
Quantum Speedups in Regret Analysis of Infinite Horizon Average-Reward Markov Decision Processes

Bhargav Ganguly^{*1} Yang Xu^{*1} Vaneet Aggarwal¹

Abstract

This paper investigates the potential of quantum acceleration in addressing infinite horizon Markov Decision Processes (MDPs) to enhance average reward outcomes. We introduce an innovative quantum framework for the agent’s engagement with an unknown MDP, extending the conventional interaction paradigm. Our approach involves the design of an optimism-driven tabular Reinforcement Learning algorithm that harnesses quantum signals acquired by the agent through efficient quantum mean estimation techniques. Through thorough theoretical analysis, we demonstrate that the quantum advantage in mean estimation leads to exponential advancements in regret guarantees for infinite horizon Reinforcement Learning. Specifically, the proposed Quantum algorithm achieves a regret bound of $\tilde{O}(1)$ ¹, a significant improvement over the $\tilde{O}(\sqrt{T})$ bound exhibited by classical counterparts.

1. Introduction

Quantum Machine Learning (QML) has garnered remarkable interest in contemporary research, predominantly attributed to the pronounced speedups achievable with quantum computers as opposed to their classical analogs (Biamonte et al., 2017; Bouland et al., 2023). The fundamental edge of quantum computing arises from the unique nature of its fundamental computing element, termed a qubit, which can exist simultaneously in states of 0 and 1, unlike classical bits that are restricted to either 0 or 1. This inherent distinction underpins the exponential advancements that quantum computers bring to specific computational tasks, surpassing

the capabilities of classical computers.

In the realm of Reinforcement Learning (RL), an agent embarks on the task of finding an efficient policy for Markov Decision Process (MDP) environment through repetitive interactions (Sutton & Barto, 2018). RL has found notable success in diverse applications, including but not limited to autonomous driving, ridesharing platforms, online ad recommendation systems, and proficient gameplay agents (Silver et al., 2017; Al-Abbasi et al., 2019; Chen et al., 2021; Bonjour et al., 2022). Most of these setups require decision making over infinite horizon, with an objective of average rewards. This setup has been studied in classical reinforcement learning (Auer et al., 2008; Agarwal et al., 2022b; Fruit et al., 2018), where $\tilde{O}(\sqrt{T})$ regret across T rounds has been found meeting the theoretical lower bound (Jaksch et al., 2010). This paper embarks on an inquiry into the potential utilization of quantum statistical estimation techniques to augment the theoretical convergence speeds of tabular RL algorithms within the context of infinite horizon learning settings.

A diverse array of quantum statistical estimation primitives has emerged, showcasing significant enhancements in convergence speeds and gaining traction within Quantum Machine Learning (QML) frameworks (Brassard et al., 2002; Harrow et al., 2009; Gilyén et al., 2019). Notably, (Hamoudi, 2021) introduces a quantum mean estimation algorithm that yields a quadratic acceleration in sample complexity when contrasted with classical mean estimation techniques. In this work, we emphasize the adoption of this particular mean estimation method as a pivotal component of our methodology. It strategically refines the signals garnered by the RL agent during its interaction with the enigmatic quantum-classical hybrid MDP environment.

It is pertinent to highlight a fundamental element in the analysis of conventional Reinforcement Learning (RL), which involves the utilization of martingale convergence theorems. These theorems play a crucial role in delineating the inherent stochastic process governing the evolution of states within the MDP. Conversely, this essential aspect lacks parallelism within the quantum setting, where comparable martingale convergence results remain absent. To address this disparity, we introduce an innovative approach to regret bound

^{*}Equal contribution ¹Purdue University, West Lafayette, IN, USA. Correspondence to: Bhargav Ganguly <bganguly@purdue.edu>, Yang Xu <xu1720@purdue.edu>, Vaneet Aggarwal <vaneet@purdue.edu>.

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

¹ $\tilde{O}(\cdot)$ conceals logarithmic terms of T .

analysis for quantum RL. Remarkably, our methodology circumvents the reliance on martingale concentration bounds, a staple in classical RL analysis, thus navigating uncharted territories in the quantum realm.

Moreover, it is important to highlight that mean estimation in quantum mechanics results in state collapse, which makes it challenging to perform estimation of state transition probabilities over multiple epochs. In this study, we present a novel approach to estimating state transition probabilities that explicitly considers the quantum state collapsing upon measurement.

To this end, the following are the major contributions of this work:

1. We introduce Quantum-UCRL (Q-UCRL), a model-based, infinite horizon, optimism-driven quantum Reinforcement Learning (QRL) algorithm. Drawing inspiration from classical predecessors like the UCRL and UC-CURL algorithms (Auer et al., 2008; Agarwal et al., 2022b; Agarwal & Aggarwal, 2023), Q-UCRL is designed to integrate an agent’s optimistic policy acquisition and an adept quantum mean estimator, and is the first quantum algorithm for infinite horizon average reward RL with provable guarantees.
2. Employing meticulous theoretical analysis, we show that the Q-UCRL algorithm achieves an exponential improvement in the regret analysis. Specifically, it attains a regret bound of $\tilde{O}(1)$, breaking through the theoretical classical lower bound of $\Omega(\sqrt{T})$ across T online rounds. The analysis is based on the novel quantum Bellman error based analysis (introduced in classical RL in (Agarwal et al., 2022b) for optimism based algorithm), where the difference between the performance of a policy on two different MDPs is bounded by long-term averaged Bellman error and the quantum mean estimation is used for improved Bellman error. Further, the analysis avoids dependence on martingale concentration bounds and incorporates the phenomenon of state collapse following measurements, a crucial aspect in quantum computing analysis.
3. We design a momentum-based estimator in equations (21)-(25) that fuses each epoch’s fresh quantum mean estimate with all past estimates using counts-based weights. The resulting algorithm preserves information that would otherwise be lost to quantum collapse, enabling accuracy to accumulate epoch-by-epoch and marking the first reuse-of-information technique fully compatible with quantum-measurement constraints. This method is essential for integrating quantum speedups into model-based RL frameworks and represents a core technical novelty of our work.

To the best of our knowledge, these are the first results for quantum speedups for infinite horizon MDPs with average reward objective.

2. Related Work and Preliminaries

Infinite Horizon Reinforcement Learning: Regret analysis of infinite horizon RL in the classical setting has been widely studied with average reward objective in both model-based settings and model-free settings (Wei et al., 2020; Bai et al., 2024; Hong et al., 2025; Ganesh et al., 2025a;b). In model-based methods, a prominent principle that underpins algorithms tailored for this scenario is the concept of “optimism in the face of uncertainty” (OFU). In this approach, the RL agent nurtures optimistic estimations of value functions and, during online iterations, selects policies aligned with the highest value estimates (Fruit et al., 2018; Auer et al., 2008; Agarwal et al., 2022b). Additionally, it’s noteworthy to acknowledge that several methodologies are rooted in the realm of posterior sampling, where the RL agent samples an MDP from a Bayesian Distribution and subsequently enacts the optimal policy (Osband et al., 2013; Agrawal & Jia, 2017; Agarwal et al., 2022a; Agarwal & Aggarwal, 2023). In our study, we follow the model-based approach and embrace the OFU-based algorithmic framework introduced in (Agarwal et al., 2022b), and we extend its scope to an augmented landscape where the RL agent gains access to supplementary quantum information. Furthermore, we render a mathematical characterization of regret, revealing a $\tilde{O}(1)$ bound, which in turn underscores the merits of astutely processing the quantum signals within our framework.

Quantum Mean Estimation: The realm of mean estimation revolves around the identification of the average value of samples stemming from an unspecified distribution. Of paramount importance is the revelation that quantum mean estimators yield a quadratic enhancement when juxtaposed against their classical counterparts (Montanaro, 2015; Hamoudi, 2021). The key reason for this improvement is based on the quantum amplitude amplification, which allows for suppressing certain quantum states w.r.t. the states that are desired to be extracted (Brassard et al., 2002). In Appendix B, we present a discussion around Quantum Amplitude Estimation and its applications in QML.

In the following, we introduce the definition of key elements and results pertaining to quantum mean estimation that are critical to our setup and analysis. First, we present the definition of a classical random variable and the quantum sampling oracle for performing quantum experiments.

Definition 1 (Random Variable, Definition 2.2 of (Cornelissen et al., 2022)). A finite random variable can be represented as $X : \Omega \rightarrow E$ for some probability space (Ω, \mathbb{P}) , where Ω is a finite sample set, $\mathbb{P} : \Omega \rightarrow [0, 1]$ is a probability mass function and $E \subset \mathbb{R}$ is the support of X . (Ω, \mathbb{P}) is

frequently omitted when referring to the random variable X .

To perform quantum mean estimation, we provide the definition of quantum experiment. This is analogous to the classical random experiments

Definition 2 (Quantum Experiment). Consider a random variable X on a probability space $(\Omega, 2^\Omega, \mathbb{P})$. Let \mathcal{H}_Ω be a Hilbert space with basis states $\{|\omega\rangle\}_{\omega \in \Omega}$ and fix a unitary $\mathcal{U}_\mathbb{P}$ acting on \mathcal{H}_Ω such that

$$\mathcal{U}_\mathbb{P} : |0\rangle \mapsto \sum_{\omega \in \Omega} \sqrt{\mathbb{P}(\omega)} |\omega\rangle$$

assuming $0 \in \Omega$. We define a *quantum experiment* as the process of applying the unitary $\mathcal{U}_\mathbb{P}$ or its inverse $\mathcal{U}_\mathbb{P}^{-1}$ on any state in \mathcal{H}_Ω .

