

Primal-Dual π Learning: Sample Complexity and Sublinear Run Time for Ergodic Markov Decision Problems

Mengdi Wang

Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ
email: mengdiw@princeton.edu

October 17, 2017

Abstract

Consider the problem of approximating the optimal policy of a Markov decision process (MDP) by sampling state transitions. In contrast to existing reinforcement learning methods that are based on successive approximations to the nonlinear Bellman equation, we propose a Primal-Dual π Learning method in light of the linear duality between the value and policy. The π learning method is model-free and makes primal-dual updates to the policy and value vectors as new data are revealed. For infinite-horizon undiscounted Markov decision process with finite state space \mathcal{S} and finite action space \mathcal{A} , the π learning method finds an ϵ -optimal policy using the following number of sample transitions

$$\tilde{O}\left(\frac{(\tau \cdot t_{mix}^*)^2 |\mathcal{S}| |\mathcal{A}|}{\epsilon^2}\right),$$

where t_{mix}^* is an upper bound of mixing times across all policies and τ is a parameter characterizing the range of stationary distributions across policies. The π learning method also applies to the computational problem of MDP where the transition probabilities and rewards are explicitly given as the input. In the case where each state transition can be sampled in $\tilde{O}(1)$ time, the π learning method gives a sublinear-time algorithm for solving the averaged-reward MDP.

Keywords: Markov decision process, reinforcement learning, sample complexity, run-time complexity, duality, primal-dual method, mixing time

1 Introduction

Consider the reinforcement learning problem in which a planner makes decisions in an unknown (sometimes stochastic) dynamic environment with the goal of maximizing the reward collected in this process. This can be modeled as a Markov decision process (MDP). MDP refers to a controlled random walk in which the planner chooses one from a number of actions at each state of the random walk and moves to another state according to some transition probability distribution. In the context of reinforcement learning, one wants to learn the optimal decision rule by using an algorithmic trial-and-error approach, without explicitly knowing the transition probabilities.

We focus on the infinite-horizon Average-reward Markov Decision Problem (AMDP) in which one aims to make an infinite sequence of decisions and optimize the average-per-time-step reward. An instance of the AMDP can be described by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$, where \mathcal{S} is a finite state space of size $|\mathcal{S}|$, \mathcal{A} is a finite action space of size $|\mathcal{A}|$, \mathcal{P} is the collection of state-to-state transition probabilities $\mathcal{P} = \{p_{ij}(a) \mid i, j \in \mathcal{S}, a \in \mathcal{A}\}$, \mathbf{r} is the collection of state-transitional

rewards $\mathbf{r} = \{r_{ij}(a) \mid i, j \in \mathcal{S}, a \in \mathcal{A}\}$ where $r_{ij}(a) \in [0, 1]$. We also denote by \mathbf{r}_a the vector of expected state-transition rewards under action a , where $\mathbf{r}_{a,i} = \sum_{j \in \mathcal{S}} p_{ij}(a)r_{ij}(a)$. Suppose that the decision process is in state i , if action a is selected, the process moves to a next state j with probability $p_{ij}(a)$ and generates a reward $r_{ij}(a)$.

We want to find a stationary policy that specifies which action to choose at each state (regardless of the time step). A stationary and randomized policy can be represented by a collection of probability distributions $\pi = \{\pi_i\}_{i \in \mathcal{S}}$, where $\pi_i : \mathcal{A} \mapsto [0, 1]$ is a vector of probability distribution over actions. We denote by P^π the transition probability matrix of the AMDP under a fixed policy π , where $P_{ij}^\pi = \sum_{a \in \mathcal{A}} \pi_i(a)p_{ij}(a)$ for all $i, j \in \mathcal{S}$. The objective of the AMDP is to find an optimal policy π^* such that the infinite-horizon average reward is maximized:

$$\max_{\pi} \lim_{T \rightarrow \infty} \mathbf{E}^\pi \left[\frac{1}{T} \sum_{t=1}^T r_{i_t i_{t+1}}(a_t) \right],$$

where $\{i_0, a_0, i_1, a_1, \dots, i_t, a_t, \dots\}$ are state-action transitions generated by the Markov decision process under the fixed policy π , and the expectation $\mathbf{E}^\pi[\cdot]$ is taken over the entire trajectory.

