Concurrent Reinforcement Learning with Aggregated States via Randomized Least Squares Value Iteration

Yan Chen¹ Qinxun Bai² Yiteng Zhang³ Maria Dimakopoulou⁴ Shi Dong⁵ Qi Sun⁶ Zhengyuan Zhou⁶⁷

Abstract

Designing learning agents that explore efficiently in a complex environment has been widely recognized as a fundamental challenge in reinforcement learning. While a number of works have demonstrated the effectiveness of techniques based on randomized value functions on a single agent, it remains unclear, from a theoretical point of view, whether injecting randomization can help a society of agents concurently explore an environment. The theoretical results established in this work tender an affirmative answer to this question. We adapt the concurrent learning framework to randomized least-squares value iteration (RLSVI) with aggregated state representation. We demonstrate polynomial worst-case regret bounds in both finite- and infinite-horizon environments. In both setups the per-agent regret decreases at an optimal rate of $\Theta\left(\frac{1}{\sqrt{N}}\right)$, highlighting the advan-tage of concurent learning. Our algorithm exhibits significantly lower space complexity compared to (Russo, 2019) and (Agrawal et al., 2021). We reduce the space complexity by a factor of Kwhile incurring only a \sqrt{K} increase in the worstcase regret bound, compared to (Agrawal et al., 2021; Russo, 2019). Interestingly, our algorithm improves the worst-case regret bound of (Russo, 2019) by a factor of $H^{1/2}$, matching the improvement in (Agrawal et al., 2021). However, this result is achieved through a fundamentally different algorithmic enhancement and proof technique. Additionally, we conduct numerical experiments to demonstrate our theoretical findings.

1. Introduction

The field of reinforcement learning (RL) is dedicated to designing agents that interact with an unknown environment, aiming to maximize the total amount of reward accumulated throughout the interactions (Sutton & Barto, 2018). When the environment is complex yet the learning budget is limited, an agent has to efficiently explore the environment, giving rise to the well-known exploration-exploitation trade-off. A large body of works in the RL literature have addressed the challenge related to smartly balancing this trade-off. Among the many proposed methods, algorithms based on randomization are receiving growing attention, both theoretically (Russo & Van Roy, 2014; Fellows et al., 2021) and empirically (Osband et al., 2016; Janz et al., 2019; Dwaracherla et al., 2020), due to their effectiveness and potential scalability in large applications.

Randomized least-squares value iteration (RLSVI) represents one example of such randomization-based algorithms (Osband et al., 2019). On a high level, RLSVI injects Gaussian noise into the rewards in the agent's previous trajectories, and allows the agent to learn a randomized value function from the perturbed dataset. With judicious noise injection, the resultant value function approximates the agent's posterior belief of state values. By acting greedily with respect to such randomized value function, the agent effectively executes an approximated version of posterior sampling for reinforcement learning (PSRL), whose efficacy has been substantiated in previous works (Osband et al., 2013; Russo & Van Roy, 2014; Xu et al., 2022). Compared with PSRL, RLSVI circumvents the need of maintaining a model of the environment, severing huge computational cost. Since its advent, RLSVI has been studied extensively in theoretical contexts, such as (Russo, 2019) and (Ishfaq et al., 2021).

In this work, we look into RLSVI from the perspective of concurrent learning (Silver et al., 2013). Specifically, concurrent RL studies the problem where a cohort of agents interact with one common environment, yet are able to share their experience with each other in order to jointly improve their decisions. Such a collaborative setting has useful applications in a variety of realms, such as robotics (Gu et al., 2017), biology (Sinai et al., 2020), and recommendation

¹Duke University ²Horizon Robotics Inc ³UC Berkeley ⁴Uber ⁵Google DeepMind ⁶New York University ⁷Arena Technologies. Correspondence to: Zhengyuan Zhou <z@arena-ai.com>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

systems (Agarwal et al., 2016). For instance, a team of drones may be cooperatively executing a search and rescue mission in a given terrain, and cooperatively learning to identify the locations of assets requires efficient concurrent RL. Another example is a team of virtual agents conducting circuit design, proceeding in parallel the design-circuit evaluation (via a simulator)-redesign loop. Here, they are cooperatively exploring the vast and complex design space, aiming to quickly identify an optimal design that meets all the sepc requirements¹. Importantly, although all agents share the identical goal, coordination between agents is nontrivial. In fact, as shown in (Dimakopoulou & Van Roy, 2018; Dimakopoulou et al., 2018), a poorly coordinated multi-agent algorithm can drastically undermine the overall learning performance. Results in the existing theoretical literature on concurrent RL have demonstrated that PSRL (Osband & Van Roy, 2017; Kim, 2017; Osband & Van Roy, 2014) in a coordinated style is provably efficient (Chen et al., 2022), compared to the earlier cooperative UCRL type of algorithms². As pointed out in (Chen et al., 2022), "these sample-complexity guarantees notwithstanding, concurrent UCRL algorithms suffer from the critical disadvantage of no coordinated exploration: since the upper confidence bounds computed by aggregating all agents' data are the same the entire team, each agent would follow the exact same policy, thereby yielding no diversity in exploration." However, despite the exploration advantage offered by concurrent Thompson sampling given in a very recent work (Chen et al., 2022), it remains unclear whether the same level of efficiency can be extended to cooperative model-free agents, which are computationally more scalable.

Our Contributions We model the environment as a Markov decision process (MDP) with Γ aggregated states, and show theoretically that a cohort of N agents, concurrently running RLSVI and sharing interaction trajectories with each other, are able to efficiently improve their joint

performance towards the optimal policy in the environment. The efficiency is established through regret analysis. Additionally, for aggregated state representation setup, the performance typically will not converge to optimal solution if $\epsilon > 0$, and our result establishes that the cumulative performance loss is no greater than $O(\epsilon)$ per period per agent. This is consistent with the findings of (Dong et al., 2019b) for single-agent scenario. Our worst-case regret bound improves upon (Russo, 2019) by a factor of $H^{1/2}$, matching the improvement in (Agrawal et al., 2021) but achieved through a fundamentally different technique and proof approach.

One major difference between our algorithm and those from current literature of RLSVI in tabular setup is that, the algorithms in (Russo, 2019; Agrawal et al., 2021) require storing the historical trajectories from the very beginning while our algorithm only stores the trajectories from the last episode for the finite-horizon case or the last pseudoepisode (defined in Section 4) for the infinite-horizon case. This is because for concurrent setup, the computational cost of storing all historical data scales at least linear in KHNfor finite horizon and TN for infinite horizon, which is infeasible for problems of practical scale.

In both cases, the regret dependence on the total number of samples is optimal, signaling well-coordinated information sharing among the agents. To our best knowledge, this work presents the first theoretical analysis of a model-free concurrent RL algorithm with aggregated states.

Besides theoretical contributions, our analysis also sheds light on the empirical role of discount factor in RL. In practical applications, a discount factor is usually not subsumed under the environment definition. Rather than prescribing a specific learning target that involves a decaying sequence of rewards, discount factors often simply function as a tuning parameter of the algorithm that keeps the value updates stable. In this work, similar as in (Xu et al., 2022), when the decision horizon is infinite, the discount factor only shows up in the algorithm and does not appear in the learning target. The sole purpose of introducing a discount factor is to decompose the single stream of interactions into "pseudoepisodes" that facilitate decision-making. Such view aligns with the authentic role of discount factor in agent design, and distinguishes this work from earlier RL theory literature on discounted regret such as (Dong et al., 2019a). While (Xu et al., 2022) focuses on PSRL, this paper is the first to extend the same view to a model-free algorithm that enjoys empirical successes.

Related Work *Multi-agent Reinforcement Learning* (MARL) has been widely studied to address problems in a variety of applications like robotics, telecommunications and e-commerce (Buşoniu et al., 2010). Scenarios of MARL

¹For instance, Arena, a world leading startup that specializes in applying AI to shorten the hardware design loop, utilizes RL to speed up the circuit design, test and debugging cycles.

²It is perhaps not surprising that UCRL algorithms are the first family of algorithms that have been adapted to the concurrent RL settings (Guo & Brunskill, 2015) and (Pazis & Parr, 2016). In particular, concurrent UCRL has been analyzed for the sample complexity performance measure (i.e. how many samples are needed to learn an ϵ -optimal policy) under both finite action space setting and infinite action space setting. More specifically, (Guo & Brunskill, 2015) provided a high-probability bound of $\tilde{O}(\frac{S^2A}{\epsilon^3} + \frac{SAn}{\epsilon})$ for the sample complexity with n agents interacting in the environment. Their algorithm was extended from MBIE(see (Strehl & Littman, 2008)), with a single agent performing concurrent RL across a set of n infinite-horizon MDPs. The results there show that with sharing samples from copies of the same MDP a linear speedup in the sample complexity of learning can be achieved. But no regret bound was derived there for concurrent RL.

include fully cooperative (Abed-alguni et al., 2015; Zhang et al., 2010), fully competitive (Bansal et al., 2017; Wang & Klabjan, 2018) and other more general settings (Ryu et al., 2021; Lowe et al., 2017). Under these settings, a group of agents share and interact with a common random environment (Shoham & Leyton-Brown, 2008; Vlassis, 2022; Weiss, 1999). There are certain challenges of multi-agent systems as pointed out by (Zhang et al., 2021). For example, the agents may have non-unique learning goals (Shoham et al., 2003). Besides, the concurrent learning structure of MARL problem can cause the environment faced by each individual to be non-stationary. Particularly in some settings, the action taken by one agent affects the reward of other opponent agents and the evolution of the state (Zhang et al., 2021). Furthermore, to handle non-stationarity of the problems, the agents may need to form the joint action space, and the dimension the action spaces can increase exponentially with the number of agents (Kulkarni & Tai, 2010).

We study the MARL problem under a concurrent learning framework involving homogeneous agents with a common learning goal. The agents interact with the unknown environment independently in parallel, and draw actions from a commonly shared action space. They communicate and share information according to some strategically devised schedule. All the agents share the same reward functions according to the states and actions they take. This is different from the Markov game setup (Szepesvári & Littman, 1999; Littman, 2001) where the reward is influenced by the joint action of all the agents in the system.

We apply concurrent learning (Min et al., 2023; Dubey & Pentland, 2021) concept with the Randomized Least-Squares Value Iteration (RLSVI) learning framework with aggregated state representations (Dong et al., 2019b; Van Roy, 2006). On the one hand, RLSVI leverages random perturbations to approximate the posterior, applying frequentist regret analysis in the tabular MDP setting (Osband et al., 2016), inspiring works that focus on theoretical analyses to improve worst-case regret in tabular MDPs (Russo, 2019; Agrawal et al., 2021) and linear settings (Zanette et al., 2020; Ishfaq et al., 2023; Dann et al., 2021). On the other hand, while a large body of literature is established on tabular representations (Agarwal et al., 2020; Azar et al., 2011; Auer et al., 2008; Jiang et al., 2017; Jin et al., 2018; Osband et al., 2013; 2016; Strehl & Littman, 2008), aggregated state representation offers an approach to reduce statistical complexity given the fact that the data requirement and learning time scales with the number of state-action pairs in tabular representations. In this paper, we present a concurrent version of the randomized least-squared value iteration algorithm with aggregated states in finite-horizon and infinite-horizon settings, and we provide their corresponding worst-case regret bounds and numerical performances.

As an outline of the rest of the paper, in Section 2, we define the finite-horizon case and introduce a concurrent learning framework with aggregated states. In Section 3, we summarize the finite-horizon concurrent RLSVI algorithm in Algorithm 3 and provide its worst-case regret bound in Theorem 2. Section 4 focuses on the infinite-horizon concurrent learning framework with aggregated states, with the infinite-horizon concurrent RLSVI algorithm summarized in Algorithm 4. The corresponding worst-case regret bound is provided in Theorem 7 of Section 5. Numerical results of both algorithms are reported in Section 6. We also include Table 1 in Appendix A comparing our approach with prior work.

2. Finite-Horizon Concurent Learning

In this section, we consider a finite-horizon Markov Decision Process (MDP) M = (H, S, A, P, R). There are N agents interacting with the same environment across Kepisodes. Each episode contains H periods. For episode $k \in [K]$, period $h \in [H]$, and agent $p \in [N]$, we use $s_{k,h}^p$ to denote the state that the agent resides in, $a_{k,h}^p$ the action that the agent takes, and $r_{k,h}^p = r(s_{k,h}^p, a_{k,h}^p)$ the reward that it receives, where $r : S \times A \mapsto [0,1]$ is a deterministic reward function. Let the information set $\mathcal{H}_k = \{(s_{k,h}^p, a_{k,h}^p, r_{k,h}^p) : h \in [H]\}$ be the trajectory during episode k for all the agents. The transition kernel P is defined as $P_{h,s,a}(s') = \mathbb{P}\left(s_{h+1} = s' | a_h = a, s_h = s\right)$. The expected reward that an agent receives in state *s* when it follows policy π at step h is represented by $R_{h,s,\pi(s)} =$ $\mathbb{E}\left[\sum_{a\in\mathcal{A}}\pi_h(a|s)\cdot r(s,a)\right]$. We assume that all agents start from a deterministic initial state $\mathbf{s}_1 = \{s_1^p\}_{p \in [N]}$ with $s_{k,1}^p \equiv s_1^p, \forall k, p$. In this work we consider deterministic rewards, which can be viewed as mappings from S to A, but all our results apply to the environments with bounded rewards without change. We say that agent $p \in [N]$ follows policy π if for all $h \in [H]$, $a_h^p = \pi_h(s_h^p)$. We use $V_h^{\pi} \in \mathbb{R}^S$ to denote the value function associated with policy π in period $h \in [H]$, such that

$$V_h^{\pi}(s) = \mathbb{E}\left[\sum_{j=h}^H R_{j,s_j,\pi_j(s_j)} \Big| s_h = s\right],$$

where the expectation is taken over all possible transitions, and we set $V_{H+1}^{\pi}(s) = 0$ for all $s \in S$. The optimal value function is denoted as $V_h^*(s) = \max_{\pi \in \Pi} V_h^{\pi}(s)$, which is the value function associated with the optimal policy. For all $s \in S, h \in [H]$, and policy π , the value function is the unique solution to the Bellman equations

$$V_h^{\pi}(s) = R_{h,s,\pi(s)} + \sum_{s \in \mathcal{S}} P_{s,h,\pi(s)}(s') V_{h+1}^{\pi}(s')$$

When π is the optimal policy π^* , there should be $V_h^{\pi^*}(s) = V_h^*(s)$, and we have

$$V_h^*(s) = R_{h,s,\pi^*(s)} + \sum_{s \in \mathcal{S}} P_{s,h,\pi^*(s)}(s') V_{h+1}^*(s')$$

For each policy π , we also define the *state-action value function* of *Q-function* of state-action pair (s, a) as the expected return when agent takes action a at state s, and then follows policy π , so that

$$Q_{h}^{\pi}(s,a) = R_{h,s,a} + \mathbb{E}\left[\sum_{j=h}^{H} R_{j,s_{j},\pi_{j}}(s_{j}) \left| s_{h} = s, a_{h} = a\right].$$

Correspondingly, we define

$$Q_h^*(s,a) = R_{h,s,a} + P_h V_{h+1}^*(s,a),$$

where we use the notation $P_hV(s,a) = \mathbb{E}_{s' \sim P_h^{s,a}}[V(s')]$. Thus by definition $Q_h^*(s,a)$ is the maximum realizable expected return when the agent starts from state s and takes action a at period h. From the optimality of V^* , we have

$$V_h^*(s) = \max_{a \in \mathcal{A}} Q_h^*(s, a), \ \forall h \in [H], s \in \mathcal{S}.$$

Furthermore, under the assumption that the reward is bounded between 0 and 1, we have

$$0 \le V_h^{\pi} \le V_h^* \le H, \ \forall h \in [H], \pi \in \Pi.$$

Regret under Finite Horizon Case The goal of an RL algorithm is for the agents to learn a good policy through consecutively interacting with the random environment, without prior knowledge about the transition probability P and the reward R. Formally, given $\pi = {\pi^{kp}}_{k \in [K], p \in [N]}$, with each agent $p \in [N]$ taking policy π^{kp} during each episode $k \in [K]$, the cumulative expected regret incurred over Kperiods and N agents is defined as

$$\operatorname{Regret}(M, K, H, N, \pi) = \sum_{p=1}^{N} \sum_{k=1}^{K} V_1^*(s_1^p) - V_{\pi^{kp}}(s_1^p).$$
(1)

Empirical Estimation We examine two types of empirical estimation methods below. The first method stores data from a single episode only, meaning the empirical counts used by the agents to evaluate policies are derived from just one episode. We take this limitation into account because, as the number of agents N increases, storing all historical data becomes practically infeasible. For the second method, we store all historical data; specifically, at each episode k, the buffer retains data from all episodes prior to k. This results in a space complexity of O(KHN). For the first empirical estimation method, define $n_{k,h}(s, a)$ to be the number of times action a has been sampled in state s, period h during episode k by all the agents $p \in [N]$:

$$n_{k,h}(s,a) = \sum_{p=1}^{N} \mathbf{1} \left\{ (s_{k,h}^{p}, a_{k,h}^{p}) = (s,a) \right\}.$$

Define the empirical mean reward for period h during episode k by

$$\hat{R}_{h,s,a}^{k} = \frac{\sum_{p=1}^{N} \mathbf{1}\left\{ (s_{k-1,h}^{p}, a_{k-1,h}^{p}) = (s,a) \right\} r_{k-1,h}^{p}}{n_{k-1,h}(s,a)},$$
(2)

and $\forall s' \in S$, define the empirical transition probabilities for period h during episode k as

$$=\frac{\hat{P}_{h,s,a}^{k}(s')}{\sum_{p=1}^{N}\mathbf{1}\left\{\left(s_{k-1,h}^{p},a_{k-1,h}^{p},s_{k-1,h+1}^{p}\right)=\left(s,a,s'\right)\right\}}{n_{k-1,h}(s,a)}.$$
(3)

If (h, s, a) is never sampled during episode k - 1, we set $\hat{R}_{h,s,a}^k = 0 \in \mathbb{R}$ and $\hat{P}_{h,s,a}^k = 0 \in \mathbb{R}^S$. Note that \hat{R}^k and \hat{P}^k are computed from the trajectory from episode k - 1.

The second method stores all historical data up to the current episode, with the reward and transition probability estimations defined as:

$$\begin{split} \hat{R}_{h,s,a}^{k,\text{full}} &= \frac{\sum_{i=0}^{k-1} \sum_{p=1}^{N} \mathbf{1} \left\{ (s_{i,h}^{p}, a_{i,h}^{p}) = (s,a) \right\} r_{i,h}^{p}}{\sum_{i=0}^{k-1} n_{i,h}(s,a)}, \\ &= \frac{\hat{P}_{h,s,a}^{k,\text{full}}(s')}{\sum_{i=0}^{k-1} \sum_{p=1}^{N} \mathbf{1} \left\{ (s_{i,h}^{p}, a_{i,h}^{p}, s_{i,h+1}^{p}) = (s,a,s') \right\}}{\sum_{i=0}^{k-1} n_{i,h}(s,a)}. \end{split}$$

2.1. Aggregated-state Representations

Many RL algorithms aim to estimate the value of each stateaction pair (e.g. under a tabular representation), but this can be infeasible in some setup where SA is large, since both the required sample size and computational cost will scale up at least linearly in SA. One alternative approach is to consider aggregated-state representations (Dong et al., 2019b; Wen & Van Roy, 2017; Jiang et al., 2017), which reduces complexity and can accelerate learning by focusing on aggregated state-action pairs. This method partitions the space of state-action pairs into Γ blocks, each block can be viewed as an aggregate state, so that the value function representation only needs to maintain one value estimate per aggregated state. Formally, let Φ be the set of all aggregated states, and let $\phi_h : S \times A \to \Phi$ be the mapping from stateaction pairs to aggregated states at period h. Without loss of generality, we let $\Phi = [\Gamma]$. We define the aggregated representation as follows:

Definition 1. We say that $\{\phi_h\}_{h=1}^H$ is an ϵ -error aggregated state-representation (or ϵ -error aggregation) of an MDP, if for all $s, s' \in S$, $a, a' \in A$ and $h \in [H]$ such that $\phi_h(s, a) = \phi_h(s', a')$, we have

$$|Q_h^*(s,a) - Q_h^*(s',a')| \le \epsilon.$$

When $\epsilon = 0$ in Definition 1, we say that the aggregation is sufficient, and one can guarantee that an algorithm finds the optimal policy as $K \to \infty$. When $\epsilon > 0$, there exists an MDP such that no RL algorithm with aggregated state representation can find the optimal policy (Van Roy, 2006). In this case, the best we can do is to approximate the optimal policy with the suboptimality bounded by a function of ϵ .