The unitary $\mathcal{U}_\mathbb{P}$ provides the ability to query samples of the random variable in superpositioned states, which is the essence for speedups in quantum algorithms. To perform quantum mean estimation for values of random variables (Cornelissen et al., 2022; Hamoudi, 2021; Montanaro, 2015), an additional quantum evaluation oracle would be needed.

Definition 3 (Quantum Evaluation Oracle). Consider a finite random variable $X : \Omega \rightarrow E$ on a probability space $(\Omega, 2^\Omega, \mathbb{P})$. Let \mathcal{H}_Ω and \mathcal{H}_E be two Hilbert spaces with basis states $\{|\omega\rangle\}_{\omega \in \Omega}$ and $\{|x\rangle\}_{x \in E}$ respectively. We say that a unitary \mathcal{U}_X acting on $\mathcal{H}_\Omega \otimes \mathcal{H}_E$ is a quantum evaluation oracle for X if

$$\mathcal{U}_X : |\omega\rangle|0\rangle \mapsto |\omega\rangle|X(\omega)\rangle$$

for all $\omega \in \Omega$, assuming $0 \in E$.

Unlike a classical sample, which discloses only one successor state, the quantum oracle stores the entire distribution in coherent superposition, offering far greater information richness per query. We note that this above form of quantum oracle, where the agent performs a quantum experiment to query the oracle and obtain a superpositioned quantum sample, is widely adopted in theoretical quantum learning literature, including optimizations (Sidford & Zhang, 2023), searching algorithms (Grover, 1996), and episodic RL (Zhong et al., 2024). These quantum experiment frameworks are key mechanisms that enable quantum speedups in sample complexity and regret analysis. We now present a key quantum mean estimation result that will be carefully employed in our algorithmic framework to utilize the quantum superpositioned states collected by the RL agent. One of the crucial aspects of Lemma 1 is that quantum mean estimation converges at the rate $\mathcal{O}(\frac{1}{n})$ as opposed to classical benchmark convergence rate $\mathcal{O}(\frac{1}{\sqrt{n}})$ for n number of samples, therefore estimation efficiency quadratically.

Lemma 1 (Quantum multivariate bounded estimator, Theorem 3.3 of (Cornelissen et al., 2022)). *Let X be a d -dimensional bounded random variable such that $\|X\|_2 \leq 1$. Given three reals $L_2 \in (0, 1]$, $\delta \in (0, 1)$ and $n \geq 1$ such that $\mathbb{E}[\|X\|_2] \leq L_2$, the multivariate bounded estimator $\text{QBounded}_d(X, L_2, n, \delta)$ obtained by Algorithm 1 performs $f(n, d) = \mathcal{O}(n \log^{1/2}(n\sqrt{d}))$ quantum experiments and outputs a mean estimate $\hat{\mu}$ of $\mu = \mathbb{E}[X]$ such that,*

$$\mathbb{P} \left[\|\hat{\mu} - \mu\|_\infty \leq \frac{\sqrt{L_2} \log(d/\delta)}{n} \right] \geq 1 - \delta. \quad (1)$$

Algorithm 1 $\text{QBounded}_d(X, L_2, n, \delta)$ (Algorithm 1 in (Cornelissen et al., 2022))

- 1: **if** $n \leq \frac{\log(d/\delta)}{\sqrt{L_2}}$ **then**
- 2: Output $\hat{\mu} = 0$.
- 3: **end if**
- 4: Set $\alpha = \frac{1}{\sqrt{\log(400\pi n\sqrt{d})}}$ and $m = \left\lceil \log \left(\frac{8\pi n}{\alpha \sqrt{L_2} \log(d/\delta)} \right) \right\rceil$.
- 5: Set $G = \left\{ \frac{j}{m} - \frac{1}{2} + \frac{1}{2m} : j \in \{0, \dots, m-1\} \right\}^d \subseteq \left(-\frac{1}{2}, \frac{1}{2}\right)^d$.
- 6: **for** $k = 1, \dots, \lceil 18 \log(d/\delta) \rceil$ **do**
- 7: Compute the uniform superposition $|G\rangle := \frac{1}{\sqrt{md/2}} \sum_{u \in G} |u\rangle$ over G .
- 8: Compute the state $|\psi\rangle := \tilde{P}_{X, L_2, m, \alpha, \epsilon} |G\rangle$ in $H_G \otimes H_{aux}$, where $\tilde{P}_{X, L_2, m, \alpha, \epsilon}$ is the directional mean oracle constructed by the quantum evaluation oracle in Proposition 3.2 (Cornelissen et al., 2022) with $\epsilon = 1/25$.
- 9: Compute the state $|\phi\rangle := (QFT_G^{-1} \otimes I_{aux})|\psi\rangle$ where the unitary $QFT_G : |u\rangle \mapsto \frac{1}{\sqrt{md/2}} \sum_{v \in G} e^{2\pi i \langle u, v \rangle} |v\rangle$ is the quantum Fourier transform over G .
- 10: Measure the H_G register of $|\phi\rangle$ in the computational basis and let $v^{(k)} \in G$ denote the obtained result. Set $\mu^{(k)} = \frac{2\pi v^{(k)}}{\alpha}$.
- 11: **end for**
- 12: Output the coordinate-wise median $\hat{\mu} = \text{median}(\mu^{(1)}, \dots, \mu^{(\lceil 18 \log(d/\delta) \rceil)})$.

Quantum Reinforcement Learning: Within the realm of QRL, a prominent strand of previous research showcases exponential enhancements in regret through amplitude amplification-based methodologies applied to the Quantum Multi-Armed Bandits (Q-MAB) problem (Wang et al., 2021; Casalé et al., 2020; Wan et al., 2023; Wu et al., 2023). Nevertheless, the theoretical underpinnings of the aforementioned Q-MAB approaches do not seamlessly extend to the QRL context since there is no state evolution in bandits.

A recent surge of research interest has been directed towards Quantum Reinforcement Learning (QRL) (Jerbi et al., 2021; Dunjko et al., 2017; Paparo et al., 2014). It is noteworthy to emphasize that the aforementioned studies do not provide an exhaustive mathematical characterization of regret. (Zhong et al., 2024; Ganguly et al., 2023) demonstrated that QRL can provide logarithmic regret in the episodic MDP settings. Ours is the first study of infinite horizon MDP with average reward objective.

3. Problem Formulation

We consider the problem of Reinforcement Learning (RL) in an infinite horizon Markov Decision Process characterized by $\mathcal{M} \triangleq (\mathcal{S}, \mathcal{A}, P, r, D)$, wherein \mathcal{S} and \mathcal{A} represent finite collections of states and actions respectively with $|\mathcal{S}| = S$ and $|\mathcal{A}| = A$, pertaining to RL agent's interaction with the unknown MDP environment. $P(s'|s, a) \in [0, 1]$ denotes the transition probability for next state $s' \in \mathcal{S}$ for a given pair of previous state and RL agent's action, i.e., $(s, a) \in \mathcal{S} \times \mathcal{A}$. Further, $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ represents the reward collected by the RL agent for state-action pair (s, a) . D is the diameter of the MDP \mathcal{M} . In the following, we first present the additional quantum transition oracle that the agent could access during its interaction with the unknown MDP environment at every RL round. Here, we would like to emphasize that unknown MDP environment implies that matrix P which encapsulates the transition dynamics of the underlying RL environment is not known beforehand.

Quantum Computing Basics: In this section, we give a brief introduction to the concepts that are most related to our work based on (Nielsen & Chuang, 2010) and (Wu et al., 2023). Consider an m dimensional Hilbert space \mathbb{C}^m , a quantum state $|x\rangle = (x_1, \dots, x_m)^T$ can be seen as a vector inside the Hilbert space with $\sum_i |x_i|^2 = 1$. Furthermore, for a finite set with m elements $\xi = \{\xi_1, \dots, \xi_m\}$, we can assign each $\xi_i \in \xi$ with a member of an orthonormal basis of the Hilbert space \mathbb{C}^m by

$$\xi_i \mapsto |\xi_i\rangle \equiv e_i \quad (2)$$

where e_i is the i th unit vector for \mathbb{C}^m . Using the above notation, we can express any arbitrary quantum state $|x\rangle = (x_1, \dots, x_m)^T$ by elements in ξ as:

$$|x\rangle = \sum_{n=1}^m x_n |\xi_n\rangle \quad (3)$$

where $|x\rangle$ is the quantum superposition of the basis $|\xi_1\rangle, \dots, |\xi_m\rangle$ and we denote x_n as the amplitude of $|\xi_n\rangle$. To obtain a classical information from the quantum state, we perform a measurement and the quantum state

would collapse to any basis $|\xi_i\rangle$ with probability $|x_i|^2$. In quantum computing, the quantum states are represented by input or output registers made of qubits that could be in superpositions.