Let us emphasize our focus on the *undiscounted average-reward* MDP. This is contrary to the majority of existing literatures that focus on the discounted cumulative reward problems, i.e., $\max_{\pi} \mathbf{E}^\pi \left[\sum_{t=1}^{\infty} \gamma^t r_{i_t i_{t+1}}(a_t) \right]$ where $\gamma \in (0, 1)$ is a pre-specified discount factor. The discount factor γ is imposed artificially for analytical purposes. It ensures contractive properties of the Bellman operator and geometric convergence of value and policy iterations. It also plays an important role in the sample and run-time complexity analysis for MDP algorithms and reinforcement learning methods. However, discounted MDP are indeed approximations to infinite-horizon undiscounted MDPs. In this paper, we attempt to obsolete the discount factor. Instead of assuming that future rewards are discounted, we focus on the undiscounted MDP that satisfies certain fast mixing property and stationary properties. The lack of a discount factor significantly complicates our analysis.

Let us focus on sampling-based methods for the AMDP. Suppose that $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$, is not explicitly given. Instead, it is possible to interact with the real-time decision process (or a simulated process) by trying different controls and observing states transitions and rewards. In particular, suppose that we are given a Sampling Oracle (\mathcal{SO}), which takes a state-action pair (i, a) as input and outputs a random future state j and reward $r_{ij}(a)$ with probability $p_{ij}(a)$. Such a \mathcal{SO} is known as the generative model in the literatures of reinforcement learning [16, 15].

In this paper, we propose a model-free policy learning method for solving the AMDP, which we refer to as Primal-Dual π Learning (π learning for short). It is motivated by a recently developed randomized primal-dual method for solving the discounted MDP [25]. The π learning method maintains a randomized policy for controlling the MDP and dynamically updates the policy and an auxiliary value vector as new observations are revealed. The method is based on a primal-dual iteration which is crafted to take advantage of the linear algebraic structures of the nonlinear Bellman equation. The π learning method is remarkably computational efficient - it uses $\mathcal{O}(|\mathcal{S}||\mathcal{A}|)$ space and $\tilde{\mathcal{O}}(1)$ arithmetic operations per update.¹ It is *model-free* in the sense that it directly updates the policy and value vectors without estimating the transition probabilities of the MDP model. We show that the π learning method finds an ϵ -optimal policy with probability $1 - \delta$ using the following sample complexity (number of queries to the \mathcal{SO}):

$$\tilde{\mathcal{O}} \left(\frac{(\tau \cdot t_{mix}^*)^2 |\mathcal{S}||\mathcal{A}|}{\epsilon^2} \log \left(\frac{1}{\delta} \right) \right),$$

¹We use $\mathcal{O}(1)$ to hide constant factors and use $\tilde{\mathcal{O}}(1)$ to hide polylog factors of $|\mathcal{S}|, |\mathcal{A}|, \epsilon$.

where τ is parameter that characterizes the range of stationary distributions across policies, and t_{mix}^* is an uniform upper bound of the mixing times of the Markov decision process under any stationary policy. This sample complexity is optimal in its dependence on $\frac{|\mathcal{S}||\mathcal{A}|}{\epsilon^2}$.

When the MDP model $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$ is explicitly given, the proposed π learning method can be used as a randomized algorithm to compute an ϵ -optimal policy. Given $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$ as the input, one can implement \mathcal{SO} using binary-tree data structures using $\mathcal{O}(|\mathcal{S}|^2|\mathcal{A}|)$ preprocessing time, such that each query to the \mathcal{SO} takes $\mathcal{O}(1)$ time [25]. In this setting, the π -learning method outputs an ϵ -optimal policy with probability $1 - \delta$ in run time $\tilde{\mathcal{O}}\left(\frac{(\tau \cdot t_{mix}^*)^2 |\mathcal{S}||\mathcal{A}|}{\epsilon^2} \log\left(\frac{1}{\delta}\right)\right)$. This is a sublinear run time in comparison with the input size $\mathcal{O}(|\mathcal{S}|^2|\mathcal{A}|)$, as long as $\epsilon \ll \sqrt{S}$.

To the author’s best knowledge, this is the first model-free learning method for infinite-horizon average-reward MDP problems that is based on a primal-dual iteration. Our sample complexity result is a first result that characterizes the role of the mixing time and range of stationary distributions, without assuming any discount factor or finite horizon. We also provide the first sublinear run-time result for approximately solving AMDP using randomization.

Outline Section 2 surveys existing model-free learning methods for MDP and their sample and run-time complexity guarantees. Section 3 states the main assumptions on the ergodic Markov decision processes, the Bellman equation and its linear programming formulations. Section 4 develops the Primal-Dual π Learning method from a saddle point formulation of the Bellman equation. Section 5 establishes the convergence analysis and sample complexity of exploration of the Primal-Dual π Learning method. Section 6 gives a summary.