3. Finite-horizon Algorithm and Regret Bound

The concurrent RLSVI algorithms for the finite-horizon case are Algorithm 1 and Algorithm 3.

Tradeoff: Regret Reduction vs. Sample Complexity Algorithm 1 and Algorithm 3 reflect a trade-off between reducing the worst-case regret and the consideration for the sample complexity in practice. Algorithm 1 stores all historical data, while Algorithm 3 retains only the data from the previous episode. Although Algorithm 1 achieves lower regret due to its larger data storage, it becomes impractical when N is large. Therefore, for practical reasons, we use Algorithm 3 in our numerical experiments. We first describe Algorithm 3, with Algorithm 1 following a similar structure, differing only in the buffer size (i.e. the amount of data stored for policy evaluation).

We initialize the all the aggregated state values as H, i.e. $\hat{Q}_{h}^{p}(\gamma) = H$ for all $h \in [H]$. At the beginning of each episode, all the agents restart at initial states $\mathbf{s}_{1} = \{s_{1}^{p}\}$, with agent p starting from state s_{1}^{p} . During pre-round (episode 0), each agent randomly samples their initial trajectory $\{s_{0,1}^{p}, a_{0,1}^{p}, r_{0,1}^{p}, \ldots, s_{0,H}^{p}, a_{0,H}^{p}, r_{0,H}^{p}\}_{p=1}^{N}$, with $s_{0,1}^{p} = s_{1}^{p}$. During each episode $k \in [K]$, each agent samples a random vector with independent components $w^{kp} \in \mathbb{R}^{HSA}$, where $w^{kp}(h, s, a) \sim \mathcal{N}(0, \sigma_{k}^{2}(h, s, a))$ and $\sigma_{k}(h, s, a) =$ $\sqrt{\frac{\beta_{k}}{N_{k-1,h}(\phi_{h}(s, a))+1}}$, where β_{k} is a tuning parameter, $N_{k-1,h}(\phi_{h}(s, a))$ is the total number of times that aggregated state $\phi_{h}(s, a)$ is reached at period h across all agents during episode k-1. Given w^{kp} , we construct a randomized perturbation of the empirical MDP for agent p as

$$\overline{M}^{kp} = (H, \mathcal{S}, \mathcal{A}, \hat{P}^k, \hat{R}^k + w^{kp}), \qquad (4)$$

where the empirical distributions \hat{R}^k , \hat{P}^k are computed as in (2) and (3). During each episode $k \in [K]$, the data set \tilde{D}_{kh}^p contains perturbation of samples from episode k-1for time period h used by agent p. Each agent computes the values for aggregated states using a backward approximation, where during episode k, the uncapped value of aggregated state γ during period h computed by agent p is

$$\begin{split} \bar{Q}_{k,h}^{p}(\gamma) &= \underset{Q \in \mathbb{R}}{\arg\min} \mathcal{L}(Q|\hat{Q}_{k,h+1}^{p}, \tilde{D}_{kh}^{p}, \alpha_{N_{k-1,h}(\gamma)}, \xi_{N_{k-1,h}(\gamma)}) \\ &+ \|Q - \alpha_{N_{k-1,h}(\gamma)} \tilde{Q}_{kh}^{p}\|_{2}^{2}, \end{split}$$

where $\xi_{N_{k-1,h}(\gamma)}$ is defined as (7) for $n = N_{k-1,h}(\gamma)$, and we set terminal values as $\hat{Q}_{k,H+1}^p(\gamma) = 0$, and the regularization noise $\tilde{Q}_{kh}^p \in \mathbb{R}^{S \times A}$ is an independently sampled random vector, such that for each $(s, a) \in S \times A$, $\tilde{Q}_{kh}^p(s, a) \sim \mathcal{N}\left(0, \frac{\beta_k}{1+N_{k,h}(\phi_h(s,a))}\right)$, β_k is a tuning parameter, and

$$\mathcal{L}(Q \mid Q_{\text{next}}, \mathcal{D}, \alpha, \xi) = \sum_{(s,a,r,s')\in\mathcal{D}} \{Q - \xi - (1 - \alpha)Q_{k-1,h}(\phi_h(s,a)) -\alpha(r + \max_{a'\in\mathcal{A}} Q_{\text{next}}(s',a'))\}^2.$$
(5)

Here the regularized loss function defined in (5) is such that $Q(\gamma)$ as the computed value of aggregated state $\gamma = \phi_h(s, a)$ for some pair (s, a) approximates

$$(1-\alpha_{N_{k-1,h}(\gamma)})Q_{k-1,h}+\alpha_{N_{k-1,h}(\gamma)}(r+\max_{a'\in\mathcal{A}}Q_{\text{next}}(s',a'))$$

And since $\alpha_{N_{k-1,h}(\gamma)} = \frac{1}{1+N_{k-1,h}(\gamma)}$ as defined in Theorem 2, we see that when $N_{k-1,h}(\gamma)$ increases, the algorithm puts more weight on the value learned from the previous episode. At the end of episode k, after each agent interacts with the environment, the algorithm takes a weighted average of the values learned by each agent $p \in [N]$, by taking

$$\hat{Q}_{k,h}(\gamma) = \frac{1}{N_{k,h}(\gamma)} \sum_{p=1}^{N} \mathbf{1}\{\phi_h(s_{k,h}^p, a_{k,h}^p) = \gamma\}\hat{Q}_{k,h}^p(\gamma)$$
(6)

for each $\gamma \in \Gamma$, where $N_{k,h}(\gamma)$ is the total number of times that aggregated state γ appears during episode k period h.

Given tuning parameter $\beta = {\beta_k}_{k \in \mathbb{N}}$, we denote Algorithm 3 by RLSVI β . While Algorithm 3 is based on the RLSVI algorithm of (Russo, 2019), there are some notable differences. The algorithm of (Russo, 2019) is with single-agent, and at each time period *h* during episode *k*, the agent needs to keep all the trajectory prior to episode *k*, which can be infeasible for multiple agents, because the space can grow very fast. In our algorithm, we only keep the historical data of the last episode, to make the algorithm computationally feasible. This leverages the fact that with *N* agents interacting with the random environment, the information within one single episode is already rich. Algorithm 3 is also related to the aggregated-state algorithm

proposed by (Dong et al., 2019b), where the authors apply the aggregated-state idea to an optimisitic variant of Q-learning on a fixed-horizon episodic Markov decision process based on the previous UCB-type result by (Jin et al., 2020). Our work incorporates this idea in the concurrent randomized least square value iteration.

Algorithm 1 follows a similar structure, with $N_{k,h}(\phi_h(s, a))$ redefined as the total number of times the aggregated state $\phi_h(s, a)$ is visited at period h across all agents up to and including episode k. We denote it by $\text{RLSVI}_{\beta}^{\text{full}}$ to indicate that Algorithm 1 uses the full history up to the current episode.

With the employment of aggregated state and the modified loss function for the learning process, the result highlights that the per-agent regret decreases at a rate of $\Theta\left(\frac{1}{\sqrt{N}}\right)$. The concurrent learning algorithm for the finite-horizon case is Algorithm 3. The worst-case regret bound for Algorithm 3 is provided in the next section.

3.1. Worst-case Regret Bound

Let \mathcal{M} be the set of MDPs with episode number K, horizon H, state space size S, action space size A, transition probabilities P, and rewards R bounded in [0, 1]. Let N be the number of agents interacting in the same environment. We use M = (K, H, S, A, P, R) to denote an MDP in \mathcal{M} .

Suppose $\{\phi_h\}_{h\in[H]}$ is an ϵ -error aggregation (defined as in Definition 1) of the underlying MDP. For a tuning parameter sequences $\beta = \{\beta_n\}_{n\in\mathbb{N}}, \alpha = \{\alpha_n\}_{n\in\mathbb{N}}, \xi = \{\xi_n\}_{n\in\mathbb{N}}$ where $\beta_n = \frac{1}{2}H^3\log(2H\Gamma n), \alpha_n = \frac{1}{1+n}$, and

$$\xi_n = \epsilon + \frac{2\alpha_n H \sqrt{\log(2KHN/\delta)}}{\sqrt{\max\{n,1\}}} + \frac{2\alpha_n \sqrt{\beta_{k-1}\log(2KHN/\delta)}}{\sqrt{(n+1)\max\{n,1\}}}$$
(7)

We now provide our main results for the finite-horizon case. As explained, Algorithm 1 stores all historical data, while Algorithm 3 retains only the previous episode, making it straightforward that Algorithm 1 has a space complexity of O(KHN), while Algorithm 3 has a space complexity of O(HN), and their worst-case regret bounds are

Theorem 2. Algorithm 1 has a worst-case regret bound

$$\sup_{\substack{M \in \mathcal{M} \\ \leq \tilde{O}(\epsilon \sqrt{K}HN + H^{5/2}\Gamma\sqrt{KN}).}} \operatorname{Regret}(M, K, N, \pi, \operatorname{RLSVI}_{\beta,\alpha,\xi}^{\text{full}})$$
(8)

Algorithm 3 has a worst-case regret bound

$$\sup_{\substack{M \in \mathcal{M} \\ \leq \widetilde{O}(\epsilon K H N + K H^{5/2} \Gamma \sqrt{N}),}} \operatorname{Regret}(M, K, N, \pi, RLSVI_{\beta, \alpha, \xi})$$
(9)

where $\widetilde{O}(\cdot)$ hides the dependence on logarithmic factors.

Since the only difference between the algorithms is the amount of data stored, and their structures are identical, we only prove (9), with (8) following immediately by reducing a factor of \sqrt{K} from the regret bound in (9). We omit the redundant proof for Algorithm 1 and defer the proof for (9) to Appendix C.

Comparison with worst-case regret bounds from (Russo, 2019; Agrawal et al., 2021) A worst-case regret bound of $\tilde{O}(H^3S^{3/2}\sqrt{AK})$ was obtained in (Russo, 2019) for a single-agent version of RLSVI algorithm. This bound was improved later by (Agrawal et al., 2021) to $\tilde{O}(H^{5/2}S\sqrt{AK})$. For the single-agent case with N =1, Algorithm 1 results in a worst-case regret bound of $\tilde{O}(\sqrt{K}H^{5/2}\Gamma)$, which translates into $\tilde{O}(H^2\Gamma\sqrt{T})$. So our bound matches that of (Agrawal et al., 2021) if $\Gamma = S \times A$ and $S \approx \sqrt{\Gamma}$. Algorithm 3 implies a worst-case regret bound of $\tilde{O}(KH^{5/2}\Gamma)$. Here an extra \sqrt{K} compared to that of (Agrawal et al., 2021; Russo, 2019) comes from the fact we only keep the trajectories of the agents from the previous episode rather than all the episodes up to the current period as in (Russo, 2019; Agrawal et al., 2021) to reduce space complexity. The extra ϵ term for both algorithms comes from model misspecification of state-aggregation formulation, which is similar to the result in (Dong et al., 2019b).

At each time period h, each agent gets N samples of tuples (s, a, r, s'), and they share information by the aggregation of information through the computation of a weighted Q-value (6) at the end of each episode. With this trick, though agents learn their own policies concurrently in parallel within each episode, we are able to obtain a sub-linear worst-case regret bound of $O(\sqrt{N})$ with respect to the number of agents.

4. Infinite-horizon Concurrent Learning

We now turn to the infinite-horizon case. Consider an unknown fixed environment as $M = (T, \mathcal{A}, \mathcal{S}, P, r, N)$, with N agents interacting in M. Here $\mathcal{A} = [A]$ is the action space, $\mathcal{S} = [S]$ is the state space, P(s' | s, a) is the transition probability from $s \in \mathcal{S}$ to $s' \in \mathcal{S}$ given action $a \in \mathcal{A}$. After agent p selects action a_t^p at state s_t^p , the agent observes s_{t+1}^p and receives a fixed reward $r_{t+1}^p = r(s_t^p, a_t^p)$ where $r \in [0, 1]$. A stochastic policy π can be represented by a probability mass function $\pi(\cdot|s_t)$ that an agent assigns to actions in \mathcal{A} given situation state s_t . For a policy π , we denote the average reward starting at state s as

$$\lambda_{\pi}(s) = \liminf_{T \to \infty} \mathbb{E}\left[\frac{1}{T} \sum_{t=0}^{T-1} r_{t+1} \middle| s_1 = s\right].$$
(10)

For any state $s \in S$, denote the optimal average reward as $\lambda_*(s) = \sup_{\pi} \lambda_{\pi}(s)$. We consider weakly-communicating MDP, which is defined as follows:

Definition 3 (Weakly-communicating MDP). A MDP is weakly communicating if there exists a closed subset of states, where each state within is reachable from any other state within that set under some deterministic stationary reward. And there exists a transient subset of states (possibly empty) under every policy.

For any $s, s' \in S$ and $a \in A$, denote $P_{s,a,s'} = P(s'|s,a)$. For each policy π define transition probabilities under π as $P_{s,\pi,s'} = \sum_{a \in A} \pi(a|s)P_{s,a,s'}$, and reward as $r_{s,\pi} = \sum_{a \in A} \pi(a|s)r(s,a)$.

Pseudo-episodes We extend our concurrent learning framework for the finite-horizon case to the infinite-horizon case by incorporating the idea of pseudo-episodes from (Xu et al., 2022). Suppose time step t is the beginning of a pseudo-episode when we sample a random variable $H \sim \text{Geometric}(1-\eta)$, where $\text{Geometric}(1-\eta)$ is geometric distribution with parameter $1 - \eta$. In the numerical experiment (section 6), we set $\eta = 0.99$. The agents compute new policies respectively according to the collected trajectories from last pseudo-episode, and sample their own MDPs respectively at time steps $t + 1, \ldots, t + H - 1$. Now the beginning of the next pseudo-episode is set as t + H. We use $\mathcal{H}_{t_1,t_2} = \bigcup_{p=1}^N \bigcup_{i=t_1}^{t_2^1} \{s_i^p, a_i^p, r_i^p\}$ to denote the trajectories of all agents from time step t_1 to time step t_2 . For policy π , denote the η -discounted value function as $V_{\pi}^{\eta} \in \mathbb{R}^{S}$, then we have

$$V_{\pi}^{\eta} := \mathbb{E}_{H} \left[\sum_{h=0}^{H-1} P_{\pi}^{h} r_{\pi} \middle| M \right] = \mathbb{E} \left[\sum_{h=0}^{\infty} \eta^{h} P_{\pi}^{h} r_{\pi} \middle| M \right],$$
(11)

where the expectation is taken over the random episode length H. A policy is said to be optimal if $V_{\pi}^{\eta} = \sup_{\pi'} V_{\pi'}^{\eta}$. For an optimal policy, we also write $V_{*}^{\eta}(s) \equiv V_{\pi}^{\eta}(s)$ as the optimal value. Note that $V_{\pi}^{\eta} \in \mathbb{R}^{S}$ satisfies the Bellman equation $V_{\pi}^{\eta} = r_{\pi} + \eta P_{\pi} V_{\pi}^{\eta}$. For any (s, a), define $Q_{\pi}^{\eta}(s, a) = r(s, a) + \eta P V_{\pi}^{\eta}(s)$, where we use the notation that $P_{\pi}V(s) = \mathbb{E}_{s' \sim P_{s,\pi(s)}}[V(s')]$. Correspondingly, define

$$Q_*^{\eta}(s,a) = r(s,a) + \eta P V_*^{\eta}(s').$$

By definition, we have $V_*^{\eta}(s) = \max_{a \in \mathcal{A}} Q_*^{\eta}(s, a)$.

Discounted Regret To analyze the algorithm over T time steps, consider $K = \arg \max\{k : t_k \leq T\}$ as the number of pesudo-episodes until time T. We use the convention that $t_{K+1} = T+1$. Given a discount factor $\eta \in [0, 1)$, define the η -discounted regret up to time T as $\operatorname{Regret}_{\eta}(M, T, \pi) =$ $\sum_{k=1}^{K} \Delta_k$, where Δ_k is the total regret of all N agents over pseudo-episode k: $\Delta_k = \sum_{p=1}^{N} V_*^{\eta}(s_{k,1}^p) - V_{\pi^{k_p}}^{\eta}(s_{k,1}^p)$, where $V_*^{\eta} = V_{\pi^*}^{\beta}$, policy π^{k_p} is computed from the trajectory $\mathcal{H}_{t_k-1,t_k-1}$ from pseudo-episode k - 1 by agent p, and $a_t^p \sim \pi^{k_p}(\cdot \mid s_t^p), s_{t+1}^p \sim P(\cdot|s_t^p, a_t^p), r_t^p = r(s_t^p, a_t^p)$ for $t \in E_k$, and E_k denotes the time steps within pseudoepisode k. So the discounted regret is a random variable depending on the algorithm's random sampling, and the random lengths of the pseudo-episodes, and as a result,

$$\Delta_{k} = \mathbb{E}_{H_{k}} \left[\sum_{p=1}^{N} \sum_{h=0}^{H_{k}} (P_{\pi^{*}}^{h} r_{\pi^{*}} - P_{\pi^{kp}} r_{\pi^{kp}}) \mid M \right]$$
$$= \mathbb{E} \left[\sum_{p=1}^{N} \sum_{h=0}^{\infty} \eta^{h} (P_{\pi^{*}}^{h} r_{\pi^{*}} - P_{\pi^{kp}} r_{\pi^{kp}}) \mid M \right].$$

Regret The optimal average reward λ_* is stateindependent under a weakly-communicating MDP. The agent p selects a policy π^{kp} and executes it within the k^{th} pseudo-episode. The cumulative expected regret incurred by the collection of policies $\pi = {\pi^{kp}}_{k \in [K], p \in [N]}$ over Ttime steps and across N agents with the fixed environment M is

Regret
$$(M, T, N, \pi) := \mathbb{E}_K \left[\sum_{k=1}^K \Delta_k | M \right],$$
 (12)

where the expectation is taken over the random seeds used by the randomized algorithm, conditioning on the true MDP M. In the following, we denote $(s_{k,h}^p, a_{k,h}^p, r_{k,h}^p)$ as the state, action and reward for agent p during pseudo-episode k and period h.