Quantum transition oracle: We utilize the concepts and notations of quantum computing in (Wang et al., 2021; Wiedemann et al., 2022; Ganguly et al., 2023; Jerbi et al., 2022) to construct the quantum sampling oracle of RL environments and capture agent's interaction with the unknown MDP environment. We now formulate the equivalent quantum-accessible RL environments for our classical MDP \mathcal{M} . For an agent at step t in state s_t and with action a_t , we construct the quantum sampling oracles for variables of the next state $P(\cdot|s_t, a_t)$. Specifically, suppose there are two Hilbert spaces $\bar{\mathcal{S}} = \mathbb{C}^{|\mathcal{S}|}$ and $\bar{\mathcal{A}} = \mathbb{C}^{|\mathcal{A}|}$ containing the superpositions of the classical states and actions. We represent the computational bases for $\bar{\mathcal{S}}$ and $\bar{\mathcal{A}}$ as $\{|s\rangle\}_{s \in \mathcal{S}}$ and $\{|a\rangle\}_{a \in \mathcal{A}}$. We assume that we can implement the quantum sampling oracle \mathcal{U}_P of the MDP's transitions as follows:

The quantum evaluation oracle for the transition probability (quantum transition oracle) \mathcal{U}_P which at step t , returns the superposition over $s' \in \mathcal{S}$ according to $P(s'|s_t, a_t)$, the probability distribution of the next state given the current state $|s_t\rangle$ and action $|a_t\rangle$ is defined as:

$$\mathcal{U}_P : |s_t\rangle \otimes |a_t\rangle \otimes |0\rangle \rightarrow |\psi\rangle^t$$

Where $|\psi\rangle^t = |s_t\rangle \otimes |a_t\rangle \otimes \sum_{s' \in \mathcal{S}} \sqrt{P(s'|s_t, a_t)} |s'\rangle$.

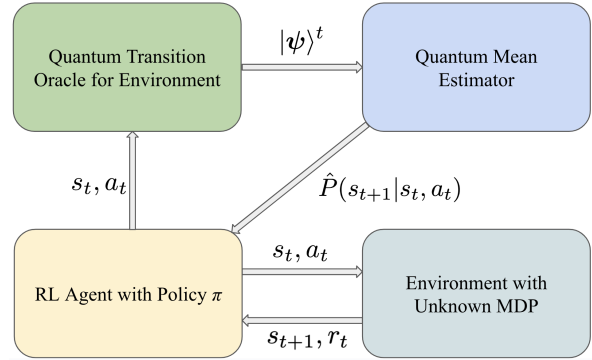


Figure 1. Agent's interaction at round t with the MDP Environment and accessible quantum transition oracle

As described in Fig 1, at step t , the agent plays an action a_t based on current state s_t according to its policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$. Consequently, the unknown MDP Q-environment shares the next state information and reward i.e., $\{s_{t+1}, r_t\}$. Additionally, the quantum transition oracle first encodes s_t and a_t into the corresponding basis of the quantum Hilbert spaces, producing basis $|s_t\rangle$ and $|a_t\rangle$, then it encodes the superpositioned quantum state $|\psi\rangle^t$ from the quantum

transition oracle $\mathcal{U}_P \cdot |\psi\rangle^t$ would be used for the Quantum Mean Estimator, QBounded, which in turn improves transition probability estimation and leads to exponentially better regret. We note that one copy of $|\psi\rangle^t$ is obtained at each (s_t, a_t) . Given that a measurement collapses the state, we can only perform one measurement on it. So, the quantum mean estimator will only perform measurements at epoch ends, as will be described in the algorithm.

RL regret formulation: Consider the RL agent following some policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ on the unknown MDP Q-environment \mathcal{M} . Then, as a result of playing policy π , the agent's long term average reward is defined next according to (Agarwal et al., 2022b).

$$\Gamma_\pi^P = \lim_{T \rightarrow \infty} \mathbb{E}_{\pi, P} \left[\frac{1}{T} \sum_{t=0}^{T-1} r(s_t, a_t) \right], \quad (4)$$

where $\mathbb{E}_{\pi, P}[\cdot]$ is the expectation taken over the trajectory $\{(s_t, a_t)\}_{t \in [0, T-1]}$ by playing policy π on \mathcal{M} at every time step. In this context, we state the definitions of value function for MDP \mathcal{M} next:

$$V_\pi^P(s; \gamma) = \mathbb{E}_{a \sim \pi} \left[Q_\pi^P(s, a; \gamma) \right], \quad (5)$$

$$= \mathbb{E}_{a \sim \pi} \left[r(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) V_\pi^P(s'; \gamma) \right] \quad (6)$$

Then, according to (Puterman, 2014), Γ_π^P can be represented as the following:

$$\Gamma_\pi^P = \lim_{\gamma \rightarrow 1} (1 - \gamma) V_\pi^P(s; \gamma), \quad \forall s \in \mathcal{S}, \quad (7)$$

$$= \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \rho_\pi^P(s, a) \bar{r}(s, a), \quad (8)$$

wherein, $V_\pi^P(\cdot; \gamma)$ is the discounted cumulative average reward by following policy π , ρ_π^P is the steady-state occupancy measure (i.e., $\rho_\pi^P(s, a) \in [0, 1]$, $\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \rho_\pi^P(s, a) = 1$), $\bar{r}(s, a)$ is the steady-state long-term average reward of pair (s, a) . In the following, we introduce key assumptions crucial from the perspective of our algorithmic framework and its analysis. In this context, we introduce the notations $\{P_{\pi, s}^t, T_{\pi, s \rightarrow s'}\}$ to denote the t time step probability distribution obtained by playing policy π on the unknown MDP \mathcal{M} starting from some arbitrary state $s \in \mathcal{S}$; and, the actual time steps to reach s' from s by playing policy π respectively. Next, we state an assumption pertaining to ergodicity of MDP \mathcal{M} .

Assumption 1 (Finite MDP mixing time). We assume that unknown MDP environment \mathcal{M} has finite mixing time, which mathematically implies:

$$T_{\text{mix}} \triangleq \max_{\pi, (s, s') \in \mathcal{S} \times \mathcal{S}} \mathbb{E}[T_{\pi, s \rightarrow s'}] < \infty, \quad (9)$$

where $T_{\pi, s \rightarrow s'}$ implies the number of time steps to reach a state s' from an initial state s by playing some policy π .

Assumption 2. The reward function $r(\cdot, \cdot)$ is known to the RL agent.

We emphasize that Assumption 2 is a commonly used assumption in RL algorithm development owing to the fact that in most setups rewards are constructed according to the underlying problem and is known beforehand (Azar et al., 2017; Agarwal et al., 2019). Next, we define the cumulative regret accumulated by the agent in expectation across T time steps as follows:

$$\mathcal{R}_{[0, T-1]} \triangleq T \cdot \Gamma_{\pi^*}^P - \mathbb{E} \left[\sum_{t=0}^{T-1} r(s_t, a_t) \right], \quad (10)$$

where the agent generates the trajectory $\{(s_t, a_t)\}_{t \in [0, T-1]}$ by playing policy π on \mathcal{M} , and π^* is the optimal policy that maximizes the long term average reward defined in Eq. (4). Finally, the RL agent's goal is to determine a policy π which minimizes $\mathcal{R}_{[0, T-1]}$.

4. Algorithmic Framework

In this section, we first delineate the key intuition behind our methodology followed by formal description of Q-UCRL algorithm. If the agent was aware of the true transition probability distribution P , then it would simply solve the following optimization problem for the optimal policy.

$$\max_{\{\rho(s, a)\}_{(s, a) \in \mathcal{S} \times \mathcal{A}}} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} r(s, a) \rho(s, a) \quad (11)$$

With the following constraints,

$$\sum_{s, a} \rho(s, a) = 1, \rho(s, a) \geq 0 \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (12)$$

$$\sum_{a \in \mathcal{A}} \rho(s', a) = \sum_{s, a} P(s'|s, a) \rho(s, a), \quad \forall s' \in \mathcal{S} \quad (13)$$

Note that the objective in Eq. (11) resembles the definition of steady-state average reward formulation in Eq. (4). Furthermore, Eq. (12), (13) preserve the properties of the transition model. Consequently, once it obtains $\{\rho^*(s, a)\}_{(s, a) \in \mathcal{S} \times \mathcal{A}}$ by solving Eq. (11), it can compute the optimal policy as:

$$\pi^*(a|s) = \frac{\rho^*(s, a)}{\sum_{\tilde{a} \in \mathcal{A}} \rho^*(s, \tilde{a})} \quad \forall s, a \quad (14)$$

However, in the absence of knowledge pertaining to true P , we propose to solve an optimistic version of Eq. (11) that will update RL agent's policy in epochs. Specifically, at the start of an epoch e , the following optimization problem will be solved:

$$\max_{\{\rho(s,a), P_e(\cdot|s,a)\}_{(s,a) \in \mathcal{S} \times \mathcal{A}}} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} r(s,a) \rho(s,a), \quad (15)$$

With the following constraints,

$$\sum_{s,a} \rho(s,a) = 1, \rho(s,a) \geq 0 \quad \forall (s,a) \in \mathcal{S} \times \mathcal{A} \quad (16)$$

$$\sum_{s' \in \mathcal{S}} P_e(s'|s,a) = 1, P_e(\cdot|s,a) \geq 0 \quad \forall (s,a) \in \mathcal{S} \times \mathcal{A} \quad (17)$$

$$\sum_{a \in \mathcal{A}} \rho(s',a) = \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} P_e(s'|s,a) \rho(s,a), \quad \forall s' \in \mathcal{S} \quad (18)$$

$$\|P_e(\cdot|s,a) - \hat{P}_e(\cdot|s,a)\|_1 \leq \frac{7S \log(S^2 A t_e)}{\max\{1, N_e(s,a)\}} \quad (19)$$

Observe that Eq. (15) increases the search space for policy in a carefully computed neighborhood of certain transition probability estimates \hat{P} via Eq. (19) at the start of epoch e . Also in Eq. (19), recall that $S = |\mathcal{S}|$ is the number of states, $A = |\mathcal{A}|$ is the number of actions, $N_e(s,a)$ is the number of visitations of state-action pair (s,a) upto the end of epoch e . Note that, solving Eq. (15) is equivalent to obtaining an optimistic policy as it extends the search space of original problem Eq. (11). Furthermore, the computation of \hat{P}_e is key to our framework design. Since performing quantum mean estimation to quantum samples would require measuring and results in quantum samples collapsing, each quantum sample can only be used once. We emphasize that the quantum mean estimator will only run at the end of the epoch, performing measurements to update the transition probability for the next epoch.