Notations All vectors are considered as column vectors. For a vector $\mathbf{x} \in \mathfrak{R}^n$, we denote by x_i or $x(i)$ its i -th component, denote by \mathbf{x}^\top its transpose, and denote by $\|\mathbf{x}\| = \sqrt{\mathbf{x}^\top \mathbf{x}}$ its Euclidean norm. We denote by $\mathbf{1} = (1, \dots, 1)^\top$ the vector with all entries equaling 1, and we denote by $\mathbf{1}_i$ the vector with its i -th entry equaling 1 and other entries equaling 0. For a positive number x , we denote by $\log x$ the natural logarithm of x . For two probability distributions p, q over a finite set X , we denote by $D_{KL}(p||q)$ their Kullback-Leibler divergence, i.e., $D_{KL}(p||q) = \sum_{x \in X} p(x) \log \frac{p(x)}{q(x)}$.

2 Related Literatures

There are two major notions of complexity for MDP: the run-time complexity and the sample complexity. The run-time complexity is critical to the computational problem where the MDP model is fully specified. It is measured by the total number of arithmetic operations performed by an algorithm. The sample complexity is critical to the reinforcement learning problem where the MDP model is unknown but a sampling oracle (\mathcal{SO}) is given. It is measured by the total number of queries to \mathcal{SO} made by an algorithm. Most existing literatures focus on either one of the two notions. They were considered as disjoint topics for years of research.

The computational complexity of MDP has been studied mainly in the setting where the MDP model is fully specified as the input. Three major deterministic approaches are the value iteration method [2, 24, 18], the policy iteration method [12, 19, 28, 22], and linear programming methods [11, 10, 28, 22, 27]. These deterministic methods inevitably require solving large linear systems. In order to compute the optimal policy exactly or to find an ϵ -optimal policy in $\tilde{\mathcal{O}}(\text{poly}(|\mathcal{S}||\mathcal{A}|) \log(\frac{1}{\epsilon}))$ time, these methods all require linear or superlinear time, i.e., the number of arithmetic operations needed is at least linear in the input size $\mathcal{O}(|\mathcal{S}|^2|\mathcal{A}|)$. For more detailed surveys on the exact solution methods for MDP, we refer the readers to the textbooks [3, 5, 21, 4] and the references therein.

Randomized versions of the classical methods have proved to achieve faster run time when $|\mathcal{S}|, |\mathcal{A}|$ are very large. Examples include the randomized primal-dual method by [25] and the variance-reduced randomized value iteration methods by [23]; both apply to the discounted MDP. These methods involve simulating the Markov decision processes and making randomized updates. As long as the input is given in suitable data structures that enable $\mathcal{O}(1)$ -time sampling, these results suggest that it is possible to compute an approximate policy for the discounted MDP in sublinear time $\tilde{\mathcal{O}}(\frac{|\mathcal{S}||\mathcal{A}|}{\epsilon^2})$ (ignoring other parameters). On the other hand, [7] recently showed that the run-time complexity for any randomized algorithm is $\Omega(|\mathcal{S}|^2|\mathcal{A}|)$ for the discounted MDP. In the case where each transition can be sampled in $\tilde{\mathcal{O}}(1)$ time, [7] showed that any randomized algorithm needs $\Omega(\frac{|\mathcal{S}||\mathcal{A}|}{\epsilon})$ run time to produce an ϵ -optimal policy with high probability. To the author’s best knowledge, existing results on randomized methods only apply to the discounted MDP. It remains unclear how to use randomized algorithms to efficiently approximate the optimal average-reward policy.

The sample complexity of MDP has been studied mainly in the setting of reinforcement learning. In this paper, we are given a \mathcal{SO} that generate state transitions from any specified by state-action pair. This is known as the generative model in reinforcement learning, which was introduced and studied in [16, 15]. In this setting, the sample complexity of the MDP is the number of queries to the \mathcal{SO} in order to find an ϵ -optimal policy (or ϵ -optimal value in some literatures) with high probability. One of the earliest reinforcement learning method is Q-learning, which are essentially sampling-based variants of value iteration. For infinite-horizon discounted MDP, [16] proved that phased Q-learning takes $\tilde{\mathcal{O}}(\frac{|\mathcal{S}||\mathcal{A}|}{\epsilon^2})$ sample transitions to compute an ϵ -optimal policy, where the dependence on γ is left unspecified. Azar, Munos and Kappen [1] considered a model-based value iteration method for the discounted MDP and showed that it takes $\tilde{\mathcal{O}}\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\epsilon^2}\right)$ samples to compute an ϵ -optimal value vector (not an ϵ -optimal policy). It also provided a matching sample complexity lower bound for estimating the value vector. It does not give explicit run-time complexity analysis.

