiscounted setting. All the exploration bonuses that were previously introduced in the RL literature explicitly depend on  $\gamma$  or  $H$  which are known to the learning agent. In the infinite-horizon undiscounted case, there is no predefined parameter informing the agent about the range of the  $Q$ -function. This makes the design of an exploration bonus very challenging. To overcome this limitation, we make the same assumption as Bartlett and Tewari (2009); Fruit et al. (2018b) i.e., we assume that the agent knows an upper-bound  $c$  on the span (i.e., range) of the optimal bias (i.e., value function). The exploration bonus used by SCAL<sup>+</sup> is thus  $b(s, a) = \tilde{\Theta}(\max\{c, r_{\max}\}/\sqrt{N(s, a)})$ . In comparison, state-of-the-art algorithms in the infinite horizon undiscounted setting like UCRL or SCAL can, to a certain extent, be interpreted as virtually using an exploration bonus of order  $\tilde{\Theta}(r_{\max}D\sqrt{\Gamma/N(s, a)})$  and  $\tilde{\Theta}(\max\{c, r_{\max}\}\sqrt{\Gamma/N(s, a)})$  respectively. This is bigger by a multiplicative factor  $\sqrt{\Gamma}$ . As a result, to the best of our knowledge, SCAL<sup>+</sup> achieves a “*tighter*” optimism than any other existing algorithm in the infinite horizon undiscounted setting and is therefore less prone to *over-exploration*.

To further illustrate the generality of the exploration bonus approach, we also present C-SCAL<sup>+</sup>, an extension of SCAL<sup>+</sup> to continuous state space –but finite action space– MDPs. As in (Ortner and Ryabko, 2012; Lakshmanan et al., 2015), we require the reward and transition functions to be Hölder continuous with parameters  $L$  and  $\alpha$ . C-SCAL<sup>+</sup> is also the first implementable algorithm in continuous problem with theoretical guarantees (existing algorithms with theoretical guarantees such as UCCRL (Ortner and Ryabko, 2012) cannot be implemented). C-SCAL<sup>+</sup> combines the idea of SCAL<sup>+</sup> with state aggregation. Compared to SCAL<sup>+</sup>, the exploration bonus contains an additional term due to the discretization: for any aggregated state  $I$ ,  $b(I, a) = \tilde{O}\left(\max\{c, r_{\max}\}(1/\sqrt{N(I, a)} + LS^{-\alpha})\right)$ .

The main result of the paper is summarized in Thm. 1:

**Theorem 1** *For any MDP with  $S$  states,  $A$  actions and  $\Gamma$  next states, the regret of SCAL<sup>+</sup> is bounded with high probability by  $\tilde{O}\left(\max\{c, r_{\max}\}\sqrt{\Gamma SAT}\right)$ . For any “smooth” MDP with*

---

1. The extended MDP is sometimes called *bounded-parameter* MDP  
2. The original bound of Jaksch et al. (2010) has  $\sqrt{S}$  instead of  $\sqrt{\Gamma}$  but  $\sqrt{\Gamma}$  can be easily achieved by replacing Hoeffding inequality by empirical Bernstein’s inequality for transition probabilities.

smoothness parameters  $L$  and  $\alpha$ , 1-dimensional state space  $\mathcal{S} = [0, 1]$  and  $A$  actions, the regret of C-SCAL<sup>+</sup> is bounded with high probability by  $\tilde{O}\left(\max\{c, r_{\max}\}L\sqrt{AT}^{(\alpha+2)/(2\alpha+2)}\right)$ .

The regret bound of SCAL<sup>+</sup> (resp. C-SCAL<sup>+</sup>) matches the one of SCAL (UCCRL). Surprisingly, the tighter optimism introduced by SCAL<sup>+</sup> compared to SCAL and UCRL is not reflected in the final regret bound with the current statistical analysis ( $\sqrt{T}$  appears in the bound although despite not being included in the bonus). We isolate and discuss where the term  $\sqrt{T}$  appears in the proof sketch of Sect. 3.4. While Azar et al. (2017); Kakade et al. (2018); Jin et al. (2018) managed to remove the  $\sqrt{T}$  term in the finite horizon setting, it remains an open question whether their result can be extended to the infinite horizon case (for example, the two definitions of regret do not match and differ by a linear term) or it is an intrinsic difficulty of the setting. Finally, SCAL<sup>+</sup> and C-SCAL<sup>+</sup> are very appealing due to their simplicity and flexibility of implementation since the planning is performed on the empirical MDP (rather than on a much more complex extended MDP). This change of paradigm results in a more computationally efficient planning compared to UCRL and SCAL, as explained in Sec. 3.1.

## 2. Preliminaries

### 2.1 Markov Decision Processes

We consider a *weakly-communicating*<sup>3</sup> MDP (Puterman, 1994, Sec. 8.3)  $M = (\mathcal{S}, \mathcal{A}, p, r)$  with a set of states  $\mathcal{S}$  and a set of actions  $\mathcal{A}$ . For sake of clarity, here we consider a finite MDP  $M$  but all the stated concepts extend to the case of continuous state space, (see e.g., Ortner and Ryabko, 2012).

Each state-action pair  $(s, a)$  is characterized by a reward distribution with mean  $r(s, a)$  and support in  $[0, r_{\max}]$  as well as a transition probability distribution  $p(\cdot|s, a)$  over next states. We denote by  $S = |\mathcal{S}|$  and  $A = |\mathcal{A}|$  the number of states and action, and by  $\Gamma = \max_{s \in \mathcal{S}, a \in \mathcal{A}} \|p(\cdot|s, a)\|_0 \leq S$  the maximum support of all transition probabilities  $p(\cdot|s, a)$ . A stationary Markov randomized policy  $\pi : \mathcal{S} \rightarrow P(\mathcal{A})$  maps states to distributions over actions. The set of stationary randomized (resp. deterministic) policies is denoted by  $\Pi^{\text{SR}}$  (resp.  $\Pi^{\text{SD}}$ ). Any policy  $\pi \in \Pi^{\text{SR}}(M)$  has an associated *long-term average reward* (or gain) and a *bias function* defined as

$$g^\pi(s) := \lim_{T \rightarrow +\infty} \mathbb{E}_{\mathbb{Q}} \left[ \frac{1}{T} \sum_{t=1}^T r(s_t, a_t) \right], \quad h^\pi(s) := C\text{-}\lim_{T \rightarrow +\infty} \mathbb{E}_{\mathbb{Q}} \left[ \sum_{t=1}^T (r(s_t, a_t) - g^\pi(s_t)) \right],$$

where  $\mathbb{Q} := \mathbb{P}(\cdot|a_t \sim \pi(s_t); s_0 = s; M)$  and the bias  $h^\pi(s)$  measures the expected total difference between the reward and the stationary reward in *Cesaro-limit*<sup>4</sup> (denoted  $C\text{-lim}$ ). Accordingly, the difference of bias values  $h^\pi(s) - h^\pi(s')$  quantifies the (dis-)advantage of starting in state  $s$  rather than  $s'$  and we denote by  $sp\{h^\pi\} := \max_s h^\pi(s) - \min_s h^\pi(s)$  the *span* of the bias function. In weakly communicating MDPs, any optimal policy  $\pi^* \in$

---

3. In a weakly-communicating MDP, the set  $\mathcal{S}$  can be decomposed into two subsets: a *communicating* set in which for any pair of states  $s, s'$  there exists a policy that has a non-zero probability to reach  $s'$  starting from  $s$ , and a set of states that are *transient* under all policies.  
 4. For policies with an aperiodic chain, the standard limit exists.

$\arg \max_{\pi} g^{\pi}(s)$  has *constant* gain, i.e.,  $g^{\pi^*}(s) = g^*$  for all  $s \in \mathcal{S}$ . Moreover, there exists a policy  $\pi^* \in \arg \max_{\pi} g^{\pi}(s)$  for which  $(g^*, h^*) = (g^{\pi^*}, h^{\pi^*})$  satisfy the *optimality equation*

$$h^* = Lh^* - g^*e, \quad \text{where } e = (1, \dots, 1)^{\top}. \quad (1)$$

where  $L$  is the *optimal* Bellman operator:  $\forall v \in \mathbb{R}^S, s \in \mathcal{S}$ ,

$$Lv(s) := \max_{a \in \mathcal{A}} \{r(s, a) + p(\cdot|s, a)^{\top} v\} \quad (2)$$

Note that  $h^*$  is finite, i.e.,  $sp\{h^*\} < +\infty$ . Finally,  $D := \max_{(s, s') \in \mathcal{S} \times \mathcal{S}, s \neq s'} \{\tau(s \rightarrow s')\}$  denotes the diameter of  $M$ , where  $\tau(s \rightarrow s')$  is the minimal expected number of steps needed to reach  $s'$  from  $s$  in  $M$  (under any policy).

## 2.2 Planning under span constraint

In this section we introduce and analyse the problem of planning under bias span constraint, i.e., by imposing that  $sp\{h^{\pi}\} \leq c$ , for any policy  $\pi$ . This problem is at the core of the proposed algorithms (SCAL<sup>+</sup> and C-SCAL<sup>+</sup>) for exploration-exploitation. Formally, we define the optimization problem:

$$g_c^*(M) := \sup_{\pi \in \Pi_c(M)} \{g^{\pi}\}, \quad (3)$$

where  $M$  is any MDP (with discrete or continuous state space) s.t.  $\Pi_c(M) := \{\pi \in \Pi^{\text{SR}} : sp\{h^{\pi}\} \leq c \wedge sp\{g^{\pi}\} = 0\} \neq \emptyset$ .<sup>5</sup> This problem is a slight variation of the bias-span constrained problem considered by (Bartlett and Tewari, 2009; Ortner and Ryabko, 2012; Lakshmanan et al., 2015), for which no known-solution is available. On the other hand, problem 3 has been widely analysed by Fruit et al. (2018b).

Problem 3 can be solved using ScOPT (Fruit et al., 2018b), a version of (relative) value iteration (Puterman, 1994; Bertsekas, 1995), where the optimal Bellman operator is modified to return value functions with span bounded by  $c$ , and the stopping condition is tailored to return a constrained-greedy policy with near-optimal gain. Given  $v \in \mathbb{R}^S$  and  $c \geq 0$ , we define the value operator  $T_c : \mathbb{R}^S \rightarrow \mathbb{R}^S$  as

$$T_c v = \Gamma_c L v = \begin{cases} Lv(s) & \forall s \in \bar{\mathcal{S}}(c, v) \\ c + \min_s \{Lv(s)\} & \forall s \in \mathcal{S} \setminus \bar{\mathcal{S}}(c, v) \end{cases} \quad (4)$$

where  $\bar{\mathcal{S}}(c, v) = \{s \in \mathcal{S} | Lv(s) \leq \min_s \{Lv(s)\} + c\}$  and  $\Gamma_c$  is the span constrain projection operator (see (Fruit et al., 2018b, App. D) for details). In other words, operator  $T_c$  applies a *span truncation* to the one-step application of  $L$ , which guarantees that  $sp\{T_c v\} \leq c$ . Given a vector  $v_0 \in \mathbb{R}^S$  and a reference state  $\bar{s}$  ScOPT implements relative value iteration where  $L$  is replaced by  $T_c$ :  $v_{n+1} = T_c v_n - T_c v_n(\bar{s})e$ . We can now state the convergence guarantees of ScOPT (see Fruit et al., 2018b, Lem. 8 and Thm. 10).

**Proposition 2** *Let's assume that I) the optimal Bellman operator  $L$  is a  $\gamma$ -span-contraction; II) all policies are unichain; III) operator  $T_c$  is globally feasible at any vector  $v \in \mathbb{R}^S$  such that  $sp\{v\} \leq c$  i.e., for all  $s \in \mathcal{S}$ ,  $\min_{a \in \mathcal{A}} \{r(s, a) + p(\cdot|s, a)^{\top} v\} \leq \min_{s'} \{Lv(s')\} + c$ . Then:*

5. Fruit et al. (2018b, Lem. 2) showed that there may not exist a deterministic optimal policy for problem 3.

- (a) Optimality equation: *there exists a solution  $(g^+, h^+) \in \mathbb{R} \times \mathbb{R}^S$  to the optimality equation  $T_c h^+ = h^+ + g^+ e$ . Moreover, any solution  $(g^+, h^+)$  satisfies  $g^+ = g_c^*$ .*
- (b) Convergence: *for any initial vector  $v_0 \in \mathbb{R}^S$ ,  $\text{SCOPT}$  converges to a solution  $h^+$  of the optimality equation, and  $\lim_{n \rightarrow +\infty} T_c^{n+1} v_0 - T_c^n v_0 = g^+ e$ .*

### 2.3 Learning Problem

Let  $M^*$  be the true *unknown* MDP. We consider the learning problem where  $\mathcal{S}$ ,  $\mathcal{A}$  and  $r_{\max}$  are *known*, while rewards  $r$  and transition probabilities  $p$  are *unknown* and need to be estimated on-line. We evaluate the performance of a learning algorithm  $\mathfrak{A}$  after  $T$  time steps by its cumulative *regret*:  $\Delta(\mathfrak{A}, T) = Tg^* - \sum_{t=1}^T r_t(s_t, a_t)$ . Finally, we assume that the algorithm is provided with the knowledge of a constant  $c > 0$  such that  $sp\{h^*\} \leq c$ . This assumption has been widely used in the literature (see e.g., Ortner, 2008; Ortner and Ryabko, 2012; Fruit et al., 2018b) and, as shown by (Fruit et al., 2018a), it is necessary in order to achieve a logarithmic regret bound in weakly-communicating MDPs.

## 3. SCAL<sup>+</sup>: SCAL with exploration bonus

In this section, we introduce SCAL<sup>+</sup>, the first online RL algorithm –in the infinite horizon undiscounted setting– that leverages on an exploration bonus to achieve provable good regret guarantees. Similarly to SCAL (Fruit et al., 2018b), SCAL<sup>+</sup> takes advantage of the prior knowledge on the optimal bias span  $sp\{h^*\} \leq c$  through the use of SCOPT. In Sec. 3.1 we present the details of SCAL<sup>+</sup> and we give an explicit formula for the exploration bonus. We then show that all the conditions of Prop. 2 are satisfied for SCAL<sup>+</sup>, meaning that SCOPT can be used. Finally, we justify the choice of the bonus by showing that SCAL<sup>+</sup> is gain-optimistic (Sec. 3.3) and we conclude this section with the regret guarantees of SCAL<sup>+</sup> (Thm. 6) and a sketch of the regret proof.

### 3.1 The algorithm

SCAL<sup>+</sup> is a variant of SCAL that uses SCOPT to (approximately) solve (3) on MDP  $\widehat{M}_k^+$  at the beginning of each episode  $k$  (see Fig. 1).<sup>6</sup> Before defining  $\widehat{M}_k^+$  we need to introduce some notations and an intermediate MDP  $\widehat{M}_k$ .

Denote by  $t_k$  the starting time of episode  $k$ ,  $N_k(s, a, s')$  the number of observations of 3-tuples  $(s, a, s')$  before episode  $k$  ( $k$  excluded) and  $N_k(s, a) := \sum_{s'} N_k(s, a, s')$ . As in UCRL, we define the empirical averages  $\bar{r}_k(s, a)$  and  $\bar{p}_k(\cdot|s, a)$  by:

$$\bar{p}_k(s'|s, a) := \frac{N_k(s, a, s')}{N_k(s, a)} \quad \text{and} \quad \bar{r}_k(s, a) := \frac{1}{N_k(s, a)} \sum_{t=1}^{t_k-1} r_t(s_t, a_t) \mathbb{1}(s_t, a_t = s, a).$$

---

6. The algorithm is reported in its general form, which applies to both finite and continuous MDPs.

The exploration bonus is defined by aggregating the uncertainty on the reward and transition functions:

$$b_k(s, a) := (c + r_{\max}) \underbrace{\sqrt{\frac{7 \ln(2SA_t k / \delta)}{\max\{1, N_k(s, a)\}}}}_{:= \beta_k^{sa}} + \frac{c}{N_k(s, a) + 1}, \quad (5)$$

where  $\beta_k^{sa}$  is derived from Hoeffding-Azuma inequality. The application of SCOPT to the MDP defined by  $(\mathcal{S}, \mathcal{A}, \bar{p}_k, \bar{r}_k + b_k)$  will not lead to a solution of problem 3 in general since none of the three assumptions of Prop. 2 is met. To satisfy the first and second assumptions, we introduce MDP  $\widehat{M}_k := (\mathcal{S}, \mathcal{A}, \widehat{p}_k, \widehat{r}_k)$  where  $\widehat{r}_k(s, a) = \bar{r}_k(s, a) + b_k(s, a)$  for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ ,  $\bar{s} \in \mathcal{S}$  is an arbitrary reference state and

$$\widehat{p}_k(s'|s, a) = \frac{N_k(s, a)\bar{p}_k(s'|s, a)}{N_k(s, a) + 1} + \frac{\mathbb{1}(s' = \bar{s})}{N_k(s, a) + 1} \quad (6)$$

is a biased (but asymptotically consistent) estimator of the probability of transition  $(s, a) \mapsto s'$ . To satisfy the third assumption, we define the *augmented* MDP  $\widehat{M}_k^+ = (\mathcal{S}, \mathcal{A}^+, \widehat{p}_k^+, \widehat{r}_k^+)$  obtained by duplicating every action in  $\widehat{M}_k$  with transition probability unchanged and reward set to 0. Formally,  $\mathcal{A}^+ = \mathcal{A} \times \{1, 2\}$  and for the sake of clarity, any pair  $(a, i) \in \mathcal{A} \times \{1, 2\}$  is denoted by  $a_i$ . We then define  $\widehat{p}_k^+(s'|s, a_i) := \widehat{p}_k(s'|s, a)$  and  $\widehat{r}_k^+(s, a_i) := \widehat{r}_k(s, a) \cdot \mathbb{1}(i = 1)$ . In the next section we will verify that  $\widehat{M}_k^+$  satisfies all the assumptions of Prop. 2. Note that the policy  $\pi_k$  returned by SCOPT takes action in the *augmented* set  $\mathcal{A}^+$ . The projection on  $\mathcal{A}$  is simply  $\pi_k(s, a) \leftarrow \pi_k(s, a_1) + \pi_k(s, a_2)$ , for all  $s \in \mathcal{S}, a \in \mathcal{A}$  (we use the same notation for the two policies).  $\pi_k$  is executed until the episode ends i.e., until the number of visits in at least one state-action pair has doubled (see Fig. 1).