Empirical Estimation We let $n_k(s, a)$ be the total number of times that (s, a)-pair appears during the k^{th} pseudo-episode, such that

$$n_k(s,a) = \sum_{p=1}^N \sum_{t=t_k}^{t_{k+1}-1} \mathbf{1} \left\{ (s_t^p, a_t^p) = (s,a) \right\}.$$

Then $\forall s'$, the empirical estimate $\hat{P}(s'|s, a)$ of the transition probability for pseudo-episode k is

$$=\frac{\sum_{p=1}^{N}\sum_{t\in E_{k-1}}\mathbf{1}\left\{(s_{t}^{p},a_{t}^{p},s_{t+1}^{p})=(s,a,s')\right\}}{n_{k-1}(s,a)}.$$
(13)

The empirical estimate of the corresponding reward is

$$\frac{P_k(s,a)}{\sum_{p=1}^N \mathbf{1}\{(s_t^p, a_t^p) = (s,a)\} \sum_{t \in E_{k-1}} r(s_t^p, a_t^p)}{n_{k-1}(s,a)}.$$
(14)

For the second empirical estimation method, we utilize the full historical data, and similar to the finite-horizon case, the empirical estimates for transition probability and reward are

$$= \frac{\hat{P}_{k}^{\text{full}}(s'|s,a)}{\sum_{i=0}^{k-1}\sum_{p=1}^{N}\sum_{t\in E_{i}}\mathbf{1}\left\{(s_{t}^{p},a_{t}^{p},s_{t+1}^{p})=(s,a,s')\right\}}{\sum_{i=0}^{k-1}n_{i}(s,a)}$$

=

$$= \frac{\hat{R}_{k}^{\text{full}}(s,a)}{\sum_{p=1}^{N} \sum_{i=0}^{k-1} \mathbf{1}\{(s_{t}^{p}, a_{t}^{p}) = (s,a)\} \sum_{t \in E_{i}} r(s_{t}^{p}, a_{t}^{p})}{\sum_{i=0}^{k-1} n_{i}(s,a)}.$$

Aggregated States We extend the aggregated states in the finite-horizon case to the infinite horizon case. Let Φ be the set of all aggregated states, and let $\phi : S \times A \rightarrow \Phi$ be the mapping from state-action pairs to aggregated states. We let $\Phi = [\Gamma]$. The aggregated representation for the infinite-horizon case is defined as follows:

Definition 4. We say that ϕ is an ϵ -error aggregated staterepresentation (or ϵ -error aggregation) of an MDP, if for all $s, s' \in S$, $a, a' \in A$ such that $\phi(s, a) = \phi(s', a')$, we have $|Q_*^{\eta}(s, a) - Q_*^{\eta}(s', a')| \le \epsilon$.

We are now ready to present the concurrent learning algorithm for the infinite-horizon case and the theoretical guarantee, as detailed in the next section.

5. Infinite-Horizon Algorithm and Regret Bound

The concurrent learning algorithm for the infinite-horizon MDP is summarized as Algorithm 4. Our result is based on the following definition of reward averaging time proposed by (Dong et al., 2022).

Definition 5. The reward averaging time τ_{π} of a policy π is the smallest value $\tau \in [0, \infty)$ such that $\forall T \ge 0, s \in S$, $\left|\mathbb{E}_{\pi}\left[\sum_{t=0}^{T-1} r_{t+1} \middle| s_0 = s\right] - T \cdot \lambda_{\pi}(s)\right| \le \tau$.

Typically the regret bounds established in the literature requires assessing an optimal policy within bounded time. Examples include episode duration (Osband et al., 2013; Jin et al., 2018), diameter (Auer et al., 2008), or span (Bartlett & Tewari, 2012). Policies that require intractably large amount of time are infeasible in practice. So we impose the following assumption:

Assumption 6. For any weakly communicating MDP M with state space S and action space A, $\exists \tau < \infty$ such that $\tau_* \leq \tau$.

When π^* is an optimal policy for M, $\tau_* := \tau_{\pi^*}$ is equivalent to the notion of span in (Bartlett & Tewari, 2012). Let \mathcal{M} be the set of infinite-horizon weakly-communicating MDPs with state space size S, action space size A, rewards bounded in [0, 1] that satisfy Assumption 6. Let N be the number of agents interacting in the same environment. Recall that τ is given by Assumption 6.

Suppose $\{\phi\}$ is an ϵ -error aggregation (defined as in Definition 4) of the underlying MDP. For a tuning parameter sequences $\beta = \{\beta_n\}_{n \in \mathbb{N}}, \alpha = \{\alpha_n\}_{n \in \mathbb{N}}, \xi = \{\xi_n\}_{n \in \mathbb{N}}$, where for k as the index of pseudo-episode,

$$\beta_{k} = \frac{1}{2}\tau^{3}\log(2\tau\Gamma k), \alpha_{n} = \frac{1}{1+n}, \text{ and}$$

$$\xi_{n} = \epsilon + \frac{2\alpha_{n}\sqrt{\log(2TN/\delta)}}{(1-\eta)\sqrt{\max\{n,1\}}}$$

$$+ \frac{2\alpha_{n}\sqrt{\beta_{k-1}\log(2TN/\delta)}}{\sqrt{(n+1)\max\{n,1\}}}$$
(15)

Theorem 7 (Infinite-horizon Worst-case Regret Bound). *Algorithm 2 has a worst-case regret bound*

$$\sup_{M \in \mathcal{M}} \operatorname{Regret}(M, T, N, \operatorname{RLSVI}_{\beta, \alpha, \xi}^{full})$$

$$\leq \widetilde{O}(\epsilon \sqrt{T}N + \tau^{3/2} \Gamma \sqrt{TN}).$$
(16)

Algorithm 4 has a worst-case regret bound

$$\sup_{M \in \mathcal{M}} \operatorname{Regret}(M, T, N, \operatorname{RLSVI}_{\beta, \alpha, \xi})$$

$$\leq \widetilde{O}(\epsilon TN + \tau^{3/2} T \Gamma \sqrt{N}).$$
(17)

Note that the bound in (17) matches that of the finitehorizion case (9) by noting that T = KH for the finitehorizon case. And $\tau^{3/2}$ corresponds to the $H^{3/2}$ factor in the finite-horizon case bound by noting that taking $\tau = H$ makes the condition holds in Definition 5 in the finitehorizon case with T = KH. Similar to the finite-horizon case, (16) follows directly from (17) with a \sqrt{T} reduction, as Algorithm 2 utilizes the full history, whereas Algorithm 4 retains only the last pseudo-episode. The intuition and worst-case regret comparison follow the discussion after Theorem 2.

6. Numerical Experiments

We present numerical results for both finite-horizon and infinite-horizon cases in Figure 1. For the finite-horizon case (S, A, K, H) or the infinite-horizon case (S, A, T), the transition probabilities are drawn from a Dirichlet distribution, and rewards, fixed as deterministic, are uniformly distributed on [0, 1], forming inherent features of the MDP class. The finite-horizon case settings are (i) K = 20, H = 30, S = 5, A = 5; (ii) K = 25, H =40, S = 10, A = 10; (iii) K = 30, H = 50, S =20, A = 20. The infinite-horizon case settings are with T = 300 and (i) S = 5, A = 5; (ii) S = 20, A = 20; (iii) S = 30, A = 30, where $\eta = 0.99$ in the pseudo-episode sampling. Under each setting, we compare the results for N = 1, 3, 5, 7, 10, 15, 20, 30, 40, 50, with $\epsilon = 0$. We set $\eta = 0.99$.

For each agent number N in the finite-horizon setting, we sample 500 MDPs from the defined class. For each sampled MDP, we compute the cumulative regret over time and then identify the maximum regret across all 500 instances, representing the worst-case regret in our analysis.

For the infinite-horizon setting, we estimate regret by averaging over 50 geometric segmentations of [T] per MDP,

consistent with the definition of infinite-horizon regret based on pseudo-episodes in (12). The worst-case regret is then obtained by taking the maximum across 500 sampled MDPs.

Figure 1 illustrates a $1/\sqrt{N}$ decreasing trend in per-agent regret for both settings, consistent with our theoretical predictions. The replication code is available at https://github.com/yz2/rlsvi_code.



Figure 1. Per-agent regret vs number of agents for finite-horizon (top panel) and infinite-horizon (bottom panel) settings. The solid curves represent the per-agent worst-case regret computed from 500 random environments. The dashed ones are the reference curves of constant/ \sqrt{N} fitting the $\Theta(1/\sqrt{N})$ trend as we show in our theoretical results.

Acknowledgements

We gratefully acknowledge the support from the NSF grant CCF-2312205 and the ONR grant ONR 13983263. We would like to thank Pratap Ranade and Engin Ural for inspiring discussions on society of agents that have helped shape and push an ambitious vision of this research agenda, for which this work is only an initial step.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Abed-alguni, B. H., Chalup, S. K., Henskens, F. A., and Paul, D. J. A multi-agent cooperative reinforcement learning model using a hierarchy of consultants, tutors and workers. *Vietnam Journal of Computer Science*, 2:213– 226, 2015.
- Agarwal, A., Bird, S., Cozowicz, M., Hoang, L., Langford, J., Lee, S., Li, J., Melamed, D., Oshri, G., Ribas, O., et al. Making contextual decisions with low technical debt. arXiv preprint arXiv:1606.03966, 2016.
- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. Optimality and approximation with policy gradient methods in markov decision processes. In *Conference on Learning Theory*, pp. 64–66. PMLR, 2020.
- Agrawal, P., Chen, J., and Jiang, N. Improved worst-case regret bounds for randomized least-squares value iteration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 6566–6573, 2021.
- Auer, P., Jaksch, T., and Ortner, R. Near-optimal regret bounds for reinforcement learning. *Advances in neural information processing systems*, 21, 2008.
- Azar, M. G., Munos, R., Ghavamzadeh, M., and Kappen, H. Speedy q-learning. In *Advances in neural information* processing systems, 2011.
- Bansal, T., Pachocki, J., Sidor, S., Sutskever, I., and Mordatch, I. Emergent complexity via multi-agent competition. arXiv preprint arXiv:1710.03748, 2017.
- Bartlett, P. L. and Tewari, A. Regal: A regularization based algorithm for reinforcement learning in weakly communicating mdps. arXiv preprint arXiv:1205.2661, 2012.
- Buşoniu, L., Babuška, R., and De Schutter, B. Multi-agent reinforcement learning: An overview. *Innovations in multi-agent systems and applications-1*, pp. 183–221, 2010.
- Chen, Y., Dong, P., Bai, Q., Dimakopoulou, M., Xu, W., and Zhou, Z. Society of agents: Regret bounds of concurrent thompson sampling. In Advances in Neural Information Processing Systems, 2022.
- Dann, C., Mohri, M., Zhang, T., and Zimmert, J. A provably efficient model-free posterior sampling method for episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 34:12040–12051, 2021.

- Desai, N., Critch, A., and Russell, S. J. Negotiable reinforcement learning for pareto optimal sequential decisionmaking. *Advances in Neural Information Processing Systems*, 31, 2018.
- Dimakopoulou, M. and Van Roy, B. Coordinated exploration in concurrent reinforcement learning. In *International Conference on Machine Learning*, pp. 1271–1279. PMLR, 2018.
- Dimakopoulou, M., Osband, I., and Van Roy, B. Scalable coordinated exploration in concurrent reinforcement learning. *Advances in Neural Information Processing Systems*, 31, 2018.
- Dong, K., Wang, Y., Chen, X., and Wang, L. Q-learning with UCB exploration is sample efficient for infinite-horizon MDP. arXiv preprint arXiv:1901.09311, 2019a.
- Dong, S., Van Roy, B., and Zhou, Z. Provably efficient reinforcement learning with aggregated states. *arXiv preprint arXiv:1912.06366*, 2019b.
- Dong, S., Van Roy, B., and Zhou, Z. Simple agent, complex environment: Efficient reinforcement learning with agent states. *Journal of Machine Learning Research*, 23(255): 1–54, 2022.
- Dubey, A. and Pentland, A. Provably efficient cooperative multi-agent reinforcement learning with function approximation. *arXiv preprint arXiv:2103.04972*, 2021.
- Dwaracherla, V., Lu, X., Ibrahimi, M., Osband, I., Wen, Z., and Van Roy, B. Hypermodels for exploration. arXiv preprint arXiv:2006.07464, 2020.
- Fellows, M., Hartikainen, K., and Whiteson, S. Bayesian bellman operators. *Advances in Neural Information Processing Systems*, 34:13641–13656, 2021.
- Gu, S., Holly, E., Lillicrap, T., and Levine, S. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In 2017 IEEE international conference on robotics and automation (ICRA), pp. 3389– 3396. IEEE, 2017.
- Guo, Z. and Brunskill, E. Concurrent pac rl. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2015.
- Ishfaq, H., Cui, Q., Nguyen, V., Ayoub, A., Yang, Z., Wang, Z., Precup, D., and Yang, L. Randomized exploration in reinforcement learning with general value function approximation. In *International Conference on Machine Learning*, pp. 4607–4616. PMLR, 2021.

- Ishfaq, H., Lan, Q., Xu, P., Mahmood, A. R., Precup, D., Anandkumar, A., and Azizzadenesheli, K. Provable and practical: Efficient exploration in reinforcement learning via langevin monte carlo. arXiv preprint arXiv:2305.18246, 2023.
- Janz, D., Hron, J., Mazur, P., Hofmann, K., Hernández-Lobato, J. M., and Tschiatschek, S. Successor uncertainties: exploration and uncertainty in temporal difference learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J., and Schapire, R. E. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pp. 1704–1713. PMLR, 2017.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is q-learning provably efficient? *Advances in neural information processing systems*, 31, 2018.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *Conference on learning theory*, pp. 2137–2143. PMLR, 2020.
- Kim, M. J. Thompson sampling for stochastic control: The finite parameter case. *IEEE Transactions on Automatic Control*, 62(12):6415–6422, 2017.
- Kulkarni, A. J. and Tai, K. Probability collectives: a multiagent approach for solving combinatorial optimization problems. *Applied Soft Computing*, 10(3):759–771, 2010.
- Littman, M. L. Value-function reinforcement learning in markov games. *Cognitive systems research*, 2(1):55–66, 2001.
- Lowe, R., Wu, Y. I., Tamar, A., Harb, J., Pieter Abbeel, O., and Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30, 2017.
- Min, Y., He, J., Wang, T., and Gu, Q. Cooperative multiagent reinforcement learning: Asynchronous communication and linear function approximation. In *International Conference on Machine Learning*, pp. 24785– 24811. PMLR, 2023.
- Osband, I. and Van Roy, B. Model-based reinforcement learning and the eluder dimension. *Advances in Neural Information Processing Systems*, 27, 2014.
- Osband, I. and Van Roy, B. On optimistic versus randomized exploration in reinforcement learning. *arXiv preprint arXiv:1706.04241*, 2017.

- Osband, I., Russo, D., and Van Roy, B. (More) efficient reinforcement learning via posterior sampling. *Advances in Neural Information Processing Systems*, 26, 2013.
- Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. Deep exploration via bootstrapped DQN. *Advances in neural information processing systems*, 29, 2016.
- Osband, I., Van Roy, B., Russo, D. J., Wen, Z., et al. Deep exploration via randomized value functions. *J. Mach. Learn. Res.*, 20(124):1–62, 2019.
- Pazis, J. and Parr, R. Efficient pac-optimal exploration in concurrent, continuous state mdps with delayed updates. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- Russo, D. Worst-case regret bounds for exploration via randomized value functions. Advances in Neural Information Processing Systems, 32, 2019.
- Russo, D. and Van Roy, B. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39 (4):1221–1243, 2014.
- Ryu, H., Shin, H., and Park, J. Cooperative and competitive biases for multi-agent reinforcement learning. arXiv preprint arXiv:2101.06890, 2021.
- Shoham, Y. and Leyton-Brown, K. Multiagent systems: Algorithmic, game-theoretic, and logical foundations. Cambridge University Press, 2008.
- Shoham, Y., Powers, R., and Grenager, T. Multi-agent reinforcement learning: a critical survey. Technical report, Citeseer, 2003.
- Silver, D., Newnham, L., Barker, D., Weller, S., and McFall, J. Concurrent reinforcement learning from customer interactions. In *International conference on machine learning*, pp. 924–932. PMLR, 2013.
- Sinai, S., Wang, R., Whatley, A., Slocum, S., Locane, E., and Kelsic, E. D. Adalead: A simple and robust adaptive greedy search algorithm for sequence design. *arXiv* preprint arXiv:2010.02141, 2020.
- Strehl, A. L. and Littman, M. L. An analysis of modelbased interval estimation for markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309– 1331, 2008.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Szepesvári, C. and Littman, M. L. A unified analysis of value-function-based reinforcement-learning algorithms. *Neural computation*, 11(8):2017–2060, 1999.

- Taiga, A. A., Courville, A., and Bellemare, M. G. Introducing coordination in concurrent reinforcement learning. In *ICLR 2022 Workshop on Gamification and Multiagent Solutions*, 2022.
- Van Roy, B. Performance loss bounds for approximate value iteration with state aggregation. *Mathematics of Operations Research*, 31(2):234–244, 2006.
- Vlassis, N. A concise introduction to multiagent systems and distributed artificial intelligence. Springer Nature, 2022.
- Wang, X. and Klabjan, D. Competitive multi-agent inverse reinforcement learning with sub-optimal demonstrations. In *International conference on machine learning*, pp. 5143–5151. PMLR, 2018.
- Weiss, G. Multiagent systems: a modern approach to distributed artificial intelligence. MIT press, 1999.
- Wen, Z. and Van Roy, B. Efficient reinforcement learning in deterministic systems with value function generalization. *Mathematics of Operations Research*, 42(3): 762–782, 2017.
- Xu, W., Dong, S., and Van Roy, B. Posterior sampling for continuing environments. *ArXiv preprint*, 2022. URL https://arxiv.org/abs/2211.15931.
- Zanette, A., Brandfonbrener, D., Brunskill, E., Pirotta, M., and Lazaric, A. Frequentist regret bounds for randomized least-squares value iteration. In *International Conference* on Artificial Intelligence and Statistics, pp. 1954–1964. PMLR, 2020.
- Zhang, K., Yang, Z., and Başar, T. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of reinforcement learning and control*, pp. 321–384, 2021.
- Zhang, P., Ma, X., Pan, Z., Li, X., and Xie, K. Multi-agent cooperative reinforcement learning in 3d virtual world. In Advances in Swarm Intelligence: First International Conference, ICSI 2010, Beijing, China, June 12-15, 2010, Proceedings, Part I 1, pp. 731–739. Springer, 2010.

A. Comparison Table: Existing Work vs. Our Approach

Agent	Setup	Algorithm	Regret Bound	Regret-Type	Data Stored	Numerical
Single	Tabular	RLSVI (Russo, 2019)	$\tilde{O}(H^3S^{3/2}\sqrt{AK})$)Worst-case	All-history	N/A
Single	Tabular	RLSVI (Agrawal et al., 2021)	$O(H^{5/2}S\sqrt{AK})$	Worst-case	All-history	N/A
Multi	Tabular	Concurrent RLSVI (Taiga et al., 2022)	N/A	Bayes	All-history	Synthetic
Multi	Linear Functional Approximation	Concurrent LSVI (Desai et al., 2018)	$\tilde{O}\!\!\left(H^2\sqrt{d^3KN}\right)$	Worst-case	All-history	N/A
Multi	Linear Functional Approximation	Concurrent LSVI (Min et al., 2023)	$\tilde{O}\!\!\left(H\sqrt{dKN}\right)$	Worst-case	All-history	N/A
Multi	Tabular	Concurrent RLSVI (ours-1)	$\tilde{O}(H^{5/2}\Gamma\sqrt{KN})$	Worst-case	All-history	N/A
Multi	Tabular	Concurrent RLSVI (ours-2)	$\widetilde{O}(H^{5/2}\Gamma K\sqrt{N})$	Worst-case	One episode	Synthetic

Table 1. Comparison of regret bounds for various RLSVI/LSVI algorithms

We outline several key observations for Table 1 below:

- The methods marked as "N/A" in the "numerical" column store all agents' trajectories at every step, making them computationally infeasible as N grows large. These approaches (including our first finite-horizon algorithm, second-to-last row in the table) provide only theoretical analyses without numerical results.
- Although Taiga et al. (2022) stores all data empirically, their RLSVI algorithm assumes a known parametric feature matrix, making it simpler to implement than ours. Also they evaluate only Bayes regret—less stringent than our worst-case regret—and provide empirical results exclusively on synthetic data without theoretical guarantees.
- The last row corresponds to storing only the latest episode data, increasing the regret bound by a factor of \sqrt{K} but reduces space complexity by a factor of K, making it computationally feasible.

B. Algorithms

Algorithm 1 and Algorithm 3 are finite-horizon algorithms, where the former stores all historical data, and the latter only keeps one episode in the buffer. Algorithm 2 and Algorithm 4 are infinite-horizon algorithms, where the former stores all historical data and the latter only keeps one pseudo-episode.