We summarize existing model-free sampling-based methods for MDP and their complexity results in Table 1. Note that the settings and assumptions in this works vary from one to another. We also note that there is a large body of works on the *sample complexity of exploration* for reinforcement learning, which is the number of suboptimal time steps an algorithm performs on a single infinite-long path of the decision process before it reaches ϵ optimality; see [14]. This differs from our notion of sample complexity under the \mathcal{SO} , which is beyond our current scope. As a result, we do not include these results for comparison in Table 1.

Our proposed algorithm and analysis was partly motivated by the stochastic mirror-prox methods for solving convex-concave saddle point problems [20] and variational inequalities [13]. The idea of stochastic primal-dual update has been used to solve a specific class of minimax bilinear programs in sublinear run time [8]. For the discounted MDP, the work [26] proposed a basic stochastic primal-dual iteration without explicit complexity analysis and later [6] established a sample complexity upper bound $\mathcal{O}(\frac{|\mathcal{S}|^{4.5}|\mathcal{A}|}{\epsilon^2})$. A most relevant prior work is the author’s recent paper [25], which focused on the discounted MDP. The work [25] proposed a randomized mirror-prox method using adaptive transition sampling, which applies to a special saddle point formulation of the Bellman equation. For discounted MDP, it achieved a total runtime/sample complexity of $\tilde{\mathcal{O}}(\frac{|\mathcal{S}|^3|\mathcal{A}|}{(1-\gamma)^6\epsilon^2})$ for finding a policy π such that $\|v^\pi - v^*\|_\infty \leq \epsilon$. For discounted MDP such that the stationary distribution satisfies τ -stationarity (see Assumption 1 in the current paper), it finds an approximate policy achieving ϵ reward from a particular initial distribution with sample size/run time $\tilde{\mathcal{O}}(\tau^2 \frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^4\epsilon^2})$.

In this work, we develop the π learning method for the case of undiscounted average-reward MDP. Our approach follows from that of [25], however, our analysis is much more streamlined and

Method	Setting	Sample Complexity	Run-Time Complexity	Space Complexity	Reference
Phased Q-Learning	γ discount factor, ϵ -optimal value	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^3\epsilon^2} \ln \frac{1}{\delta}$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^3\epsilon^2} \ln \frac{1}{\delta}$	$ \mathcal{S} \mathcal{A} $	[17]
Model-Based Q-Learning	γ discount factor, ϵ -optimal value	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^3\epsilon^2} \ln \frac{ \mathcal{S} \mathcal{A} }{\delta}$	NA	$ \mathcal{S} ^2 \mathcal{A} $	[1]
Randomized P-D	γ discount factor, ϵ -optimal policy	$\frac{ \mathcal{S} ^3 \mathcal{A} }{(1-\gamma)^6\epsilon^2}$	$\frac{ \mathcal{S} ^3 \mathcal{A} }{(1-\gamma)^6\epsilon^2}$	$ \mathcal{S} \mathcal{A} $	[25]
Randomized P-D	γ discount factor, τ -stationary, ϵ -optimal policy	$\tau^4 \frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^4\epsilon^2}$	$\tau^4 \frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^4\epsilon^2}$	$ \mathcal{S} \mathcal{A} $	[25]
Randomized VI	γ discount factor, ϵ -optimal policy	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^4\epsilon^2}$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^4\epsilon^2}$	$ \mathcal{S} \mathcal{A} $	[23]
Primal-Dual π Learning	τ -stationary, t_{mix}^* -mixing, ϵ -optimal policy	$\frac{(\tau \cdot t_{mix}^*)^2 \mathcal{S} \mathcal{A} }{\epsilon^2}$	$\frac{(\tau \cdot t_{mix}^*)^2 \mathcal{S} \mathcal{A} }{\epsilon^2}$	$ \mathcal{S} \mathcal{A} $	This Paper

Table 1: Complexity Results for Sampling-Based Methods for MDP. The sample complexity is measured by the number of queries to the \mathcal{SO} . The run-time complexity is measured by the total run-time complexity under the assumption that each query takes $\tilde{O}(1)$ time. The space complexity is the additional space needed by the algorithm in addition to the input.

applies to the more general undiscounted problems. Without assuming any discount factor, we are able to characterize the complexity upperbound for infinite-horizon MDP using its mixing and stationary properties. Comparing to [25], the complexity results achieved in the current paper are much sharper, mainly due to the natural simplicity of average-reward Markov processes. To the author’s best knowledge, our results provide the first sublinear run time for solving infinite-horizon average-reward MDP without any assumption on discount factor or finite horizon.