Details on Quantum Mean Estimator: At the end of the epoch, we use all (s_i, a_i) observed in the samples, and perform quantum mean estimation on all the samples $|\psi\rangle^t$ obtained for that (s_i, a_i) . More precisely, $\hat{P}_e(\cdot|s_i, a_i)$ is the quantum mean estimation for the random variable $P_e(\cdot|s_i, a_i)$ performed by `QBounded` (Algorithm 1) using the quantum samples $\{|\psi\rangle\}_i^e$ observed in the previous epoch with the total number being $\nu_{e-1}(s_i, a_i)$:

$$\{|\psi\rangle\}_i^e = \{|\psi\rangle^t \mid t_{e-1}^{start} \leq t \leq t_e^{start} - 1, (s_t, a_t) = (s_i, a_i)\} \quad (20)$$

where t_e^{start} is the starting step of epoch e . Using these samples, the transition probability is obtained as

$$\tilde{P}_e(\cdot|s_i, a_i) = \text{QBounded}_S(\{P_e(\cdot|s_i, a_i), L_2, n_e^{s_i, a_i}, \delta_e\}) \quad (21)$$

For a specific state-action pair (s_i, a_i) , we construct an estimator $\hat{P}_e(\cdot|s_i, a_i)$ used in epoch e that consists of a weighted average of the estimator $\hat{P}_{e-1}(\cdot|s_i, a_i)$ and the

estimator obtained using the quantum samples in the latest completed epoch $\tilde{P}_e(\cdot|s_i, a_i)$ as follows:

$$\hat{P}_e(\cdot|s_i, a_i) = \begin{cases} 0 & \text{if } e = 1 \\ \tilde{P}_e(\cdot|s_i, a_i) & \text{if } e = 2 \\ \hat{w} \hat{P}_{e-1}(\cdot|s_i, a_i) + \tilde{w} \tilde{P}_e(\cdot|s_i, a_i) & \text{if } e \geq 3 \end{cases} \quad (22)$$

where

$$\hat{w} = \frac{N_{e-1}(s_i, a_i)}{N_{e-1}(s_i, a_i) + \nu_{e-1}(s_i, a_i)} \quad (23)$$

$$\tilde{w} = \frac{\nu_{e-1}(s_i, a_i)}{N_{e-1}(s_i, a_i) + \nu_{e-1}(s_i, a_i)} \quad (24)$$

In particular, given that $\nu_{e-1}(s_i, a_i)$ is the maximum number of quantum experiments we can perform in Eq. (21), we set

$$n_e^{s_i, a_i} = \left\lfloor \frac{\nu_{e-1}(s_i, a_i)}{c \log^{1/2}(T\sqrt{S})} \right\rfloor \quad (25)$$

for $c \in \mathbb{R}$ satisfying $cn \log^{1/2}(n\sqrt{S}) \geq f(n, S) \quad \forall n$ where $f(n, S)$ is defined in Lemma 1. Such choice of c would allow $\nu_e(s, a)$ samples to be used for analysis of our proposed algorithm. Note that the choice of c is achievable because $cn \log^{1/2}(n\sqrt{S})$ and $f(n, S)$ are the same order. Since $\mathbb{E}[\|P_e(\cdot|s_i, a_i)\|_2] \leq 1$, we set $L_2 = 1$. Different from the works in classical approaches based on (Jaksch et al., 2010), our algorithm utilizes each quantum sample only once for estimating all \hat{P}_e throughout the entire process. This strategy fits with our quantum framework since quantum samples would collapse once being measured. In Appendix A, we show that Eq. (15) is a convex optimization problem, therefore standard optimization solvers can be employed to solve it. Finally, the agent's policy at the start of epoch e is given by:

$$\pi_e(a|s) = \frac{\rho_e(s, a)}{\sum_{\tilde{a} \in \mathcal{A}} \rho_e(s, \tilde{a})} \quad (26)$$

Next, we formally present Q-UCRL in Algorithm 2.

Algorithm 2 proceeds in epochs, and each epoch e contains multiple rounds of RL agent's interaction with the environment. At each t , by playing action a_t for instantaneous state s_t through current epoch policy π_e , the agent collects next state configuration consisting of classical signals $s_{t+1}, r(s_t, a_t)$ and a quantum sample $|\psi\rangle^{(t)}$. Additionally, the agent keeps track of its visitations to the different states via the variables $\{\nu_{s,a}, \mu_{s,a,s'}\}$. Note that a new epoch is triggered at the "if" block of Line 15 in algorithm 2, which in turn is aligned to the widely used concept of *doubling trick* used in model-based RL algorithm design. Finally, the

Algorithm 2 Q-UCRL

```

1: Inputs:  $\mathcal{S}, \mathcal{A}, r(\cdot, \cdot)$ .
2: Set  $t \leftarrow 1$ ,  $e \leftarrow 1$ ,  $\delta_e \leftarrow 0$ ,  $t_e \leftarrow 0$ ,  $t_e^{start} \leftarrow 1$ 
3: Set  $\mu(s, a, s') \leftarrow 0 \forall (s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ .
4: Set  $\hat{P}_e(\cdot|s, a) \leftarrow 0 \forall (s, a) \in \mathcal{S} \times \mathcal{A}$ .
5: for  $(s, a) \in \mathcal{S} \times \mathcal{A}$  do
6:   Set  $\nu_e(s, a) \leftarrow 0$ ,  $N_e(s, a) \leftarrow 0$ .
7: end for
8: Obtain  $\pi_e$  by solving Eq. (15) and Eq. (26).
9: for  $t = 1, 2, \dots$  do
10:  Observe state  $s_t$  and play action  $a_t \sim \pi_e(\cdot|s_t)$ .
11:  Observe  $s_{t+1}, r(s_t, a_t)$  and quantum samples  $|\psi\rangle^t$ .
12:  Update  $\nu_e(s_t, a_t) \leftarrow \nu_e(s_t, a_t) + 1$ .
13:  Update  $\mu(s_t, a_t, s_{t+1}) \leftarrow \nu_e(s_t, a_t, s_{t+1}) + 1$ .
14:  Set  $t_e \leftarrow t_e + 1$ .
15:  if  $\nu_e(s_t, a_t) = \max\{1, N_e(s_t, a_t)\}$  then
16:    for  $(s, a) \in \mathcal{S} \times \mathcal{A}$  do
17:       $N_{e+1}(s, a) \leftarrow N_e(s, a) + \nu_e(s, a)$ .
18:    end for
19:    Set  $e \leftarrow e + 1$ ,  $\nu_e(s, a) \leftarrow 0$ ,  $\delta_e \leftarrow \frac{1}{S^2 A t_e^7}$ .
20:    Set  $t_e \leftarrow 0$ ,  $t_e^{start} \leftarrow t + 1$ .
21:    Obtain  $\hat{P}_e(\cdot|s, a)$  by Eq. (22)  $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$ 
22:    Obtain  $\pi_e$  by solving Eq. (15) and Eq. (26).
23:  end if
24: end for
    
```

empirical visitation counts and probabilities up to epoch e , i.e., $\{N_{e+1}(s, a), \hat{P}_{e+1}(s'|s, a)\}$ attributes are updated (line 15-21), and the policy is updated via re-solving from Eq. (15) - (19) and using Eq. (26).

5. Theoretical Results

In this section, we characterize the cumulative regret for T rounds by running Q-UCRL in a unknown MDP hybrid Q-Environment. In the following, we first present a crucial result bounding the error accumulated till e epochs between the true transition probabilities $\{P(\cdot|s, a)\}$ and the agent's estimates $\{\hat{P}_e(\cdot|s, a)\}$.

Lemma 2. *Execution of Q-UCRL (Algorithm 2) up to the beginning of epoch $e + 1$ with total e completed epochs comprising $t = 1, 2, \dots, t_e$ rounds, guarantees that the following holds for \hat{P}_{e+1} :*

$$\mathbb{P}\left[\left|\tilde{P}_{e+1}(s'|s, a) - P(s'|s, a)\right| \leq \frac{C \log(S/\delta)}{\nu_e(s, a)}\right] \geq 1 - \delta \quad (27)$$

$\forall (s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$, where $C = c \log^{1/2}(T\sqrt{S})$ for some $c \in \mathbb{R}$ defined in Eq. (25) and $\{\nu_e(s, a)\}$ are the state-action pair visitation counts up in epoch e , as depicted in Line 12 of Algorithm 2, and $\{P(s'|s, a), \tilde{P}_{e+1}(s'|s, a)\}$ are the actual and estimated transition probabilities by Eq.

(21).

Proof. For all $(s, a) \in \mathcal{S} \times \mathcal{A}$, since \tilde{P}_{e+1} is obtained by Eq. (21), we obtain using Lemma 1:

$$\mathbb{P}\left[\left\|\tilde{P}_{e+1}(\cdot|s, a) - P(\cdot|s, a)\right\|_\infty \leq \frac{\log(S/\delta)}{n_{e+1}^{s_i, a_i}}\right] \geq 1 - \delta, \quad (28)$$

By switching $n_{e+1}^{s_i, a_i}$ with $\nu_e(s, a)/C$, Eq. (28) gives:

$$\mathbb{P}\left[\left\|\tilde{P}_{e+1}(\cdot|s, a) - P(\cdot|s, a)\right\|_\infty \leq \frac{C \log(S/\delta)}{\nu_e(s, a)}\right] \geq 1 - \delta \quad (29)$$

Eq. (29) can be transformed into Eq. (27) by taking the entry-wise expression. \square

Note that the choice of C in Eq. (27) is to ensure that the samples $\nu_e(s, a)$ are sufficient for performing Algorithm 1.