Algorithm 1 Concurent RLSVI: Finite-Horizon Storing All Historical Data **Data:** $K, H, S, \mathcal{A}, N, \mathbf{s}_1, \{\phi_h\}_{h=1}^H$, Tuning parameters $\{\beta_k\}_{k \in \mathbb{N}}$ Define constants $\alpha_t \leftarrow 1/(1+t), t = 1, 2, \dots$ /* Define squared temporal difference error */ $\mathcal{L}(Q \mid Q_{k-1,h}, Q_{\text{next}}, \mathcal{D}, \xi, \alpha) = \sum_{(s,a,r,s') \in \mathcal{D}} (Q - \xi - (1 - \alpha)Q_{k-1,h}(\phi_h(s,a)) - \alpha(r + \max_{a' \in \mathcal{A}} Q_{\text{next}}(s',a')))^2$ **Initialize:** $\hat{Q}_{0h}^{p}(\gamma) = H, \forall h \in [H], \gamma \in [\Gamma], p \in [N]$ Each agent randomly samples the initial trajectory $\{s_{0,1}^p, a_{0,1}^p, r_{0,1}^p, \dots, s_{0,H}^p, a_{0,H}^p, r_{0,1}^H\}_{p=1}^N$, with $s_{0,1}^p = s_1^p$ $N_{0,h}(\gamma) = \sum_{p=1}^{N} \mathbf{1}\{\phi_h(s_{0,h}^p, a_{0,h}^p) = \gamma\}, \forall \gamma \in [\Gamma], h \in [H]$ compute $\hat{Q}_{0,h}$ by (6) for *episode* k = 1, 2, ... do /*Each agent rollouts in the evironment*/ for p = 1, ..., N do /*Executed in parallel*/ for period $h = 1, \ldots, H$ do
$$\begin{split} \hat{a}_{k,h}^{p} &\leftarrow \arg\max_{a \in \mathcal{A}} \hat{Q}_{k-1,h}^{p}(\phi_{h}(s_{k,h}^{p},a)) \\ \text{observe reward } r_{k,h}^{p} \text{ and next state } s_{k,h+1}^{p} \\ \mathcal{D}_{h} &\leftarrow \mathcal{D}_{h} \cup \{(s_{k,h}^{p}, a_{k,h}^{p}, r_{k,h}^{p}, s_{k,h+1}^{p})\} \end{split}$$
end end $\begin{array}{l} \texttt{'*Visitation of aggregated-states*/} \\ N_{k,h}(\gamma) \leftarrow \sum_{i=1}^{k-1} \sum_{p=1}^{N} \mathbf{1}\{\phi_h(s_{i,h}^p, a_{i,h}^p) = \gamma\}, \forall \gamma \in [\Gamma], h \in [H] \end{array}$ /*Construct perturbed data sets and sample regularization noise $\hat{Q}\star/$ for $p \in [N]$ and $h \in [H]$ do /*Executed in parallel*/ For any $(s, a) \in \mathcal{S} \times \mathcal{A}$, sample array $\tilde{Q}_{kh}^{p}(s,a) \sim \mathcal{N}\left(0, \frac{\beta_{k}}{1+N_{k,h}(\phi_{h}(s,a))}\right),$ /* Draw prior sample */ $\tilde{D}_{kh}^p \leftarrow \{\}$ for $(s, a, r, s') \in \mathcal{D}_h^k$ do Sample $w^p(s,a) \sim \mathcal{N}\left(0, \frac{\beta_k}{1+N_{k,h}(\phi_h(s,a))}\right)$ $\tilde{D}^p_{kh} \leftarrow \tilde{D}^p_{kh} \cup \{(s, a, r + w^p, s')\}$ end end /*Estimate Q on perturbed data*/ for p = 1, ..., N do /*Executed in parallel*/ Define terminal value $\hat{Q}_{k,H+1}^p(\gamma) \leftarrow H \ \forall \gamma \in [\Gamma]$ for period $h = H, \ldots, 1$ do $\bar{Q}_{k,h}^p(\gamma) \leftarrow \arg\min_{Q \in \mathbb{R}} \mathcal{L}(Q|\hat{Q}_{k-1,h}, \hat{Q}_{k,h+1}^p, \tilde{D}_{kh}^p, \xi, \alpha) + \|Q - \alpha_{N_{k-1,h}(\gamma)} \tilde{Q}_{kh}^p\|_2^2, \forall \gamma \in [\Gamma]$ $\hat{Q}_{k,h}^p(\gamma) \leftarrow \min\{\bar{Q}_{k,h}^p(\gamma), H\}, \forall \gamma \in [\Gamma]$ end $s_{k,1}^p \leftarrow s_1^p$ end Update $\hat{Q}_{k,h}(\gamma), \forall \gamma \in [\Gamma]$ by (6) end

Algorithm 2 Concurent RLSVI: Infinite-Horizon Storing All Historical Data **Data:** Discount factor η , $t_0 = 1$, t = 1, k = 0, $X_1 = 0$, S, A, N, T, ϕ , tuning parameters $\{\beta_k\}_{k \in \mathbb{N}}, \xi, \eta$ Initialize $N_k(\gamma) \leftarrow 0, \forall \gamma \in [\Gamma], k \in [K]; \hat{Q}_0(\gamma) \leftarrow 0, \forall \gamma \in [\Gamma]$ Define constants $\alpha_t \leftarrow 1/(1+t), t = 1, 2, \dots$ /* Define squared temporal difference error */ $\mathcal{L}(Q \mid \hat{Q}, Q_{\text{next}}, \mathcal{D}, \xi, \eta, \alpha) = \sum_{(s, a, r, s') \in \mathcal{D}} (Q - \eta \xi - \eta (1 - \alpha) \hat{Q}(\phi(s, a)) - \alpha \eta (r + \max_{a' \in \mathcal{A}} Q_{\text{next}}(s', a')))^2$ Sample $H_0 \sim \text{Geometric}(1-\eta)$, set $H_0 \leftarrow \min\{H_0, T+1-t\}$ Each agent randomly samples the initial trajectory $\{s_{0,1}^p, a_{0,1}^p, r_{0,1}^p, \dots, s_{0,H}^p, a_{0,H}^p, r_{0,H}^p\}_{p=1}^N$, with $s_{0,1}^p = s_1^p$ $k \leftarrow k+1, t_k = 1+H_0$ $t_k \leftarrow$ the start time of pseudo-episode k while $t \leq T$ do Sample $H \sim \text{Geometric}(1 - \eta)$ $H \leftarrow \min\{H, T+1-t\}$ $t_{k+1} \leftarrow t_k + H$ (the start time of pseudo-episode k + 1) /*Each agent rollouts in the environment */ for p = 1, ..., N do /*Executed in parallel*/ for $t = t_k, \ldots, t_{k+1} - 1$ do $\begin{array}{l} a_t^p \leftarrow \arg\max_{a \in \mathcal{A}} \hat{Q}_t^p(\phi(s_t^p, a)) \\ \text{observe reward } r_t^p \text{ and next state } s_{t+1}^p \end{array}$ $\mathcal{D}_k \leftarrow \mathcal{D}_k \cup \{(s_t^p, a_t^p, r_t^p, s_{t+1}^p)\}$ end /* Visitation of aggregated-states */ $N_{k-1}(\gamma) \leftarrow \sum_{p=1}^{N} \sum_{t=0}^{t_k-1} \mathbf{1}\{\phi(s_t^p, a_t^p) = \gamma\}, \ \forall \gamma \in [\Gamma]$ end /* Construct perturbed datasets and sample regularization noise $ar{Q}$ */ for $p \in [N]$ and $t = t_k, ..., t_{k+1} - 1$ do $\begin{array}{l} \texttt{/* Executed in parallel */} \\ \texttt{Sample array } \tilde{Q}_t^p(s,a) \sim \mathcal{N}(0, \frac{\beta_{N_{k-1}(\phi(s,a))}}{N_{k-1}(\phi(s,a))+1}), \forall (s,a) \\ \texttt{/* Draw prior sample */} \end{array}$ $\mathcal{H}_k^p \leftarrow \{\}$ for $(s, a, r, s') \in \mathcal{D}_k$ do $\begin{array}{l} \texttt{ (* Randomly perturb data */} \\ \textbf{Sample } w^p(s,a) \sim \mathcal{N}(0, \frac{\beta_{N_{k-1}}(\phi(s,a))}{N_{k-1}(\phi(s,a))+1}) \\ \mathcal{H}^p_k \leftarrow \mathcal{H}^p_k \cup \{(s,a,r+w^p,s')\} \end{array}$ end end /* Estimate Q on perturbed data */ for p = 1, ..., N do /* Executed in Parallel */ Define terminal value $\hat{Q}_{t_{k+1}}^p(\gamma) \leftarrow 0 \ \forall \gamma \in [\Gamma]$ for $t = t_{k+1} - 1, \ldots, t_k$ do /* Estimate Q on noisy data */ $\hat{Q}_t^p \leftarrow \arg\min_{Q \in \mathbb{R}} \mathcal{L}(Q \mid \hat{Q}_{k-1}, \hat{Q}_{t+1}^p, \mathcal{H}_k^p, \xi, \eta, \alpha_{N_{k-1}(\gamma)}) + \|Q - \eta \alpha_{N_{k-1}(\gamma)} \tilde{Q}_t^p\|_2^2$ $\hat{Q}_t^p \leftarrow \min\{\hat{Q}_t^p, \frac{1}{1-n}\}, \forall \gamma \in [\Gamma]$ end $s_{t_k}^p = s_1^p, \forall p \in [N].$ end $\hat{Q}_{k}(\gamma) = \frac{\sum_{p=1}^{N} \sum_{t=t_{k}}^{t_{k+1}-1} \mathbf{1}\{\phi(s_{t}^{p}, a_{t}^{p}) = \gamma\} \hat{Q}_{t_{k}}^{p}(\gamma)}{N_{k}(\gamma)}, \forall \gamma \in \Gamma$ $t \leftarrow t_{k+1}, k \leftarrow k+1$ end

Algorithm 3 Concurent RLSVI: Finite-Horizon (Storing the Data of One Episode) **Data:** $K, H, S, \mathcal{A}, N, \mathbf{s}_1, \{\phi_h\}_{h=1}^H$, Tuning parameters $\{\beta_k\}_{k \in \mathbb{N}}$ Define constants $\alpha_t \leftarrow 1/(1+t), t = 1, 2, \dots$ /* Define squared temporal difference error */ $\mathcal{L}(Q \mid Q_{k-1,h}, Q_{\text{next}}, \mathcal{D}, \xi, \alpha) = \sum_{(s,a,r,s') \in \mathcal{D}} (Q - \xi - (1 - \alpha)Q_{k-1,h}(\phi_h(s,a)) - \alpha(r + \max_{a' \in \mathcal{A}} Q_{\text{next}}(s',a')))^2$ **Initialize:** $\hat{Q}_{0h}^{p}(\gamma) = H, \forall h \in [H], \gamma \in [\Gamma], p \in [N]$ Each agent randomly samples the initial trajectory $\{s_{0,1}^p, a_{0,1}^p, r_{0,1}^p, \dots, s_{0,H}^p, a_{0,H}^p, r_{0,1}^H\}_{p=1}^N$, with $s_{0,1}^p = s_1^p$ $N_{0,h}(\gamma) = \sum_{p=1}^{N} \mathbf{1}\{\phi_h(s_{0,h}^p, a_{0,h}^p) = \gamma\}, \forall \gamma \in [\Gamma], h \in [H]$ compute $Q_{0,h}$ by (6) for *episode* k = 1, 2, ... do /*Each agent rollouts in the evironment*/ for p = 1, ..., N do /*Executed in parallel*/ for period $h = 1, \ldots, H$ do $\begin{array}{l} \hat{a}_{k,h}^{p} \leftarrow \arg\max_{a \in \mathcal{A}} \hat{Q}_{k-1,h}^{p}(\phi_{h}(s_{k,h}^{p},a)) \\ \text{observe reward } r_{k,h}^{p} \text{ and next state } s_{k,h+1}^{p} \\ \mathcal{D}_{h}^{k} \leftarrow \mathcal{D}_{h}^{k} \cup \{(s_{k,h}^{p}, a_{k,h}^{p}, r_{k,h}^{p}, s_{k,h+1}^{p})\} \end{array}$ end end /*Visitation of aggregated-states*/ $N_{k,h}(\gamma) \leftarrow \sum_{p=1}^{N} \mathbf{1}\{\phi_h(s_{k,h}^p, a_{k,h}^p) = \gamma\}, \forall \gamma \in [\Gamma], h \in [H]$ /*Construct perturbed data sets and sample regularization noise $\hat{Q}\star/$ for $p \in [N]$ and $h \in [H]$ do /*Executed in parallel*/ For any $(s, a) \in \mathcal{S} \times \mathcal{A}$, sample array $\tilde{Q}_{kh}^{p}(s,a) \sim \mathcal{N}\left(0, \frac{\beta_{k}}{1+N_{k,h}(\phi_{h}(s,a))}\right),$ /* Draw prior sample */ $\tilde{D}_{kh}^{p} \leftarrow \{\}$ for $(s, a, r, s') \in \mathcal{D}_{h}^{k}$ do Sample $w^p(s,a) \sim \mathcal{N}\left(0, \frac{\beta_k}{1+N_{k,h}(\phi_h(s,a))}\right)$ $\tilde{D}^p_{kh} \leftarrow \tilde{D}^p_{kh} \cup \{(s, a, r + w^p, s')\}$ end end /*Estimate Q on perturbed data*/ for p = 1, ..., N do /*Executed in parallel*/ Define terminal value $\hat{Q}_{k,H+1}^p(\gamma) \leftarrow H \ \forall \gamma \in [\Gamma]$ for period $h = H, \ldots, 1$ do $\bar{Q}_{k,h}^p(\gamma) \leftarrow \arg\min_{Q \in \mathbb{R}} \mathcal{L}(Q|\hat{Q}_{k-1,h}, \hat{Q}_{k,h+1}^p, \tilde{D}_{kh}^p, \xi, \alpha) + \|Q - \alpha_{N_{k-1,h}(\gamma)} \tilde{Q}_{kh}^p\|_2^2, \forall \gamma \in [\Gamma]$ $\hat{Q}_{k,h}^p(\gamma) \leftarrow \min\{\bar{Q}_{k,h}^p(\gamma), H\}, \forall \gamma \in [\Gamma]$ end $s_{k,1}^p \leftarrow s_1^p, \mathcal{D}_h^k = \emptyset$ end Update $\hat{Q}_{k,h}(\gamma), \forall \gamma \in [\Gamma]$ by (6) end

Algorithm 4 Concurrent RLSVI: Infinite-Horizon (Storing the Data of One Pseudo-episode) **Data:** Discount factor η , $t_0 = 1$, t = 1, k = 0, $X_1 = 0$, S, A, N, T, ϕ , tuning parameters $\{\beta_k\}_{k \in \mathbb{N}}, \xi, \eta$ **Initialize** $N_k(\gamma) \leftarrow 0, \forall \gamma \in [\Gamma], k \in [K]; \ddot{Q}_0(\gamma) \leftarrow 0, \forall \gamma \in [\Gamma]$ Define constants $\alpha_t \leftarrow 1/(1+t), t = 1, 2, \dots$ /* Define squared temporal difference error */ $\mathcal{L}(Q \mid Q, Q_{\text{next}}, \mathcal{D}, \xi, \eta, \alpha) = \sum_{(s, a, r, s') \in \mathcal{D}} (Q - \eta \xi - \eta (1 - \alpha) \hat{Q}(\phi(s, a)) - \alpha \eta (r + \max_{a' \in \mathcal{A}} Q_{\text{next}}(s', a')))^2$ Sample $H_0 \sim \text{Geometric}(1-\eta)$, set $H_0 \leftarrow \min\{H_0, T+1-t\}$ Each agent randomly samples the initial trajectory $\{s_{0,1}^p, a_{0,1}^p, r_{0,1}^p, \dots, s_{0,H}^p, a_{0,H}^p, r_{0,H}^p\}_{p=1}^N$, with $s_{0,1}^p = s_1^p$ $k \leftarrow k+1, t_k = 1+H_0$ $t_k \leftarrow$ the start time of pseudo-episode k while t < T do Sample $H \sim \text{Geometric}(1 - \eta)$ $H \leftarrow \min\{H, T+1-t\}$ $t_{k+1} \leftarrow t_k + H$ (the start time of pseudo-episode k + 1) /*Each agent rollouts in the environment */ for p = 1, ..., N do /*Executed in parallel*/ for $t = t_k, \ldots, t_{k+1} - 1$ do $\begin{array}{l} a_t^p \leftarrow \arg\max_{a \in \mathcal{A}} \hat{Q}_t^p(\phi(s_t^p, a)) \\ \text{observe reward } r_t^p \text{ and next state } s_{t+1}^p \end{array}$ $\mathcal{D}_k \leftarrow \mathcal{D}_k \cup \{(s_t^p, a_t^p, r_t^p, s_{t+1}^p)\}$ end /* Visitation of aggregated-states */ $N_{k-1}(\gamma) \leftarrow \sum_{p=1}^{N} \sum_{t=t_{k-1}}^{t_k-1} \mathbf{1}\{\phi(s_t^p, a_t^p) = \gamma\}, \ \forall \gamma \in [\Gamma]$ end /* Construct perturbed datasets and sample regularization noise \hat{Q} */ for $p \in [N]$ and $t = t_k, ..., t_{k+1} - 1$ do $\begin{array}{l} & \text{Figure 1} \text{ for } k_{k}, \dots, k_{k+1} \text{ for } k_{k} \text{ for } k_{k}, \dots, k_{k+1} \text{ for } k_{k} \text{ fo$ $\mathcal{H}_k^p \leftarrow \{\}$ for $(s, a, r, s') \in \mathcal{D}_k$ do /* Randomly perturb data */ Sample $w^p(s, a) \sim \mathcal{N}(0, \frac{\beta_{N_{k-1}(\phi(s, a))}}{N_{k-1}(\phi(s, a))+1})$ $\mathcal{H}_k^p \leftarrow \mathcal{H}_k^p \cup \{(s, a, r + w^p, s')\}$ end end /* Estimate Q on perturbed data */ for p = 1, ..., N do /* Executed in Parallel */ Define terminal value $\hat{Q}_{t_{k+1}}^p(\gamma) \leftarrow 0 \ \forall \gamma \in [\Gamma]$ for $t = t_{k+1} - 1, \ldots, t_k$ do /* Estimate Q on noisy data */ $\hat{Q}_t^p \leftarrow \arg\min_{Q \in \mathbb{R}} \mathcal{L}(Q \mid \hat{Q}_{k-1}, \hat{Q}_{t+1}^p, \mathcal{H}_k^p, \xi, \eta, \alpha_{N_{k-1}(\gamma)}) + \|Q - \eta \alpha_{N_{k-1}(\gamma)} \tilde{Q}_t^p\|_2^2$ $\hat{Q}_t^p \leftarrow \min\{\hat{Q}_t^p, \frac{1}{1-\eta}\}, \forall \gamma \in [\Gamma]$ end $s_{t_k}^p = s_1^p, \forall p \in [N]; \mathcal{D}_k \leftarrow \emptyset.$ end $\hat{Q}_{k}(\gamma) = \frac{\sum_{p=1}^{N} \sum_{t=t_{k}}^{t_{k+1}-1} \mathbf{1}\{\phi(s_{t}^{p}, a_{t}^{p}) = \gamma\} \hat{Q}_{t_{k}}^{p}(\gamma)}{N_{k}(\gamma)}$ $\forall \gamma \in \Gamma$ $t \leftarrow t_{k+1}, k \leftarrow k+1$ end

C. Proofs for the Finite-Horizon Case

In this section, we prove the worst-case regret bound for the finite-horizon case. We use the following notations:

- $n_h^k(\gamma)$: the number of visits to aggregate state γ at period h from episodes 0 to k-1.
- $N_{k,h}(\gamma)$: total number of agents who visit aggregate state γ during episode k and period h.
- π^{kp} : the greedy policy with respect to $\hat{Q}_{k,h}^p$, i.e. the policy that the agent p follows to produce the trajectory $s_{k,1}^p, a_{k,1}^p, r_{k,1}^p, \dots, s_{k,H}^p, a_{k,H}^p, r_{k,H}^p$.
- $\hat{V}^p_{k,h}(s)$: the state value function estimate at period h, induced by $\hat{Q}^p_{k,h}(\gamma)$ through

$$\hat{V}_{k,h}^p(s) = \max_{a \in \mathcal{A}} \hat{Q}_{k,h}^p(\phi_h(s,a)).$$

- $V^p(M, \pi)$: the value function corresponding to policy π from the initial state s_1^p , where s_1^p is the initial state of agent p at the beginning of each episode. For the true MDP M we have $V_1^{\pi}(s_1^p) := V^p(M, \pi)$.
- $\hat{V}_{k,1}^{\pi^{k_p}}(s_1^p) := V^p(\overline{M}^{k_p}, \pi^{k_p})$: the value function corresponding to MDP \overline{M}^{k_p} with initial state s_1^p and policy π^{k_p} , where MDP \overline{M}^{k_p} is defined as (4).