**Remark.** SCAL<sup>+</sup> only requires to plan on an empirical MDP with exploration bonus rather than an extended MDP (with continuous action space). This removes the burden of computing the best probability in a confidence interval which has a worst-case computational complexity linear in the number of states  $S$  (Jaksch et al., 2010, Sec. 3.1.2). Therefore, SCAL<sup>+</sup> is not only simpler to implement but also less computationally demanding. Furthermore, removing the optimistic step on the transition probabilities allows the exploration bonus scheme to be easily adapted to any MDP that can be efficiently solved (e.g., continuous smooth MDPs).

### 3.2 Requirements for ScOpt

We show that the three assumptions of Prop. 2 required from SCOPT to solve (3) for  $\widehat{M}_k^+$  are satisfied. The arguments are similar to those used by Fruit et al. (2018b, Sec. 6) for SCAL. We denote by  $\widehat{L}^+, \widehat{L}$  and  $L$  the optimal Bellman operators of  $\widehat{M}_k^+, \widehat{M}_k$  and  $M^*$  respectively. Similarly, we denote by  $\widehat{T}_c^+, \widehat{T}_c$  and  $T_c$  the truncated Bellman operators (Eq. 4) of  $\widehat{M}_k^+, \widehat{M}_k$  and  $M^*$  respectively.

**Contraction.** The small bias in the definition of  $\widehat{p}_k$  ensures that the “attractive” state  $\bar{s}$  is reached with non-zero probability from any state-action pair  $(s, a_i)$  implying that the *ergodic coefficient* of  $\widehat{M}_k^+$  defined as  $\gamma_k = 1 - \min_{a, b \in \mathcal{A}^+} \left\{ \sum_{j \in \mathcal{S}} \min\{\widehat{p}_k(j|s, a), \widehat{p}_k(j|u, b)\} \right\}$




Figure 1: Structure of SCAL<sup>+</sup> and C-SCAL<sup>+</sup>. For SCAL<sup>+</sup> by definition we have  $I(s) = s$ .

is smaller than  $1 - \min_{s,a} \left\{ \frac{N_k(s,a,\bar{s})+1}{N_k(s,a)+1} \right\} < 1$  and thus  $\widehat{L}^+$  (the Bellman operator of  $\widehat{M}_k^+$ ) is  $\gamma_k$ -contractive (Puterman, 1994, Thm. 6.6.6).

**Unichain.** By construction, the attractive state  $\bar{s}$  necessarily belongs to all *recurrent classes* of all policies implying that  $\widehat{M}_k^+$  is unichain (i.e., all policies are unichain).

**Global feasibility.** Let  $v \in \mathbb{R}^{\mathcal{S}}$  such that  $sp\{v\} \leq c$  and let  $(s^*, a_i^*) \in \mathcal{S} \times \mathcal{A}^+$  be such that  $\widehat{r}_k^+(s^*, a_i^*) + \widehat{p}_k^+(\cdot|s^*, a_i^*)^\top v = \min_{s \in \mathcal{S}} \{ \max_{a \in \mathcal{A}^+} \{ \widehat{r}_k^+(s, a) + \widehat{p}_k^+(\cdot|s, a)^\top v \} \}$ . For all  $(s, a_2) \in \mathcal{S} \times \mathcal{A}^+$  we have:

$$\widehat{p}_k^+(\cdot|s, a_2)^\top v - \widehat{p}_k^+(\cdot|s^*, a_i^*)^\top v \leq \max_{s \in \mathcal{S}} \{v(s)\} - \min_{s \in \mathcal{S}} \{v(s)\} = sp\{v\} \leq c$$

and  $\widehat{r}_k^+(s, a_2) = 0 \leq \widehat{r}_k^+(s^*, a_i^*)$ . Therefore, for all  $s \in \mathcal{S}$ ,  $\min_{a_j \in \mathcal{A}^+} \{ \widehat{r}_k^+(s, a_j) + \widehat{p}_k^+(\cdot|s, a_j)^\top v \} \leq \min_{s'} \{ \widehat{L}^+ v(s') \} + c$  implying that  $\widehat{T}_c^+$  is globally feasible at  $v$ .

### 3.3 Optimistic Exploration Bonus

All algorithms relying on the OFU principle (e.g., UCRL, OPT-PSRL, SCAL, etc.) have the property that the optimal gain of the MDP used for planning is an upper bound on the optimal gain of the true MDP  $g^*$ . This is a key step in deriving regret guarantees. If we want to use the same proof technique for SCAL<sup>+</sup>, we also have to ensure that the policy  $\pi_k$  is *gain-optimistic* (up to an  $\varepsilon_k$ -accuracy), i.e.,  $\widehat{g}_k^+ := g_c^*(\widehat{M}_k^+) \geq g^*$ . The exploration bonus was tailored to enforce this property. To prove gain-optimism we rely on the following proposition which is a direct consequence of Fruit et al. (2018b, Lem. 8):

**Proposition 3 (Dominance)** *If there exists  $(g, h)$  satisfying  $\widehat{T}_c^+ h \geq h + ge$  then  $\widehat{g}_k^+ \geq g$ .*

**Proof** By induction, using the monotonicity and linearity of  $\widehat{T}_c^+$  (Fruit et al., 2018b, Lemma 16 (a) & (c)), we have that  $\forall n \in \mathbb{N}$ ,  $(\widehat{T}_c^+)^{n+1} h \geq (\widehat{T}_c^+)^n h + ge$ . By Prop. 2,  $\lim_{n \rightarrow +\infty} (\widehat{T}_c^+)^{n+1} h - (\widehat{T}_c^+)^n h = \widehat{g}_k^+$ . Taking the limit when  $n$  tends to infinity in the



previous inequality yields:  $\widehat{g}_k^+ \geq g$ .  $\blacksquare$

Recall that the optimal gain and bias of the true MDP  $(g^*, h^*)$  satisfy the optimality equation  $Lh^* = h^* + g^*e$  (Sec. 2.1). Since in addition  $sp\{h^*\} \leq c$  (by assumption), we also have  $sp\{Lh^*\} = sp\{h^* + g^*e\} = sp\{h^*\} \leq c$  and so  $T_ch^* = Lh^*$ . According to Prop. 3, it is sufficient to show that  $\widehat{T}_c^+ h^* \geq h^* + g^*e = T_ch^*$  to prove optimism. Fruit et al. (2018b, Lemma 15) also showed that the span projection  $\Gamma_c$  (see Eq. 4) is monotone implying that a sufficient condition for  $\Gamma_c \widehat{L}^+ h^* = \widehat{T}_c^+ h^* \geq T_ch^* = \Gamma_c \widehat{L} h^*$  to hold is to have  $\widehat{L}^+ h^* \geq Lh^*$ . With our choice of bonus, this inequality holds with high probability (w.h.p) as a consequence of the following lemma:

**Lemma 4** *For all  $T \geq 1$  and  $k \geq 1$ , with probability at least  $1 - \frac{\delta}{15t_k^6}$ , for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  we have:  $\bar{r}_k(s, a) + b_k(s, a) + \widehat{p}_k(\cdot|s, a)^\top h^* \geq r(s, a) + p(\cdot|s, a)^\top h^*$ .*

**Proof** Hoeffding-Azuma inequality implies that with probability at least  $1 - \frac{\delta}{15t_k^6}$ , for all  $k \geq 1$  and for all pairs  $(s, a) \in \mathcal{S} \times \mathcal{A}$ ,  $|\bar{r}_k(s, a) - r(s, a)| \leq r_{\max} \beta_k^{sa}$  and  $|(\bar{p}_k(\cdot|s, a) - p(\cdot|s, a))^\top h^*| \leq c \beta_k^{sa}$ . Finally, we also need to take into account the small bias introduced by  $\widehat{p}_k(\cdot|s, a)$  compared to  $\bar{p}_k(\cdot|s, a)$  which is not bigger than  $c/(N_k(s, a) + 1)$  by definition.  $\blacksquare$

Denote by  $\widehat{L}$  the optimal Bellman operator of  $\widehat{M}_k$ . A direct implication of Lem. 4 is that  $\widehat{L}h^* \geq Lh^*$  w.h.p. Since by definition  $\widehat{p}_k^+(s'|s, a_1) = \widehat{p}_k^+(s'|s, a_2) = \widehat{p}_k(s'|s, a)$  and  $\widehat{r}_k^+(s, a_2) \leq \widehat{r}_k^+(s, a_1)$  it is immediate to see that  $\widehat{L}^+ h^* = \widehat{L}h^*$  implying that  $\widehat{L}^+ h^* \geq Lh^*$  w.h.p. As a result, we have the following desired property:

**Lemma 5** *For all  $T \geq 1$  and  $k \geq 1$ , with probability at least  $1 - \frac{\delta}{15t_k^6}$ ,  $\widehat{g}_k^+ = g_c^*(\widehat{M}_k) \geq g^*$ .*

**Remark.** Note that the argument used in this section to prove optimism (Lem. 5) significantly differs from the one used by Jaksch et al. (2010, UCRL) and Fruit et al. (2018b, SCAL). UCRL and SCAL compute a (nearly) optimal policy of an *extended* MDP that “contains” the true MDP  $M^*$  (w.h.p.). This immediately implies that the gain of the extended MDP is bigger than  $g^*$  (analogue property of Lem. 5). The main advantage of our argument compared to theirs is that it allows for a *tighter* optimism. To see why, note that the exploration bonus quantifies by how much  $\widehat{L}^+ h^*$  is bigger than  $Lh^*$  and approximately scales as  $b_k(s, a) = \widetilde{\Theta}\left(\max\{r_{\max}, c\}/\sqrt{N_k(s, a)}\right)$ . In contrast, UCRL and SCAL use an optimistic Bellman operator  $\widetilde{L}$  such that  $\widetilde{L}h^*$  is bigger than  $Lh^*$  by respectively  $\widetilde{\Theta}\left(r_{\max} D \sqrt{\Gamma/N_k(s, a)}\right)$  (UCRL) and  $\widetilde{\Theta}\left(\max\{r_{\max}, c\} \sqrt{\Gamma/N_k(s, a)}\right)$  (SCAL). In other words, the optimism in SCAL<sup>+</sup> is tighter by a multiplicative factor  $\sqrt{\Gamma}$ . A natural next step would be to investigate whether our argument could be extended to UCRL and SCAL in order to save  $\sqrt{\Gamma}$  for the optimism. We keep this open question for future work.

### 3.4 Regret Analysis of SCAL<sup>+</sup>

We now give the main result of this section:

**Theorem 6** *For any weakly communicating MDP  $M$  such that  $sp\{h_M^*\} \leq c$ , with probability at least  $1 - \delta$  it holds that for any  $T \geq 1$ , the regret of  $\text{SCAL}^+$  is bounded as*

$$\Delta(\text{SCAL}^+, T) = O\left(\max\{r_{\max}, c\} \left(\sqrt{\Gamma SAT \ln\left(\frac{T}{\delta}\right)} + S^2 A \ln^2\left(\frac{T}{\delta}\right)\right)\right).$$

Unlike SCAL,  $\text{SCAL}^+$  does not have a regret scaling with  $\min\{r_{\max}D, c\}$  implying that whenever  $c > D$ ,  $\text{SCAL}^+$  performs worse than UCRL. SCAL builds an extended MDP that contains the true MDP and therefore the shortest path in the extended MDP is shorter than the shortest path in the true MDP implying that  $sp\{\tilde{v}_k\} \leq r_{\max}D$  with  $\tilde{v}_k$  being the solution returned by extended value iteration (Thm. 4 of Bartlett and Tewari (2009)). Let  $v_k$  be the solution returned by  $\text{ScOPT}$  on  $\widehat{M}_k^+$ , it is not clear how to bound  $sp\{v_k\}$  other than using the prior knowledge  $c$ . This open question seems a lot related to the one of Sec. 3.3 (i.e., how to have a tighter optimism in UCRL) and we also keep it for future work.

**Proof sketch.** We now provide a sketch of the main steps of the proof of Thm. 6 (the full proof is reported in App. 3.4). In order to preserve readability, in the following, all inequalities should be interpreted up to minor approximations and in high probability. Let  $\nu_k(s, a)$  be the number of visits in  $(s, a)$  during episode  $k$  and  $m$  be the total number of episodes. By using the optimism proved in Sec. 3.3, we can decompose the regret as:

$$\Delta(\text{SCAL}^+, T) \lesssim \sum_{k=1}^m \sum_{s,a} \nu_k(s, a) \left( g_k - \sum_a r(s, a) \pi_k(s, a) \right) \quad (7)$$

where  $g_k = 1/2(\max\{\widehat{T}_k^+ v_k - v_k\} + \min\{\widehat{T}_k^+ v_k - v_k\})$  and  $(v_k, \pi_k)$  is the solution of  $\text{ScOPT}$ . The stopping condition of  $\text{ScOPT}$  applied to  $\widehat{M}_k^+$  is such that (after manipulation):  $g_k \leq \sum_a \pi_k(s, a) (\widehat{r}_k(s, a) + \widehat{p}_k(\cdot|s, a)^\top v_k) - v_k(s) + \varepsilon_k$ . By plugging this inequality into (7) we obtain two terms:  $\bar{r}_k(s, a) - r(s, a) + b_k(s, a)$  and  $(\widehat{p}_k(\cdot|s, a) - e_s)^\top v_k$ . We can further decompose the scalar product as  $(\widehat{p}_k(\cdot|s, a) - p(\cdot|s, a))^\top v_k + (p(\cdot|s, a) - e_s)^\top v_k$ . The second terms is negligible in the final regret since it is of order  $\widetilde{O}(c\sqrt{T} + cSA)$  when summed over  $\mathcal{S}, \mathcal{A}$  and episodes (Fruit et al., 2018b, Eq. 56). On the other hand, the term  $(\widehat{p}_k(\cdot|s, a) - p(\cdot|s, a))^\top v_k$  is the dominant term of the regret and represents the error of using the estimated  $\widehat{p}_k$  in place of  $p$  in a step of value iteration. As shown in Sec. 3.3, we can start bounding the error of using  $\widehat{p}_k$  in place of  $\bar{p}_k$  by  $c/(N_k(s, a) + 1)$ . The remaining term is thus  $(\bar{p}_k - p)^\top v_k$ . Since  $v_k$  depends on  $\bar{p}_k$ , we cannot apply Hoeffding-Azuma inequality as done in Sec. 3.3 for the design of  $b_k$ . Instead we use a worst-case approach and bound separately  $\|\bar{p}_k(\cdot|s, a) - p(\cdot|s, a)\|_1 \lesssim \sqrt{\Gamma} \beta_k^{sa}$  and  $sp\{v_k\} \leq c$  which will introduce a  $\sqrt{\Gamma}$  factor (by using Bernstein-Freedman inequality instead of Hoeffding-Azuma inequality). It is worth pointing out that  $\Gamma$  only appears due to statistical fluctuations that we cannot control, and not from the optimism (i.e., exploration bonus) that is explicitly encoded in the algorithm. Concerning the reward, as shown in Sec. 3.3, we have that  $|\bar{r}(s, a) - r(s, a)| \leq r_{\max} \beta_k^{sa}$ . As a consequence, we can approximately write that:

$$\Delta(\text{SCAL}^+, T) \lesssim \sum_{k=1}^m \sum_{s,a} \nu_k(s, a) \pi_k(s, a) \left( b_k(s, a) + \underbrace{(c\sqrt{\Gamma} + r_{\max}) \beta_k^{sa} + c/(N_k(s, a) + 1)}_{:=d_k(s, a)} \right)$$

The proof follows by noticing that  $d_k(s, a) + b_k(s, a) \leq 2d_k(s, a)$ , thus all the remaining terms can be bounded as in (Fruit et al., 2018b).

**Remarks.** Given the fact that the optimism in SCAL<sup>+</sup> is tighter than in SCAL by a factor  $\sqrt{\Gamma}$ , one might have expected to get a regret bound scaling as  $c\sqrt{SAT}$  instead of  $c\sqrt{STAT}$ , thus matching the lower bound of Jaksch et al. (2010) as for the dependency in  $S$ .<sup>7</sup> Unfortunately, such a bound seems difficult to achieve with SCAL<sup>+</sup> (and even SCAL) for the reason explained above.