Lemma 3. *Execution of Q-UCRL (Algorithm 2) up to the beginning of epoch $e + 1$ with total e completed epochs comprising $t = 1, 2, \dots, t_e$ rounds, guarantees that the following holds for \hat{P}_{e+1} :*

$$\mathbb{P}\left[\left|\hat{P}_{e+1}(s'|s, a) - P(s'|s, a)\right| \leq \frac{eC \log(eS/\delta)}{N_{e+1}(s, a)}\right] \geq 1 - \delta \quad (30)$$

$\forall (s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$, where $C = c \log^{1/2}(T\sqrt{S})$ for some $c \in \mathbb{R}$ and $\{\nu_e(s, a)\}$ are the state-action pair visitation counts up in epoch e , as depicted in Line 12 of Algorithm 2, and $\{P(s'|s, a), \hat{P}_{e+1}(s'|s, a)\}$ are the actual and weighted average of transition probabilities by Eq. (22).

Proof. Please refer to Appendix C. \square

Lemma 2 and Lemma 3 provide the acceleration for mean estimation, which further improve global convergence rate.

Lemma 4. *Execution of Q-UCRL (Algorithm 2) up to total e epochs comprising $t = 1, 2, \dots, t_e$ rounds, guarantees that the following holds for confidence parameter $\delta_e = \frac{1}{S A t_e^7}$ with probability at least $1 - 1/t_e^6$:*

$$\begin{aligned} & \|P_e(\cdot|s, a) - P(\cdot|s, a)\|_1, \\ & \leq \frac{7S(1 + Ce) \log(S^2 A T) + S C e \log(eS)}{\max\{1, N_e(s, a)\}}, \end{aligned} \quad (31)$$

$\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, where $\{N_{e-1}^{(s, a)}\}$ is the cumulative state-action pair visitation counts up to epoch e , as depicted in Line 17 of Algorithm 2, and $\{P(\cdot|s, a), P_e(\cdot|s, a)\}$ are the actual and estimated transition probabilities.

Proof. Please refer to Appendix D. \square

Remark 1. Note that for $P_e(\cdot|s, a), P(\cdot|s, a)$ we have the bound: $\|P_e(\cdot|s, a) - P(\cdot|s, a)\|_1 \leq 2S$, in the worst case scenario. However, the probability that a *bad event* (i.e., complementary of Eq. (31)) happens is at most $\frac{1}{t_e^6}$.

Lemma 5 (Regret Decomposition). *Regret of Q-UCRL defined in Eq. (10) can be decomposed as follows:*

$$\mathcal{R}_{[0, T-1]} \leq \sum_{e=1}^E T_e \left(\Gamma_{\pi_e}^{\bar{P}_e} - \Gamma_{\pi_e}^P \right). \quad (32)$$

Proof. Please refer to Appendix E. \square

Regret in terms of Bellman Errors: Here, we introduce the notion of Bellman errors $B_{\pi_e}^{P_e}(s, a)$ at epoch e as follows:

$$B_{\pi_e}^{P_e}(s, a) \triangleq \lim_{\gamma \rightarrow 1} [Q_{\pi_e}^{P_e}(s, a; \gamma) - r(s, a) - \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) V_{\pi_e}^{P_e}(s'; \gamma)], \quad (33)$$

which essentially captures the gap in the average rewards by playing action a in state s accordingly to some policy π in one step of the optimistic transition model P_e w.r.t. the actual probabilities P . This leads us to the next result that ties the regret expression in Eq. (32) and the Bellman errors.

Lemma 6 (Lemma 5.2, (Agarwal et al., 2022b)). *The difference in long-term expected reward by playing optimistic policy of epoch e , i.e., π_e on the optimistic transition model estimates, i.e., $\Gamma_{\pi_e}^{P_e}$, and the expected reward collected by playing π_e on the actual model P , i.e., $\Gamma_{\pi_e}^P$ is the long-term average of the Bellman errors. Mathematically, we have:*

$$\Gamma_{\pi_e}^{\bar{P}_e} - \Gamma_{\pi_e}^P = \sum_{s, a} \rho_{\pi_e}^P(s, a) B_{\pi_e}^{P_e}(s, a), \quad (34)$$

where $B_{\pi_e}^{P_e}(s, a)$ is defined in Eq. (33).

With Eq. (34), it necessary to mathematically bound the Bellman errors. To this end, we leverage the following result from (Agarwal et al., 2022b) which is stated as Lemma 10 from Appendix F characterizing the Bellman errors in terms of $\tilde{h}(\cdot)$, a quantity that measures the bias contained in the optimistic MDP model.

$$B_{\pi_e, P_e}(s, a) \leq \left\| P_e(\cdot|s, a) - P(\cdot|s, a) \right\|_1 \|\tilde{h}(\cdot)\|_{\infty}, \quad (35)$$

where $\tilde{h}(\cdot)$ (defined in (Puterman, 2014)) is the bias of the optimistic MDP P_e for which the optimistic policy π_e obtains the maximum expected reward in the confidence interval. Using the bound of the error arising from estimation of

transition model in Lemma 4 into Eq. (35), we obtain the following result with probability $1 - 1/t_e^6$:

$$\begin{aligned} B_{\pi_e, P_e}(s, a) &\leq \frac{7S(1 + Ce) \log(S^2 AT) + SCe \log(eS)}{\max\{1, N_e(s, a)\}} \|\tilde{h}(\cdot)\|_{\infty}, \end{aligned} \quad (36)$$

Consequently, it is necessary to bound the model bias error term in Eq. (36). We recall that the optimistic MDP P_e corresponding to epoch e maximizes the average reward in the confidence set as described in the formulation of Eq. (15). In other words, $\Gamma_{\pi_e}^{P_e} \geq \Gamma_{\pi_e}^{P'}$ for every P' in the confidence interval. Specifically, the model bias for optimistic MDP can be explicitly defined in terms of Bellman equation (Puterman, 2014), as follows:

$$\tilde{h}(s) \triangleq r_{\pi_e}(s) - \lambda_{\pi_e}^{P_e} + \left(P_{e, \pi_e}(\cdot|s) \right)^T \tilde{h}(\cdot), \quad (37)$$

where the following simplified notations are used: $r_{\pi_e}(s) = \sum_{a \in \mathcal{A}} \pi_e(a|s) r(s, a)$, $P_{e, \pi_e}(s'|s) = \sum_{a \in \mathcal{A}} \pi_e(a|s) P_e(s'|s, a)$. Next, we incorporate the following result from (Agarwal et al., 2022b) stated as Lemma 11 in Appendix F correlating model bias $\tilde{h}(\cdot)$ to the mixing time T_{mix} of MDP \mathcal{M} .

$$\tilde{h}(s) - \tilde{h}(s') \leq T_{\text{mix}}, \forall s, s' \in \mathcal{S}. \quad (38)$$

Using bound of bias term $\tilde{h}(\cdot)$ into Eq. (36), we get the final bound for Bellman error with probability $1 - 1/T^6$ as:

$$B_{\pi_e, P_e}(s, a) \leq \frac{7S(1 + Ce) \log(S^2 AT) + SCe \log(eS)}{\max\{1, N_e(s, a)\}} T_{\text{mix}}. \quad (39)$$

Note that Q-UCRL obtains a quadratic improvement for the Bellman error over the classical UC-UCRL (Agarwal et al., 2022b). We next present the final regret bound of Q-UCRL.

Theorem 1. *In an unknown MDP environment $\mathcal{M} \triangleq (\mathcal{S}, \mathcal{A}, P, r, D)$, the regret incurred by Q-URL (Algorithm 2) across T rounds is bounded as follows:*

$$\mathcal{R}_{[0, T-1]} = \mathcal{O} \left(S^5 A^4 T_{\text{mix}} \log^3 \left(\frac{T}{SA} \right) \log^{\frac{1}{2}}(T\sqrt{S}) \log(S^2 AT) \right). \quad (40)$$

Proof. Please refer to Appendix G. \square

Remark 2. We reiterate that our characterization of regret incurred by Q-UCRL in Theorem 1 is a martingale-free analysis and thus deviates from its classical counterparts (Fruit et al., 2018; Auer et al., 2008; Agarwal et al., 2022b).

6. Conclusions

This paper demonstrates that quantum computing helps provide significant reduction in the regret bounds for infinite horizon reinforcement learning with average rewards. This paper not only unveils the quantum potential for bolstering reinforcement learning but also beckons further exploration into novel avenues that bridge quantum and classical paradigms, thereby advancing the field to new horizons.

A promising future direction is parameterized quantum RL, which has recently shown speedups for discounted settings (Xu & Aggarwal, 2025); achieving similar gains for average-reward problems, where the state-of-the-art in non-quantum RL is given by (Ganesh et al., 2025b), remains an open challenge.