For any $\gamma \in [\Gamma]$, we define the following events:

$$\mathcal{E}(\gamma) := \left\{ \left| \frac{1}{N_{k-1,h}(\gamma)} \sum_{j=1}^{N_{k-1,h}(\gamma)} \mathbf{1}\{\phi_h(s_{k-1,h}^j, a_{k-1,h}^j) = \gamma\}\{V_{h+1}^*(s_{k-1,h}^j) - P_h V_{h+1}^*(s_{k-1,h}^j)\} \right| \\
\leq \frac{2H\sqrt{\log(2KHN/\delta)}}{\sqrt{N_{k-1,h}(\gamma)}} \right\}.$$
(18)

$$\mathcal{G}(\gamma) := \left\{ \left| \frac{1}{N_{k-1,h}(\gamma)} \sum_{j=1}^{N_{k-1,h}(\gamma)} \mathbf{1}\{\phi_h(s_{k-1,h}^j, a_{k-1,h}^j) = \gamma\} \tilde{Q}_{k-1,h}^j(s_{k-1,h}^j, a_{k-1,h}^j) \right| \\
\leq 2 \frac{\sqrt{\beta_{k-1} \log(2KHN/\delta)}}{\sqrt{(N_{k-1,h}(\gamma) + 1)N_{k-1,h}(\gamma)}} \right\}.$$
(19)

C.1. Proof of Finite-Horizon Main Result: Theorem 2

Proof of Theorem 2. Following the derivation of (37) and (38) in Lemma 11, we have

$$\sum_{\gamma \in [\Gamma]} \sum_{h=1}^{H} \sum_{k=1}^{K} \frac{1}{\sqrt{N_{k-1,h}(\gamma)+1}} = \sum_{h=1}^{H} \sum_{k=1}^{K} \sum_{\gamma \in [\Gamma]} \sum_{j=1}^{N_{k-1,h}(\gamma)} \frac{1}{\sqrt{j}}$$

$$\leq \sum_{k=1}^{K} \sum_{h=1}^{H} \sum_{\gamma \in [\Gamma]} 2\sqrt{N_{k,h}(\gamma)} \leq 2\sqrt{KH\Gamma\sum_{k=1}^{K} \sum_{h=1}^{H} \sum_{\gamma \in [\Gamma]} N_{k,h}(\gamma)} = 2KH\sqrt{\Gamma N}$$
(20)

and

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \sum_{h=1}^{H} \frac{1}{\sqrt{N_{K-1,h}(\gamma_{kh}^{p})}} = \sum_{h=1}^{H} \sum_{k=1}^{K} \sum_{\gamma \in [\Gamma]} \sum_{j=1}^{N_{K-1,h}(\gamma)} \frac{1}{\sqrt{j}}$$
$$\leq \sum_{h=1}^{H} \sum_{k=1}^{K} \sum_{\gamma \in [\Gamma]} 2\sqrt{N_{k-1,h}(\gamma)}$$
$$\leq 2\sqrt{HK\Gamma} \sum_{h=1}^{H} \sum_{k=1}^{K} \sum_{\gamma \in [\Gamma]} N_{k-1,h}(\gamma)$$
$$= 2HK\sqrt{\Gamma N}.$$

Under events $\mathcal{E}(\gamma_{kh}^p), \mathcal{G}(\gamma_{kh}^p), \forall \gamma_{kh}^p \in [\Gamma]$, we have

$$\operatorname{Regret}(M, K, N, \pi, \operatorname{RLSVI}_{\beta, \alpha, \xi}) = \sum_{k=1}^{K} \sum_{p=1}^{N} V_1^*(s_1^p) - V_{\pi^{k_p}}^{\eta}(s_1^p) \le \sum_{k=1}^{K} \sum_{p=1}^{N} \hat{V}_{k,1}(s_1^p) - V_{\pi^{k_p}}^{\eta}(s_1^p).$$

Recall from Lemma 11 that when $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ hold for all $\gamma_{kh}^p \in [\Gamma]$, with probability $1 - 2\delta$, we have

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \hat{V}_{k,1}^{p}(s_{1}^{p}) - V_{\pi^{kp}}^{\eta}(s_{1}^{p}) \\ \leq 2\epsilon K H N + 8K H^{2} \sqrt{\Gamma N} \sqrt{\log(2KHN/\delta)} + 32H^{2} \sqrt{K\Gamma N} \sqrt{\log(HKN/\delta)} \\ + 2K H^{5/2} \Gamma \sqrt{N} \sqrt{\log(2KH\Gamma)} \sqrt{\log(2KHN/\delta)}.$$

Note that by (18), (19) and the first statement of Lemma 10, we have

$$\mathbb{P}(\mathcal{E}(\gamma_{kh}^{p}), \mathcal{G}(\gamma_{kh}^{p}), \gamma_{kh}^{p}, \forall k \in [K], h \in [H], p \in [N]) \\
\geq 1 - \sum_{k \in [K], h \in [H], p \in [N]} (\mathbb{P}(\mathcal{E}(\gamma_{kh}^{p})^{c}) + \mathbb{P}(\mathcal{G}(\gamma_{kh}^{p})^{c})) \\
\geq 1 - 2NKH\delta/(2KHN) \geq 1 - \delta.$$
(21)

Hence with probability $1 - 3\delta$, we have

$$\begin{aligned} &\operatorname{Regret}(M, K, N, \pi, \operatorname{RLSVI}_{\beta, \alpha, \xi}) \\ &\leq 2\epsilon K H N + 8 K H^2 \sqrt{\Gamma N} \sqrt{\log(2KHN/\delta)} + 32 H^2 \sqrt{K\Gamma N} \sqrt{\log(HKN/\delta)} \\ &+ 2K H^{5/2} \Gamma \sqrt{N} \sqrt{\log(2KH\Gamma)} \sqrt{\log(2KHN/\delta)}. \end{aligned}$$

So we obtain (9).

C.2. Lemmas for State-Aggregation Results

Lemma 8 (Concentration Bound). Conditioning on $\mathcal{D}_k = \bigcup_{p \in [N]} \sum_{h \in [H]} \{s_{k-1,h}^p, a_{k-1,h}^p\}$ as the trajectories of episode k - 1 across all agents, for every possible $\gamma_{kh}^p = \phi_h(s_{k,h}^p, a_{k,h}^p)$ in episode k, event $\mathcal{E}(\gamma_{kh}^p)$ defined in (18) and event $\mathcal{G}(\gamma_{kh}^p)$ both hold with probability $1 - \delta/(2KHN)$.

Proof of Lemma 8. By Höeffding's inequality, conditional on the trajectory during episode k - 1, with probability $1 - \delta/(2KHN)$, we have

$$\left|\frac{1}{N_{k-1,h}(\gamma_{kh}^{p})}\sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})}\left\{V_{h+1}^{*}(s_{k-1,h+1}^{j}) - P_{h}V_{h+1}^{*}(s_{k-1,h}^{j})\right\}\right| \le \frac{2H\sqrt{\log(2KHN/\delta)}}{\sqrt{N_{k-1,h}(\gamma_{kh}^{p})}}.$$
(22)

Additionally, recall from Algorithm 3 that for $k \ge 2$, the random perturbation during episode k-1 is drawn from normal distribution $\tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \sim \mathcal{N}(0, \frac{\beta_{k-1}}{N_{k-2,h}(\phi_h(s_{k-1,h}^{j}, a_{k-1,h}^{j}))+1})$, and since $\phi_h(s_{k-1,h}^{j}, a_{k-1,h}^{j}) = \gamma_{kh}^{p}$ in (24), we have $\tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \sim \mathcal{N}(0, \frac{\beta_{k-1}}{N_{k-2,h}(\gamma_{kh}^{p})+1})$. Also note that the random perturbations $\tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})$ are i.i.d. across $j \in [N]$, hence by Höeffding bound, conditioning on \mathcal{D}_k , with probability $1 - \delta/(2KHN)$, we have

$$\left|\frac{1}{N_{k-1,h}(\gamma_{kh}^p)}\sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^p)}\tilde{Q}_{k-1,h}^j(s_{k-1,h}^j,a_{k-1,h}^j)\right| \le 2\sqrt{\frac{\beta_{k-1}}{N_{k-2,h}(\gamma_{kh}^p)+1}}\sqrt{\frac{\log(2KHN/\delta)}{N_{k-1,h}(\gamma_{kh}^p)}}.$$

Thus the result follows.

Lemma 9 (Optimism). When events $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ hold for all γ_{kh}^p in episode k and for all $k \in [K], p \in [N]$, the on-policy error $\hat{V}_{k,h}^p(s) - V_h^*(s)$ is lower bounded by zero for any $s \in S$, $k \in [K]$, $p \in [N]$.

Proof of Lemma 9. Recall from Algorithm 3 that the unclipped value function estimates $\bar{Q}_{k,h}^{p}(\cdot)$ in episode k can be written as

$$\bar{Q}_{k,h}^{p}(\gamma) = \arg\min_{Q \in \mathbb{R}} \sum_{(s,a):\phi_{h}(s,a)=\gamma} (Q - \xi_{N_{k-1,h}(\gamma)} - (1 - \alpha_{N_{k-1,h}(\gamma)})\hat{Q}_{k-1,h}(\gamma) - \alpha_{N_{k-1,h}(\gamma)} \{r(s,a) + \max_{a' \in \mathcal{A}} \hat{Q}_{k,h+1}^{p}(s',a')\})^{2} + \left\|Q - \alpha_{N_{k-1,h}(\gamma)}\tilde{Q}_{k,h}^{p}\right\|_{2}^{2}.$$

By first-order condition, we have

$$0 = 2 \sum_{(s,a):\phi_h(s,a)=\gamma} (\bar{Q}^p_{k,h}(\gamma) - \xi_{N_{k-1,h}(\gamma)} - (1 - \alpha_{N_{k-1,h}(\gamma)}) \hat{Q}_{k-1,h}(\gamma) - \alpha_{N_{k-1,h}(\gamma)} \{r(s,a) + \max_{a' \in \mathcal{A}} \hat{Q}^p_{k,h+1}(s',a')\}) + 2 \sum_{(s,a)\in \mathcal{D}^k_h:\phi_h(s,a)=\gamma} (\bar{Q}^p_{k,h}(\gamma) - \alpha_{N_{k-1,h}(\gamma)} \tilde{Q}^p_{k,h}(s,a)),$$

so by calculation we have

$$\begin{split} \bar{Q}_{k,h}^{p}(\gamma) &= \xi_{N_{k-1,h}(\gamma)} + \frac{1}{N_{k-1,h}(\gamma)} \sum_{(s,a) \in \mathcal{D}_{h}^{k}:\phi_{h}(s,a) = \gamma} (1 - \alpha_{N_{k-1,h}(\gamma)}) \hat{Q}_{k-1,h}(\gamma) \\ &+ \alpha_{N_{k-1,h}(\gamma)}(r(s,a) + \max_{a' \in \mathcal{A}} \hat{Q}_{k,h+1}^{p}(s',a')) + \alpha_{N_{k-1,h}(\gamma)} \tilde{Q}_{k,h}^{p}(s,a)) \\ &= \xi_{N_{k-1,h}(\gamma)} + \frac{1}{N_{k-1,h}(\gamma)} \sum_{p=1}^{N} \sum_{\substack{(s_{k-1,h}^{p}, a_{k-1,h}^{p}) \in \mathcal{D}_{kh}^{p}:\phi_{h}(s,a) = \gamma}} (1 - \alpha_{N_{k-1,h}(\gamma)}) \hat{Q}_{k-1,h}(\gamma) \\ &+ \alpha_{N_{k-1,h}(\gamma)}(r(s_{k-1,h}^{p}, a_{k-1,h}^{p}) + \hat{V}_{k,h+1}^{p}(s_{k-1,h+1}^{p})) + \alpha_{N_{k-1,h}(\gamma)} \tilde{Q}_{k,h}^{p}(s_{k-1,h}^{p}, a_{k-1,h}^{p}) \end{split}$$

$$=\xi_{N_{k-1,h}(\gamma)} + (1 - \alpha_{N_{k-1,h}(\gamma)})Q_{k-1,h}(\gamma) + \frac{\alpha_{N_{k-1,h}(\gamma)}}{N_{k-1,h}(\gamma)}\sum_{p=1}^{N}\sum_{\substack{(s_{k-1,h}^{p}, a_{k-1,h}^{p}) \in \mathcal{D}_{k,h}^{p}: \phi_{h}(s,a) = \gamma}} (r(s_{k-1,h}^{p}, a_{k-1,h}^{p}) + \hat{V}_{k,h+1}^{p}(s_{k-1,h+1}^{p}) + \tilde{Q}_{k,h}^{p}(s_{k-1,h}^{p}, a_{k-1,h}^{p}))$$

$$=\xi_{N_{k-1,h}(\gamma)} + (1 - \alpha_{N_{k-1,h}(\gamma)})\hat{Q}_{k-1,h}(\gamma) + \frac{\alpha_{N_{k-1,h}(\gamma)}}{N_{k-1,h}(\gamma)} \sum_{p=1}^{N_{k-1,h}(\gamma)} (r(s_{k-1,h}^{p}, a_{k-1,h}^{p}) + \hat{V}_{k,h+1}^{p}(s_{k-1,h+1}^{p}) + \tilde{Q}_{k,h}^{p}(s_{k-1,h}^{p}, a_{k-1,h}^{p}))$$

$$=\xi_{N_{k-1,h}(\gamma)} + \frac{1}{N_{k-1,h}(\gamma)} \sum_{p=1}^{N_{k-1,h}(\gamma)} (1 - \alpha_{N_{k-1,h}(\gamma)}) \hat{Q}_{k-1,h}^{p}(\gamma) + \alpha_{N_{k-1,h}(\gamma)} (r(s_{k-1,h}^{p}, a_{k-1,h}^{p}) + \hat{V}_{k,h+1}^{p}(s_{k-1,h+1}^{p}) + \tilde{Q}_{k,h}^{p}(s_{k-1,h}^{p}, a_{k-1,h}^{p})),$$

where in the above derivation we used the fact that $\sum_{p=1}^{N} \mathbf{1}\{\phi_h(s_{k-1,h}^p, a_{k-1,h}^p) = \gamma\} = N_{k-1,h}(\gamma).$ Thus for $\gamma_{kh}^p = \phi_h(s_{k,h}^p, a_{k,h}^p)$, with $\phi_h(s_{k-1,h}^j, a_{k-1,h}^j) = \gamma_{kh}^p$ in the following, we have

$$\begin{split} &\tilde{Q}_{k,h}^{p}(s_{k,h}^{p}, a_{k,h}^{p}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}) \\ &= \xi_{N_{k-1,h}(\gamma)} + \frac{1}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} (1 - \alpha_{N_{k-1,h}(\gamma_{kh}^{p})}) (\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p})) \\ &+ \alpha_{N_{k-1,h}(\gamma_{kh}^{p})} [\{r_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) + \hat{V}_{k-1,h+1}(s_{k-1,h+1}^{j}) + \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\} \\ &= \xi_{N_{k-1,h}(\gamma)} + \frac{1}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} (1 - \alpha_{N_{k-1,h}(\gamma_{kh}^{p})}) (\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p})] \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{p=1}^{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j})} (1 - \alpha_{N_{k-1,h}(\gamma_{k}^{p})}) (\hat{Q}_{k-1,h+1}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k,h}^{j}, a_{k,h}^{j})) \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{p=1}^{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j})} (1 - \alpha_{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j}) + \hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})) \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{k}^{p})}}{N_{k-1,h}(\gamma_{k}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j})} (2q_{h}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k,h}^{j}, a_{k,h}^{j})) \\ &= 2 - \epsilon + \xi_{N_{k-1,h}(\gamma)} + \frac{1}{N_{k-1,h}(\gamma_{k}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k}^{p})} (1 - \alpha_{N_{k-1,h}(\gamma_{k}^{p}, a_{k,h}^{j})) \\ &= 2 - \epsilon + \xi_{N_{k-1,h}(\gamma)} + \frac{1}{N_{k-1,h}(\gamma_{k}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k}^{p})} (1 - \alpha_{N_{k-1,h}(\gamma_{k}^{p}, a_{k,h}^{j})) \\ &= 2 - \epsilon + \xi_{N_{k-1,h}(\gamma)} + \frac{1}{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j})} (1 - \alpha_{N_{k-1,h}(\gamma_{k}^{p}, a_{k,h}^{j})) \\ &= - \epsilon - \epsilon + \xi_{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j}) \sum_{j=1}^{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j})} (1 - \alpha_{N_{k-1,h}(\gamma_{k}^{p}, a_{k,h}^{j})) \\ &= - \epsilon - \epsilon + \xi_{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j}) \sum_{j=1}^{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k}^{p}, a_{k-1,h}^{j})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k}^{p},$$

where we use the definition of ϵ -error aggregation as in Definition 1 in the last inequality above, and the right hand side of (23) is equal to

$$\begin{aligned} \text{RHS of } (23) &=_{(1)} -\epsilon + \xi_{N_{k-1,h}(\gamma)} + \frac{\left(1 - \alpha_{N_{k-1,h}(\gamma_{kh}^{p})}\right)}{N_{k-1,h}(\gamma)} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left(\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\right) \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left(\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) + \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left\{\hat{V}_{k-1,h+1}^{j}(s_{k-1,h+1}^{j}) - V_{h+1}^{*}(s_{k-1,h+1}^{j})\right\} \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left\{V_{h+1}^{*}(s_{k-1,h+1}^{j}) - P_{h}V_{h+1}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\right\} \\ &= (2) - \epsilon + \xi_{N_{k-1,h}(\gamma)} + \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left\{\hat{V}_{k-1,h+1}^{j}(s_{k-1,h+1}^{j}) - V_{h+1}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\right\} \\ &+ \frac{1}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left\{V_{h+1}^{*}(s_{k-1,h+1}^{j}) - V_{h+1}^{*}(s_{k-1,h+1}^{j})\right\} \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left\{V_{h+1}^{*}(s_{k-1,h+1}^{j}) - P_{h}V_{h+1}^{*}(s_{k-1,h}^{j})\right\} \end{aligned}$$

where the equalities (1) and (2) use the fact that

$$Q_h(s',a') = r_h(s',a') + P_h V_{h+1}^*(s',a'), \ \forall (s',a') \in \mathcal{S} \times \mathcal{A}.$$

(24)

Suppose events $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ hold for all γ_{kh}^p in episode k, then by (18),(19), we have

$$\frac{\alpha_{N_{k-1,h}(\gamma_{kh}^p)}}{N_{k-1,h}(\gamma_{kh}^p)} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^p)} \{V_{h+1}^*(s_{k-1,h+1}^j) - P_h V_{h+1}^*(s_{k-1,h}^j)\} \ge -\frac{2\alpha_{N_{k-1,h}(\gamma_{kh}^p)}H\sqrt{\log(2KHN/\delta)}}{\sqrt{\max\{N_{k-1,h}(\gamma_{kh}^p),1\}}},$$

and

$$\frac{\alpha_{N_{k-1,h}(\gamma_{kh}^p)}}{N_{k-1,h}(\gamma_{kh}^p)} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^p)} \tilde{Q}_{k-1,h}^j(s_{k-1,h}^j, a_{k-1,h}^j) \ge -\frac{2\alpha_{N_{k-1,h}(\gamma_{kh}^p)}\sqrt{\beta_{k-1}\log(2KHN/\delta)}}{\sqrt{(N_{k-1,h}(\gamma_{kh}^p)+1)\max\{N_{k-1,h}(\gamma_{kh}^p),1\}}}$$