On the other side, (Azar et al., 2017; Kakade et al., 2018) achieved such an optimal dependence in finite-horizon problems. The main issue in extending such results is the different definition of the regret: their regret is defined as the difference between the value function at episode  $k$  and the optimal one. It is not clear how to map their definition to ours without introducing a linear term in  $T$ . Concerning infinite-horizon undiscounted problems, Agrawal and Jia (2017) claimed to have obtained the optimal dependence in their optimistic posterior sampling approach. To achieve such goal, they exploited the fact that  $|(\bar{p}_k(\cdot|s, a) - p(\cdot|s, a))^T \tilde{v}_k| \lesssim r_{\max} D \beta_k^{s_a}$ . Unfortunately, as explained above, it is not possible to achieve such tight concentration by using a worst-case argument, as they do. As a result, optimistic PSRL would have a regret scaling as  $D\sqrt{STAT}$ , while the improved bound in (Agrawal and Jia, 2017) should be rather considered as a conjecture.<sup>8</sup>

#### 4. C-SCAL<sup>+</sup>: SCAL<sup>+</sup> for continuous state space

We now consider an MDP with continuous state space  $\mathcal{S} = [0, 1]$  and discrete action space  $\mathcal{A}$ . In general, it is impossible to approximate an arbitrary function with only a finite number of samples. As a result, we introduce the same smoothness assumption as Ortner and Ryabko (2012) (Hölder continuity):

**Assumption 7** *There exist  $L, \alpha > 0$  s.t. for any two states  $s, s' \in \mathcal{S}$  and any action  $a \in \mathcal{A}$ :*

$$|r(s, a) - r(s', a)| \leq r_{\max} L |s - s'|^\alpha \quad \text{and} \quad \|p(\cdot|s, a) - p(\cdot|s', a)\|_1 \leq L |s - s'|^\alpha$$

Similarly to Sec. 3 we start presenting C-SCAL<sup>+</sup>, the variant of SCAL<sup>+</sup> for continuous state space, and then we provide its theoretical guarantees (see Sec. 4.2).

##### 4.1 The algorithm

In order to apply SCAL<sup>+</sup> to a continuous problem, a natural idea is to discretize the state space as is done by Ortner and Ryabko (2012). We therefore partition  $\mathcal{S}$  into  $S$  intervals defined as  $I_1 := [0, \frac{1}{S}]$  and  $I_k = (\frac{k-1}{S}, \frac{k}{S}]$  for  $k = 2, \dots, S$ . The set of “aggregated” states is then  $\mathcal{I} := \{I_1, \dots, I_S\}$  ( $|\mathcal{I}| = S$ ). As can be expected, we will see that the number of intervals  $S$  will play a central role in the regret. Note that the terms  $N_k(s, a, s')$  and  $N_k(s, a)$  defined in Sec. 3 are still well-defined for  $s$  and  $s'$  lying in  $[0, 1]$  but are 0 except for a finite number of  $s$  and  $s'$  (see Def. 9). For any subset  $I \subseteq \mathcal{S}$ , the sum  $\sum_{s \in I} u_s$  is also well-defined as long as the collection  $(u_s)_{s \in I}$  contains only a finite number of non-zero elements. We can

7. From an algorithmic perspective we achieve the optimal dependence on  $S$ , although this is not reflecting in the regret bound.

8. The problem has been acknowledged by the authors via personal communication.

therefore define the *aggregated* counts, rewards and transition probabilities for all  $I, J \in \mathcal{I}$  as:  $N_k(I, a) := \sum_{s \in I} N_k(s, a)$ ,

$$\bar{r}_k^{ag}(I, a) := \frac{1}{N_k(s, a)} \sum_{t=1}^{t_k-1} r_t(s_t, a_t) \mathbb{1}(s_t \in I, a_t = a), \quad \bar{p}_k^{ag}(J|I, a) := \frac{\sum_{s' \in J} \sum_{s \in I} N_k(s, a, s')}{\sum_{s \in I} N_k(s, a)}.$$

Similarly to the discrete case, we define the exploration bonus of an aggregated state as

$$b_k(I, a) := (c + r_{\max}) (\beta_k^{Ia} + LS^{-\alpha}) + \frac{c}{N_k(I, a) + 1} \quad (8)$$

While  $\beta_k^{Ia}$  is defined as in (5) on the discrete “aggregated” MDP, the main difference with the discrete bonus (5) is an additional  $O(cLS^{-\alpha})$  term that accounts for the fact that the states that we aggregate are not completely identical but have parameters that differ by at most  $LS^{-\alpha}$ . We pick an arbitrary reference aggregated state  $\bar{I}$  and define  $\widehat{M}_k^{ag} = (\mathcal{I}, \mathcal{A}, \widehat{p}_k^{ag}, \widehat{r}_k^{ag})$  the “aggregated” (discrete) analogue of  $\widehat{M}_k$  defined in Sec. 3, where  $\widehat{r}_k^{ag} = \bar{r}_k^{ag} + b_k$  and

$$\widehat{p}_k^{ag}(J|I, a) := \frac{N_k(I, a) \bar{p}_k^{ag}(J|I, a)}{N_k(I, a) + 1} + \frac{\mathbb{1}(J = \bar{I})}{N_k(I, a) + 1},$$

Similarly we “augment”  $\widehat{M}_k^{ag}$  into  $\widehat{M}_k^{ag+} = (\mathcal{I}, \mathcal{A}^+, \widehat{p}_k^{ag+}, \widehat{r}_k^{ag+})$  (analogue of  $\widehat{M}_k^+$  in Sec. 3) by duplicating each transition in  $\widehat{M}_k^{ag}$  with the transition probability unchanged and the reward set to 0.

At each episode  $k$ , C-SCAL<sup>+</sup> uses SCOPT (with the same parameters as in Sec. 3) to solve optimization problem (3) on  $\widehat{M}_k^{ag+}$ . This is possible because although the state space of  $M^*$  is infinite,  $\widehat{M}_k^{ag+}$  has only  $S < +\infty$  states. SCOPT returns an *optimistic* (nearly) optimal policy  $\pi_k$  satisfying the span constraint. This policy is defined in the aggregated discrete state space but can easily be extended to the continuous case as  $\pi_k(s, a) = \pi_k(I(s), a)$  for any  $(s, a)$ . Policy  $\pi_k$  is executed until the end of the episode (see Alg. 1).

## 4.2 Regret Analysis of C-SCAL<sup>+</sup>

This section is devoted to the analysis of C-SCAL<sup>+</sup>. We start providing the regret bound:

**Theorem 8** *For any continuous MDP  $M$  with state space  $\mathcal{S} \in [0, 1]$  and  $A$  actions such that  $sp\{h_M^*\} \leq c$ , with probability at least  $1 - \delta$  it holds that for any  $T \geq 1$ , the regret of C-SCAL<sup>+</sup> is bounded as*

$$\Delta(\text{C-SCAL}^+, T) = O \left( \max\{r_{\max}, c\} \left( S \sqrt{AT \ln \left( \frac{T}{\delta} \right)} + S^2 A \ln^2 \left( \frac{T}{\delta} \right) + LS^{-\alpha} T \right) \right)$$

For  $T \geq L^{2/\alpha} A$  and by setting  $S = \left( \alpha L \sqrt{\frac{T}{A}} \right)^{1/(\alpha+1)}$ , the regret is bounded w.h.p. as

$$\Delta(\text{C-SCAL}^+, T) = \widetilde{O} \left( \max\{r_{\max}, c\} L^{1/(\alpha+1)} A^{\alpha/(2\alpha+2)} T^{(\alpha+2)/(2\alpha+2)} \right)$$

Thm. 8 shows that C-SCAL<sup>+</sup> achieves the same regret of UCCRL (Ortner and Ryabko, 2012) while being the only implementable algorithm for regret minimization in this setting. It is worth to mention that it is possible to exploit the recent advances in the literature (see Fruit et al., 2018b) in order to derive an implementable variant of UCCRL (based on SCAL). However, this new algorithm will still require to plan on an extended MDP making it less computation efficient than C-SCAL<sup>+</sup>, while having the same regret bound. Moreover, note that the analysis can be extended to the general  $d$ -dimensional case. As pointed out by (Ortner and Ryabko, 2012),  $S^d$  intervals are used for the discretization leading to a regret bound of  $\tilde{O}(T^{(2d+\alpha)/(2d+2\alpha)})$  when  $S = T^{1/(2d+2\alpha)}$ . Finally, we believe that C-SCAL<sup>+</sup> can be extended to the setting considered by (Lakshmanan et al., 2015) where, in addition to Hölder conditions, the transition function is assumed to be  $\kappa$ -times smoothly differentiable. In the case of Lipschitz model, i.e.,  $\alpha = 1$ , this means that it is possible obtain an asymptotic regret (as  $\kappa \rightarrow \infty$ ) of  $\tilde{O}(T^{2/3})$  while C-SCAL<sup>+</sup> is achieving  $\tilde{O}(T^{3/4})$ . We leave the derivation of this variant for future work.

**Proof sketch.** The continuous case considered in this section can be interpreted as a generalization of the discrete case, thus presenting more challenges. The *main technical challenge* is to be able to compare the solution of problem (3) on  $\widehat{M}_k^{ag+}$  (discrete state space MDP) with the solution of (3) on  $M^*$  (continuous state space MDP) and thus prove the optimism. We start introducing an intermediate empirical continuous MDP  $\widehat{M}_k$  that will be used in the rest of the proof.

**Definition 9 (Estimated continuous MDP)** Let  $\widehat{M}_k = (\mathcal{S}, \mathcal{A}, \widehat{p}_k, \widehat{r}_k)$  be the continuous state space MDP s.t. for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$

$$\begin{aligned}\widehat{r}_k(s, a) &:= \bar{r}_k(s, a) + b_k(s, a) = \bar{r}_k^{ag}(I(s), a) + b_k(I(s), a) \\ \widehat{p}_k(s'|s, a) &:= \frac{N_k(I(s), a)\bar{p}_k(s'|s, a)}{N_k(I(s), a) + 1} + \frac{S \cdot \mathbb{1}(s' \in I(\bar{s}))}{N_k(I(s), a) + 1}.\end{aligned}$$

where  $I : \mathcal{S} \rightarrow \mathcal{I}$  is the function mapping a state  $s$  to the interval containing  $s$  and the term  $\bar{p}_k(s'|s, a)$  is the Radon-Nikodym derivative of the cumulative density function  $F(s) = \sum_{s' \leq s} \frac{\sum_{x \in I(s)} N_k(x, a, s')}{N_k(I(s), a)}$ , meaning that for any measurable function  $f$  and any measurable set  $Z \subseteq [0, 1]$ ,  $\int_Z \bar{p}_k(s'|s, a) f(s') ds' = \sum_{s' \in Z} \frac{\sum_{x \in I(s)} N_k(x, a, s')}{\sum_{x \in I(s)} N_k(x, a)} f(s')$ .

This MDP is one of the possible instances of continuous MDP that, when aggregated over the interval set  $\mathcal{I}$ , matches the discrete MDP  $\widehat{M}_k^{ag}$ .<sup>9</sup> In particular, by definition,  $\forall J \in \mathcal{I}$ ,  $\int_J \bar{p}_k(s'|s, a) ds' = \bar{p}_k^{ag}(J|s, a) := \bar{p}_k^{ag}(J|I(s), a)$  and  $\forall (s, J) \in \mathcal{S} \times \mathcal{I}$ :

$$\int_J \widehat{p}_k(s'|s, a) ds' = \int_J \frac{N_k(I(s), a)\bar{p}_k(s'|s, a)}{N_k(I(s), a) + 1} ds' + \frac{S \int_J \mathbb{1}(s' \in I(\bar{s})) ds'}{N_k(I(s), a) + 1} = \widehat{p}_k^{ag}(J|I(s), a) \quad (9)$$

We leverage this definition to prove an analogous of Lem. 4 for the continuous case.

9. It seems that an intermediate (*extended*) continuous MDP is also used in the proof of UCCRL but never formally defined, leaving to the reader the need of interpreting the properties of this MDP. Due to the lack of rigorous definition, few steps in the regret proof of UCCRL are not completely clear.

**Lemma 10** For all  $T \geq 1$  and  $k \geq 1$ , with probability at least  $1 - \frac{\delta}{15t_k^6}$ , for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  we have:  $\underbrace{\bar{r}_k(s, a) + b_k(I(s), a) + \int_{\mathcal{S}} \hat{p}_k(s'|s, a)h^*(s')ds'}_{:=\hat{r}_k(s, a)} \geq r(s, a) + \int_{\mathcal{S}} p(s'|s, a)h^*(s')ds'$

**Proof** The main and crucial difference in the proof is that due to the aggregation of states,  $\bar{p}_k$  and  $\bar{r}_k$  do not statistically concentrate around the true values  $p$  and  $r$ . To overcome this problem we decompose  $\hat{p}_k - p$  into the sum of three terms  $(\hat{p}_k - \bar{p}_k) + (\bar{p}_k - \tilde{p}_k) + (\tilde{p}_k - p)$  with  $\tilde{p}_k(s'|s, a) := \frac{1}{N_k(I(s), a)} \sum_{x \in I(s)} N_k(x, a)p(s'|x, a)$ . We show that  $\int_{\mathcal{S}} (\tilde{p}_k(s') - p(s'))h^*(s')ds' = O(LS^{-\alpha})$  (Asm. 7: smoothness assumption) while  $\int_{\mathcal{S}} (\hat{p}_k(s') - \bar{p}_k(s'))h^*(s')ds' = O(1/N_k)$  (biased estimator). The term  $\int_{\mathcal{S}} (\bar{p}_k(s') - \tilde{p}_k(s'))h^*(s')ds'$  can be bounded using concentration inequalities but requires more work than in the discrete case. In the discrete case, for a given state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , the difference  $(\bar{p}_k(\cdot|s, a) - \tilde{p}_k(\cdot|s, a))^T h^*$  is usually interpreted as the deviation of a sum of independent random variables from its expectation (Lattimore and Szepesvári, 2018, Section 4.4) and can be bounded using Hoeffding inequality. Since there is only a finite number of possible  $(s, a)$ , it is possible to take a union bound over state-action pairs. In the continuous case, the difference  $\int_{\mathcal{S}} (\bar{p}_k(s') - \tilde{p}_k(s'))h^*(s')ds'$  does not just depend on a single  $s$  but on the (random) set of states belonging to a given interval  $I \in \mathcal{I}$  that have been visited. There is an uncountable number of possible such sequences of states and so we can not use a union bound argument. Instead, we rely on a martingale argument and Azuma inequality for the proof. We decompose  $\hat{r}_k - r$  as  $(\hat{r}_k - \tilde{r}_k) + (\tilde{r}_k - r)$  with  $\tilde{r}_k(s, a) := \frac{1}{N_k(I(s), a)} \sum_{x \in I(s)} N_k(x, a)r(x, a)$  and proceed similarly for the reward. The detailed proof can be found in App. A (Lem. 12).  $\blacksquare$

Note that, as a consequence of Lem. 10, we have that  $\hat{L}h^* \geq Lh^*$  w.h.p. where  $\hat{L}$  is the optimal Bellman operator of the continuous empirical MDP  $\widehat{M}_k$ . This, together with Prop. 3, is sufficient to prove that the exploration bonus in (8) makes  $\widehat{M}_k$  optimistic w.r.t.  $M^*$ :  $g_c^*(\widehat{M}_k) \geq g^*$  w.h.p. We cannot directly extend this argument to prove optimism for  $\widehat{M}_k^{ag+}$  since the aggregated MDPs lie in a different state space. The key property used in this setting is that the  $n$ -times application of  $\widehat{L}^{ag}$  (optimal Bellman operator of  $\widehat{M}_k^{ag}$ ) and  $\widehat{L}$  to a constant vector are identical.<sup>10</sup> As a consequence, we can prove that  $\widehat{M}_k^{ag+}$  is optimistic:

**Lemma 11** For all  $T \geq 1$  and  $k \geq 1$ , with probability at least  $1 - \frac{\delta}{15t_k^6}$ ,  $\widehat{g}_k^+ = g_c^*(\widehat{M}_k^{ag+}) \geq g^*$ .

**Proof** We start noticing that, starting from a value  $v_0(s) = 0$  ( $\forall s$ ), the application of the Bellman operator of  $\widehat{M}_k^{ag}$  and  $\widehat{M}_k$  is such that  $(\widehat{L}^{ag})^n v_0(s) = (\widehat{L})^n v_0(s)$ ,  $\forall s \in \mathcal{S}$ ,  $n > 0$ . This is due to the fact that  $v_n$  is constant over any interval  $I$  (for any  $n$ ) and (9) holds:

$$\forall n \geq 0, \quad \int_{s \in \mathcal{S}} \hat{p}_k(s'|s, a)v_n(s')ds' = \sum_{J \in \mathcal{I}} v_n(J) \int_J \hat{p}_k(s'|s, a)ds' = \sum_{J \in \mathcal{I}} v_n(J) \hat{p}_k^{ag}(J|I(s), a)$$

Then,  $\forall n, s$ ,  $\widehat{T}_c^{ag} v_n(s) = \widehat{T}_c v_n(s)$  implying  $g_c^*(\widehat{M}_k^{ag}) = g_c^*(\widehat{M}_k)$ . The optimism of  $\widehat{M}_k^{ag+}$  follows by the fact that the ‘‘augmentation’’ does not impact the gain:  $g_c^*(\widehat{M}_k^{ag+}) =$

<sup>10</sup>. In the App. A we show that  $(\widehat{L}^{ag})^n v_0 = (\widehat{L})^n v_0$  for any piecewise constant vector over intervals in  $\mathcal{I}$ .