3 Ergodic MDP, Bellman Equation, and Duality

Consider an AMDP that is described by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$. In this paper, we focus on AMDP that is ergodic (aperiodic and recurrent) under any stationary policy. For a stationary policy π , we denote by ν^π the stationary distribution of the Markov decision process which satisfies $(P^\pi)^\top \nu^\pi = \nu^\pi$. We make the following assumptions on the stationary distributions and mixing times:

Assumption 1 (Ergodic Decision Process). *The Markov decision process specified by $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$ is τ -stationary in the sense that it is ergodic under any stationary policy π and there exists $\tau > 1$ such that*

$$\frac{1}{\sqrt{\tau}|\mathcal{S}|} \mathbf{1} \leq \nu^\pi \leq \frac{\sqrt{\tau}}{|\mathcal{S}|} \mathbf{1}.$$

Assumption 1 characterizes a form of complexity of MDP in terms of the range of its stationary distributions. The factor τ characterizes a notion of complexity of ergodic MDP, i.e., the variation of stationary distributions associated with different policies. Suppose that some policies induce transient states (so the stationary distribution is not bounded away from zero). In this case, we as long as there is some policy that leads to an ergodic process, we can restrict our attention to mixture policies in order to guarantee ergodicity. In this way, we can always guarantee that Assumption 1 holds on the restricted problem at a cost of some additional approximation error.

Assumption 2 (Fast Mixing Markov Chains). *The Markov decision process specified by $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$ is t_{mix}^* -mixing in the sense that*

$$t_{mix}^* \geq \max_{\pi} \min \left\{ t \geq 1 \mid \|(P^{\pi})^t(i, \cdot) - \nu^{\pi}\|_{TV} \leq \frac{1}{4}, \forall i \in \mathcal{S} \right\},$$

where $\|\cdot\|_{TV}$ is the total variation.

Assumption 2 requires that the Markov chains be sufficiently “rapidly mixing.” The factor t_{mix}^* characterizes how fast the Markov decision process reaches its stationary distribution from any state under any policy. Our results suggest that the π learning method would work extremely well on “rapidly mixing” decision processes where t_{mix}^* is a small constant. A typical example is autonomous driving, where the previous actions get forgotten quickly. On the other hand, the current format of π learning might work poorly for problems such as the maze in which the mixing time can be very large and most policies are non-ergodic. This is to be improved.

Consider an MDP tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$ that satisfies Assumptions 1 and 2. For a fixed policy π , the average reward $\bar{v}^{\pi} > 0$ is defined as

$$\bar{v}^{\pi} \equiv \bar{v}^{\pi}(i) = \lim_{N \rightarrow \infty} \mathbf{E}^{\pi} \left[\frac{1}{N} \sum_{t=1}^N r_{i_t i_{t+1}}(a_t) \mid i_1 = i \right], \quad i \in \mathcal{S},$$

where $\mathbf{E}^{\pi}[\cdot]$ is taken over the random state-action trajectory $\{i_1, a_1, i_2, a_2, \dots\}$ generated by the Markov decision process under policy π . Note that the average reward $\bar{v}^{\pi} \equiv \bar{v}^{\pi}(i)$ is state-invariant, so we treat it as a scalar.

Bellman Equation According to the theory of dynamic programming [21, 3], the value \bar{v}^* is the optimal average reward to the AMDP \mathcal{M} if and only if it satisfies the following $|\mathcal{S}| \times (|\mathcal{S}| + 1)$ system of equations, known as the *Bellman equation*, given by

$$\bar{v}^* + h_i^* = \max_{a \in \mathcal{A}} \left\{ \sum_{j \in \mathcal{S}} p_{ij}(a) h_j^* + \sum_{j \in \mathcal{S}} p_{ij}(a) r_{ij}(a) \right\}, \quad \forall i \in \mathcal{S},$$

for some vector $\mathbf{h}^* \in \mathfrak{R}^{|\mathcal{S}|}$. A stationary policy π^* is an optimal policy of the AMDP if it attains the elementwise maximization in the Bellman equation (Theorem 8.4.5 [21]). For finite-state AMDP, there always exists at least one optimal policy π^* . If the optimal policy is unique, it is also a deterministic policy. If there are multiple optimal policies, there exist infinitely many optimal randomized policies.