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

- Agarwal, A., Jiang, N., Kakade, S. M., and Sun, W. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep.*, 32, 2019.
- Agarwal, M. and Aggarwal, V. Reinforcement learning for joint optimization of multiple rewards. *Journal of Machine Learning Research*, 24(49):1–41, 2023.
- Agarwal, M., Aggarwal, V., and Lan, T. Multi-objective reinforcement learning with non-linear scalarization. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pp. 9–17, 2022a.
- Agarwal, M., Bai, Q., and Aggarwal, V. Concave utility reinforcement learning with zero-constraint violations. *Transactions on Machine Learning Research*, 2022b.
- Agrawal, S. and Jia, R. Optimistic posterior sampling for reinforcement learning: worst-case regret bounds. *Advances in Neural Information Processing Systems*, 30, 2017.
- Al-Abbasi, A. O., Ghosh, A., and Aggarwal, V. Deepool: Distributed model-free algorithm for ride-sharing using deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 20(12):4714–4727, 2019.
- 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., Osband, I., and Munos, R. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pp. 263–272. PMLR, 2017.
- Bai, Q., Mondal, W. U., and Aggarwal, V. Regret analysis of policy gradient algorithm for infinite horizon average reward markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 10980–10988, 2024.
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., and Lloyd, S. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.
- Bonjour, T., Haliem, M., Alsalem, A., Thomas, S., Li, H., Aggarwal, V., Kejriwal, M., and Bhargava, B. Decision making in monopoly using a hybrid deep reinforcement learning approach. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 6(6):1335–1344, 2022.
- Boulund, A., Getachew, Y. M., Jin, Y., Sidford, A., and Tian, K. Quantum speedups for zero-sum games via improved dynamic gibbs sampling. In *International Conference on Machine Learning*, pp. 2932–2952. PMLR, 2023.
- Brassard, G., Hoyer, P., Mosca, M., and Tapp, A. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002.
- Casalé, B., Di Molfetta, G., Kadri, H., and Ralaivola, L. Quantum bandits. *Quantum Machine Intelligence*, 2:1–7, 2020.
- Chen, J., Umrawal, A. K., Lan, T., and Aggarwal, V. Deep-freight: A model-free deep-reinforcement-learning-based algorithm for multi-transfer freight delivery. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 31, pp. 510–518, 2021.
- Cornelissen, A., Hamoudi, Y., and Jerbi, S. Near-optimal quantum algorithms for multivariate mean estimation. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 33–43, 2022.
- Dunjko, V., Taylor, J. M., and Briegel, H. J. Advances in quantum reinforcement learning. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 282–287. IEEE, 2017.
- Farhi, E., Goldstone, J., and Gutmann, S. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.
- Fruit, R., Pirota, M., Lazaric, A., and Ortner, R. Efficient bias-span-constrained exploration-exploitation in reinforcement learning. In *International Conference on Machine Learning*, pp. 1578–1586. PMLR, 2018.

- Ganesh, S., Mondal, W. U., and Aggarwal, V. Order-optimal regret with novel policy gradient approaches in infinite horizon average reward mdps. In *The 28th International Conference on Artificial Intelligence and Statistics*, 2025a.
- Ganesh, S., Mondal, W. U., and Aggarwal, V. A sharper global convergence analysis for average reward reinforcement learning via an actor-critic approach. In *Forty-second International Conference on Machine Learning*, 2025b.
- Ganguly, B., Wu, Y., Wang, D., and Aggarwal, V. Quantum computing provides exponential regret improvement in episodic reinforcement learning. *arXiv preprint arXiv:2302.08617*, 2023.
- Gilyén, A., Su, Y., Low, G. H., and Wiebe, N. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 193–204, 2019.
- Grover, L. K. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pp. 212–219, 1996.
- Hamoudi, Y. Quantum sub-gaussian mean estimator. In *29th Annual European Symposium on Algorithms (ESA 2021)*, pp. 50–1, 2021.
- Harrow, A. W., Hassidim, A., and Lloyd, S. Quantum algorithm for linear systems of equations. *Physical review letters*, 103(15):150502, 2009.
- Hong, K., Chae, W., Zhang, Y., Lee, D., and Tewari, A. Reinforcement learning for infinite-horizon average-reward linear mdps via approximation by discounted-reward mdps. In *The 28th International Conference on Artificial Intelligence and Statistics*, 2025.
- Jaksch, T., Ortner, R., and Auer, P. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11:1563–1600, 2010.
- Jerbi, S., Trenkwalder, L. M., Nautrup, H. P., Briegel, H. J., and Dunjko, V. Quantum enhancements for deep reinforcement learning in large spaces. *PRX Quantum*, 2(1): 010328, 2021.
- Jerbi, S., Cornelissen, A., Ozols, M., and Dunjko, V. Quantum policy gradient algorithms. *arXiv e-prints*, pp. arXiv–2212, 2022.
- Kitaev, A. Y. Quantum measurements and the abelian stabilizer problem. *arXiv preprint quant-ph/9511026*, 1995.
- Lloyd, S., Mohseni, M., and Rebentrost, P. Quantum algorithms for supervised and unsupervised machine learning. *arXiv preprint arXiv:1307.0411*, 2013.
- Montanaro, A. Quantum speedup of monte carlo methods. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2181):20150301, 2015.
- Nielsen, M. A. and Chuang, I. L. *Quantum computation and quantum information*. Cambridge university press, 2010.
- Osband, I., Russo, D., and Van Roy, B. (more) efficient reinforcement learning via posterior sampling. *Advances in Neural Information Processing Systems*, 26, 2013.
- Paparo, G. D., Dunjko, V., Makmal, A., Martin-Delgado, M. A., and Briegel, H. J. Quantum speedup for active learning agents. *Physical Review X*, 4(3):031002, 2014.
- Puterman, M. L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Rosenberg, A. and Mansour, Y. Online convex optimization in adversarial markov decision processes. In *International Conference on Machine Learning*, pp. 5478–5486. PMLR, 2019.
- Sidford, A. and Zhang, C. Quantum speedups for stochastic optimization. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Wan, Z., Zhang, Z., Li, T., Zhang, J., and Sun, X. Quantum multi-armed bandits and stochastic linear bandits enjoy logarithmic regrets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 10087–10094, 2023.
- Wang, D., You, X., Li, T., and Childs, A. M. Quantum exploration algorithms for multi-armed bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 10102–10110, 2021.
- Wei, C.-Y., Jahromi, M. J., Luo, H., Sharma, H., and Jain, R. Model-free reinforcement learning in infinite-horizon average-reward markov decision processes. In *International conference on machine learning*, pp. 10170–10180. PMLR, 2020.

- Weissman, T., Ordentlich, E., Seroussi, G., Verdu, S., and Weinberger, M. J. Inequalities for the ℓ_1 deviation of the empirical distribution. *Hewlett-Packard Labs, Tech. Rep*, 2003.
- Wiedemann, S., Hein, D., Udfluft, S., and Mendl, C. B. Quantum policy iteration via amplitude estimation and grover search—towards quantum advantage for reinforcement learning. *Transactions on Machine Learning Research*, 2022.
- Wu, Y., Guan, C., Aggarwal, V., and Wang, D. Quantum heavy-tailed bandits. *arXiv preprint arXiv:2301.09680*, 2023.
- Xu, Y. and Aggarwal, V. Accelerating quantum reinforcement learning with a quantum natural policy gradient based approach. In *Forty-second International Conference on Machine Learning*, 2025.
- Zhong, H., Hu, J., Xue, Y., Li, T., and Wang, L. Provably efficient exploration in quantum reinforcement learning with logarithmic worst-case regret. In *Forty-first International Conference on Machine Learning*, 2024.

A. Appendix A: Convexity of Optimization Problem

In order to solve the optimistic optimization problem in Eq. (15) - (19) for epoch e , we adopt the approach proposed in (Agarwal et al., 2022b; Rosenberg & Mansour, 2019). The key is that the constraint in (18) seems non-convex. However, that can be made convex with an inequality as

$$\sum_{a \in \mathcal{A}} \rho(s', a) \leq \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} P_e(s' | s, a) \rho(s, a), \quad \forall s' \in \mathcal{S} \quad (41)$$

This makes the problem convex since $xy \geq c$ is convex region in the first quadrant.

However, since the constraint region is made bigger, we need to show that the optimal solution satisfies (18) with equality. This follows since

$$\sum_{s' \in \mathcal{S}} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} P_e(s' | s, a) \rho(s, a) = 1 = \sum_{s' \in \mathcal{S}} \sum_{a \in \mathcal{A}} \rho(s', a) \quad (42)$$

Thus, since we have point-wise inequality in (41) with the sum of the two sides being the same, thus, the equality holds for all s' .

B. Appendix B: A brief overview of Quantum Amplitude Amplification

Quantum amplitude amplification stands as a critical outcome within the realm of quantum mean estimation. This pivotal concept empowers the augmentation of the amplitude associated with a targeted state, all the while suppressing the amplitudes linked to less desirable states. The pivotal operator, denoted as Q , embodies the essence of amplitude amplification: $Q = 2|\psi\rangle\langle\psi| - I$, where $|\psi\rangle$ signifies the desired state and I represents the identity matrix. This operator orchestrates a double reflection of the state $|\psi\rangle$ - initially about the origin and subsequently around the hyperplane perpendicular to $|\psi\rangle$ - culminating in a significant augmentation of the amplitude of $|\psi\rangle$.

Moreover, the application of this operator can be iterated to achieve a repeated amplification of the desired state's amplitude, effectively minimizing the amplitudes of undesired states. Upon applying the amplitude amplification operator a total of t times, the outcome is Q^t , which duly enhances the amplitude of the desired state by a scaling factor of \sqrt{N}/M , where N corresponds to the overall count of states and M signifies the number of solutions.