$g_c^*(\widehat{M}_k^{ag})$  (Fruit et al., 2018b, Lem. 20). The detailed proof can be found in App. A (Lem. 13). ■

**Remark.** The proof of optimism does not seem as straightforward as suggested by Ortner and Ryabko (2012) (regret proof of UCCRL). They use an informal “inclusion” argument (i.e.,  $M^*$  is included in the discretized extended MDP used for planning) which seems not easy to formally prove since the true and “optimistic” MDPs are of different nature (the true MDP has a continuous state space unlike the optimistic one which is discretized). Overall, we believe that an additional contribution of this paper is to provide a more rigorous analysis of the continuous case compared to the existing literature.

We have now all the key properties to apply the same regret analysis stated for the discrete case (with several technical arrangements to deal with the continuous case) in order to prove Thm. 8. The complete proof can be found in App. A.3.

## 5. Conclusion

In this paper we provide the first analysis of exploration bonus in infinite-horizon undiscounted problems, a more challenging setting than the finite-horizon (see e.g., Azar et al., 2017; Kakade et al., 2018; Jin et al., 2018) and discounted (Strehl and Littman, 2008). Compared to these other settings, we also extended the analysis to the continuous case and we provided the first implementable and efficient (no need to plan on an extend MDP) algorithm. We finally showed, through a formal derivation of the exploration bonus, that the empirical count-based exploration bonuses are in general not sufficient to provide optimism and thus prone to under-exploration. In particular, the knowledge of the span of the optimal bias function is required in order to properly scale the bonus. Moreover, even in the finite-horizon case, the mentioned approaches use the knowledge of the horizon to scale the bonus. The planning horizon is in turn an upper-bound to the span of the optimal value function, thus they exploit the same prior knowledge required by SCAL<sup>+</sup> and C-SCAL<sup>+</sup>.

We also provide the tightest level of optimism for OFU algorithms by achieving the optimal dependence in the bonus w.r.t. the state dimensionality (it cannot further reduced while preserving theoretical guarantees). Unfortunately, this tighter optimism does not imply a tighter bound leaving open the quest for closing the gap between lower and upper bound in infinite-horizon undiscounted settings. Moreover, it is unclear to us if the exploration bonus can be extended to settings where no-prior knowledge of the span of the optimal bias function is available, e.g., in communicating (see UCRL (Jaksch et al., 2010)) or weakly-communicating MDPs (see TUCRL (Fruit et al., 2018a)). We leave this question for future work. Finally, C-SCAL<sup>+</sup> requires to know the smoothness parameters in order to define the discretization of the state space. We believe that some effort should be spent in the direction of removing such prior knowledge, making the algorithm more adaptive.

## References

- Shipra Agrawal and Randy Jia. Optimistic posterior sampling for reinforcement learning: worst-case regret bounds. In *NIPS*, pages 1184–1194, 2017.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 263–272. PMLR, 2017.
- Peter L. Bartlett and Ambuj Tewari. REGAL: A regularization based algorithm for reinforcement learning in weakly communicating MDPs. In *UAI*, pages 35–42. AUAI Press, 2009.
- Marc G. Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Rémi Munos. Unifying count-based exploration and intrinsic motivation. In *NIPS*, pages 1471–1479, 2016.
- Dimitri P Bertsekas. *Dynamic programming and optimal control. Vol II*. Number 2. Athena scientific Belmont, MA, 1995.
- Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1): 1–122, 2012. ISSN 1935-8237. doi: 10.1561/22000000024.
- Nicolò Cesa-Bianchi and Claudio Gentile. Improved risk tail bounds for on-line algorithms. In *NIPS*, pages 195–202, 2005.
- Y.S. Chow and H. Teicher. *Probability Theory: Independence, Interchangeability, Martingales*. Springer texts in statistics. World Publishing Company, 1988. ISBN 9780387966953.
- David A. Freedman. On tail probabilities for martingales. *Ann. Probab.*, 3(1):100–118, 02 1975. doi: 10.1214/aop/1176996452.
- Ronan Fruit, Matteo Pirotta, Alessandro Lazaric, and Emma Brunskill. Regret minimization in mdps with options without prior knowledge. In *NIPS*, pages 3169–3179, 2017.
- Ronan Fruit, Matteo Pirotta, and Alessandro Lazaric. Near optimal exploration-exploitation in non-communicating markov decision processes. In *NIPS*, 2018a.
- Ronan Fruit, Matteo Pirotta, Alessandro Lazaric, and Ronald Ortner. Efficient bias-span-constrained exploration-exploitation in reinforcement learning. In *ICML*, Proceedings of Machine Learning Research. PMLR, 2018b.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11:1563–1600, 2010.
- Chi Jin, Zeyuan Allen-Zhu, Sébastien Bubeck, and Michael I. Jordan. Is q-learning provably efficient? *CoRR*, abs/1807.03765, 2018.
- Sham Kakade, Mengdi Wang, and Lin F. Yang. Variance reduction methods for sublinear reinforcement learning. *CoRR*, abs/1802.09184, 2018.



- A. Klenke and M. Loève. *Probability Theory: A Comprehensive Course*. Graduate texts in mathematics. Springer, 2013. ISBN 9781447153627.
- K. Lakshmanan, Ronald Ortner, and Daniil Ryabko. Improved regret bounds for undiscounted continuous reinforcement learning. In *ICML*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 524–532. JMLR.org, 2015.
- Tor Lattimore and Csaba Szepesvári. Bandit algorithms. Pre-publication version, 2018. URL <http://downloads.tor-lattimore.com/banditbook/book.pdf>.
- Jarryd Martin, Suraj Narayanan Sasikumar, Tom Everitt, and Marcus Hutter. Count-based exploration in feature space for reinforcement learning. *CoRR*, abs/1706.08090, 2017.
- Ronald Ortner. Optimism in the face of uncertainty should be refutable. *Minds and Machines*, 18(4):521–526, 2008.
- Ronald Ortner and Daniil Ryabko. Online regret bounds for undiscounted continuous reinforcement learning. In *NIPS*, pages 1772–1780, 2012.
- Georg Ostrovski, Marc G. Bellemare, Aäron van den Oord, and Rémi Munos. Count-based exploration with neural density models. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 2721–2730. PMLR, 2017.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1994. ISBN 0471619779.
- Alexander L Strehl and Michael L Littman. An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008.
- Mohammad Sadegh Talebi and Odalric-Ambrym Maillard. Variance-aware regret bounds for undiscounted reinforcement learning in mdps. In *ALT*, volume 83 of *Proceedings of Machine Learning Research*, pages 770–805. PMLR, 2018.
- Haoran Tang, Rein Houthoofd, Davis Foote, Adam Stooke, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. #exploration: A study of count-based exploration for deep reinforcement learning. In *NIPS*, pages 2750–2759, 2017.

**Structure of the appendix.** We first present the proofs for the continuous case (App. A) and then for the discrete case (App. B) since the latter can be viewed as a special case of the former. That way, we only need to highlight the main differences in the discrete case. Because of the continuous nature of the state space and the aggregation of states in the continuous case, extra care needs to be taken while using concentration inequalities compared to standard regret proofs in RL (see App. A.1.3 for more details). In App. C we recall/prove all the necessary results from probability theory that we use in the regret proofs.

## Appendix A. Continuous state MDPs: the analysis of C-SCAL<sup>+</sup>

In all this section we say that a function  $v : s \in \mathcal{S} \mapsto \mathbb{R}$  is *piecewise constant on  $\mathcal{I}$*  when  $\forall J \in \mathcal{I}, \forall s, s' \in J$  we have  $v(s) = v(s')$  and we denote  $v(J)$  the joint value.

### A.1 High probability bound using the exploration bonus (proof of Lem. 4)

To begin with, we introduce two variants of the exploration bonus that will be used for the regret proof:

$$b_k(J, a) := c \cdot \min \left\{ \beta_{p,k}^{Ja} + \frac{1}{N_k(J, a) + 1}; 2 \right\} + \min \{ \beta_{r,k}^{Ja}; r_{\max} \} + (c + r_{\max})LS^{-\alpha} \quad (10)$$

$$d_k(J, a) := c \cdot \min \left\{ \phi_{p,k}^{Ja} + \frac{1}{N_k(J, a) + 1}; 2 \right\} + \min \{ \beta_{r,k}^{Ja}; r_{\max} \} + (c + r_{\max})LS^{-\alpha} \quad (11)$$

where  $\beta_{p,k}^{Ja} = \beta_k^{Ja}$  (see Eq. 5 for the definition of  $\beta_k^{sa}$ ),  $\beta_{r,k}^{Ja} = r_{\max}\beta_k^{Ja}$  and

$$\phi_{p,k}^{Ja} := \sqrt{\frac{7S \ln \left( \frac{3SA_t k}{\delta} \right)}{N_k^+(J, a)}} + \frac{14S}{N_k^+(J, a)} \ln \left( \frac{3SA_t k}{\delta} \right) \geq \beta_{p,k}^{Ja} \quad (12)$$

with  $N_k^+(J, a) := \max\{1, N_k(J, a)\}$ . Note that the Eq. 10 is a slightly tighter version of the exploration bonus considered in Eq. 5. We define for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ ,  $\bar{r}_k(s, a) := \bar{r}_k^{ag}(I(s), a)$ . We state a slightly more general result than Lem. 4:

**Lemma 12** *Consider the estimated continuous MDP  $\widehat{M}_k$  defined in Def. 9. Let  $(g^*, h^*)$  be a solution of the optimality equation  $Lh^* = h^* + g^*$  such that  $sp\{h^*\} \leq c$ . For all  $T \geq 1$  and  $k \geq 1$ , with probability at least  $1 - \frac{\delta}{15t_k^6}$ , for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and for any function  $v$  piecewise constant on  $\mathcal{I}$  s.t.  $sp\{v\} \leq c$  we have:*

$$(a) \quad b_k(s, a) \geq \left| \bar{r}_k(s, a) - r(s, a) + \int_{\mathcal{S}} (\widehat{p}_k(s'|s, a) - p(s'|s, a)) h^*(s') ds' \right|$$

$$(b) \quad d_k(s, a) \geq \left| \bar{r}_k(s, a) - r(s, a) + \int_{\mathcal{S}} (\widehat{p}_k(s'|s, a) - p(s'|s, a)) v(s') ds' \right|$$

where  $b_k$  and  $d_k$  are defined as in Eq. 10 and Eq. 11, respectively.

The rest of section A.1 is devoted to proving Lem. 12. To do this, we introduce an intermediate continuous state-space MDP  $\widetilde{M}_k := (\mathcal{S}, \mathcal{A}, \widetilde{r}_k, \widetilde{p}_k)$  defined for all pairs  $(s, a) \in \mathcal{S} \times \mathcal{A}$  by:

$$\begin{aligned}\widetilde{r}_k(s, a) &:= \frac{1}{N_k(I(s), a)} \sum_{x \in I(s)} N_k(x, a) r(x, a) \\ \widetilde{p}_k(s'|s, a) &:= \frac{1}{N_k(I(s), a)} \sum_{x \in I(s)} N_k(x, a) p(s'|x, a)\end{aligned}$$

We decompose  $\widehat{p}_k - p$  and  $\bar{r}_k - r$  as

$$\widehat{p}_k - p = (\widehat{p}_k - \bar{p}_k) + (\bar{p}_k - \widetilde{p}_k) + (\widetilde{p}_k - p) \quad \text{and} \quad \bar{r}_k - r = (\bar{r}_k - \widetilde{r}_k) + (\widetilde{r}_k - r) \quad (13)$$

and bound separately all the terms. Similarly, we decompose  $\bar{r}_k - r$  as  $(\bar{r}_k - \widetilde{r}_k) + (\widetilde{r}_k - r)$ . We also define  $w^*(s) := h^*(s) - (\inf\{h^*(s)\} + \sup\{h^*(s)\})/2$  implying that for all  $s \in \mathcal{S}$ ,  $w^*(s) \in [-c/2, c/2]$ .

#### A.1.1 BOUNDING THE DIFFERENCE BETWEEN $\widetilde{r}_k/\widetilde{p}_k$ AND $r/p$

To bound the differences  $\widetilde{r}_k(s'|s, a) - r(s, a)$  and  $\int (\widetilde{p}_k(s'|s, a) - p(s'|s, a)) w^*(s') ds'$  we simply use the smoothness assumption on the reward and transition model (see Asm. 7). For all  $(s, a) \in \mathcal{S} \times \mathcal{A}$  (using the triangle inequality):

$$|\widetilde{r}_k(s, a) - r(s, a)| \leq \frac{1}{N_k(I(s), a)} \sum_{x \in I(s)} N_k(x, a) \underbrace{|r(x, a) - r(s, a)|}_{\leq r_{\max} LS^{-\alpha} \text{ since } x, s \in I(s)} \leq r_{\max} LS^{-\alpha} \quad (14)$$

For the transition probability we have that for all  $J \in \mathcal{I}$  (using the triangle inequality):

$$\begin{aligned}\sum_{J \in \mathcal{I}} \left| \int_J (\widetilde{p}_k(s'|s, a) - p(s'|s, a)) ds' \right| &\leq \sum_{J \in \mathcal{I}} \int_J |\widetilde{p}_k(s'|s, a) - p(s'|s, a)| ds' \\ &= \frac{1}{N_k(I(s), a)} \sum_{x \in I(s)} N_k(x, a) \int_{\mathcal{S}} \underbrace{|p(s'|x, a) - p(s'|s, a)|}_{\leq LS^{-\alpha} \text{ since } x, s \in I(s)} ds' \\ &\leq LS^{-\alpha}\end{aligned} \quad (15)$$

and similarly:

$$\begin{aligned}\left| \int_{\mathcal{S}} (\widetilde{p}_k(s'|s, a) - p(s'|s, a)) w^*(s') ds' \right| &\leq c \int_{\mathcal{S}} |\widetilde{p}_k(s'|s, a) - p(s'|s, a)| ds' \\ &= \frac{c}{N_k(I(s), a)} \sum_{x \in I(s)} N_k(x, a) \int_{\mathcal{S}} \underbrace{|p(s'|x, a) - p(s'|s, a)|}_{\leq LS^{-\alpha} \text{ since } x, s \in I(s)} ds' \\ &\leq cLS^{-\alpha}\end{aligned} \quad (16)$$

A.1.2 BOUNDING THE DIFFERENCE BETWEEN  $\widehat{p}_k$  AND  $\bar{p}_k$

Using the triangle inequality and the fact that  $\int_{\mathcal{S}} \mathbb{1}(s' \in I(\bar{s})) ds' = \int_{I(\bar{s})} 1 ds' = |I(\bar{s})| = 1/S$  we have that for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$ :

$$\begin{aligned} \left| \int_{\mathcal{S}} (\widehat{p}_k(s'|s, a) - \bar{p}_k(s'|s, a)) w^*(s') ds' \right| &\leq \int_{\mathcal{S}} |\widehat{p}_k(s'|s, a) - \bar{p}_k(s'|s, a)| \cdot |w^*(s')| ds' \\ &= \left| \frac{N_k(I(s), a)}{N_k(I(s), a) + 1} - 1 \right| \int_{\mathcal{S}} \bar{p}_k(s'|s, a) \underbrace{|w^*(s')|}_{\leq c/2} ds' \\ &\quad + S \int_{\mathcal{S}} \frac{|w^*(s')| \mathbb{1}(s' \in I(\bar{s}))}{N_k(I(s), a) + 1} ds' \\ &\leq \frac{c}{N_k(I(s), a) + 1} \end{aligned} \tag{17}$$

and similarly:

$$\sum_{J \in \mathcal{I}} \left| \int_J (\widehat{p}_k(s'|s, a) - \bar{p}_k(s'|s, a)) ds' \right| \leq \int_{\mathcal{S}} |\widehat{p}_k(s'|s, a) - \bar{p}_k(s'|s, a)| ds' \leq \frac{1}{N_k(I(s), a) + 1} \tag{18}$$

A.1.3 BOUNDING THE DIFFERENCE BETWEEN  $\tilde{r}_k/\tilde{p}_k$  AND  $\bar{r}_k/\bar{p}_k$

Let's consider a fixed pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and a fixed aggregated state  $J \in \mathcal{I}$ . Our goal is to bound the differences  $\int_{\mathcal{S}} (\tilde{p}_k(s'|s, a) - \bar{p}_k(s'|s, a)) w^*(s') ds'$ ,  $\int_J \tilde{p}_k(s'|s, a) - \bar{p}_k(s'|s, a) ds'$  and  $\tilde{r}_k(s, a) - \bar{r}_k(s, a)$ . Since  $\tilde{p}_k$  and  $\tilde{r}_k$  are somehow the expected values of  $\bar{p}_k$  and  $\bar{r}_k$ , we would like to use concentration inequalities. In the case of a finite state space  $\mathcal{S}$ , Jaksch et al. (2010, UCRL) and Fruit et al. (2018b, SCAL) use concentration inequalities that apply to *independent* random variables (r.v.). We argue that a more careful analysis is needed here since the states lie in an *uncountable* set. Indeed, the implicit assumption made about the RL model for UCRL and SCAL is that for each state-action pair  $(s, a)$ , the rewards (respectively next states) are sampled from an infinite *stack* of independent and identically distributed (i.i.d.) rewards (respectively next states). More precisely, each time the agent visits  $(s, a)$ , it receives a reward from the top of the stack of rewards associated to  $(s, a)$  and moves to the state on the top of the stack of next states associated to  $(s, a)$ . The two samples are then withdrawn from their respective stacks (meaning that they cannot be popped again). For more details about why this is a valid model refer to (Lattimore and Szepesvári, 2018, Section 4.4). In the case where  $\mathcal{S}$  and  $\mathcal{A}$  are discrete sets (finite or countable), it is possible to use any concentration inequality for i.i.d. r.v. and then take a union bound over all “stacks”  $(s, a)$  (and over rewards and next states). When  $\mathcal{S}$  is uncountable however, the same argument cannot be used (the probability of an uncountable union of events is not even always defined). Moreover, the terms  $\tilde{r}_k$  and  $\tilde{p}_k$  are obtained using sampled from different states  $x \in I(s)$  instead of a single state  $s$ . To overcome these technical problems, we use a variant of Doob's optional skipping (e.g., Chow and Teicher, 1988, Sec. 5.3, Lem. 4) and concentration inequalities for martingales (Azuma and Freedman inequalities). The theorem that we use (Thm. 29) is formally proved in App. C.