Note that the preceding Bellman equation has one unique optimal solution \bar{v}^* but infinitely many solutions \mathbf{h}^* . In the remainder of this paper, we augment the Bellman equation with an additional linear equality constraint

$$(\nu^{\pi^*})^{\top} \mathbf{h}^* = 0,$$

where ν^{π^*} is the stationary distribution under policy π^* . Now the augmented Bellman equation has an unique optimal solution. We refer to such a unique solution as the *difference-of-value* vector, and we denote it by \mathbf{h}^* throughout the rest of the paper. The difference-of-value vector \mathbf{h}^* can be informally defined as

$$h_i^* = \lim_{N \rightarrow \infty} \mathbf{E}^{\pi^*} \left[\sum_{t=1}^N r_{i_t i_{t+1}}(a_t) - N \bar{v}^* \mid i_1 = i \right], \quad \forall i \in \mathcal{S}.$$

It characterizes the transient effect of each initial state under the optimal policy.

Linear Duality Of The Bellman Equation The nonlinear Bellman equation is equivalent to the following $(|\mathcal{S}| + 1) \times (|\mathcal{S}||\mathcal{A}|)$ linear programming problem (see [21] Section 8.8):

$$\begin{aligned} & \text{minimize}_{\bar{v}, \mathbf{h}} \bar{v} \\ & \text{subject to } \bar{v} \cdot \mathbf{1} + (I - P_a) \mathbf{h} - \mathbf{r}_a \geq 0, \quad \forall a \in \mathcal{A}, \end{aligned} \quad (1)$$

where $P_a \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ is the matrix whose (i, j) -th entry equals to $p_{ij}(a)$, I is the identity matrix with dimension $|\mathcal{S}| \times |\mathcal{S}|$, and $\mathbf{r}_a \in \mathbb{R}^{|\mathcal{S}|}$ is the expected state-transition reward vector under action a , i.e.,

$$r_a(i) = \sum_{j \in \mathcal{S}} p_{ij}(a) r_{ij}(a), \quad \forall i \in \mathcal{S}.$$

We associate each constraint of the primal program (1) with a dual variable $\mu_a \in \mathbb{R}^{|\mathcal{S}|}$, $a \in \mathcal{A}$. The dual linear program of (1) is

$$\begin{aligned} & \text{maximize } \sum_{a \in \mathcal{A}} \mu_a^\top \mathbf{r}_a \\ & \text{subject to } \sum_{a \in \mathcal{A}} (I - P_a^\top) \mu_a = 0, \\ & \sum_{a \in \mathcal{A}} \sum_{i \in \mathcal{S}} \mu_{a,i} = 1, \quad \mu_a \geq 0, \quad \forall a \in \mathcal{A}. \end{aligned} \quad (2)$$

It is well known that each deterministic policy of the AMDP corresponds to a basic feasible solution to the dual linear program (2). A randomized policy is a mixture of deterministic policies, so it corresponds to a feasible solution of program (2). We denote by $\mu^* = (\mu_a^*)_{a \in \mathcal{A}} \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$ the optimal solution to the dual linear program (2). If there is a unique optimal dual solution, it must be a basic feasible solution. In this case, the basis of μ^* corresponds to an optimal deterministic policy.

4 Primal-Dual π Learning

In this section, we develop the Primal-Dual π Learning Method (π learning for short). Our first step is to examine the nonlinear Bellman equation and formulate it into a bilinear saddle point problem with specially chosen primal and dual constraints. Our second step is to develop the Primal-Dual π Learning method and discuss its implementation and run-time complexity per iteration.

4.1 Saddle Point Formulation of Bellman Equation

In light of linear duality, we formulate the linear programs (1)-(2) into an equivalent minimax problem, given by

$$\min_{\bar{v}, \mathbf{h}} \max_{\mu \geq 0} \bar{v} + \sum_{a \in \mathcal{A}} \mu_a^\top (-\bar{v} \cdot \mathbf{1} + (P_a - I) \mathbf{h} + \mathbf{r}_a).$$

The minimax formulation is more preferable to the linear program formulation because it has much simpler constraints. We construct the sets \mathcal{H} and \mathcal{U} to be the *search spaces* for the value and the policy, respectively, given by

$$\mathcal{H} = \left\{ \mathbf{h} \in \mathbb{R}^{|\mathcal{S}|} \mid \|\mathbf{h}\|_\infty \leq 2t_{mix}^* \right\},$$

and

$$\mathcal{U} = \left\{ \mu = (\mu_a)_{a \in \mathcal{A}} \mid \mathbf{1}^\top \mu = 1, \mu \geq 0, \sum_{a \in \mathcal{A}} \mu_a \geq \frac{1}{\sqrt{\tau} |\mathcal{S}|} \mathbf{1} \right\}.$$

Since $\sum_{a \in \mathcal{A}} \mu_a^\top \mathbf{1} = \mu^\top \mathbf{1} = 1$, we simplify the minimax problem to

$$\min_{\mathbf{h} \in \mathcal{H}} \max_{\mu \in \mathcal{U}} \sum_{a \in \mathcal{A}} \mu_a^\top ((P_a - I)\mathbf{h} + \mathbf{r}_a). \quad (3)$$