The implications of quantum amplitude amplification extend across a diverse spectrum of applications within quantum algorithms. By bolstering their efficacy and hastening the resolution of intricate problems, quantum amplitude amplification carves a substantial niche. Noteworthy applications encompass:

1. **Quantum Search:** In the quantum algorithm for unstructured search, known as Grover's algorithm (Grover, 1996), the technique of amplitude amplification is employed to enhance the amplitude of the target state and decrease the amplitude of the non-target states. As a result, this technique provides a quadratic improvement over classical search algorithms.
2. **Quantum Optimization:** Quantum amplification can be used in quantum optimization algorithms, such as Quantum Approximate Optimization Algorithm (QAOA) (Farhi et al., 2014), to amplify the amplitude of the optimal solution and reduce the amplitude of sub-optimal solutions. This can lead to an exponential speedup in solving combinatorial optimization problems.
3. **Quantum Simulation:** Quantum amplification has also been used in quantum simulation algorithms, such as Quantum Phase Estimation (QPE) (Kitaev, 1995), to amplify the amplitude of the eigenstate of interest and reduce the amplitude of other states. This can lead to an efficient simulation of quantum systems, which is an intractable problem in the classical domain.
4. **Quantum Machine Learning:** The utilization of quantum amplification is also evident in quantum machine learning algorithms, including Quantum Support Vector Machines (QSVM) (Lloyd et al., 2013). The primary objective is to amplify the amplitude of support vectors while decreasing the amplitude of non-support vectors, which may result in an exponential improvement in resolving particular classification problems.

In the domain of quantum Monte Carlo methods, implementing quantum amplitude amplification can enhance the algorithm's efficiency by decreasing the amount of samples needed to compute the integral or solve the optimization problem. By amplifying the amplitude of the desired state, a substantial reduction in the number of required samples can be achieved while still maintaining a certain level of accuracy, leading to significant improvements in efficiency when compared to classical methods. This technique has been particularly useful in (Hamoudi, 2021) for achieving efficient convergence rates in quantum Monte Carlo simulations.

C. Appendix C: Proof of Lemma 3

Lemma 7 (Lemma 3 in the main paper). *Execution of Q -UCRL (Algorithm 2) up to the beginning of epoch $e + 1$ with total e completed epochs comprising $t = 1, 2, \dots, t_e$ rounds, guarantees that the following holds for \hat{P}_{e+1} :*

$$\mathbb{P}[|\hat{P}_{e+1}(s'|s, a) - P(s'|s, a)| \leq \frac{eC \log(eS/\delta)}{N_{e+1}(s, a)}] \geq 1 - \delta \quad (43)$$

$\forall (s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$, where $C = c \log^{1/2}(T\sqrt{S})$ for some $c \in \mathbb{R}$ and $\{\nu_e(s, a)\}$ are the state-action pair visitation counts up in epoch e , as depicted in Line 12 of Algorithm 2, and $\{P(s'|s, a), \hat{P}_{e+1}(s'|s, a)\}$ are the actual and weighted average of transition probabilities by Eq. (22).

Proof. For all $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$, $\hat{P}_{e+1}(s'|s, a)$ could be equivalently expressed as:

$$\hat{P}_{e+1}(s'|s, a) = \frac{\sum_{i=1}^e \nu_i(s, a) \tilde{P}_i(s'|s, a)}{\sum_{j=1}^e \nu_j(s, a)} \quad (44)$$

Denote event $|\tilde{P}_i(s'|s, a) - P(s'|s, a)| \leq \frac{C \log(S/\delta)}{\nu_e(s, a)}$ as event E_i . Thus if events E_1, \dots, E_e all occur, we would have:

$$|\hat{P}_{e+1}(s'|s, a) - P(s'|s, a)| \quad (45)$$

$$= \left| \frac{\sum_{i=1}^e \nu_i(s, a) (\tilde{P}_i(s'|s, a) - P(s'|s, a))}{\sum_{j=1}^e \nu_j(s, a)} \right|, \quad (46)$$

$$\leq \frac{\sum_{i=1}^e \nu_i(s, a) |\tilde{P}_i(s'|s, a) - P(s'|s, a)|}{\sum_{j=1}^e \nu_j(s, a)}, \quad (47)$$

$$\leq \frac{\sum_{i=1}^e C \log(S/\delta)}{\sum_{j=1}^e \nu_j(s, a)}, \quad (48)$$

$$= \frac{eC \log(S/\delta)}{\sum_{j=1}^e \nu_j(s, a)}, \quad (49)$$

$$= \frac{eC \log(S/\delta)}{N_{e+1}(s, a)}, \quad (50)$$

Where Eq. (47) is the result of triangle inequality and Eq. (48) is by Lemma 2. The probability of events E_1, \dots, E_e all occurring is:

$$\mathbb{P}[E_1 \cap \dots \cap E_e], \quad (51)$$

$$= 1 - \mathbb{P}[E_1^C \cup \dots \cup E_e^C], \quad (52)$$

$$\geq 1 - \sum_{i=1}^e \mathbb{P}[E_i^C], \quad (53)$$

$$\geq 1 - e\delta \quad (54)$$

Where Eq. (54) is by Lemma 2. By combining Eq. (50) and Eq. (54) and substituting $e\delta$ with δ , we would obtain Eq. (43) \square

D. Appendix D: Proof of Lemma 4

Lemma 8 (Lemma 4 in the main paper). *Execution of Q -UCRL (Algorithm 2) up to total e epochs comprising $t = 1, 2, \dots, t_e$ rounds, guarantees that the following holds for confidence parameter $\delta_e = \frac{1}{S^2 A t_e^7}$ with probability at least $1 - 1/t_e^6$:*

$$\begin{aligned} & \|P_e(\cdot|s, a) - P(\cdot|s, a)\|_1, \\ & \leq \frac{7S(1 + Ce) \log(S^2 A t_e) + S C e \log(eS)}{\max\{1, N_e(s, a)\}}, \end{aligned} \quad (55)$$

$\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, where $\{N_{e-1}^{(s,a)}\}$ is the cumulative state-action pair visitation counts up to epoch e , as depicted in Line 17 of Algorithm 2, and $\{P(\cdot|s, a), P_e(\cdot|s, a)\}$ are the actual and estimated transition probabilities.

Proof. To prove the claim in Eq. (55), note that for an arbitrary pair (s, a) we have:

$$\begin{aligned} & \|P_e(\cdot|s, a) - P(\cdot|s, a)\|_1 \\ & \leq \|P_e(\cdot|s, a) - \hat{P}_e(\cdot|s, a)\|_1 \\ & \quad + \|\hat{P}_e(\cdot|s, a) - P(\cdot|s, a)\|_1, \end{aligned} \quad (56)$$

$$\leq \frac{7S \log(S^2 A t_e)}{\max\{1, N_e(s, a)\}} + \|\hat{P}_e(\cdot|s, a) - P(\cdot|s, a)\|_1, \quad (57)$$

$$\begin{aligned} & = \frac{7S \log(S^2 A t_e)}{\max\{1, N_e(s, a)\}} \\ & \quad + \sum_{s' \in \mathcal{S}} \underbrace{|\hat{P}_e(s'|s, a) - P(s'|s, a)|}_{(a)}, \end{aligned} \quad (58)$$

where Eq. (57) is a direct consequence of constraint Eq. (19) of optimization problem Eq. (15). Next, in order to analyze (a) in Eq. (58), consider the following event:

$$\begin{aligned} \mathcal{E} & = \left\{ |\hat{P}_e(s'|s, a) - P(s'|s, a)| \right. \\ & \quad \left. \leq \frac{7 \log(S^2 A t_e)}{\max\{1, N_e(s, a)\}}, \forall (s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \right\}, \end{aligned} \quad (59)$$

For a tuple (s, a, s') , by replacing $\delta = e^{-\epsilon \cdot N_e(s, a)}$ into Eq. (1) from Lemma 1, we get:

$$\begin{aligned} & \mathbb{P} \left[|\hat{P}_e(s'|s, a) - P(s'|s, a)| \leq \frac{C e \log(eS)}{N_e(s, a)} + C e \epsilon \right] \\ & \geq 1 - e^{-\epsilon \cdot N_e(s, a)}, \end{aligned} \quad (60)$$

By plugging in $\epsilon = \frac{\log(S^2 A t_e^7)}{N_e(s, a)}$ into Eq. (60), we get:

$$\begin{aligned} & \mathbb{P} \left[|\hat{P}_e(s'|s, a) - P(s'|s, a)| \geq \frac{C e \log(eS) + C e \log(S^2 A t_e^7)}{N_e(s, a)} \right] \\ & \leq \frac{1}{S^2 A t_e^7}, \end{aligned} \quad (61)$$

Finally, we sum over all possible values of $N_e(s, a)$ as well as tuples (s, a, s') to bound the probability in the RHS of (61) that implies failure of event \mathcal{E} :

$$\sum_{(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}} \sum_{N_e(s, a)=1}^{t_e} \frac{1}{S^2 A t_e^7} \leq \frac{1}{t_e^6} \quad (62)$$