For any  $t \geq 0$ , the  $\sigma$ -algebra induced by the past history of state-action pairs and rewards up to time  $t$  is denoted  $\mathcal{F}_t := \sigma(s_1, a_1, r_1, \dots, s_t, a_t)$  where by convention  $\mathcal{F}_0 = \sigma(\emptyset)$  and  $\mathcal{F}_\infty := \cup_{t \geq 0} \mathcal{F}_t$ . Let  $\mathbb{F}$  denote the filtration  $(\mathcal{F}_t)_{t \geq 0}$ . We define the following adapted sequences and stopping times:

**1) Adapted sequences:**

We consider the following stochastic processes adapted to  $\mathbb{F}$ :  $(w^*(s_t))_{t \geq 0}$ ,  $(\mathbb{1}(s_t \in J))_{t \geq 0}$  and  $(r_{t-1}(s_{t-1}, a_{t-1}))_{t \geq 0}$  (with the conventions  $r_{-1}(s_{-1}, a_{-1}) = r_0(s_0, a_0) = r_\infty(s_\infty, a_\infty) = 0$ ,  $w^*(s_0) = w^*(s_\infty) = 0$  and  $\mathbb{1}(s_0 \in J) = \mathbb{1}(s_\infty \in J) = 0$ ). These processes are bounded as  $|w^*(s_t)| \leq 2 \times \|w^*\|_\infty \leq c$ ,  $|\mathbb{1}(s_t \in J)| \leq 1$  and  $|r_{t-1}(s_{t-1}, a_{t-1})| \leq r_{\max}$  for all  $t \geq 0$ .

**2) Stopping times:**

We define  $\tau := (\tau_l)_{l \geq 0}$  s.t.  $\tau_0 := 0$  and  $\inf\{t_k > t > \tau_l : s_t \in I(s), a_t = a\}$ . For all  $l \geq 0$  and for all  $t \geq 0$ ,  $\tau_l := \{\tau_l = t\} \in \mathcal{F}_t$  and so  $\tau_l$  is a stopping time w.r.t.  $\mathbb{F}$  (see Def. 18 in App. C). By definition for any  $l \geq 0$ ,  $\tau_l < \tau_{l+1}$  a.s. (i.e.,  $\tau$  is strictly increasing, see Lem. 26). We denote  $\mathcal{G}_l := \mathcal{F}_{\tau_{l+1}}$  the  $\sigma$ -algebra at stopping time  $\tau_{l+1}$  (see Def. 19 in App. C).

All the assumptions of Thm. 29 are satisfied and so by taking a union bound and using the fact that  $N_k(I(s), a) \leq t_k$  a.s., we obtain that with probability at least  $1 - 3\delta$  (simultaneously):

$$\begin{aligned} \left| \sum_{l=1}^{N_k(I(s), a)} \left( r_{\tau_l}(s_{\tau_l}, a_{\tau_l}) - \mathbb{E}[r_{\tau_l}(s_{\tau_l}, a_{\tau_l}) | \mathcal{G}_{l-1}] \right) \right| &\leq r_{\max} \sqrt{N_k(I(s), a) \ln \left( \frac{2t_k}{\delta} \right)} \\ \left| \sum_{l=1}^{N_k(I(s), a)} \left( \mathbb{1}(s_{\tau_{l+1}} \in J) - \mathbb{E}[\mathbb{1}(s_{\tau_{l+1}} \in J) | \mathcal{G}_{l-1}] \right) \right| &\leq 2 \sqrt{V_k(J) \ln \left( \frac{4t_k}{\delta} \right)} + 4 \ln \left( \frac{4t_k}{\delta} \right) \end{aligned} \quad (19)$$

$$\text{and } \left| \sum_{l=1}^{N_k(I(s), a)} \left( w^*(s_{\tau_{l+1}}) - \mathbb{E}[w^*(s_{\tau_{l+1}}) | \mathcal{G}_{l-1}] \right) \right| \leq c \sqrt{N_k(I(s), a) \ln \left( \frac{2t_k}{\delta} \right)}$$

where  $V_k(J) := \sum_{l=1}^{N_k(I(s), a)} \mathbb{V}(\mathbb{1}(s_{\tau_{l+1}} \in J) | \mathcal{G}_{l-1})$ . We now need to relate the above sums to  $\int_S (\tilde{p}_k(s' | s, a) - \bar{p}_k(s' | s, a)) w^*(s') ds'$ ,  $\int_J \tilde{p}_k(s' | s, a) - \bar{p}_k(s' | s, a) ds'$  and  $\tilde{r}_k(s, a) - \bar{r}_k(s, a)$ . We also need to give an explicit formula for  $V_k(J)$ . By definition of  $\tau$ , we can rewrite  $\tilde{r}_k$  and  $\bar{p}_k$  as follows:

$$\begin{aligned} \bar{r}_k(s, a) &= \frac{1}{N_k(I(s), a)} \sum_{l=1}^{N_k(I(s), a)} r_{\tau_l}(s_{\tau_l}, a_{\tau_l}) \\ \int_J \bar{p}_k(s' | s, a) ds' &= \frac{1}{N_k(I(s), a)} \sum_{l=1}^{N_k(I(s), a)} \mathbb{1}(s_{\tau_{l+1}} \in J) \\ \int_S \bar{p}_k(s' | s, a) w^*(s') ds' &= \frac{1}{N_k(I(s), a)} \sum_{l=1}^{N_k(I(s), a)} w^*(s_{\tau_{l+1}}) \end{aligned}$$

It is also easy to verify that the following holds:  $\mathbb{E}[w^*(s_{\tau_{l+1}}) | \mathcal{G}_{l-1}] = \int_S p(s' | s_{\tau_l}, a_{\tau_l}) w^*(s') ds'$ ,  $\mathbb{E}[\mathbb{1}(s_{\tau_{l+1}} \in J) | \mathcal{G}_{l-1}] = \int_J p(s' | s_{\tau_l}, a_{\tau_l}) ds'$  and  $\mathbb{E}[r_{\tau_l}(s_{\tau_l}, a_{\tau_l}) | \mathcal{G}_{l-1}] = r(s_{\tau_l}, a_{\tau_l})$  (see Lem. 30

in App. C for a formal proof). As a result, we can rewrite  $\tilde{r}_k$  and  $\tilde{p}_k$  as follows:

$$\tilde{r}_k(s, a) = \frac{1}{N_k(I(s), a)} \sum_{x \in I(s)} N_k(x, a) r(x, a) = \frac{1}{N_k(I(s), a)} \sum_{l=1}^{N_k(I(s), a)} \mathbb{E}[r_{\tau_l}(s_{\tau_l}, a_{\tau_l}) | \mathcal{G}_{l-1}]$$

$$\begin{aligned} \text{and } \int_{\mathcal{S}} \tilde{p}_k(s' | s, a) w^*(s') ds' &= \frac{1}{N_k(I(s), a)} \sum_{x \in I(s)} N_k(x, a) \int_{\mathcal{S}} p(s' | x, a) w^*(s') ds' \\ &= \frac{1}{N_k(I(s), a)} \sum_{l=1}^{N_k(I(s), a)} \mathbb{E}[w^*(s_{\tau_l+1}) | \mathcal{G}_{l-1}] \end{aligned}$$

$$\text{and similarly } \int_J \tilde{p}_k(s' | s, a) ds' = \frac{1}{N_k(I(s), a)} \sum_{l=1}^{N_k(I(s), a)} \mathbb{E}[\mathbb{1}(s_{\tau_l+1} \in J) | \mathcal{G}_{l-1}]$$

We can also give a more explicit expression for  $V_k$ :

$$\begin{aligned} \mathbb{V}(\mathbb{1}(s_{\tau_l+1} \in J) | \mathcal{G}_{l-1}) &:= \mathbb{E}[\underbrace{\mathbb{1}(s_{\tau_l+1} \in J)^2}_{= \mathbb{1}(s_{\tau_l+1} \in J)} | \mathcal{G}_{l-1}] - \mathbb{E}[\mathbb{1}(s_{\tau_l+1} \in J) | \mathcal{G}_{l-1}]^2 \\ &= \int_J p(s' | s_{\tau_l}, a_{\tau_l}) ds' - \left( \int_J p(s' | s_{\tau_l}, a_{\tau_l}) ds' \right)^2 \end{aligned}$$

implying:

$$V_k(J) = \sum_{l=1}^{N_k(I(s), a)} \underbrace{\left(1 - \int_J p(s' | s_{\tau_l}, a_{\tau_l}) ds'\right)}_{\leq 1} \underbrace{\int_J p(s' | s_{\tau_l}, a_{\tau_l}) ds'}_{\geq 0} \leq \sum_{x \in I(s)} N_k(x, a) \int_J p(s' | x, a) ds'$$

Using Cauchy-Schwartz inequality:

$$\sum_{J \in \mathcal{I}} \sqrt{V_k(J)} \leq \sqrt{S \sum_{J \in \mathcal{I}} V_k(J)} \leq \sqrt{S \sum_{x \in I(s)} N_k(x, a) \sum_{J \in \mathcal{I}} \int_J p(s' | x, a) ds'} = \sqrt{S N_k(I(s), a)}$$

To conclude, we take a union bound over all possible  $(I(s), a) \in \mathcal{I} \times \mathcal{A}$  and  $J \in \mathcal{I}$ . Note that we only need to take a union bound over  $I(s) \in \mathcal{I}$  (and not  $\mathcal{S}$ ) because  $s \mapsto \tilde{p}_k(\cdot | s, a)$  and  $s \mapsto \tilde{r}_k(s, a)$  are piecewise constant on  $\mathcal{I}$  (and similarly for  $\bar{p}_k$  and  $\bar{r}_k$ ). With probability at least  $1 - \frac{\delta}{15l_k^6}$ , for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and for all  $J \in \mathcal{I}$ :

$$|\bar{r}_k(s, a) - \tilde{r}_k(s, a)| \leq r_{\max} \sqrt{\frac{\ln\left(\frac{90S^2 At_k^7}{\delta}\right)}{N_k(I(s), a)}} \leq r_{\max} \sqrt{\frac{7 \ln\left(\frac{2S At_k}{\delta}\right)}{N_k(I(s), a)}} \quad (20)$$

$$\sum_{J \in \mathcal{I}} \left| \int_J \bar{p}_k(s' | s, a) - \tilde{p}_k(s' | s, a) ds' \right| \leq 2 \sqrt{\frac{7S \ln\left(\frac{3S At_k}{\delta}\right)}{N_k(I(s), a)}} + \frac{28S}{N_k(I(s), a)} \ln\left(\frac{3S At_k}{\delta}\right) \quad (21)$$

$$\left| \int_{\mathcal{S}} (\bar{p}_k(s' | s, a) - \tilde{p}_k(s' | s, a)) w^*(s') ds' \right| \leq c \sqrt{\frac{7 \ln\left(\frac{2S At_k}{\delta}\right)}{N_k(I(s), a)}} \quad (22)$$

Since by definition  $N_k^+(I(s), a) := \max\{1, N_k(I(s), a)\}$ , the above inequalities also hold with  $N_k(I(s), a)$  replaced by  $N_k^+(I(s), a)$ .

#### A.1.4 GATHERING ALL THE TERMS

We first notice that  $\int_{\mathcal{S}} (\widehat{p}_k(s'|s, a) - p(s'|s, a)) h^*(s') ds' = \int_{\mathcal{S}} (\widehat{p}_k(s'|s, a) - p(s'|s, a)) w^*(s') ds'$  since  $w^*$  and  $h^*$  are equal up to a constant shift and  $\int_{\mathcal{S}} \widehat{p}_k(s'|s, a) ds' = \int_{\mathcal{S}} p(s'|s, a) ds' = 1$ . Gathering equations (22), (16) and (17) we have:

$$\left| \int_{\mathcal{S}} (\widehat{p}_k(s'|s, a) - p(s'|s, a)) h^*(s') ds' \right| \leq c \sqrt{\frac{7 \ln \left( \frac{2SA t_k}{\delta} \right)}{N_k^+(I(s), a)}} + cLS^{-\alpha} + \frac{c}{N_k(I(s), a) + 1}$$

Gathering equations (20) and (14) we have:

$$|\bar{r}_k(s, a) - r(s, a)| \leq r_{\max} \sqrt{\frac{7 \ln \left( \frac{2SA t_k}{\delta} \right)}{N_k^+(I(s), a)}} + r_{\max} LS^{-\alpha}$$

Let  $v$  be a piecewise constant function on  $\mathcal{I}$  s.t.  $sp\{v\} \leq c$  and define  $w(s) := v(s) - (\inf\{v(s)\} + \sup\{v(s)\})/2$ .  $w$  is also piecewise constant on  $\mathcal{I}$  and for all  $J \in \mathcal{I}$ ,  $w(J) \in [-c/2, c/2]$ . Gathering equations (21), (15) and (18) we have that :

$$\begin{aligned} \sum_{J \in \mathcal{I}} \left| \int_J (\widehat{p}_k(s'|s, a) - p(s'|s, a)) v(s') ds' \right| &= \sum_{J \in \mathcal{I}} \left| w(J) \int_J (\widehat{p}_k(s'|s, a) - p(s'|s, a)) ds' \right| \\ &\leq \frac{c}{2} \sum_{J \in \mathcal{I}} \left| \int_J (\widehat{p}_k(s'|s, a) - p(s'|s, a)) ds' \right| \\ &\leq c \sqrt{\frac{7S \ln \left( \frac{3SA t_k}{\delta} \right)}{N_k^+(I(s), a)}} + \frac{14cS}{N_k^+(I(s), a)} \ln \left( \frac{3SA t_k}{\delta} \right) \\ &\quad + cLS^{-\alpha} + \frac{c}{N_k(I(s), a) + 1} \end{aligned}$$

Properties (a) and (b) of Lem. 14 follow by definition of the exploration bonuses and application of the triangle inequality.

## A.2 Optimism (Proof of Lem. 5)

Let  $\widehat{g}_k^{ag+}$  denote the solution of optimisation problem (3) on  $\widehat{M}_k^{ag+}$ . In this section we prove that:

**Lemma 13** *Consider the MDP  $\widehat{M}_k^{ag+}$  defined in Sec. 4. Then for any  $k > 0$ , with probability at least  $1 - \frac{\delta}{15t_k^6}$ ,  $\widehat{g}_k^{ag+} \geq g^*$ .*

$\widehat{M}_k^{ag+}$  only has a finite number of states while the true MDP  $M^*$  has an uncountable state-space. Thus, it is difficult to compare directly  $\widehat{g}_k^{ag+}$  with  $g^*$ . To overcome this difficulty, we first compare  $g^*$  with the gain of  $\widehat{M}_k$  and then compare the latter to  $\widehat{g}_k^{ag+}$ .

**1. Optimism of  $\widehat{M}_k$ .** Let  $\widehat{g}_k$  denote the solution of optimisation problem (3) on  $\widehat{M}_k$ . To prove that  $\widehat{g}_k \geq g^*$  we can use Prop. 3 which –as explained in the main body of the paper– only requires to show that  $\widehat{L}h^* \geq Lh^*$  where  $\widehat{L}$  is the optimal Bellman operator of  $\widehat{M}_k$ . By applying property (a) of Lem. 12, we have that with probability at least  $1 - \frac{\delta}{15l_k^6}$ :

$$\begin{aligned} \forall s \in \mathcal{S}, \widehat{L}h^*(s) &:= \max_{a \in \mathcal{A}} \left\{ \widehat{r}_k(s, a) + b_k(s, a) + \int_{\mathcal{S}} \widehat{p}_k(s'|s, a) h^*(s') ds' \right\} \\ &\geq \max_{a \in \mathcal{A}} \left\{ r(s, a) + \int_{\mathcal{S}} p(s'|s, a) h^*(s') ds' \right\} = Lh^*(s) \end{aligned}$$

Therefore,  $\widehat{g}_k \geq g^*$  with probability at least  $1 - \frac{\delta}{15l_k^6}$ .