The search space for the dual vector given by \mathcal{U} essentially reflects Assumption 1. Recall that Assumption 1 suggests that the stationary distribution of any policy belongs to a certain range, therefore it is sufficient to search for the dual variable within that range. The search space for the difference-of-value vector given by \mathcal{H} essentially reflects Assumption 2 on the fast mixing property of the MDP. The fast mixing condition implies that one can move from any state to any state within a bounded number of steps, therefore the relative difference in their values is bounded by the expected traverse time. In what follows, we verify that $\mathbf{h}^* \in \mathcal{H}$ and $\mu^* \in \mathcal{U}$ under Assumptions 1 and 2.

Lemma 1. *Under Assumptions 1 and 2, the optimal primal and dual solutions $\bar{v}^*, \mathbf{h}^*, \mu^*$ to the linear programs (1)-(2) satisfy $\bar{v}^* \in [0, 1]$, $\mathbf{h}^* \in \mathcal{H}$ and $\mu^* \in \mathcal{U}$.*

Proof. Since \bar{v}^* is the average reward under π^* and each reward per period belongs to $[0, 1]$, we obtain that $\bar{v}^* \in [0, 1]$.

Let P^* be the transition probability matrix under π^* . Let ν^* be the stationary distribution under π^* , so the difference-of-value vector \mathbf{h}^* satisfies $(\mathbf{h}^*)^\top \nu^* = 0$. Let Π be the matrix with all rows equaling to $(\nu^*)^\top$, therefore $\Pi \mathbf{h}^* = 0$. Letting $m = t_{mix}^*$, we have $\|(P^*)^m(i, \cdot) - \pi^*\|_{TV} \leq 1/4$ for all $i \in \mathcal{S}$, therefore $\|(P^*)^m - \Pi\|_\infty \leq 1/4$. We apply the relation $\mathbf{h}^* = P^* \mathbf{h}^* + \mathbf{r}^* - \bar{v}^* \cdot \mathbf{1}$ inductively for m times, use $\Pi \mathbf{h}^* = 0$ and obtain

$$\mathbf{h}^* = \sum_{k=0}^{m-1} (P^*)^k \mathbf{r}^* + (P^*)^m \mathbf{h}^* - m \bar{v}^* \cdot \mathbf{1} = \sum_{k=0}^{m-1} \left((P^*)^k \mathbf{r}^* - \bar{v}^* \cdot \mathbf{1} \right) + ((P^*)^m - \Pi) \mathbf{h}^*.$$

We take $\|\cdot\|_\infty$ on both sides of the above, use the triangle inequality and obtain

$$\|\mathbf{h}^*\|_\infty \leq \sum_{k=0}^{m-1} \|(P^*)^k \mathbf{r}^* - \bar{v}^* \cdot \mathbf{1}\|_\infty + \|(P^*)^m - \Pi\|_\infty \|\mathbf{h}^*\|_\infty \leq m + \frac{1}{4} \|\mathbf{h}^*\|_\infty$$

It follows that $\|\mathbf{h}^*\|_\infty \leq (4/3)m \leq 2t_{mix}^*$ and $\mathbf{h}^* \in \mathcal{H}$.

Recall μ^* is the optimal dual solution to the linear programs (1)-(2). The dual feasibility of μ^* suggests that $0 = \sum_{a \in \mathcal{A}} (I - P_a^\top) \mu_a^* = (I - (P^*)^\top) \sum_{a \in \mathcal{A}} \mu_a^*$, therefore $\sum_{a \in \mathcal{A}} \mu_a^*$ is the stationary distribution corresponding to the transition matrix P^* under the optimal policy π^* . It follows from Assumption 1 that $\mu^* \in \mathcal{U}$. \blacksquare

4.2 The Primal Dual π Learning Algorithm

Motivated by the minimax formulation of the Bellman equation, we propose the Primal-Dual π Learning method as follows: The π learning method makes iterative updates to a sequence of primal and dual variables $\{\mu^t, \mathbf{h}^t\}_{t=0}^T$. At the $(t+1)$ iteration, the algorithm draws a random state-action

pair (i, a) with probability $\mu_{i,a}^t$ and query the \mathcal{SO} for a state transition to a random next state j with probability $p_{ij}(a)$. Then the π learning method updates according to