Here, it is critical to highlight that we obtain this high probability bound for model errors by careful utilization of Lemma 1 and quantum signals $|\psi\rangle^t$ via Eq. (55) - (62), while avoiding classical concentration bounds for probability norms that use martingale style analysis such as in the work (Weissman et al., 2003). Finally, plugging the bound of term (a) as described by event \mathcal{E} , we obtain the following with probability $1 - 1/t_e^6$:

$$\begin{aligned} & \|P_e(\cdot|s, a) - P(\cdot|s, a)\|_1 \\ & \leq \frac{7S \log(S^2 A t_e)}{\max\{1, N_e(s, a)\}} + \sum_{s' \in \mathcal{S}} \frac{C e \log(eS) + C e \log(S^2 A t_e^7)}{\max\{1, N_e(s, a)\}}, \end{aligned} \quad (63)$$

$$\leq \frac{7S \log(S^2 A t_e)}{\max\{1, N_e(s, a)\}} + \sum_{s' \in \mathcal{S}} \frac{C e \log(eS) + 7C e \log(S^2 A t_e)}{\max\{1, N_e(s, a)\}}, \quad (64)$$

$$= \frac{7S(1 + C e) \log(S^2 A t_e) + S C e \log(eS)}{\max\{1, N_e(s, a)\}}, \quad (65)$$

$$\leq \frac{7S(1 + C e) \log(S^2 A T) + S C e \log(eS)}{\max\{1, N_e(s, a)\}}. \quad (66)$$

This proves the claim as in the statement of the Lemma. \square

E. Appendix E: Proof of Lemma 5

Lemma 9 (Regret Decomposition, Same as Lemma 5 in the main paper). *Regret of Q -UCRL defined in Eq. (10) can be decomposed as follows:*

$$\mathcal{R}_{[0,T-1]} \leq \sum_{e=1}^E T_e \left(\Gamma_{\pi_e}^{\tilde{P}_e} - \Gamma_{\pi_e}^P \right). \quad (67)$$

Proof.

$$\mathcal{R}_{[0,T-1]} = T \times \Gamma_{\pi^*}^P - \mathbb{E} \left[\sum_{t=0}^{T-1} r(s_t, a_t) \right], \quad (68)$$

$$\begin{aligned} &= T \times \Gamma_{\pi^*}^P - \sum_{e=1}^E T_e \Gamma_{\pi_e^*}^P + \sum_{e=1}^E T_e \Gamma_{\pi_e^*}^P - \mathbb{E} \left[\sum_{t=0}^{T-1} r(s_t, a_t) \right], \\ &= \sum_{e=1}^E T_e \underbrace{\left(\Gamma_{\pi^*}^P - \Gamma_{\pi_e^*}^P \right)}_{(a)} + \sum_{e=1}^E T_e \Gamma_{\pi_e^*}^P - \mathbb{E} \left[\sum_{t=0}^{T-1} r(s_t, a_t) \right], \end{aligned} \quad (69)$$

$$= \sum_{e=1}^E T_e \Gamma_{\pi_e^*}^P - \mathbb{E} \left[\sum_{t=0}^{T-1} r(s_t, a_t) \right], \quad (70)$$

$$\leq \sum_{e=1}^E T_e \Gamma_{\pi_e}^{P_e} - \mathbb{E} \left[\sum_{t=0}^{T-1} r(s_t, a_t) \right], \quad (71)$$

$$= \sum_{e=1}^E T_e \Gamma_{\pi_e}^{P_e} - \sum_{e=1}^E T_e \Gamma_{\pi_e}^P + \sum_{e=1}^E T_e \Gamma_{\pi_e}^P - \mathbb{E} \left[\sum_{t=0}^{T-1} r(s_t, a_t) \right], \quad (72)$$

$$= \sum_{e=1}^E T_e \Gamma_{\pi_e}^{P_e} - \sum_{e=1}^E T_e \Gamma_{\pi_e}^P + \left[\sum_{e=1}^E \sum_{t=t_{e-1}+1}^{t_e} \underbrace{\left(\Gamma_{\pi_e}^P - \mathbb{E}[r(s_t, a_t)] \right)}_{(b)} \right], \quad (73)$$

$$= \sum_{e=1}^E T_e \left(\Gamma_{\pi_e}^{\tilde{P}_e} - \Gamma_{\pi_e}^P \right). \quad (74)$$

□

Firstly, (a) in Eq. (69) is 0 owing to the fact that the agent would have gotten π_e^* exactly as the true π^* in epoch e solving for original problem Eq. (11), had it known the true transition model P . Next, Eq. (71) is because $\{P_e, \pi_e\}$ are outputs from the optimistic problem Eq. (15) solved at round e , which extends the search space of Eq. (11) and therefore always upper bounds true long-term rewards. Finally, term (b) in Eq. (73) is 0 in expectation, because the RL agent actually collects rewards by playing policy π_e against the true P during epoch e .

F. Appendix F: Some Auxiliary Lemmas used in this Paper

Lemma 10 (Lemma 5.3, (Agarwal et al., 2022b)). *The Bellman errors for any arbitrary state-action pair (s, a) corresponding to epoch e is bounded as follows:*

$$B^{\pi_e, P_e}(s, a) \leq \left\| P_e(\cdot | s, a) - P(\cdot | s, a) \right\|_1 \|\tilde{h}(\cdot)\|_\infty, \quad (75)$$

where $\tilde{h}(\cdot)$ (defined in (Puterman, 2014)) is the bias of the optimistic MDP P_e for which the optimistic policy π_e obtains the maximum expected reward in the confidence interval.

Lemma 11. [Lemma D.5, (Agarwal et al., 2022b)] *For an MDP with the transition model P_e that generates rewards $r(s, a)$ on playing policy π_e , the difference of biases for states s and s' are bounded as:*

$$\tilde{h}(s) - \tilde{h}(s') \leq T_{\max}, \forall s, s' \in \mathcal{S}. \quad (76)$$

G. Appendix G: Proof of Theorem 1

Theorem 2 (Theorem 1 in main paper). *In an unknown MDP environment $\mathcal{M} \triangleq (\mathcal{S}, \mathcal{A}, P, r, D)$, the regret incurred by Q -URL (Algorithm 2) across T rounds is bounded as follows:*

$$\mathcal{R}_{[0, T-1]} = \mathcal{O}\left(S^5 A^4 T_{\text{mix}} \log^3\left(\frac{T}{SA}\right) \log^{1/2}(T\sqrt{S}) \log(S^2 AT)\right). \quad (77)$$

Proof. Using the final expression from regret decomposition, i.e., Eq. (74), we have:

$$\mathcal{R}_{[0, T-1]} = \sum_{e=1}^E T_e \left(\Gamma_{\pi_e}^{\tilde{P}_e} - \Gamma_{\pi_e}^P \right), \quad (78)$$

$$= \frac{1}{T} \sum_{e=1}^E \sum_{t=t_{e-1}+1}^{t_e} \sum_{s,a} \rho_{\pi_e}^P B^{\pi_e, P_e}(s, a), \quad (79)$$

$$\leq \sum_{e=1}^E \sum_{t=t_{e-1}+1}^{t_e} \frac{7S(1+Ce) \log(S^2 AT) + SCe \log(eS)}{\max\{1, N_e(s_t, a_t)\}} T_{\text{mix}}, \quad (80)$$

$$= \sum_{e=1}^E \sum_{t=t_{e-1}+1}^{t_e} \sum_{s,a} \left(\mathbf{1}[s_t = s, a_t = a] \cdot \frac{7S(1+Ce) \log(S^2 AT) + SCe \log(eS)}{\max\{1, N_e(s, a)\}} T_{\text{mix}} \right), \quad (81)$$

$$= \sum_{e=1}^E \sum_{s,a} \nu_e(s, a) \cdot \frac{7S(1+Ce) \log(S^2 AT) + SCe \log(eS)}{\max\{1, N_e(s, a)\}} T_{\text{mix}}, \quad (82)$$

$$= \sum_{s,a} T_{\text{mix}} (7S(E + CE^2) \log(S^2 AT) + SCE^2 \log(ES)) \cdot \sum_{e=1}^E \frac{\nu_e(s, a)}{\max\{1, N_e(s, a)\}}, \quad (83)$$

$$\begin{aligned} &\leq \sum_{s,a} T_{\text{mix}} (7S(E + CE^2) \log(S^2 AT) + SCE^2 \log(ES)) E, \\ &= T_{\text{mix}} E (7S^2 A (E + CE^2) \log(S^2 AT) + S^2 ACE^2 \log(ES)), \\ &= \mathcal{O}\left(S^5 A^4 T_{\text{mix}} \log^3\left(\frac{T}{SA}\right) \log^{1/2}(T\sqrt{S}) \log(S^2 AT)\right). \end{aligned} \quad (84)$$

where in Eq. (79), we directly incorporate the result of Lemma 6. Next, we obtain Eq. (80) by using the Bellman error bound in Eq. (39), as well as unit sum property of occupancy measures (Puterman, 2014). Furthermore, in Eq. (80) note

that we have omitted the regret contribution of *bad event* (Remark 1) which is atmost $\tilde{\mathcal{O}}\left(\frac{2ST_{\text{mix}}}{T^4}\right)$ as obtained below:

$$\sum_{e=1}^E \sum_{t=t_{e-1}+1}^{t_e} \frac{2ST_{\text{mix}}}{t_e^6} \leq \sum_{e=1}^E \frac{2ST_{\text{mix}}}{t_e^5}, \quad (85)$$

$$\leq \sum_{t_e=1}^T \frac{2ST_{\text{mix}}}{t_e^5} = \tilde{\mathcal{O}}\left(\frac{2ST_{\text{mix}}}{T^4}\right). \quad (86)$$

Subsequently, Eq. (82) - (83) are obtaining by using the definition of epoch-wise state-action pair visitation frequencies $\nu_e(s, a)$ and the epoch termination trigger condition in Line 13 of Algorithm 2. Finally, we obtain Eq. (84) by the using Proposition 18 in Appendix C.2 of (Auer et al., 2008). \square