So recall from (7)

$$\xi_{N_{k-1,h}(\gamma_{kh}^p)} = \epsilon + \frac{2\alpha_{N_{k-1,h}(\gamma_{kh}^p)}H\sqrt{\log(2KHN/\delta)}}{\sqrt{\max\{N_{k-1,h}(\gamma_{kh}^p),1\}}} + \frac{2\alpha_{N_{k-1,h}(\gamma_{kh}^p)}\sqrt{\beta_{k-1}\log(2KHN/\delta)}}{\sqrt{(N_{k-1,h}(\gamma_{kh}^p)+1)\max\{N_{k-1,h}(\gamma_{kh}^p),1\}}},$$

thus we have

$$\bar{Q}_{k,h}^{p}(s_{k,h}^{p}, a_{k,h}^{p}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}) \geq \frac{1}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \{\hat{V}_{k-1,h+1}^{j}(s_{k-1,h+1}^{j}) - V_{h+1}^{*}(s_{k-1,h+1}^{j})\}.$$
(25)

Note that when $N_{k-1,h}(\gamma) = 0$, then $\bar{Q}_{k,h}^p(\gamma) = \hat{Q}_{k,h}^p(\gamma) = H$, and we define terminal values as H with $\bar{Q}_{k,h+1}^p(\gamma) = \hat{Q}_{k,h+1}^p(\gamma) = H$, so by plugging in h = H in (25), for any possible $\gamma \in [\Gamma]$, we have $\bar{Q}_{k,H}^p(\gamma) \ge Q_H^*(\gamma)$. Thus

$$\hat{V}_{k,H}^p(s) \ge V_H^*(s), \, \forall s \in \mathcal{S}, k \in [K], p \in [N].$$

Furthermore, suppose that at period h, we have

$$\hat{V}_{k,h}^p(s) \ge V_h^*(s), \ \forall s \in \mathcal{S}, k \in [K], p \in [N],$$

then from (25) we have

$$\bar{Q}_{k,h-1}^{p}(s,a) - Q_{h}^{*}(s,a) \geq \frac{1}{N_{k-1,h}(\phi_{h}(s,a))} \sum_{j=1}^{N_{k-1,h}(\phi_{h}(s,a))} \{\hat{V}_{k-1,h+1}^{j}(s_{k-1,h+1}^{j}) - V_{h+1}^{*}(s_{k-1,h+1}^{j})\} \geq 0,$$

which implies that

$$\bar{Q}_{k,h-1}^p(s,a) \ge Q_h^*(s,a), \; \forall (s,a) \in \mathcal{S} \times \mathcal{A}, k \in [K], p \in [N]$$

By maximizing over $a \in \mathcal{A}$ on both sides above, we have

$$\hat{V}_{k-1,h-1}^p \geq V_{h-1}^*(s), \ \forall s \in \mathcal{S}, k \in [K].$$

Thus by induction, we conclude that when events $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ hold for all γ_{kh}^p in episode k for all $k \in [K]$, we have

$$\hat{V}_{k,h}^p(s) \ge V_h^*(s), \ \forall s \in \mathcal{S}, k \in [K], p \in [N].$$

Lemma 10. Suppose events $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ hold for all γ_{kh}^p , for all $k \in [K]$ and $p \in [N]$, then we have

$$\begin{split} \sum_{p=1}^{N} \sum_{k=1}^{K} \left[\hat{V}_{k,h}^{p}(s_{k,h}^{p}) - V_{h}^{*}(s_{k,h}^{p}) \right] &\leq 2\epsilon K N (H-h+1) \\ &+ 4H \sqrt{\log(2KHN/\delta)} \sum_{k=1}^{K} \sum_{p=1}^{N} \sum_{\ell=h}^{M} \frac{1}{\sqrt{N_{k-1,\ell}(\gamma_{k\ell}^{p})}} \frac{1}{1+N_{k-1,\ell}(\gamma_{k\ell}^{p})} \\ &+ 4\sum_{\ell=h}^{H} \sum_{p=1}^{N} \sum_{k=2}^{K} \frac{\sqrt{\beta_{k-1}\log(2KHN/\delta)}}{\sqrt{\left(n_{k-1}^{h}(\gamma_{kh}^{p})+1\right)N_{k-2,h}(\gamma_{kh}^{p})}} \frac{1}{1+N_{k-1,\ell}(\gamma_{k\ell}^{p})}. \end{split}$$

Proof of Lemma 10. By (23) we have

$$\begin{split} \bar{Q}_{k,h}^{p}(s_{k,h}^{p}, a_{k,h}^{p}) &= Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}) \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} \left[(1 - \alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}) (\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}) \right] \\ &+ \alpha_{N_{k-1,h}(\gamma_{k,h}^{p})} \left[[r_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) + \hat{V}_{k-1,h+1}(s_{k-1,h+1}^{j}) + \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \right] \\ &= \xi_{N_{k-1,h}(\gamma)} + \frac{1}{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} (1 - \alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}) (\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}) \right] \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}}{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{p=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} \{ [r_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) + \hat{V}_{k-1,h+1}^{j}(s_{k-1,h+1}^{j}) + \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \right] \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}}{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} \{ Q_{h}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}) \} \\ &\leq \epsilon + \xi_{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} (1 - \alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}) (\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \right] \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}}{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} (1 - \alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}) (\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \right] \\ &\leq \epsilon + \xi_{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} (1 - \alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}) (\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \right] \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}}{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} \{ [r_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) + \hat{V}_{k-1,h+1}^{j}(s_{k-1,h+1}^{j}) + \hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \right] \\ &- Q_{h}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \}$$

where we use the definition of ϵ -error aggregation as in Definition 1 in the last inequality above, and the right hand side of (26) is equal to

$$\begin{aligned} \text{RHS of } (26) &=_{(1)} \epsilon + \xi_{N_{k-1,h}(\gamma)} + \frac{\left(1 - \alpha_{N_{k-1,h}(\gamma_{kh}^{p})}\right)}{N_{k-1,h}(\gamma)} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left(\hat{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) - Q_{h}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\right) \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left\{\hat{V}_{k-1,h+1}^{j}(s_{k-1,h+1}^{j}) - V_{h+1}^{*}(s_{k-1,h+1}^{j})\right\} \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left\{V_{h+1}^{*}(s_{k-1,h+1}^{j}) - P_{h}V_{h+1}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\right\} \\ &= _{(2)} \epsilon + \xi_{N_{k-1,h}(\gamma)} + \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}, a_{k-1,h}^{j}) \\ &+ \frac{1}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left\{\hat{V}_{k-1,h+1}^{j}(s_{k-1,h+1}^{j}) - V_{h+1}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\right\} \\ &+ \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \left\{\hat{V}_{k-1,h+1}^{j}(s_{k-1,h+1}^{j}) - V_{h+1}^{*}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\right\} \end{aligned}$$

where the equalities (1) and (2) above use the fact that

$$Q_h(s',a') = r_h(s',a') + P_h V_{h+1}^*(s',a'), \ \forall (s',a') \in \mathcal{S} \times \mathcal{A}.$$

(27)

Now suppose $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ holds for all aggregated states γ_{kh}^p during period k for all $p \in [N]$. By definition, note that

$$\hat{V}_{k,h}^{p}(s_{k,h}^{p}) - V_{h}^{*}(s_{k,h}^{p}) \le \hat{Q}_{k,h}^{p}(s_{k,h}^{p}, a_{k,h}^{p}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}) \le \bar{Q}_{k,h}^{p}(s_{k,h}^{p}, a_{k,h}^{p}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}),$$

$$(28)$$

Denote

$$\Delta_{k,h}^{p} = \hat{V}_{k,h}^{p}(s_{k,h}^{p}) - V_{h}^{*}(s_{k,h}^{p}),$$
⁽²⁹⁾

then under events $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p),$

$$\bar{Q}_{k,h}^{p}(s_{k,h}^{p}, a_{k,h}^{p}) - Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}) \leq \epsilon + \xi_{N_{k-1,h}(\gamma)} + \frac{2\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}\sqrt{\beta_{k-1}\log(2KHN/\delta)}}{\sqrt{(N_{k-2,h}(\gamma_{k,h}^{p}) + 1)N_{k-1,h}(\gamma_{k,h}^{p})}} \\
+ \frac{\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}}{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{k,h}^{p})} \Delta_{k,h+1}^{j} + \frac{2\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}H\sqrt{\log(2KHN/\delta)}}{\sqrt{N_{k-1,h}(\gamma_{k,h}^{p})}}.$$
(30)

Since $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ hold for all $\gamma_{kh}^p \in [\Gamma]$, and recall from (7)

$$\xi_{N_{k-1,h}(\gamma_{kh}^p)} = \epsilon + \frac{2\alpha_{N_{k-1,h}(\gamma_{kh}^p)}H\sqrt{\log(2KHN/\delta)}}{\sqrt{\max\{N_{k-1,h}(\gamma_{kh}^p),1\}}} + \frac{2\alpha_{N_{k-1,h}(\gamma_{kh}^p)}\sqrt{\beta_{k-1}\log(2KHN/\delta)}}{\sqrt{(N_{k-1,h}(\gamma_{kh}^p)+1)\max\{N_{k-1,h}(\gamma_{kh}^p),1\}}},$$

so we have

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \Delta_{k,h}^{p}(s_{k,h}^{p}) \leq \sum_{p=1}^{N} \sum_{k=1}^{K} Q_{h}^{*}(s_{k,h}^{p}, a_{k,h}^{p}) - \bar{Q}_{k,h}^{p}(s_{k,h}^{p}, a_{k,h}^{p}) \\ \leq 2KN\epsilon + 4 \sum_{p=1}^{N} \sum_{k=1}^{K} \frac{\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})} H \sqrt{\log(2KHN/\delta)}}{\sqrt{N_{k-1,h}(\gamma_{k,h}^{p})}} \\ + 4 \sum_{p=1}^{N} \sum_{k=2}^{K} \frac{\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})} \sqrt{\beta_{k-1}\log(2KHN/\delta)}}{\sqrt{(N_{k-2,h}(\gamma_{k,h}^{p}) + 1)N_{k-1,h}(\gamma_{k,h}^{p})}} \\ + \sum_{p=1}^{N} \sum_{k=1}^{K} \frac{\alpha_{N_{k-1,h}(\gamma_{k,h}^{p})}}{N_{k-1,h}(\gamma_{k,h}^{p})} \sum_{j=1}^{N} \Delta_{k,h+1}^{j}.$$
(31)

For any $\gamma \in [\Gamma]$, for $s_{k,h}^p, a_{k,h}^p$ such that $\phi_h(s_{k,h}^p, a_{k,h}^p) = \gamma$, denote

$$\bar{\Delta}_{k,h+1}(\gamma) = \frac{1}{N_{k-1,h}(\gamma)} \sum_{p=1}^{N_{k-1,h}(\gamma)} \Delta_{k,h+1}^p.$$

Note that

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \frac{\alpha_{N_{k-1,h}}(\gamma_{kh}^{p})}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \Delta_{k,h+1}^{j}$$

$$= \sum_{k=1}^{K} N_{k-1,h}(\gamma_{kh}^{p}) \alpha_{N_{k-1,h}}(\gamma_{kh}^{p}) \bar{\Delta}_{k,h+1} = \sum_{k=1}^{K} \alpha_{N_{k-1,h}}(\gamma_{kh}^{p}) \sum_{p=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \Delta_{k,h+1}^{p}$$

$$\leq \sum_{k=1}^{K} \sum_{p=1}^{N} \Delta_{k,h+1}^{p},$$
(32)

where the last inequality above follows from Lemma 9.

Hence by above recursion and (31) we have that when $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ hold for all $\gamma_{kh}^p \in [\Gamma]$, we have

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \Delta_{k,h}^{p} \leq 2\epsilon K N (H-h+1) + 4H \sqrt{\log(2KHN/\delta)} \sum_{k=1}^{K} \sum_{p=1}^{N} \sum_{\ell=h}^{H} \frac{1}{\sqrt{N_{k-1,\ell}(\gamma_{k,\ell}^{p})}} \frac{1}{1+N_{k-1,\ell}(\gamma_{k,\ell}^{p})} + 4\sum_{\ell=h}^{H} \sum_{p=1}^{N} \sum_{k=2}^{K} \frac{\sqrt{\beta_{k-1}\log(2KHN/\delta)}}{\sqrt{(N_{k-2,h}(\gamma_{k,h}^{p})+1)N_{k-1,h}(\gamma_{k,h}^{p})}} \frac{1}{1+N_{k-1,\ell}(\gamma_{k,\ell}^{p})}.$$
(33)

Lemma 11. Suppose $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ hold for all $\gamma_{kh}^p \in [\Gamma]$, then with probability $1 - 2\delta$, we have

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \hat{V}_{k,1}^{p}(s_{1}^{p}) - V_{\pi^{k_{p}}}^{\eta}(s_{1}^{p}) \\ \leq 2\epsilon KHN + 8KH^{2} \sqrt{\Gamma N} \sqrt{\log(2KHN/\delta)} + 32H^{2} \sqrt{K\Gamma N} \sqrt{\log(HKN/\delta)} \\ + 2KH^{5/2} \Gamma \sqrt{N} \sqrt{\log(2KH\Gamma)} \sqrt{\log(2KHN/\delta)}.$$

Proof of Lemma 11. Note that

Summing over k = 1, 2, ..., K, by (28), (31), (32) and (33) of Lemma 10 we have

$$\begin{split} &\sum_{p=1}^{N} \sum_{k=1}^{K} \hat{V}_{k,h}^{p}(s_{k,h}^{p}) - V_{h}^{\pi^{kp}}(s_{k,h}^{p}) \\ &\leq \sum_{p=1}^{N} \sum_{k=1}^{K} [\hat{V}_{k,h+1}^{p}(s_{k,h+1}^{p}) - V_{h+1}^{\pi^{kp}}(s_{k,h+1}^{p})] - \sum_{p=1}^{N} \sum_{k=1}^{K} \Delta_{k,h+1}^{p} \\ &+ 2KN\epsilon + 4 \sum_{p=1}^{N} \sum_{k=1}^{K} \frac{\alpha_{N_{k-1,h}}(\gamma_{k,h}^{p}) H \sqrt{\log(2KHN/\delta)}}{\sqrt{N_{k-1,h}}(\gamma_{k,h}^{p})} \\ &+ 2 \sum_{p=1}^{N} \sum_{k=2}^{K} \frac{\alpha_{N_{k-1,h}}(\gamma_{k,h}^{p}) \sqrt{\beta_{k-1}} \log(2KHN/\delta)}{\sqrt{(N_{k-2,h}}(\gamma_{k,h}^{p}) + 1)N_{k-1,h}}(\gamma_{k,h}^{p})} \frac{1}{1 + N_{k-1,h}} + \sum_{k=1}^{K} \sum_{p=1}^{N} [P_{h}V_{h+1}^{*}(s_{k,h+1}^{p}) - V_{h+1}^{*}(s_{k,h+1}^{p})] \\ &+ \sum_{k=1}^{K} \sum_{p=1}^{N} [V_{h+1}^{\pi^{kp}}(s_{k,h+1}^{p}) - P_{h}V_{h+1}^{\pi^{kp}}(s_{k,h+1}^{p})]. \end{split}$$

$$(35)$$

Note that (35) is recursive. Recall that $s_{k,1}^p = s_1^p$ for any $k \in [K], p \in [N]$. Thus by Lemma 10, when $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$

hold for all $\gamma_{kh}^p,$ for all $k\in[K],p\in[N],$ we have

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \hat{V}_{k,1}^{p}(s_{1}^{p}) - V_{\pi^{kp}}^{\eta}(s_{1}^{p})$$

$$\leq 2\epsilon KHN + 4H\sqrt{\log(2KHN/\delta)} \sum_{k=1}^{K} \sum_{p=1}^{N} \sum_{h=1}^{H} \frac{1}{\sqrt{N_{k-1,h}(\gamma_{kh}^{p})}} \frac{1}{1 + N_{k-1,h}(\gamma_{kh}^{p})}$$

$$+ 4\sum_{h=1}^{H} \sum_{p=1}^{N} \sum_{k=2}^{K} \frac{\sqrt{\beta_{k-1}\log(2KHN/\delta)}}{\sqrt{(N_{k-2,h}(\gamma_{kh}^{p}) + 1)N_{k-1,h}(\gamma_{kh}^{p})}} \frac{1}{1 + N_{k-1,h}(\gamma_{kh}^{p})}$$

$$+ \sum_{h=1}^{H} \sum_{k=1}^{K} \sum_{p=1}^{N} [V_{h+1}^{\pi^{kp}}(s_{k,h+1}^{p}) - P_{h}V_{h+1}^{\pi^{kp}}(s_{k,h}^{p})].$$
(36)

Further note that

$$\sum_{k=1}^{K} \sum_{p=1}^{N} [P_h V_{h+1}^{\pi^{k_p}}(s_{k,h+1}^p) - V_{h+1}^{\pi^{k_p}}(s_{k,h+1}^p)]$$

=
$$\sum_{\gamma \in [\Gamma]} \sum_{k=1}^{K} \sum_{p=1}^{N} \mathbf{1} \{ \phi_h(s_{k,h}^p, a_{k,h}^p) = \gamma \} [P_h V_{h+1}^{\pi^{k_p}}(s_{k,h+1}^p) - V_{h+1}^{\pi^{k_p}}(s_{k,h+1}^p)]$$

and

$$\begin{split} &\sum_{k=1}^{K}\sum_{p=1}^{N}[P_{h}V_{h+1}^{*}(s_{k,h+1}^{p})-V_{h+1}^{*}(s_{k,h+1}^{p})] \\ &=\sum_{\gamma\in[\Gamma]}\sum_{k=1}^{K}\sum_{p=1}^{N}\mathbf{1}\{\phi_{h}(s_{k,h}^{p},a_{k,h}^{p})=\gamma\}[P_{h}V_{h+1}^{*}(s_{k,h+1}^{p})-V_{h+1}^{*}(s_{k,h+1}^{p})]. \end{split}$$

By Azuma-Höeffding's inequality, with probability $\geq 1-\delta/\Gamma,$

$$\sum_{k=1}^{K} \sum_{p=1}^{N} \mathbf{1}\{\phi_h(s_{k,h}^p, a_{k,h}^p) = \gamma\}[P_h V_{h+1}^{\pi^{k_p}}(s_{k,h+1}^p) - V_{h+1}^{\pi^{k_p}}(s_{k,h+1}^p)] \le 2H\sqrt{n_h^K(\gamma)}\sqrt{\log\frac{\Gamma}{\delta}},$$

and with probability $\geq 1-\delta/\Gamma,$ we have

$$\sum_{k=1}^{K} \sum_{p=1}^{N} \mathbf{1}\{\phi_h(s_{k,h}^p, a_{k,h}^p) = \gamma\}[P_h V_{h+1}^*(s_{k,h+1}^p) - V_{h+1}^*(s_{k,h+1}^p)] \le 2H\sqrt{n_h^K(\gamma)}\sqrt{\log\frac{\Gamma}{\delta}}$$

thus by summing over all the possible $\gamma^p_{kh} \in [\Gamma],$ with probability $1-2\delta,$

$$\begin{split} &\sum_{k=1}^{K} \sum_{p=1}^{N} [P_h V_{h+1}^{\pi^{kp}}(s_{k,h+1}^p) - V_{h+1}^{\pi^{kp}}(s_{k,h+1}^p)] + [P_h V_{h+1}^*(s_{k,h+1}^p) - V_{h+1}^*(s_{k,h+1}^p)] \\ &\leq 4H \sum_{\gamma \in [\Gamma]} \sqrt{n_h^K(\gamma)} \sqrt{\log \frac{\Gamma}{\delta}}. \end{split}$$

Finally, note that

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \sum_{h=1}^{H} \frac{1}{\sqrt{N_{k-1,h}(\gamma_{kh}^{p})}} \frac{1}{1+N_{k-1,h}(\gamma_{kh}^{p})}$$

$$\leq \sum_{h=1}^{H} \sum_{k=1}^{K} \sum_{\gamma \in [\Gamma]} \sum_{j=1}^{N_{k-1,h}(\gamma)} \frac{1}{j}$$

$$\leq_{(1)} \sum_{k=1}^{K} \sum_{h=1}^{H} \sum_{\gamma \in [\Gamma]} \sqrt{N_{k-1,h}(\gamma)} \sqrt{\sum_{j=1}^{\infty} 1/j^{2}} \leq_{(2)} \sum_{h=1}^{H} \sum_{k=1}^{K} \sum_{\gamma \in [\Gamma]} 2\sqrt{N_{k-1,h}(\gamma)}$$

$$\leq_{(3)} 2\sqrt{HK\Gamma} \sum_{h=1}^{H} \sum_{k=1}^{K} \sum_{\gamma \in [\Gamma]} N_{k-1,h}(\gamma)$$

$$= 2HK\sqrt{\Gamma N}$$
(37)

where (1) and (3) hold by Cauchy's inequality, and (2) holds because $\sqrt{\pi^2/6} < 2$.