**2. Relationship between  $\widehat{M}_k$  and  $\widehat{M}_k^{ag+}$ .** We now show that  $\widehat{g}_k^{ag+} = \widehat{g}_k$ . Consider a piecewise-constant function  $v_0$  on  $\mathcal{I}$  (e.g.,  $v_0 = 0$ ) and a vector  $u_0 \in \mathbb{R}^{\mathcal{S}}$  satisfying  $u_0(J) = v_0(J)$  for all  $J \in \mathcal{I}$ . We define the sequences  $v_{n+1} := \widehat{T}_c v_n$  and  $u_{n+1} := \widehat{T}_c^{ag+} u_n$ . We show by induction that  $u_n(J) = v_n(J)$  for all  $n \geq 0$  and for all  $J \in \mathcal{I}$ . By definition it is true for  $n = 0$  and for all  $n \geq 0$ :

$$\begin{aligned} \int_{s \in \mathcal{S}} \widehat{p}_k(s'|s, a) v_n(s') ds' &= \sum_{J \in \mathcal{I}} \int_J \widehat{p}_k(s'|s, a) v_n(s') ds' \\ &= \sum_{J \in \mathcal{I}} v_n(J) \int_J \widehat{p}_k(s'|s, a) ds' = \sum_{J \in \mathcal{I}} u_n(J) \widehat{p}_k^{ag+}(J|I(s), a) \end{aligned} \tag{23}$$

where the last equality follows from (9) and the induction hypothesis. In addition  $\widehat{r}_k(s, a)$  is also piecewise-constant on  $\mathcal{I}$  and  $\widehat{r}_k(s, a) = \widehat{r}_k^{ag+}(I(s), a)$  for all  $s \in \mathcal{S}$ . Therefore, we have that  $\widehat{L}^{ag+} u_n(I(s)) = \widehat{L} v_n(s)$  for any  $s \in \mathcal{S}$ . Finally, the augmentation is not impacting the optimal Bellman operator (i.e., for any  $v$ ,  $\widehat{L}^{ag+} v = \widehat{L}^{ag} v$ ) so  $\widehat{L}^{ag+} u_n(I(s)) = \widehat{L} v_n(s)$  and consequently  $\widehat{T}_c^{ag+} u_n(I(s)) = \widehat{T}_c v_n(s)$  for any  $s \in \mathcal{S}$ . This shows that  $v_{n+1}(J) = u_{n+1}(J)$  for all  $J \in \mathcal{I}$  which concludes the proof by induction.

As shown by Fruit et al. (2018b, Theorem 10),  $\lim_{n \rightarrow +\infty} v_{n+1}(J) - v_n(J) = \widehat{g}_k^{ag+}$  and  $\lim_{n \rightarrow +\infty} u_{n+1}(J) - u_n(J) = \widehat{g}_k$  so that  $\widehat{g}_k^{ag+} = \widehat{g}_k \geq g^*$  with probability at least  $1 - \frac{\delta}{15l_k^6}$ .

### A.3 Regret Proof of C-SCAL<sup>+</sup> (Proof of Thm. 8)

In this section, we provide a complete proof of the regret bound for C-SCAL<sup>+</sup>. Defining  $\Delta_k = \sum_{s \in \mathcal{S}} \nu_k(s) \left( g^* - \sum_{a \in \mathcal{A}_{s_t}} r(s, a) \widetilde{\pi}_k(s, a) \right)$  and using the arguments in (Jaksch et al., 2010; Fruit et al., 2018b), it holds with probability at least  $1 - \frac{\delta}{20T^{5/4}}$  that:  $\Delta(\text{SCAL}^+, T) \leq \sum_{k=1}^m \Delta_k + r_{\max} \sqrt{\frac{5}{2} T \ln \left( \frac{11T}{\delta} \right)}$ . Note that  $\nu_k(s)$  is the total number of observation of state  $s$  in episode  $k$  and is well-defined for  $s$  lying in  $[0, 1]$ . Finally, recall that for any subset  $I \subseteq \mathcal{S}$ , the sum  $\sum_{s \in I} u_s$  is also well-defined as long as the collection  $(u_s)_{s \in I}$  contains only a finite number of non-zero elements.

**▲** In this section we will abuse of notation and write  $p(\cdot|s, a)^\top v = \int_{\mathcal{S}} p(s'|s, a) v(s') ds'$  for any probability density function  $p$  defined on  $\mathcal{S} = [0, 1]$ .



## A.3.1 OPTIMISM AND BELLMAN EQUATION

We consider the case where  $\Delta_k \cdot \mathbb{1} \left( g_c^*(\widehat{M}_k^{ag+}) \geq g^* \right)$ , the complementary case is (cumulatively) bounded by  $r_{\max} \sqrt{T}$  as in (Jaksch et al., 2010; Fruit et al., 2018b). Denote by  $g_k := 1/2(\max\{\widehat{T}_c^+ v_k - v_k\} + \min\{\widehat{T}_c^+ v_k - v_k\})$  where  $v_k$  is the value function returned by  $\text{SCOPT}(0, \bar{s}, \gamma_k, \varepsilon_k)$ . Remember that  $v_k \in \mathbb{R}^S$  is a discrete vectors obtained by applying  $\text{SCOPT}$  on  $\widehat{M}_k^{ag+}$ . The stopping condition of  $\text{SCOPT}$  is such that (see Fruit et al., 2018b)

$$g_k \geq g_c^*(\widehat{M}_k^{ag+}) - \underbrace{\varepsilon_k}_{=r_{\max}/\sqrt{t_k}} \stackrel{\text{Lem. 5}}{\geq} g^* - \frac{r_{\max}}{\sqrt{t_k}}$$

implying:

$$\Delta_k \cdot \mathbb{1} \left( g_c^*(\widehat{M}_k^{ag+}) \geq g^* \right) \leq r_{\max} \sum_{s \in \mathcal{S}} \frac{\nu_k(s)}{\sqrt{t_k}} + \sum_{s \in \mathcal{S}} \nu_k(s) \underbrace{\left( g_k - \sum_{a \in \mathcal{A}_s} r(s, a) \tilde{\pi}_k(s, a) \right)}_{:=\Delta'_k(s)}$$

Note that we can associated a continuous piecewise-constant function to the discrete vector  $v_k$ :  $u_k(s) = v_k(I(s))$ ,  $\forall s \in \mathcal{S}$ . A consequence of Lem. 12(b) applied to vector  $u_k$  is that  $-r(s, a) \leq d_k(s, a) - \bar{r}_k(s, a) + (p(\cdot|s, a) - \hat{p}_k(\cdot|s, a))^\top u_k$ . Note that we cannot use the tighter version (a) since it does not hold for any function  $u_k$ . Moreover, by definition  $\bar{r}_k(s, a) = \bar{r}_k^{ag}(I(s), a) = \hat{r}_k(s, a) - b_k(s, a)$ . Therefore:

$$\begin{aligned} \Delta'_k(s) &\leq g_k - \sum_{a \in \mathcal{A}_s} \pi_k(s, a) \left( \underbrace{\hat{r}_k(s, a)}_{:=\hat{r}_k^{ag}(I(s), a)} + \hat{p}_k(\cdot|s, a)^\top u_k \right) \\ &\quad + \sum_{a \in \mathcal{A}_s} \pi_k(s, a) \left( \underbrace{b_k(s, a)}_{:=b_k(I(s), a)} + d_k(s, a) + p(\cdot|s, a)^\top u_k \right) \end{aligned} \quad (24)$$

A direct consequence of the stopping condition used by  $\text{SCOPT}$  (see Thm. 18 in (Fruit et al., 2018b)) is that:  $\forall J \in \mathcal{I}$ ,

$$\left| g_k - \sum_{a \in \mathcal{A}} \sum_{i \in \{1, 2\}} \hat{r}_k^{ag+}(J, a_i) \pi_k(J, a_i) + v_k(J) - \sum_{a \in \mathcal{A}} \sum_{i \in \{1, 2\}} \pi_k(J, a_i) \hat{p}_k^{ag+}(\cdot|J, a_i)^\top v_k \right| \leq \frac{r_{\max}}{\sqrt{t_k}} \quad (25)$$

Recall that by definition:  $\pi_k(J, a) = \pi_k(J, a_1) + \pi_k(J, a_2)$ ,  $\hat{r}_k^{ag+}(J, a_i) \leq \hat{r}_k^{ag}(J, a)$  (since  $\hat{r}_k^{ag+}(J, a_1) = \hat{r}_k^{ag}(J, a)$  and  $\hat{r}_k^{ag+}(J, a_2) = 0$ ) and  $\hat{p}_k^{ag}(\cdot|J, a) = \hat{p}_k^{ag+}(\cdot|J, a_i)$ . We can thus write:

$$\begin{aligned} \sum_{a \in \mathcal{A}} \hat{r}_k^{ag}(J, a) \pi_k(J, a) &= \sum_{a \in \mathcal{A}} \sum_{i \in \{1, 2\}} \hat{r}_k^{ag}(J, a) \pi_k(J, a_i) \geq \sum_{a \in \mathcal{A}} \sum_{i \in \{1, 2\}} \hat{r}_k^{ag+}(J, a_i) \pi_k(J, a_i) \\ \text{and } \sum_{a \in \mathcal{A}} \pi_k(s, a) \hat{p}_k^{ag}(\cdot|J, a) &= \sum_{a \in \mathcal{A}} \sum_{i \in \{1, 2\}} \pi_k(J, a_i) \hat{p}_k^{ag+}(\cdot|J, a_i) \end{aligned} \quad (26)$$

Plugging this last two equations into (25) and using Eq. 23 (i.e.,  $\widehat{p}_k(\cdot|s, a)^\top u_k = \widehat{p}_k^{ag}(\cdot|s, a)^\top v_k$ ), the fact  $u_k(s) = v_k(I(s))$ ,  $\widehat{r}_k(s, a) = \widehat{r}_k^{ag}(I(s), a)$ , and  $\pi_k(s, a) = \pi_k(I(s), a)$ , we obtain:

$$\forall s \in \mathcal{S}, \quad g_k - \sum_{a \in \mathcal{A}} \pi_k(s, a) \left( \widehat{r}_k(s, a) + \widehat{p}_k(\cdot|s, a)^\top u_k \right) \leq -u_k(s) + \frac{r_{\max}}{\sqrt{t_k}} \quad (27)$$

Combining (27) with (24) we have:

$$\Delta'_k(s) \leq \sum_{a \in \mathcal{A}_s} \pi_k(s, a) \left( d_k(s, a) + \underbrace{b_k(s, a)}_{\leq d_k(s, a)} + p(\cdot|s, a)^\top u_k \right) - u_k(s) + \frac{r_{\max}}{\sqrt{t_k}}$$

Note that  $d_k(s, a) \geq b_k(s, a)$  for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  since the term  $\phi_{p,k}^{Ia}$  (see Eq. 12) contains a  $\sqrt{S}$  dependence that is not present in  $\beta_{p,k}^{Ia}$ . Since the dominant term is given by  $d_k(s, a)$ , we will consider the following loser upper-bound  $d_k(s, a) + b_k(s, a) \leq 2d_k(s, a)$  in the remaining of the proof. We can now state that

$$\begin{aligned} \Delta_k &\leq \underbrace{\sum_s \nu_k(s) \left( \sum_a \left( \pi_k(s, a) p(\cdot|s, a)^\top w_k \right) - w_k(s) \right)}_{\xi_k} \\ &\quad + 2 \sum_{s, a} \nu_k(s) \pi_k(s, a) d_k(s, a) + 2r_{\max} \sum_{s \in \mathcal{S}} \frac{\nu_k(s)}{\sqrt{t_k}} \end{aligned} \quad (28)$$

where  $w_k = u_k - (\inf_s \{u_k(s)\} + \sup_s \{u_k(s)\})/2$  is obtained by “recentering”  $u_k$  around 0 so that  $\|w_k\|_\infty = sp \{w_k\} / 2 \leq c/2$  (see Fruit et al., 2018b, App. F4). Then, similarly to what is done in (Jaksch et al., 2010, Sec. 4.3.2) and (Fruit et al., 2018b, App. F.7, pg. 32), we have

$$\begin{aligned} \xi_k &= \sum_{t=t_k}^{t_{k+1}-1} \underbrace{\sum_a \int p(s'|s_t, a_t) \pi_k(s_t, a_t) w_k(s') ds' - w_k(s_{t+1})}_{:=X_t} + \sum_{t=t_k}^{t_{k+1}-1} w_k(s_{t+1}) - w_k(s_t) \\ &= \sum_{t=t_k}^{t_{k+1}-1} X_t + \underbrace{w_k(s_{t_{k+1}}) - w_k(s_{t_k})}_{\leq sp \{w_k\} \leq c} \end{aligned}$$

Given the filtration  $\mathcal{F}_t = \sigma(s_1, a_1, r_1, \dots, s_{t+1})$ ,  $X_t$  is an MDS since  $|X_t| \leq c$  and  $\mathbb{E}[X_t | \mathcal{F}_{t-1}] = 0$  since  $\pi_{k_t}$  is  $\mathcal{F}_{t-1}$ -measurable. By using Azuma inequality we have that with probability at least  $1 - \frac{\delta}{20T^{5/4}}$

$$\sum_{k=1}^m \xi_k \leq c \sqrt{\frac{5}{2} T \ln \left( \frac{11T}{\delta} \right)} + mc \quad (29)$$

with  $m \leq SA \log_2 \left( \frac{8T}{SA} \right)$  when  $T \geq SA$  (see App. C.2 in (Jaksch et al., 2010)).

## A.3.2 BOUNDING THE EXPLORATION BONUS

Using the same argument in App. F.6 in (Fruit et al., 2018b) and by noticing that  $d_k(s, a) \leq 2c + r_{\max} \leq 2 \max\{c, r_{\max}\}$ , we obtain with probability at least  $1 - \frac{\delta}{20T^{5/4}}$ :

$$\sum_{k=1}^m \sum_{s,a} \nu_k(s) \pi_k(s, a) d_k(s, a) \leq \sum_{k=1}^m \sum_{s,a} \nu_k(s, a) d_k(s, a) + 2 \max\{c, r_{\max}\} \sqrt{\frac{5}{2} T \ln \left( \frac{11T}{\delta} \right)} \quad (30)$$

We now gather inequalities (30), (29) and a result in (Fruit et al., 2018b, App. F.7, pg. 33) into inequality (28) summed over all the episodes  $k$  which yields (after taking a union bound) that with probability at least  $1 - \frac{2\delta}{20T^{5/4}} \geq 1 - \frac{3\delta}{20T^{5/4}}$  (for  $T \geq SA$ ):

$$\begin{aligned} \sum_{k=1}^m \Delta_k \mathbb{1}\{g_c^*(\widehat{M}_k) \geq g^*\} &\leq 3 \max\{c, r_{\max}\} \sqrt{\frac{5}{2} T \ln \left( \frac{11T}{\delta} \right)} + cSA \log_2 \left( \frac{8T}{SA} \right) \\ &\quad + 2r_{\max} (\sqrt{2} + 1) \sqrt{SAT} + 2 \sum_{k=1}^m \sum_{s,a} \nu_k(s, a) d_k(s, a) \end{aligned} \quad (31)$$

Let  $\phi_{p,k}^{sa}$  as defined in Eq. 12, then

$$\begin{aligned} \sum_{k=1}^m \sum_{s,a} \nu_k(s, a) d_k(s, a) &\leq \underbrace{\sum_{k=1}^m \sum_{s,a} \nu_k(s, a) \beta_{r,k}^{sa}}_{\text{see p. 33 (Fruit et al., 2018b)}} + c \underbrace{\sum_{k=1}^m \sum_{s,a} \nu_k(s, a) \phi_{p,k}^{sa}}_{\text{see p. 33 (Fruit et al., 2018b)}} \\ &\quad + 2c \sum_{k=1}^m \sum_{s,a} \frac{\nu_k(s, a)}{N_k(s, a) + 1} + \underbrace{(c + r_{\max})}_{\leq 2 \max\{c, r_{\max}\}} LS^{-\alpha} T \end{aligned} \quad (32)$$

We recall that

$$\sum_{k=1}^m \sum_{s,a} \nu_k(s, a) \beta_{r,k}^{sa} = \widetilde{O}(r_{\max} \sqrt{SAT}) \quad \text{and} \quad \sum_{k=1}^m \sum_{s,a} \nu_k(s, a) \phi_{p,k}^{sa} = \widetilde{O}(cS\sqrt{AT} + cS^2A).$$

Similarly to what done in (Fruit et al., 2018b, Eq. 58-60), we can write

$$\sum_{k=1}^m \sum_{s,a} \frac{\nu_k(s, a)}{N_k(s, a) + 1} \leq 2 \sum_{s,a} \sum_{t=1}^T \frac{\mathbb{1}_{\{(s_t, a_t) = (s, a)\}}}{N_k(s_t, a_t) + 1} = 2 \sum_{s,a} \sum_{j=1}^{N_{T+1}(s, a)} \frac{1}{j+1} \leq 2SA \ln(T) \quad (33)$$

## A.3.3 COMPLETING THE PROOF

Summing up all the contributions and taking a union bound over all possible values of  $T$  and use the fact that  $\sum_{T=2}^{+\infty} \frac{\delta}{4T^{5/4}} < \delta$ , we write that there exists a numerical constant  $\chi$  such that at least with probability  $1 - \delta$  our algorithm C-SCAL<sup>+</sup> has a regret bounded by

$$\Delta(\text{C-SCAL}^+, T) \leq \chi \left( \max\{r_{\max}, c\} \left( S \sqrt{AT \ln \left( \frac{T}{\delta} \right)} + S^2 A \ln^2 \left( \frac{T}{\delta} \right) + LS^{-\alpha} T \right) \right)$$

We now set  $S = \left(\alpha L \sqrt{\frac{T}{A}}\right)^{1/(\alpha+1)}$  so that

$$\begin{aligned} \Delta(\text{C-SCAL}^+, T) = \tilde{O} & \left( \max\{r_{\max}, c\} \underbrace{\left( \max\left\{ \alpha^{1/(\alpha+1)}, \alpha^{-\alpha/(1+\alpha)} \right\} \right)}_{\leq 2, \forall \alpha \geq 0} \times \right. \\ & \left. \times L^{1/(\alpha+1)} A^{\alpha/(2\alpha+2)} T^{(\alpha+2)/(2\alpha+2)} + \underbrace{\alpha^{2/(1+\alpha)} L^{2/(1+\alpha)} A^{\alpha/(1+\alpha)} T^{1/(1+\alpha)}}_{\leq 2, \forall \alpha > 0} \right) \end{aligned}$$

Finally, when  $T \geq L^{2/\alpha} A$ , the regret of C-SCAL<sup>+</sup> is bounded with probability at least  $1 - \delta$  by

$$\Delta(\text{C-SCAL}^+, T) = \tilde{O} \left( \max\{r_{\max}, c\} L^{1/(\alpha+1)} A^{\alpha/(2\alpha+2)} T^{(\alpha+2)/(2\alpha+2)} \right).$$

## Appendix B. Finite MDPs: the analysis of SCAL<sup>+</sup>

In this section we analyse SCAL<sup>+</sup> by leveraging the results provided for the continuous state case. We start presenting the derivation of the bonus  $b_k$  and an analogous of Lem. 12 which implies SCAL<sup>+</sup> is optimistic at each episode  $k$ . Finally, we provide the proof of the regret bound stated in Thm. 8.

### B.1 High probability bound using the exploration bonus (proof of Lem. 4)

To begin with, we introduce two variants of the exploration bonus that will be used for the regret proof:

$$\begin{aligned} b_k(s, a) &:= c \cdot \min \left\{ \beta_{p,k}^{sa} + \frac{1}{N_k(s, a) + 1}; 2 \right\} + \min \{ \beta_{r,k}^{sa}; r_{\max} \} \\ d_k(s, a) &:= c \cdot \min \left\{ \phi_{p,k}^{sa} + \frac{1}{N_k(s, a) + 1}; 2 \right\} + \min \{ \beta_{r,k}^{sa}; r_{\max} \} \end{aligned} \tag{34}$$

where  $\beta_{p,k}^{sa} = \beta_k^{sa}$  (see Eq. 5),  $\beta_{r,k}^{sa} = r_{\max} \beta_k^{sa}$  and

$$\phi_{p,k}^{sa} := \sqrt{\frac{7(\Gamma - 1) \ln \left( \frac{3SA t_k}{\delta} \right)}{\max\{1, N_k(s, a)\}}} + \frac{14S}{\max\{1, N_k(s, a)\}} \ln \left( \frac{3SA t_k}{\delta} \right)$$

Notice that compared to the bonus  $b_k$ ,  $d_k$  explicitly depends on the number of states (linearly in  $S$ ) and next states (sublinearly in  $\Gamma$ ). As a consequence,  $d_k(s, a) \geq b_k(s, a)$  for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$ . In the continuous case we might consider the number of next states in the (true) aggregated MDP. However, this quantity is not very informative so we have decided (for sake of clarity) to upper-bound it by the number of intervals.

**Lemma 14** *Let  $(g^*, h^*)$  be a solution of the optimality equation  $Lh^* = h^* + g^*$  such that  $sp\{h^*\} \leq c$ . For all  $T \geq 1$  and  $k \geq 1$ , with probability at least  $1 - \frac{\delta}{15t_k^6}$ , for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$*

and for any  $v \in \mathbb{R}^S$  s.t.  $\text{sp}\{v\} \leq c$  we have:

$$\begin{aligned} (a) \quad b_k(s, a) &\geq \left| \bar{r}_k(s, a) - r(s, a) + (\widehat{p}_k(\cdot|s, a) - p(\cdot|s, a))^\top h^* \right| \\ (b) \quad d_k(s, a) &\geq \left| \bar{r}_k(s, a) - r(s, a) + (\widehat{p}_k(\cdot|s, a) - p(\cdot|s, a))^\top v \right| \end{aligned}$$

where  $b_k$  and  $d_k$  are defined as in Eq. 34.

**Proof** We consider the discrete case as a special sub-case of the continuous one considered in Lem. 12. As explained in Sec. A.1.3, for the discrete case we can even use an independence argument based on “stack of samples” idea as done for bandits (Lattimore and Szepesvári, 2018, Sec. 4.4). However, for sake of clarity we use the same MDS argument exploited in the continuous case. The main difference is that in the discrete case we do not need state aggregation and thus we replace every interval with a singleton function, i.e.,  $I(s) = s$ ,  $\forall s \in \mathcal{S}$ . Define  $w := h^* - (\min\{h^*\} + \max\{h^*\})/2$  such that  $w \in [-c/2, c/2]$ . We decompose  $\widehat{p}_k - p$  into  $(\widehat{p}_k - \bar{p}_k) + (\bar{p}_k - p)$ . As done in Eq. 17 (Sec. A.1.2), we can write that

$$\begin{aligned} |(\widehat{p}_k(\cdot|s, a) - \bar{p}_k(\cdot|s, a))^\top w| &\leq \left| \frac{N_k(s, a)}{N_k(s, a) + 1} - 1 \right| \underbrace{\|\bar{p}_k(\cdot|s, a)\|_1}_{=1} \underbrace{\|w\|_\infty}_{\leq c/2} + \frac{|w(\bar{s})|}{N_k(s, a) + 1} \\ &\leq \frac{c}{2} \left( 1 - \frac{N_k(s, a)}{N_k(s, a) + 1} + \frac{1}{N_k(s, a) + 1} \right) = \frac{c}{N_k(s, a) + 1} \end{aligned} \quad (35)$$

In order to bound the term depending on  $(\bar{p}_k - p)$  we use the same MDS argument as in Sec. A.1.3. You can consider  $p$  equivalent to  $\tilde{p}$  defined in the continuous case since:

$$\tilde{r}_k(s, a) = \frac{1}{N_k(\underbrace{I(s)}_{:=s}, a)} \sum_{x \in \underbrace{I(s)}_{:=s}} N_k(x, a) r(x, a) = r(s, a).$$

Similarly, we can prove that  $\tilde{p}_k(s'|s, a) = p(s'|s, a)$ . Then, we consider the same adapted sequences, stopping times and predictable processes except from the fact that intervals are replaced by singletons (i.e., discrete states). As a consequence, (an analogous of) Lem. 30 holds. By following the same steps in Sec. A.1.3, we can prove that with probability at least  $1 - \delta$ , for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$

$$\begin{aligned} |\bar{r}_k(s, a) - r_k(s, a)| &\leq r_{\max} \sqrt{\frac{7 \ln \left( \frac{2SA_t k}{\delta} \right)}{N_k^+(s, a)}} := \beta_{r,k}^{sa} \\ \left| (\bar{p}_k(\cdot|s, a) - p(\cdot|s, a))^\top h^* \right| &\leq c \sqrt{\frac{7 \ln \left( \frac{2SA_t k}{\delta} \right)}{N_k^+(s, a)}} := c\beta_{p,k}^{sa} \end{aligned}$$

where we recall the  $N_k^+(s, a) := \max\{1, N_k(s, a)\}$ . We now consider the concentration of  $(\widehat{p}_k - p)^\top v$  for which we need to use Freedman’s inequality (see Thm. 29). Similarly to what done before, let  $z = v - (\min\{v\} + \max\{v\})/2$  such that  $(\widehat{p}_k - p)^\top v = (\widehat{p}_k - p)^\top z$ . We start noticing that, Eq. 19 holds for the discrete case where we replace

the adapted sequence  $\mathbb{1}(s_{\tau_l+1} \in I)$  by  $\mathbb{1}(s_{\tau_l+1} = s')$  and the conditional variance  $V_k(J)$  by  $V_k(s') = \sum_{l=1}^{N_k(s,a)} \mathbb{1}(\tau_l < t_k) \mathbb{V}(\mathbb{1}(s_{\tau_l+1} = s') | \mathcal{G}_{l-1})$ . Furthermore,  $\mathbb{V}(\mathbb{1}(s_{\tau_l+1} = s') | \mathcal{G}_{l-1}) = (1 - p(s'|s, a))p(s'|s, a)$  and

$$V_k(s') = \sum_{l=1}^{N_k(s,a)} (1 - p(s'|s_{\tau_l}, a_{\tau_l}))p(s'|s_{\tau_l}, a_{\tau_l}) \leq N_k(s, a)(1 - p(s'|s, a))p(s'|s, a)$$

As done in (Fruit et al., 2018b, App. F.7) we use Cauchy-Schwartz inequality to write that

$$\sum_{s' \in \mathcal{S}} \sqrt{p(s'|s, a)(1 - p(s'|s, a))} = \sum_{s' \in \mathcal{S}: p(s'|s, a) > 0} \sqrt{p(s'|s, a)(1 - p(s'|s, a))} \leq \sqrt{\Gamma - 1}$$

where we recall that  $\Gamma := \max_{s,a} \|p(\cdot|s, a)\|_0$  is the maximum support of  $p$ . Then, for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and for any vector  $z \in [-c/2, c/2]$ , we have that

$$\begin{aligned} \left| (\bar{p}_k(\cdot|s, a) - p(\cdot|s, a))^\top z \right| &\leq \|z\|_\infty \sum_{s' \in \mathcal{S}} |\bar{p}_k(s'|s, a) - p(s'|s, a)| \\ &\leq \frac{c}{2N_k^+(s, a)} \sum_{s' \in \mathcal{S}} \left( 2\sqrt{V_k(s') \ln\left(\frac{4t_k}{\delta}\right)} + 4\ln\left(\frac{4t_k}{\delta}\right) \right) \\ &\leq c \left( \sqrt{\frac{7(\Gamma - 1) \ln\left(\frac{3SA t_k}{\delta}\right)}{N_k^+(s, a)}} + \frac{14S}{N_k^+(s, a)} \ln\left(\frac{3SA t_k}{\delta}\right) \right) := c \phi_{p,k}^{sa} \end{aligned}$$

We can also write with probability 1 that:

$$\left| (\hat{p}_k(\cdot|s, a) - p(\cdot|s, a))^\top w \right| \leq \hat{p}_k(\cdot|s, a)^\top w + p(\cdot|s, a)^\top w \leq 2c$$

So we can take the minimum between the two upper-bounds. We also know that the difference in reward is bound by  $r_{\max}$ . ■

In order to prove optimism we start noticing that the bonus  $b_k(s, a)$  (see Lem. 14) implies that  $\hat{L}_k h^* \geq Lh^*$ . As a consequence, we can use Prop. 3 to show that  $g_k^+ \geq g^*$ .

## B.2 Regret Proof of SCAL<sup>+</sup>

The regret proof follows the same steps of the one for C-SCAL<sup>+</sup>. The main difference resides in the fact that there is no need of state aggregation, thus simplifying the proof.

By using the optimism of  $\widehat{M}_k^+$ , the stopping condition of SCOPT and the relationships between  $\widehat{M}_k^+$  and  $\widehat{M}_k$  (see Eq. 26), we can prove Eq. 28 for the discrete case. Note that the analysis of the cumulative contribution of the term  $d_k(s, a)$  and  $b_k(s, a)$  will lead to the following terms  $\tilde{O}(c\sqrt{\Gamma SAT})$  and  $\tilde{O}(c\sqrt{SAT})$ , respectively. Since the dominant term is the one associated to  $d_k$ , even in this case we upper-bound  $b_k$  by  $d_k$ .

From this point, we follow the same steps as in Sec. A.3. The only difference resides in Eq. 31 where the term  $(c + r_{\max})LS^{-\alpha}T$  disappears since it depends on aggregation and/or

smoothness. Finally, the regret bound in Thm. 6 follows by noticing that the order of the term  $\sum_{k=1}^m \sum_{s,a} \nu_k(s,a) \phi_{p,k}^{sa}$  is  $\tilde{O}(\sqrt{\Gamma SAT} + S^2 A)$ .

As a consequence, there exists a numerical constant  $\chi$  such that at least with probability  $1 - \delta$  our algorithm SCAL<sup>+</sup> has a regret bounded by

$$\Delta(\text{SCAL}^+, T) \leq \chi \left( \max\{r_{\max}, c\} \left( \sqrt{\Gamma SAT \ln \left( \frac{T}{\delta} \right)} + S^2 A \ln^2 \left( \frac{T}{\delta} \right) \right) \right)$$

## Appendix C. Results of probability theory

### C.1 Reminder

We start by recalling some well-known properties of filtrations, stopping times and martingales (Klenke and Loève, 2013, Chapter 2). For simplicity, we use “a.s.” to denote “almost surely” (i.e., with probability 1). In this section, we consider a probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ . We call *filtration* any *increasing* (for the inclusion) sequence of sub- $\sigma$ -algebras of  $\mathcal{F}$  i.e.,  $(\mathcal{F}_n)_{n \in \mathbb{N}}$  where  $\forall n \in \mathbb{N}, \mathcal{F}_n \subseteq \mathcal{F}_{n+1} \subseteq \mathcal{F}$ . We denote by  $\mathcal{F}_\infty := \cup_{n \in \mathbb{N}} \mathcal{F}_n$ . For any sub- $\sigma$ -algebra  $\mathcal{G} \subseteq \mathcal{F}$ , we say that a real-valued random variable (r.v.)  $X : \Omega \rightarrow \mathbb{R}^d$  is  $\mathcal{G}$ -measurable if for all borel sets  $B \in \mathcal{B}(\mathbb{R}^d)$ ,  $X^{-1}(B) \in \mathcal{G}$ . We say that  $X$  is  $\mathcal{G}$ -integrable if it is  $\mathcal{G}$ -measurable and  $\mathbb{E}[|X|] < +\infty$  (componentwise). We call *stochastic process* any sequence of r.v.  $(X_n)_{n \in \mathbb{N}}$ . We say that the stochastic process  $(X_n)_{n \in \mathbb{N}}$  is *adapted* to the filtration  $(\mathcal{F}_n)_{n \in \mathbb{N}}$  if for all  $n \in \mathbb{N}$ ,  $X_n$  is  $\mathcal{F}_n$ -measurable. In this case, the sequence  $(X_n, \mathcal{F}_n)_{n \in \mathbb{N}}$  is called an *adapted sequence*. If in addition,  $X_n$  is integrable for all  $n \in \mathbb{N}$  then we say that  $(X_n, \mathcal{F}_n)_{n \in \mathbb{N}}$  is an *integrable adapted sequence*. We say that a stochastic process  $(X_n)_{n \in \mathbb{N}}$  is almost surely:

1. *increasing* (resp. *strictly increasing*) if for all  $n \geq N$ ,  $\mathbb{P}(X_n \leq X_{n+1}) = 1$  (resp.  $\mathbb{P}(X_n < X_{n+1}) = 1$ ),
2. *bounded* if there exists a universal constant  $K$  such that for all  $n \in \mathbb{N}$ ,  $\mathbb{P}(X_n < K) = 1$ ,

**Definition 15 (Conditional expectation)** *Let  $X$  be an  $\mathcal{F}$ -integrable r.v. with values in  $\mathbb{R}^d$ . Let  $\mathcal{G} \subseteq \mathcal{F}$  be a sub- $\sigma$ -algebra of  $\mathcal{F}$ . The conditional expectation of  $X$  given  $\mathcal{G}$  (denoted  $\mathbb{E}[X|\mathcal{G}]$ ) is the (a.s. unique) r.v. that is  $\mathcal{G}$ -integrable and satisfies:*

$$\forall A \in \mathcal{G}, \quad \mathbb{E}[\mathbb{1}(A) \cdot \mathbb{E}[X|\mathcal{G}]] = \mathbb{E}[\mathbb{1}(A) \cdot X]$$

**Proposition 16 (Law of total expeptations)** *Let  $X$  be an  $\mathcal{F}$ -integrable r.v. with values in  $\mathbb{R}^d$ . For any sub- $\sigma$ -algebra  $\mathcal{G} \subseteq \mathcal{F}$ ,  $\mathbb{E}[\mathbb{E}[X|\mathcal{G}]] = \mathbb{E}[X]$ .*

**Proposition 17** *Let  $X$  be an  $\mathcal{F}$ -integrable real-valued r.v. and  $\mathcal{G} \subseteq \mathcal{F}$  a sub- $\sigma$ -algebra. For any  $\mathcal{G}$ -integrable real-valued r.v.  $Y$  s.t.  $YX$  is also integrable we have  $\mathbb{E}[YX|\mathcal{G}] = Y\mathbb{E}[X|\mathcal{G}]$ .*

**Definition 18 (Stopping time)** *A random variable  $\tau : \Omega \rightarrow \mathbb{N} \cup \{+\infty\}$  is called stopping time w.r.t. a filtration  $(\mathcal{F}_n)_{n \in \mathbb{N}}$  if for all  $n \in \mathbb{N}$ ,  $\{\tau = n\} \in \mathcal{F}_n$ .*

**Definition 19 ( $\sigma$ -algebra at stopping time)** *Let  $\tau$  be a stopping time. An event prior to  $\tau$  is any event  $A \in \mathcal{F}_\infty$  s.t.  $A \cap \{\tau = n\} \in \mathcal{F}_n$  for all  $n \in \mathbb{N}$ . The set of events prior to  $\tau$  is a  $\sigma$ -algebra denoted  $\mathcal{F}_\tau$  and called  $\sigma$ -algebra at time  $\tau$ :*

$$\mathcal{F}_\tau := \{A \in \mathcal{F}_\infty : \forall n \in \mathbb{N}, A \cap \{\tau = n\} \in \mathcal{F}_n\}$$