Furthermore, we also have

$$\sum_{h=1}^{H} \sum_{\gamma \in [\Gamma]} \sqrt{n_h^K(\gamma)} \le \sqrt{H\Gamma \sum_{h=1}^{H} \sum_{\gamma \in [\Gamma]} n_h^K(\gamma)} = \sqrt{H^2 K \Gamma N}.$$
(38)

Additionally, note that whenever $N_{k-2,h}(\gamma_{kh}^p) \ge 1$, we have

$$\sum_{h=1}^{H} \sum_{p=1}^{N} \sum_{k=2}^{K} \frac{\sqrt{\beta_{k-1} \log(2KHN/\delta)}}{\sqrt{(N_{k-2,h}(\gamma_{kh}^{p})+1)N_{k-1,h}(\gamma_{kh}^{p})}} \frac{1}{1+N_{k-1,h}(\gamma_{kh}^{p})} \\ \leq \sum_{h=1}^{H} \sum_{p=1}^{N} \sum_{k=2}^{K} \frac{\sqrt{\beta_{k-1} \log(2KHN/\delta)}}{\sqrt{N_{k-2,h}(\gamma_{kh}^{p})+1}} \frac{1}{1+N_{k-1,h}(\gamma_{kh}^{p})},$$
(39)

and note that

$$\sum_{k=1}^{H} \sum_{p=1}^{N} \sum_{k=2}^{K} \frac{1}{\sqrt{N_{k-2,h}(\gamma_{k+1,h}^{p}) + 1}} \frac{1}{1 + N_{k-1,h}(\gamma_{k+1,h}^{p})} = (1) \sum_{k=1}^{K} \sum_{h=1}^{H} \sum_{\gamma \in [\Gamma]} \frac{N_{k,h}(\gamma)}{\sqrt{N_{k-1,h}(\gamma) + 1}} \frac{1}{1 + N_{k,h}(\gamma)} \leq \sum_{k=1}^{K} \sum_{h=1}^{H} \sum_{\gamma \in [\Gamma]} \frac{1}{\sqrt{N_{k-1,h}(\gamma)}} \leq \sqrt{KH\Gamma\sum_{k=1}^{K} \sum_{h=1}^{H} \sum_{\gamma \in [\Gamma]} N_{k,h}(\gamma)} = KH\Gamma\sqrt{N},$$

$$(40)$$

where equality (1) holds because $N_{k,h}(\gamma)$ is the number of agents that reach aggregated state γ at period h during episode k. And the last inequality holds by Cauchy's inequality. Recall that $\beta_k = \frac{1}{2}H^3 \log(2kH\Gamma)$, thus by (39),

$$\sum_{h=1}^{H} \sum_{p=1}^{N} \sum_{k=2}^{K} \frac{\sqrt{\beta_{k-1} \log(2KHN/\delta)}}{\sqrt{\left(N_{k-2}, h(\gamma_{kh}^{p})+1)N_{k-1,h}(\gamma_{kh}^{p}\right)}} \frac{1}{1+N_{k-1,h}(\gamma_{kh}^{p})} \leq KH^{5/2} \Gamma \sqrt{N} \sqrt{\log(2KH\Gamma)} \sqrt{\log(2KHN/\delta)}.$$
(41)

Thus by (36), when events $\mathcal{E}(\gamma_{kh}^p)$ and $\mathcal{G}(\gamma_{kh}^p)$ hold for all $\gamma_{kh}^p \in [\Gamma]$, with probability $1 - 2\delta$, we have

$$\begin{split} &\sum_{p=1}^{N}\sum_{k=1}^{K}\hat{V}_{k,1}^{p}(s_{1}^{p})-V_{\pi^{kp}}^{\eta}(s_{1}^{p})\\ &\leq 2\epsilon KHN+4KH^{2}\sqrt{\Gamma N}\sqrt{\log(2KHN/\delta)}+32H^{2}\sqrt{K\Gamma N}\sqrt{\log(HKN/\delta)}\\ &+2KH^{5/2}\Gamma\sqrt{N}\sqrt{\log(2KH\Gamma)}\sqrt{\log(2KHN/\delta)}. \end{split}$$

D. Proofs for the Infinite-Horizon Case

Additional notations For true MDP M and policy π we denote the discounted value function under π at state s as. $V_{\pi,1}^{\eta}(s) = V_{\pi}^{\eta}(s)$. We denote V_*^{η} as the discounted value function under the optimal policy π^* .

During each pseudo-episode $k \in [K]$, each agent samples a random vector with independent components $w^{kp} \in \mathbb{R}^{H_kSA}$, where $w^{kp}(h, s, a) \sim \mathcal{N}(0, \sigma_k^2(h, s, a))$ and $\sigma_k(h, s, a) = \sqrt{\frac{\beta_k}{N_{k-1}(\phi(s, a))+1}}$, where β_k is a tuning parameter, $N_{k-1,h}(\phi(s, a))$ is the total number of times that aggregated state $\phi(s, a)$ is reached across all agents during episode k-1. Given w^{kp} , we construct a randomized perturbation of the empirical MDP for agent p as

$$\overline{M}^{kp} = (T, \mathcal{S}, \mathcal{A}, \hat{P}^k, \hat{R}^k + w^{kp}, N),$$
(42)

where the empirical distributions \hat{P}^k and empirical rewards \hat{R}^k are computed as in (13) and (14). During each pseudo-episode $k \in [K]$, the data set \tilde{D}_k^p contains perturbation of samples from pseudo-episode k - 1 used by agent p.

We define the following events:

$$\mathcal{E}^{I}(\gamma) := \left\{ \left| \frac{1}{N_{k-1}(\gamma)} \sum_{j=1}^{N} \sum_{h \in [H_{k-1}]} \mathbf{1} \{ \phi(s_{k-1,h}^{j}, a_{k-1,h}^{j}) = \gamma \} \{ V_{*}^{\eta}(s_{k-1,h}^{j}) - PV_{*}^{\eta}(s_{k-1,h}^{j}) \} \right|$$

$$\leq \frac{2H_{k-1}\sqrt{\log(TN/\delta)}}{\sqrt{N_{k-1,h}(\gamma)}} \right\}.$$
(43)

$$\mathcal{G}^{I}(\gamma) := \left\{ \left| \frac{1}{N_{k-1}(\gamma)} \sum_{j=1}^{N} \sum_{h \in [H_{k-1}]} \mathbf{1} \{ \phi(s_{k-1,h}^{j}, a_{k-1,h}^{j}) = \gamma \} \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \right| \\
\leq \frac{2\sqrt{\beta_{k-1} \log(TN/\delta)}}{\sqrt{(N_{k-1}(\gamma)+1)N_{k-1}(\gamma)}} \right\}.$$
(44)

Lemma 12 (Lemma 2 of (Dong et al., 2022)). For all $\pi, s \in S$ and $\eta \in [0, 1)$, $\left| V_{\pi}^{\eta}(s) - \frac{\lambda_{\pi}(s)}{1-\eta} \right| \leq \tau_{\pi}$.

Remark 13. For weakly communicating M, the optimal average reward is state independent, so under Assumption 6, almost surely for any $s, s' \in S$, we have $|V_*^{\eta}(s) - V_*^{\eta}(s')| \le 2\tau_* \le 2\tau$.

Regret Decomposition Recall that $\operatorname{Regret}(M, T, N, \operatorname{RLSVI}_{\beta,\alpha,\xi})$ denotes the regret under infinite-horizon case for MDP M. Here $K = \arg \max\{k : t_k \leq T\}$. We put K explicitly here to derive the regret bound in a way dependent on the random K. Let H_k be the length of pseudo-episode k. So we can decompose the regret as

$$\operatorname{Regret}(M, T, N, \operatorname{RLSVI}_{\beta, \alpha, \xi}) = \mathbb{E}_{K, H_{k}} \left[\sum_{p=1}^{N} \sum_{k=1}^{K} \sum_{t=1}^{H_{k}} (\lambda_{*} - R_{t}^{p}) | M \right] \\ = \mathbb{E}_{K, H_{k}} \left[\sum_{p=1}^{N} \sum_{k=1}^{K} H_{k} \lambda_{*} - \sum_{p=1}^{N} \sum_{k=1}^{K} \sum_{t=1}^{H_{k}} R_{t}^{p} | M \right] \\ = \underbrace{\mathbb{E}_{K, H_{k}} \left[\sum_{p=1}^{N} \sum_{k=1}^{K} \{H_{k} \lambda_{*} - V_{*}^{\eta}(s_{k,1}^{p})\} | M \right]}_{(a)} + \underbrace{\sum_{p=1}^{N} \sum_{k=1}^{K} \{V_{*}^{\eta}(s_{k,1}^{p}) - V_{\pi^{kp}}^{\eta}(s_{k,1}^{p})\}}_{(b)} \right]$$

$$(45)$$

Recall from (11) that $V_*^{\eta}, V_{\pi^{kp}}^{\eta} \leq \frac{1}{1-\eta}$. Note that (a) in (45) is the difference between the optimal average reward weighted by the pseudo-horizons and the discounted reward, and part (b) is the sum of the differences between the optimal discounted value and cumulative the discounted value of the employed policies throughout K pseudo-episodes by all agents. We provide an upper bound for the worst-case regret by bounding (a) and (b) respectively.

D.1. Proof for Infinite-Horizon Main Result: Theorem 7

Proof of Theorem 7. By Azuma-Höeffding inequality, conditional on the trajectory during pseudo-episode k - 1, with probability $1 - \delta/(2TN)$, we have

$$\left|\frac{1}{N_{k-1}(\gamma)}\sum_{j=1}^{N_{k-1}(\gamma)} \{V_*^{\eta}(s_{k-1,h+1}^j) - P_h V_*^{\eta}(s_{k-1,h}^j)\}\right| \le \frac{2\sqrt{\log(2TN/\delta)}}{(1-\eta)\sqrt{N_{k-1}(\gamma)}}.$$

Additionally, recall from Algorithm 4 that $\tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \sim \mathcal{N}\left(0, \frac{\beta_{k-1}}{N_{k-1}(\gamma)+1}\right)$. Note that conditional on the trajectories during pseudo-episode k-1, the random perturbations are i.i.d. across $j \in [p], h \in [H_{k-1}]$, thus by Höefdding's inequality, conditional on the trajectory during pseudo-period k-1, with probability $1-\delta/(2TN)$, we have

$$\left|\frac{1}{N_{k-1}(\gamma)}\sum_{j=1}^{N}\sum_{h\in[H_{k-1}]}\mathbf{1}\{\phi_h(s_{k-1,h}^p, a_{k-1,h}^p) = \gamma\}\tilde{Q}_{k-1,h}^j(s_{k-1,h}^p, a_{k-1,h}^p)\right| \le 2\sqrt{\frac{\beta_{k-1}}{N_{k-1}(\gamma)+1}}\sqrt{\frac{\log(2TN/\delta)}{N_{k-1}(\gamma)}}$$

So we have

$$\mathbb{P}(\mathcal{E}(\gamma_{kh}^{p}), \mathcal{G}(\gamma_{kh}^{p}), \gamma_{kh}^{p}, \forall k \in [K], h \in [H_{k}], p \in [N]) = \mathbb{P}(\mathcal{E}(\gamma_{t}^{p}), \mathcal{G}(\gamma_{t}^{p}), \gamma_{t}^{p}, \forall t \in [T], p \in [N])$$

$$\geq 1 - \sum_{t \in [T], p \in [N]} (\mathbb{P}(\mathcal{E}(\gamma_{t}^{p})^{c}) + \mathbb{P}(\mathcal{G}(\gamma_{t}^{p})^{c})) \geq 1 - 2NT\delta/(2NT) \geq 1 - \delta.$$
(46)

By Lemma 15 and Lemma 14, with probability $1 - 3\delta$, we have

$$\operatorname{Regret}(M, T, N, \operatorname{RLSVI}_{\beta, \alpha, \xi}) \leq \frac{\tau N[(1-\eta)T+1]}{\sqrt{NT}} + [(1-\eta)T+1]\tau \left(1+\sqrt{N\log(NT)}\right) \\ + \frac{8T\sqrt{\Gamma N\log(2TN/\delta)}}{(1-\eta)^2} + \frac{8\Gamma\tau^{3/2}T\sqrt{N\log(2\tau\Gamma T)\log(2TN/\delta)}}{(1-\eta)^2} + \frac{4K\sqrt{N\Gamma}\sqrt{\log(\Gamma/\delta)}}{1-\eta} + 2\eta\epsilon TN.$$

$$(47)$$

Recall that $\tau = \max{\{\tau, \tau\}}$, so with probability $1 - \delta$, we have

$$\begin{aligned} \operatorname{Regret}(M,T,N,\operatorname{RLSVI}_{\beta,\alpha,\xi}) &\leq \frac{\tau N [(1-\eta)T+1]}{\sqrt{NT}} + [(1-\eta)T+1]\tau \left(1+\sqrt{N\log(NT)}\right) \\ &+ \frac{8T\sqrt{\Gamma N \log(6TN/\delta)}}{(1-\eta)^2} + \frac{8\tau^{3/2}T\Gamma\sqrt{N\log(2\tau\Gamma T)\log(6TN/\delta)}}{(1-\eta)^2} \\ &+ \frac{4K\sqrt{N\Gamma}\sqrt{\log(3\Gamma/\delta)}}{1-\eta} + 2\eta\epsilon TN. \end{aligned}$$

When $1 - \eta$ such that $\frac{1}{(1-\eta)^2} \leq C$ for some constant C, is bounded from below, we have

$$\begin{array}{ll} \operatorname{Regret}(M,T,N,\operatorname{RLSVI}_{\beta,\alpha,\xi}) &\leq 2\epsilon TN + 2\tau\sqrt{NT} + 4T\tau\sqrt{N\log(NT)} \\ &\quad +16C\max\{\tau^{3/2},1\}T\Gamma\sqrt{\Gamma N\log(2\tau T)\log(6TN/\delta)} \end{array}$$

D.2. Lemmas for bounding (a) and (b) in (45)

Lemma 14 (Bound for (a) of (45)). $\mathbb{E}_{H_k,K}\left[\sum_{p=1}^N \sum_{k=1}^K H_k \lambda_* - V_*^{\eta}(s_{k,1}^p)\right] \leq \frac{\tau N[(1-\eta)T+1]}{\sqrt{NT}} + [(1-\eta)T + 1]\tau \left(1 + \sqrt{N\log(NT)}\right).$

Proof of Lemma 14. Note that the expected length of each pseudo-episode is independent of the policy and is equal to $\frac{1}{1-\eta}$. Thus for K fixed, $\mathbb{E}_{H_k} \left[\sum_{p=1}^N \sum_{k=1}^K H_k \lambda_* - V_*^{\eta}(s_{k,1}^p) \right] = \sum_{p=1}^N \sum_{k=1}^K \left(\frac{\lambda_*}{1-\eta} - V_*^{\eta}(s_{k,1}^p) \right)$, thus $\mathbb{E} \left[\sum_{p=1}^N \sum_{k=1}^K H_k \lambda_* - V_*^{\eta}(s_{k,1}^p) \right] \leq \mathbb{E} \left[\sum_{p=1}^N \sum_{k=1}^K \left| \frac{\lambda_*}{1-\eta} - V_*^{\eta}(s_{k,1}^p) \right| \right].$

By Lemma 12, we know that $\left|\frac{\lambda_*}{1-\eta} - V^{\eta}_*(s^p_{k,1})\right| \leq \tau$. So that for any $p \in [N]$, for any fixed K,

$$\sum_{k=1}^{K} \left| \frac{\lambda_*}{1-\eta} - V_*^{\eta}(s_{k,1}^p) \right| \le K\tau.$$

Note that $s_{k,1}^p$ are sampled i.i.d. across p at the beginning of each pseudo-episode k. Höeffding's inequality shows that for any $\epsilon > 0$,

$$\mathbb{P}\left(\sum_{p=1}^{N}\sum_{k=1}^{K}\left|\frac{\lambda_{*}}{1-\eta} - V_{*}^{\eta}(s_{k,1}^{p})\right| \ge \epsilon + K\tau\right) \le \exp\left(-\frac{2\epsilon^{2}}{4NK^{2}\tau^{2}}\right).$$

Take

$$\epsilon = K\tau \sqrt{N\log(NT)},$$

we then have

$$\mathbb{P}\left(\sum_{p=1}^{N}\sum_{k=1}^{K}\left|\frac{\lambda_{*}}{1-\eta}-V_{*}^{\eta}(s_{k,1}^{p})\right| \geq K\tau\left(1+\sqrt{N\log(NT)}\right)\right) \leq \frac{1}{\sqrt{NT}}.$$

Thus conditioning on the total number of pseudo-episodes K,

$$\mathbb{E}_{K}\left[\sum_{p=1}^{N}\sum_{k=1}^{K}\left|\frac{\lambda_{*}}{1-\eta}-V_{*}^{\eta}(s_{k,1}^{p})\right|\right] \leq \frac{\tau KN}{\sqrt{NT}}+K\tau\left(1+\sqrt{N\log(NT)}\right).$$

Further note that since $H_k \sim \text{Geometric}(1 - \eta)$ and are i.i.d. across k, so

$$\mathbb{E}[K] \le (1 - \eta)T + 1.$$

Hence by taking expectation over K, we have

$$\mathbb{E}\left[\sum_{p=1}^{N}\sum_{k=1}^{K}\left|\frac{\lambda_{*}}{1-\eta}-V_{*}^{\eta}(s_{k,1}^{p})\right|\right] \leq \frac{\tau\mathbb{E}[K]N}{\sqrt{NT}} + \mathbb{E}[K]\tau\left(1+\sqrt{N\log(NT)}\right)$$
$$\leq \frac{\tau N[(1-\eta)T+1]}{\sqrt{NT}} + [(1-\eta)T+1]\tau\left(1+\sqrt{N\log(NT)}\right).$$

Lemma 15 (Bound for (b) of (45)). Suppose events $\mathcal{E}^{I}(\gamma)$ and $\mathcal{G}^{I}(\gamma)$ hold for any $\gamma \in [\Gamma]$, then with probability $1 - 2\delta$, we have

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \left\{ V_*^{\eta}(s_{k,1}^p) - V_{\pi^{kp},M}^{\eta}(s_{k,1}^p) \right\} \leq \frac{8T\sqrt{\Gamma N \log(2TN/\delta)}}{(1-\eta)^2} + \frac{8\tau^{3/2}T\Gamma\sqrt{N \log(2\tau\Gamma T) \log(2TN/\delta)}}{1-\eta} + \frac{4T\sqrt{N\Gamma}\sqrt{\log(\Gamma/\delta)}}{1-\eta} + 2\eta\epsilon TN.$$