**Proposition 20** *Let  $\tau_1$  and  $\tau_2$  be two stopping times w.r.t. the same filtration  $(\mathcal{F}_n)_{n \in \mathbb{N}}$  s.t.  $\tau_1 \leq \tau_2$  a.s. Then  $\mathcal{F}_{\tau_1} \subseteq \mathcal{F}_{\tau_2}$ .*

**Definition 21 (Stopped Process)** *Let  $(X_n, \mathcal{F}_n)_{n \in \mathbb{N}}$  be an adapted sequence with values in  $\mathbb{R}^d$ . If  $\tau$  is a stopping time w.r.t. the filtration  $(\mathcal{F}_n)_{n \in \mathbb{N}}$ , then the process stopped at time  $\tau$  (denoted by  $X_\tau$ ) is the r.v. defined as:*

$$\forall \omega \in \Omega, \quad X_\tau(\omega) := \sum_{n \in \mathbb{N}} X_n(\omega) \cdot \mathbb{1}(\tau(\omega) = n) \quad (\text{i.e., } X_\infty(\omega) = 0 \text{ by convention})$$

**Proposition 22**  *$X_\tau$  -the process stopped at time  $\tau$ - is  $\mathcal{F}_\tau$ -measurable.*

**Definition 23 (Martingale difference sequence)** *An adapted sequence  $(X_n, \mathcal{F}_n)_{n \in \mathbb{N}}$  is a martingale difference sequence (MDS for short) if for all  $n \in \mathbb{N}$ ,  $X_n$  is  $\mathcal{F}_n$ -integrable and  $\mathbb{E}[X_{n+1} | \mathcal{F}_n] = 0$  a.s.*

**Proposition 24 (Azuma's inequality)** *Let  $(X_n, \mathcal{F}_n)_{n \in \mathbb{N}}$  be an MDS such that  $|X_n| \leq a$  a.s. for all  $n \in \mathbb{N}$ . Then for all  $\delta \in ]0, 1[$ ,*

$$\mathbb{P} \left( \forall n \geq 1, \left| \sum_{i=1}^n X_i \right| \leq a \sqrt{n \ln \left( \frac{2n}{\delta} \right)} \right) \geq 1 - \delta$$

**Proof** Azuma's inequality states that:

$$\mathbb{P} \left( \left| \sum_{i=1}^n X_i \right| \leq a \sqrt{\frac{n}{2} \ln \left( \frac{2}{\delta} \right)} \right) \geq 1 - \delta$$

We can then choose  $\delta \leftarrow \frac{\delta}{2n^2}$  and take a union bound over all possible values of  $n \geq 1$ . The result follows by noting that  $\sum_{n \geq 1} \frac{\delta}{2n^2} < \delta$ . ■

**Proposition 25 (Freedman's inequality)** *Let  $(X_n, \mathcal{F}_n)_{n \in \mathbb{N}}$  be an MDS such that  $|X_n| \leq a$  a.s. for all  $n \in \mathbb{N}$ . Then for all  $\delta \in ]0, 1[$ ,*

$$\mathbb{P} \left( \forall n \geq 1, \left| \sum_{i=1}^n X_i \right| \leq 2 \sqrt{\left( \sum_{i=1}^n \mathbb{V}(X_i | \mathcal{F}_{i-1}) \right) \cdot \ln \left( \frac{4n}{\delta} \right) + 4a \ln \left( \frac{4n}{\delta} \right)} \right) \geq 1 - \delta$$

**Proof** Freedman (1975) showed that when  $a = 1$ :

$$\mathbb{P} \left( \forall n \geq 1, \sum_{i=1}^n X_i \geq \varepsilon, \sum_{i=1}^n \mathbb{V}(X_i | \mathcal{F}_{i-1}) \leq k \right) \leq \exp \left( \frac{-\varepsilon^2}{2k + 2\varepsilon/3} \right)$$

Since  $(-X_n, \mathcal{F}_n)_{n \in \mathbb{N}}$  is also an MDS, the above inequality holds also in absolute value (with a factor 2 appearing in front of the exponential term after taking a union bound). In order to reverse the inequality (i.e., replace  $\varepsilon$  by  $\delta$ ), we can use the same technique as Cesa-Bianchi and Gentile (2005, Section 2). Finally, to account for the case where  $a \neq 1$  we can simply apply the result to  $(X_n/a, \mathcal{F}_n)_{n \in \mathbb{N}}$ . ■



## C.2 A useful concentration with optional skipping

In this section we prove a very simple theorem inspired by Doob's optional skipping (e.g., Chow and Teicher, 1988, Sec. 5.3, Lem. 4). We start with useful definitions and lemmas.

**Lemma 26** *Let  $\tau_1$  and  $\tau_2$  be two stopping times w.r.t. the same filtration  $(\mathcal{F}_n)_{n \in \mathbb{N}}$ . We say that  $\tau_1 < \tau_2$  a.s. if  $\mathbb{P}(\{\tau_1 < \tau_2\} \cup \{\tau_1 = \tau_2 = +\infty\}) = 1$ . If  $\tau_1 < \tau_2$  a.s. then  $\mathcal{F}_{\tau_1+1} \subseteq \mathcal{F}_{\tau_2}$ .*

**Proof** If  $\tau_1 < \tau_2$  then  $\tau_1 + 1 \leq \tau_2$  since  $\tau_1$  is an integer-valued r.v. If  $\tau_1 = \tau_2 = +\infty$  then  $\tau_1 + 1 = +\infty$  and so  $\tau_1 + 1 = \tau_2$ . In conclusion,  $\tau_1 + 1 \leq \tau_2$  a.s. and so by Prop. 20,  $\mathcal{F}_{\tau_1+1} \subseteq \mathcal{F}_{\tau_2}$ .  $\blacksquare$

**Definition 27** *We say that a sequence of stopping times  $(\tau_m)_{m \in \mathbb{N}}$  w.r.t.  $(\mathcal{F}_n)_{n \in \mathbb{N}}$  is strictly increasing if  $\tau_m < \tau_{m+1}$  a.s. for all  $m \geq 0$ .*

**Lemma 28** *Let  $(X_n, \mathcal{F}_n)_{n \in \mathbb{N}}$  be a bounded adapted sequence and let  $(\tau_m)_{m \in \mathbb{N}}$  be a strictly increasing sequence of stopping times w.r.t.  $(\mathcal{F}_n)_{n \in \mathbb{N}}$ . For all  $m \in \mathbb{N}$ , define  $Y_m := X_{\tau_m+1} - \mathbb{E}[X_{\tau_m+1} | \mathcal{F}_{\tau_m}]$  and  $\mathcal{G}_m := \mathcal{F}_{\tau_{m+1}}$ . Then,  $(Y_m, \mathcal{G}_m)_{m \in \mathbb{N}}$  is an MDS.*

**Proof** By assumption, for any  $m \in \mathbb{N}$ ,  $\tau_m < \tau_{m+1}$  a.s. and Prop. 20 implies that  $\mathcal{F}_{\tau_m} \subseteq \mathcal{F}_{\tau_{m+1}}$ . As a consequence,  $(\mathcal{G}_m)_{m \in \mathbb{N}} = (\mathcal{F}_{\tau_{m+1}})_{m \in \mathbb{N}}$  is a filtration. By Prop. 22 we know that  $X_{\tau_m+1}$  is  $\mathcal{F}_{\tau_m+1}$ -measurable and Lem. 26 implies that  $\mathcal{F}_{\tau_m+1} \subseteq \mathcal{F}_{\tau_{m+1}} = \mathcal{G}_m$  so  $X_{\tau_m+1}$  is  $\mathcal{G}_m$ -measurable. Finally,  $\mathbb{E}[X_{\tau_m+1} | \mathcal{F}_{\tau_m}]$  is  $\mathcal{F}_{\tau_m}$ -measurable by definition (see Def. 15). Therefore,  $Y_m$  is  $\mathcal{G}_m$ -measurable.

Since by assumption  $X_n$  is a.s. bounded ( $\mathbb{P}(X_n < K) = 1$  for all  $n \geq 0$ ), we can write a.s. (see Def. 21)

$$|X_{\tau_m+1}| = \left| \sum_{n=0}^{+\infty} \mathbb{1}(\tau_m + 1 = n) \cdot X_n \right| \leq \sum_{n=0}^{+\infty} \mathbb{1}(\tau_m + 1 = n) \cdot |X_n| \leq K \sum_{n=0}^{+\infty} \mathbb{1}(\tau_m + 1 = n) = K$$

Thus,  $X_{\tau_m+1}$  is a.s. bounded hence integrable implying that  $\mathbb{E}[X_{\tau_m+1} | \mathcal{F}_{\tau_m}]$  is well-defined (see Def. 15). Therefore,  $Y_m$  is a.s. bounded and so integrable.

Finally, we can apply Prop. 17 and we obtain:

$$\mathbb{E}[Y_{m+1} | \mathcal{G}_m] = \mathbb{E}[X_{\tau_{m+1}} - \mathbb{E}[X_{\tau_{m+1}} | \mathcal{F}_{\tau_m}] | \mathcal{F}_{\tau_m}] = \mathbb{E}[X_{\tau_{m+1}} | \mathcal{F}_{\tau_m}] - \mathbb{E}[X_{\tau_{m+1}} | \mathcal{F}_{\tau_m}] = 0$$

which concludes the proof.  $\blacksquare$

**Theorem 29** *Let  $(X_n, \mathcal{F}_n)_{n \in \mathbb{N}}$  be an adapted sequence a.s. bounded by  $a_1$  and let  $(\tau_m)_{m \in \mathbb{N}}$  be a strictly increasing sequence of stopping times w.r.t.  $(\mathcal{F}_n)_{n \in \mathbb{N}}$ . If  $(Y_m, \mathcal{G}_m)_{m \in \mathbb{N}}$  is defined as in Lem. 28 then the following concentration inequalities hold:*

$$\begin{aligned} \mathbb{P} \left( \forall m \geq 1, \left| \sum_{i=1}^m Y_m \right| \leq a_1 \sqrt{m \ln \left( \frac{2m}{\delta} \right)} \right) &\geq 1 - \delta \\ \mathbb{P} \left( \forall m \geq 1, \left| \sum_{i=1}^m Y_m \right| \leq 2 \sqrt{\left( \sum_{i=1}^m \mathbb{V}(Y_i | \mathcal{G}_{i-1}) \right) \cdot \ln \left( \frac{4m}{\delta} \right) + 4a_1 \ln \left( \frac{4m}{\delta} \right)} \right) &\geq 1 - \delta \end{aligned}$$

In particular for any  $\mathcal{F}$ -measurable integer-valued r.v.  $N : \Omega \rightarrow \mathbb{N}$  the above inequalities hold true with  $m$  replaced by  $N$  e.g.,

$$\mathbb{P} \left( \left| \sum_{i=1}^N Y_m \right| \leq a_1 \sqrt{N \ln \left( \frac{2N}{\delta} \right)} \right) \geq 1 - \delta \quad \dots$$

**Proof** The concentration inequalities follow from Lem. 28 and Azuma's and Freedman's inequalities. If the results hold for all  $n \in \mathbb{N}$  and  $N$  takes values in  $\mathbb{N}$ , then the result holds for  $N$  too which concludes the proof.  $\blacksquare$

### C.3 In the regret proof

For any  $t \geq 0$ , the  $\sigma$ -algebra induced by the past history of state-action pairs and rewards up to time  $t$  is denoted  $\mathcal{F}_t := \sigma(s_1, a_1, r_1, \dots, s_t, a_t)$  where by convention  $\mathcal{F}_0 = \sigma(\emptyset)$  and  $\mathcal{F}_\infty := \cup_{t \geq 0} \mathcal{F}_t$ . Trivially, for all  $t \geq 0$ ,  $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$  and the filtration  $(\mathcal{F}_t)_{t \geq 0}$  is denoted by  $\mathbb{F}$ . We recall that the sequence  $(t_k)_{k \geq 1}$  (starting times of episodes  $k \geq 1$ ) is formally defined by  $t_1 := 1$  and for all  $k \geq 1$ ,

$$t_{k+1} := 1 + \inf \left\{ T \geq t > t_k : \sum_{u=t_k}^{t-1} \mathbb{1}(s_u \in I(s), a_u = a) \geq \sum_{u=0}^{t_k-1} \mathbb{1}(s_u \in I(s), a_u = a) \right\}.$$

where by convention  $\inf\{\emptyset\} := T$ . It is immediate to see that for all  $t \geq 0$ ,  $\{t_k = t\} \in \mathcal{F}_{t-1} \subseteq \mathcal{F}_t$  and so  $t_k$  is a *stopping time* w.r.t. filtration  $\mathbb{F}$  (see Def. 18).

The following lemma is used in App. A.1.3:

**Lemma 30** *For all  $l \geq 1$ , we have:*

1.  $\mathbb{E}[w^*(s_{\tau_l+1}) | \mathcal{G}_{l-1}] = \int_{\mathcal{S}} p(s' | s_{\tau_l}, a_{\tau_l}) w^*(s') ds'$ ,
2.  $\mathbb{E}[\mathbb{1}(s_{\tau_l+1} \in J) | \mathcal{G}_{l-1}] = \int_J p(s' | s_{\tau_l}, a_{\tau_l}) ds'$ ,
3. and  $\mathbb{E}[r_{\tau_l}(s_{\tau_l}, a_{\tau_l}) | \mathcal{G}_{l-1}] = r(s_{\tau_l}, a_{\tau_l})$ .

**Proof** To prove this result, we rely on the definition of conditional expectation (see Def. 15).

1) By Prop. 22,  $(s_{\tau_l}, a_{\tau_l})$  is  $\mathcal{G}_{l-1}$ -measurable ( $\mathcal{G}_{l-1} = \mathcal{F}_{\tau_l}$ ) and so  $\int_{\mathcal{S}} p(s' | s_{\tau_l}, a_{\tau_l}) w^*(s') ds'$  is  $\mathcal{G}_{l-1}$ -measurable too. Moreover,  $|\int_{\mathcal{S}} p(s' | s_{\tau_l}, a_{\tau_l}) w^*(s') ds'| \leq c/2$  a.s. so  $\int_{\mathcal{S}} p(s' | s_{\tau_l}, a_{\tau_l}) w^*(s') ds'$  is also integrable (and therefore  $\mathcal{G}_{l-1}$ -integrable).

2) We recall that for any stochastic process  $(X_t)_{t \geq 0}$ , we use the convention that  $X_\infty = 0$  a.s. implying that  $X_{\tau_l} = \sum_{t=0}^{+\infty} X_t \mathbb{1}(\tau_l = t)$  (see Def. 21). Using the law of total expectations (see Prop. 16) we have that  $\forall A \in \mathcal{G}_{l-1}$ ,

$$\begin{aligned} \mathbb{E}[\mathbb{1}(A) \times w^*(s_{\tau_l+1})] &= \sum_{t=0}^{+\infty} \mathbb{E}[\mathbb{1}(A \cap \{\tau_l = t\}) \times w^*(s_{\tau_l+1})] \\ &= \sum_{t=0}^{+\infty} \mathbb{E} \left[ \underbrace{\mathbb{E}[\mathbb{1}(A \cap \{\tau_l = t\})]}_{\in \mathcal{F}_t} \times \underbrace{w^*(s_{\tau_l+1})}_{=s_{t+1}} \middle| \mathcal{F}_t \right] \end{aligned}$$

In the first equality, the fact that we can move the sum outside the expectation is a direct consequence of the dominated convergence theorem (for series) since

$$\begin{aligned} \sum_{t=0}^{+\infty} \mathbb{E}[\mathbb{1}(A \cap \{\tau_l = t\}) \times |w^*(s_{\tau_l+1})|] &\leq c/2 \sum_{t=0}^{+\infty} \mathbb{E}[\mathbb{1}(A \cap \{\tau_l = t\})] \\ &= c/2 \sum_{t=0}^{+\infty} \mathbb{P}(A \cap \{\tau_l = t\}) = c/2 \cdot \mathbb{P}(A) < +\infty \end{aligned}$$

Under event  $\{\tau_l = t\}$  we have that  $s_{\tau_l+1} = s_{t+1}$  a.s. Moreover,  $A \cap \{\tau_l = t\} \in \mathcal{F}_t$  since  $\tau_l$  is a stopping time (see Def. 19) so by Prop. 17 we can move it outside the conditional expectation and we get:

$$\begin{aligned} \mathbb{E}[\mathbb{1}(A) \times w^*(s_{\tau_l+1})] &= \sum_{t=0}^{+\infty} \mathbb{E} \left[ \mathbb{1}(A \cap \{\tau_l = t\}) \times \underbrace{\mathbb{E}[w^*(s_{t+1}) | \mathcal{F}_t]}_{=\int_{\mathcal{S}} p(s'|s_t, a_t) w^*(s') ds'} \right] \\ &= \mathbb{E} \left[ \mathbb{1}(A) \times \underbrace{\sum_{t=0}^{+\infty} \mathbb{1}(\tau_l = t) \int_{\mathcal{S}} p(s'|s_t, a_t) w^*(s') ds'}_{=\int_{\mathcal{S}} p(s'|s_{\tau_l}, a_{\tau_l}) w^*(s') ds' \text{ (see Def. 21)}} \right] \\ &= \mathbb{E} \left[ \mathbb{1}(A) \times \int_{\mathcal{S}} p(s'|s_{\tau_l}, a_{\tau_l}) w^*(s') ds' \right] \end{aligned}$$

This proves the first inequality (see Def. 15). The second and third equality can be proved using the same technique. ■