Proof of Lemma 15. Recall from Algorithm 4 that the unclipped value function estimates $\bar{Q}_{k,h}^{p}(\cdot)$ during pseudo-episode k at time period $h \in [H_k]$ (recall that H_k here is random) as

$$\begin{split} \bar{Q}_{k,h}^{p}(\gamma) &= \arg\min_{Q \in \mathbb{R}} \quad \sum_{(s,a):\phi_{h}(s,a)=\gamma} (Q - \xi_{N_{k-1}(\gamma)} - (1 - \alpha_{N_{k-1}(\gamma)})\hat{Q}_{k-1}(\gamma) \\ &- \eta \alpha_{N_{k-1,h}(\gamma)} \{r(s,a) + \max_{a' \in \mathcal{A}} \hat{Q}_{k,h+1}^{p}(s',a')\})^{2} + \left\| Q - \eta \alpha_{N_{k-1,h}(\gamma)} \tilde{Q}_{k,h}^{p} \right\|_{2}^{2} \end{split}$$

So similar to the derivation in the proof of Lemma 10, we have

$$\bar{Q}_{k,h}^{p}(\gamma) = \xi_{N_{k-1}(\gamma)} + \frac{1}{N_{k-1}(\gamma)} \sum_{p=1}^{N_{k-1}(\gamma)} (1 - \alpha_{N_{k-1}(\gamma)}) \hat{Q}_{k-1}^{p}(\gamma) + \alpha_{N_{k-1}(\gamma)} \eta(r(s_{k-1}^{p}, a_{k-1}^{p}) + \hat{V}_{k}^{p}(s_{k-1,h+1}^{p}) + \tilde{Q}_{k,h}^{p}(s_{k-1,h}^{p}, a_{k-1,h}^{p}))$$

By definition, we have $\bar{Q}_{k,1}^p(\gamma) = \bar{Q}_{t_k}^p(\gamma)$. We denote $\bar{Q}_k^p(\gamma) := \bar{Q}_{k,1}^p(\gamma)$ in the following. Thus for any $\gamma = \phi(s_{k,h}^p, a_{k,h}^p)$ during pseudo-episode k, with $\phi(s_{k-1,h}^j, a_{k-1,h}^j) = \gamma$ in the following, where $\{(s_{k-1,h}^j, a_{k-1,h}^j)\}$ come from all the state-action pairs during pseudo-period k-1 (note that we don't distinguish between different pseudo periods now), where $h \in \{1, 2, \dots, H_{k-1}\}$, and H_{k-1} is the random length of pseudo-episode k-1, and recall that $N_{k-1}(\gamma) = \sum_{p=1}^{N} \sum_{h \in [H_{k-1}]} \mathbf{1}\{\phi_h(s_{k-1,h}^p, a_{k-1,h}^p) = \gamma\}$, so we have

$$\begin{split} & \bar{Q}_{k,h}^{p}(s_{k,h}^{p}, a_{k,h}^{p}) - Q_{*}^{q}(s_{k,h}^{p}, a_{k,h}^{p}) \\ &= \eta \xi_{N_{k-1}(\gamma)} + \frac{1}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} (1 - \alpha_{N_{k-1}(\gamma)}) (Q_{*}^{q}(s_{k,h}^{p}, a_{k,h}^{p}) - \hat{Q}_{k-1}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})) \\ &+ \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{p=1}^{N_{k-1}(\gamma)} \{-[r_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) + \hat{V}_{k-1}^{j}(s_{k-1,h+1}^{j}) + \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})] \\ &+ Q_{*}^{\eta}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\} \\ &+ \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} \underbrace{\{Q_{*}^{\eta}(s_{k,h}^{p}, a_{k,h}^{p}) - Q_{*}^{\eta}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\}\}}_{\leq \epsilon} \\ &\leq \epsilon + \eta \xi_{N_{k-1}(\gamma)} + \frac{\eta}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} (1 - \alpha_{N_{k-1}(\gamma)}) (Q_{*}^{\eta}(s_{k,h}^{p}, a_{k,h}^{p}) - \hat{Q}_{k-1}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})) \\ &+ \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{p=1}^{N_{k-1}(\gamma)} \{-[r_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) + \hat{V}_{k-1}^{j}(s_{k-1,h+1}^{j}) + \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})] \\ &= \epsilon + \eta \xi_{N_{k-1}(\gamma)} - \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \\ &- \eta \frac{\eta}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} \{V_{k-1}^{j}(s_{k-1,h+1}^{j}) - V_{*}^{\eta}(s_{k-1,h+1}^{j})\} \\ &- \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} \{V_{*}^{\eta}(s_{k-1,h+1}^{j}) - P_{h}V_{*}^{\eta}(s_{k-1,h+1}^{j})\}, \end{split}$$

where we used the fact that under optimal policy π^* for the true MDP, we have

$$Q_*(s',a') = r(s',a') + \eta P_* V_*(s'),$$

and under the optimal policy $\pi^{k-1,j}$ under MDP $\overline{M}^{k-1,j}$, we have

$$\hat{Q}_{k-1}^{j}(s',a') = r(s',a') + \eta \hat{P}_{\pi^{kj}} \hat{V}_{k-1}(s)$$

Now suppose that events $\mathcal{E}^{I}(\gamma_{kh}^{p})$ and $\mathcal{G}^{I}(\gamma_{kh}^{p})$ hold for all aggregated states γ_{kh}^{p} during pseudo-perido k for all $k \in [K]$. Then by similar derivation as that for the finite-horion case in Lemma 9, we have

$$V_{k,h}^p(s) - V_*(s) \ge 0, \ \forall s \in \mathcal{S}, k \in [K], p \in [N].$$

$$\tag{49}$$

We denote

$$\Delta^{p}_{k,h} = \hat{V}^{p}_{k,h}(s^{p}_{k,h}) - V^{\eta}_{*}(s^{p}_{k,h}).$$

Note that

$$\hat{V}_{k,h}(s_{k,h}^p) - V_*(s_{k,h}^p) \le \hat{Q}_{k,h}^p(s_{k,h}^p, a_{k,h}^p) - Q_*^\eta(s_{k,h}^p, a_{k,h}^p) \le \bar{Q}_{k,h}^p(s_{k,h}^p, a_{k,h}^p) - Q_*^\eta(s_{k,h}^p, a_{k,h}^p).$$

By (48) we have

$$\begin{split} & \bar{Q}_{k,h}^{p}(s_{k,h}^{p}, a_{k,h}^{p}) - Q_{*}^{q}(s_{k,h}^{p}, a_{k,h}^{p}) \\ &= \eta \xi_{N_{k-1}(\gamma)} + \frac{\eta}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} (1 - \alpha_{N_{k-1}(\gamma)}) (Q_{*}^{\eta}(s_{k,h}^{p}, a_{k,h}^{p}) - \hat{Q}_{k-1}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})) \\ &+ \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{p=1}^{N_{k-1}(\gamma)} \{-[r_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) + \hat{V}_{k-1}^{j}(s_{k-1,h+1}^{j}) + \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})] \\ &+ Q_{*}^{\eta}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\} \\ &+ \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} \underbrace{\{Q_{*}^{\eta}(s_{k,h}^{p}, a_{k-1,h}^{p}) - Q_{*}^{\eta}(s_{k-1,h}^{j}, a_{k-1,h}^{j})\}}_{\leq \epsilon} \\ &\leq \eta \epsilon + \xi_{N_{k-1}(\gamma)} + \frac{1}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} (1 - \alpha_{N_{k-1}(\gamma)}) (Q_{*}^{\eta}(s_{k,h}^{p}, a_{k,h}^{p}) - \hat{Q}_{k-1}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})) \\ &+ \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{p=1}^{N_{k-1}(\gamma)} \{-[r_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) + \hat{V}_{k-1}^{j}(s_{k-1,h+1}^{j}) + \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j})) \\ &= \eta \xi_{N_{k-1}(\gamma)} + \eta \epsilon - \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} \tilde{Q}_{k-1,h}^{j}(s_{k-1,h}^{j}, a_{k-1,h}^{j}) \\ &- \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} \{V_{*}^{\eta}(s_{k-1,h+1}^{j}) - V_{*}^{\eta}(s_{k-1,h+1}^{j})\} \\ &- \eta \frac{\alpha_{N_{k-1}(\gamma)}}{N_{k-1}(\gamma)} \sum_{j=1}^{N_{k-1}(\gamma)} \{V_{*}^{\eta}(s_{k-1,h+1}^{j}) - P_{h}V_{*}^{\eta}(s_{k-1,h}^{j})\}, \end{split}$$

where the inequality follows by definition of ϵ -error aggregated states as in Definition 4.

Then by similar derivation as for (32), by summing from the *i*-th pseudo-episode to K-th pseudo-episode, we have

$$\begin{split} \sum_{p=1}^{N} \sum_{k=i}^{K} \Delta_{k,h}^{p} &\leq 2N\eta \sum_{k=i}^{K} \epsilon + \frac{4\eta}{1-\eta} \sum_{p=1}^{N} \sum_{k=i}^{K} \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})} \sqrt{\log(2TN/\delta)}}{\sqrt{N_{k-1,h}(\gamma_{kh}^{p})}} \\ &+ \eta \sum_{p=1}^{N} \sum_{k=i+1}^{K} \frac{\alpha_{N_{k-1,h}(\gamma_{kh}^{p})}}{N_{k-1,h}(\gamma_{kh}^{p})} \sum_{j=1}^{N_{k-1,h}(\gamma_{kh}^{p})} \Delta_{k,h+1}^{j}. \end{split}$$

Then by similar derivation as in (33) under events $\mathcal{E}^{I}(\gamma_{kh}^{p})$ and $\mathcal{G}^{I}(\gamma_{kh}^{p})$ for all γ_{kh}^{p} across K pseudo-episodes and p agents, we have

$$\begin{split} \sum_{p=1}^{N} \sum_{k=i}^{K} \Delta_{k,h}^{p} &\leq 2\eta \epsilon N \sum_{k=i}^{K} H_{k} + \frac{4\eta}{1-\eta} \sqrt{\log(2TN/\delta)} \sum_{k=i}^{K} \sum_{p=1}^{N} \sum_{\ell=h}^{H} \frac{1}{\sqrt{N_{k-1,\ell}(\gamma_{k,\ell}^{p})}} \frac{1}{1+N_{k-1,\ell}(\gamma_{k,\ell}^{p})} \\ &+ 4\eta \sum_{\ell=h}^{H} \sum_{p=1}^{N} \sum_{k=i+1}^{K} \frac{\sqrt{\beta_{k-1}\log(2TN/\delta)}}{\sqrt{(N_{k-2,h}(\gamma_{k,h}^{p})+1)N_{k-1,h}(\gamma_{k,h}^{p})}} \frac{1}{1+N_{k-1,\ell}(\gamma_{k,\ell}^{p})}. \end{split}$$

Then following similar steps as in the proof of Lemma 11, we have $\Delta_{k,h}^p = \hat{V}_{k,h}^p(s_{k,h}^p) - V_*^{\eta}(s_{k,h}^p)$.

$$\hat{V}_{k,h}^{p}(s_{k,h}^{p}) - V_{\pi^{k_{p}}}^{\eta}(s_{k,h}^{p}) = \left[\hat{Q}_{h}^{\pi^{k_{p}}}(s_{k,h}^{p}, a_{k,h}^{p}) - Q_{*}^{\eta}(s_{k,h}^{p}, a_{k,h}^{p})\right] + \eta \left[\hat{V}_{k,h+1}^{\pi^{k_{p}}}(s_{k,h+1}^{p}) - V_{\pi^{k_{p}}}^{\eta}(s_{k,h+1}^{p})\right] - \eta \Delta_{k,h+1}^{p} \\ + \eta \left[V_{\pi^{k_{p}}}^{\eta}(s_{k,h+1}^{p}) - P_{h}V_{\pi^{k_{p}}}^{\eta}(s_{k,h+1}^{p})\right].$$
(51)

Note that (51) is recursive, so when events $\mathcal{E}^{I}(\gamma_{kh}^{p})$ and $\mathcal{G}^{I}(\gamma_{kh}^{p})$ hold for all aggregated states γ_{kh}^{p} during pseudo-perido k for all $k \in [K]$, we have

$$\sum_{k=1}^{K} \sum_{p=1}^{N} \hat{V}_{k,1}^{p}(s_{1}^{p}) - V_{\pi^{kp}}^{\eta}(s_{1}^{p})$$

$$\leq 2\eta\epsilon TN + \frac{4}{1-\eta}\sqrt{\log(2TN/\delta)} \sum_{k=1}^{K} \sum_{p=1}^{N} \sum_{h=1}^{M+k} \eta^{h-1} \frac{1}{\sqrt{N_{k-1}(\gamma_{kh}^{p})}} \frac{1}{1+N_{k-1}(\gamma_{kh}^{p})}$$

$$+4\sum_{k=2}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sum_{p=1}^{N} \frac{\sqrt{\beta_{k-1}\log(2TN/\delta)}}{\sqrt{(N_{k-2}(\gamma_{kh}^{p})+1)N_{k-1}(\gamma_{kh}^{p})}} \frac{1}{1+N_{k-1}(\gamma_{kh}^{p})}$$

$$+\sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sum_{p=1}^{N} [P_{h}V_{*}^{\eta}(s_{k,h+1}^{p}) - V_{*}^{\eta}(s_{k,h+1}^{p})]]$$

$$+\sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sum_{p=1}^{N} [V_{\pi^{kp}}^{\eta}(s_{k,h+1}^{p}) - P_{\pi^{kp}}^{h}V_{\pi^{kp}}^{\eta}(s_{k,h+1}^{p})]].$$
(52)

Note that

$$\frac{4}{1-\eta}\sqrt{\log(2TN/\delta)} \sum_{k=1}^{K} \sum_{p=1}^{N} \sum_{h=1}^{H_{k}} \eta^{h-1} \frac{1}{\sqrt{N_{k-1}(\gamma_{kh}^{p})}} \frac{1}{1+N_{k-1}(\gamma_{kh}^{p})} \\
\leq \frac{4}{1-\eta}\sqrt{\log(2TN/\delta)} \sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sum_{p=1}^{N} \frac{1}{\sqrt{N_{k-1}(\gamma_{kh}^{p})}} \frac{1}{1+N_{k-1}(\gamma_{kh}^{p})} \\
\leq \frac{4}{1-\eta}\sqrt{\log(2TN/\delta)} \sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sum_{\gamma \in [\Gamma]} \sum_{j=1}^{N_{k-1}(\gamma)} 1/j \\
\leq \frac{4}{1-\eta}\sqrt{\log(2TN/\delta)} \sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sum_{\gamma \in [\Gamma]} 2\sqrt{N_{k-1}(\gamma)} \\
\leq \frac{8}{(1-\eta)}\sqrt{\log(2TN/\delta)} \sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sqrt{\Gamma} \sum_{\gamma \in [\Gamma]} N_{k-1}(\gamma) \\
\leq \frac{8}{(1-\eta)}\sqrt{\log(2TN/\delta)} \sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sqrt{\Gamma NH_{k-1}} \leq \frac{8\sqrt{\Gamma N\log(2TN/\delta)}}{(1-\eta)} \sum_{h=1}^{\infty} \eta^{h-1} \sum_{k=1}^{K} \sqrt{H_{k-1}} \\
\leq \frac{8\sqrt{\Gamma N\log(2TN/\delta)}}{(1-\eta)^{2}} \sum_{k=1}^{K} H_{k-1} \leq \frac{8T\sqrt{\Gamma N\log(2TN/\delta)}}{(1-\eta)^{2}}.$$
(53)

Next, note that $\beta_k = \frac{1}{2}\tau^3 \log(2\tau\Gamma k)$, following similar steps as (53) and (40) we have

$$4\sum_{k=1}^{K}\sum_{h=1}^{H_{k}}\eta^{h-1}\sum_{p=1}^{N}\frac{\sqrt{\beta_{k-1}\log(2TN/\delta)}}{\sqrt{(N_{k-2}(\gamma_{kh}^{p})+1)N_{k-1}(\gamma_{kh}^{p})}}\frac{1}{1+N_{k-1}(\gamma_{kh}^{p})} \leq \frac{8\tau^{3/2}T\Gamma\sqrt{N\log(2\tau\Gamma T)\log(2TN/\delta)}}{(1-\eta)^{2}}.$$
 (54)

Additionally, denote $N_{k,h}(\gamma)$ as the number of times that aggregated state γ is attained during pseudo-episode k and pseudo

period h, then by Azuma-Höeffding's inequality, with probability $1 - 2\delta$, we have

$$\sum_{p=1}^{N} \{ [V_{\pi^{kp}}^{\eta}(s_{k,h+1}^{p}) - P_{\pi^{kp}}^{h}V_{\pi^{kp}}^{\eta}(s_{k,h+1}^{p})] + [P_{h}V_{*}^{\eta}(s_{k,h+1}^{p}) - V_{*}^{\eta}(s_{k,h+1}^{p})] \} \leq \frac{4}{1-\eta} \sum_{\gamma \in [\Gamma]} \sqrt{N_{k,h}(\gamma)} \sqrt{\log \frac{\Gamma}{\delta}}.$$

Hence with probability $1 - 2\delta$,

$$\sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sum_{p=1}^{N} \{ [P_{h} V_{*}^{\eta}(s_{k,h+1}^{p}) - V_{*}^{\eta}(s_{k,h+1}^{p})] + [V_{\pi^{k_{p}}}^{\eta}(s_{k,h+1}^{p}) - P_{\pi^{k_{p}}}^{h} V_{\pi^{k_{p}}}^{\eta}(s_{k,h+1}^{p})] \}$$

$$\leq \frac{4\sqrt{\log(\Gamma/\delta)}}{1-\eta} \sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1} \sum_{\gamma \in [\Gamma]} \sqrt{N_{k}(\gamma)} \leq \frac{4\sqrt{\Gamma N \log(\Gamma/\delta)}}{1-\eta} \sum_{k=1}^{K} \sum_{h=1}^{H_{k}} \eta^{h-1}$$

$$\leq \frac{4K\sqrt{\Gamma N \log(\Gamma/\delta)}}{(1-\eta)^{2}}$$
(55)

Thus by (52), (53), (54), (55), when events $\mathcal{E}^{I}(\gamma_{kh}^{p})$ and $\mathcal{G}^{I}(\gamma_{kh}^{p})$ hold for all aggregated states γ_{kh}^{p} during pseudo-perido k for all $k \in [K]$, with probability $1 - 2\delta$, we have

$$\sum_{k=1}^{K} \sum_{p=1}^{N} \hat{V}_{k,1}^{p}(s_{1}^{p}) - V_{\pi^{kp}}^{\eta}(s_{1}^{p}) \leq \frac{8T\sqrt{\Gamma N \log(2TN/\delta)}}{\frac{(1-\eta)^{2}}{(1-\eta)^{2}}} + \frac{8\tau^{3/2}T\Gamma\sqrt{N \log(2\tau\Gamma T) \log(2TN/\delta)}}{(1-\eta)^{2}}.$$
(56)

Finall, by (49), when events $\mathcal{E}^{I}(\gamma_{kh}^{p})$ and $\mathcal{G}^{I}(\gamma_{kh}^{p})$ hold for all aggregated states γ_{kh}^{p} during pseudo-perido k for all $k \in [K]$,

$$\hat{V}_{k,h}^p(s) - V_*(s) \ge 0, \ \forall s \in \mathcal{S}, k \in [K], p \in [N],$$

And with probability $1 - 2\delta$, we have

$$\sum_{p=1}^{N} \sum_{k=1}^{K} \left\{ V_*^{\eta}(s_{k,1}^p) - V_{\pi^{kp},M}^{\eta}(s_{k,1}^p) \right\} \leq \frac{8T\sqrt{\Gamma N \log(2TN/\delta)}}{\frac{(1-\eta)^2}{1-\eta}} + \frac{8\tau^{3/2}T\Gamma\sqrt{N \log(2\tau\Gamma T) \log(2TN/\delta)}}{1-\eta}.$$