$$\begin{aligned}\mu^{t+1} &= \operatorname{argmin}_{\mu \in \mathcal{U}} D_{KL}(\mu || \mu^t \cdot \exp(\Delta^{t+1})), \\ \mathbf{h}^{t+1} &= \mathbf{Proj}_{\mathcal{H}} [\mathbf{h}^t + \mathbf{d}^{t+1}],\end{aligned}\tag{4}$$

where “ \cdot ” denotes elementwise multiplication, $\mathbf{Proj}_{\mathcal{H}}$ denotes the Euclidean projection onto \mathcal{H} , $\Delta^{t+1} \in \Re^{|\mathcal{S}||\mathcal{A}|}$, $\mathbf{d}^{t+1} \in \Re^{|\mathcal{S}|}$ are random vectors generated conditioned on μ^t, \mathbf{h}^t according to

$$\begin{aligned}\Delta^{t+1} | \mathcal{F}_t &= \beta \cdot \frac{h_j^t - h_i^t + r_{ij}(a) - M}{\mu_{i,a}^t} \mathbf{e}_{i,a}, & \text{with probability } \mu_{i,a}^t, \\ \mathbf{d}^{t+1} | \mathcal{F}_t &= \alpha \cdot (\mathbf{e}_i - \mathbf{e}_j), & \text{with probability } \mu_{i,a}^t p_{i,j}(a),\end{aligned}\tag{5}$$

where we use \mathcal{F}_t to denote the collection of all random variables up to the t -th iteration. We note that Δ^{t+1} is a vector of dimension $|\mathcal{S}||\mathcal{A}|$ but it only has one single nonzero entry. Similarly, \mathbf{d}^{t+1} is a vector of dimension $|\mathcal{S}|$ but it has only two nonzero entries, whose coordinates are randomly generated by sampling a single state transition. We can easily verify that

$$\mathbf{E} [\Delta_a^{t+1} | \mathcal{F}_t] = \beta ((P_a - I)\mathbf{h}^t + \mathbf{r}_a - M \cdot \mathbf{1}), \quad a \in \mathcal{A},$$

and

$$\mathbf{E} [\mathbf{d}^{t+1} | \mathcal{F}_t] = \alpha \sum_{a \in \mathcal{A}} \mu_a^\top (I - P_a).$$

In other words, the primal and dual updates $\Delta^{t+1}, \mathbf{d}^{t+1}$ are conditionally unbiased partial derivatives of the minimax objective.

4.3 Implementations and Fast Updates

Let us consider how to implement the π learning method in order to minimize the run time per iteration. We define the auxiliary variables $\xi^t = (\xi_i^t)_{i \in \mathcal{S}}$, $\pi = (\pi_{i,a}^t)_{i \in \mathcal{S}, a \in \mathcal{A}}$ such that

$$\xi_i^t = \sum_{a \in \mathcal{A}} \mu_{i,a}^t, \quad \pi_{i,a}^t = \frac{\mu_{i,a}^t}{\xi_i^t}, \quad \mu_{i,a}^t = \xi_i^t \pi_{i,a}^t \quad \forall i \in \mathcal{S}, a \in \mathcal{A}.$$

Note that ξ^t is a vector of probability over states, and π^t is a randomized stationary policy that specifies the probability distribution for choosing actions at each given state. We implement the Primal-Dual π Learning method given by iteration (5) in Algorithm 1.

Now we analyze the computational complexity of Algorithm 1. Each iteration draws one state-action-state triplet from the \mathcal{SO} . The updates on \mathbf{h} are made to two coordinates, thus taking $\tilde{\mathcal{O}}(1)$ time. The updates on π are multiplicative, which take $\tilde{\mathcal{O}}(1)$ time if π is represented using convenient data structures like binary trees (see Prop. 1 of [25]). The updates on ξ involve information projection onto the set $\{\xi \geq \frac{1}{\sqrt{\tau}|\mathcal{S}|}\mathbf{1}, \xi^\top \mathbf{1} = 1, \xi \geq 0\}$. This can be done by maintaining and updating the shifted vector $\xi - \frac{1}{\sqrt{\tau}|\mathcal{S}|}\mathbf{1}$ using a binary-tree structure. This idea is also used in the algorithm implementation of [25]. Accordingly, Step 10 of Algorithm 1 takes $\tilde{\mathcal{O}}(1)$ run time. To sum up, each iteration of Algorithm 1 draws one sample transition and makes updates in $\tilde{\mathcal{O}}(1)$ time. The space complexity of Algorithm 1 is $\mathcal{O}(|\mathcal{S}||\mathcal{A}|)$ space, mainly to keep track of π and its running average.

Algorithm 1 Primal-Dual π Learning

- 1: **Input:** Precision level $\epsilon > 0$, \mathcal{S} , \mathcal{A} , t_{mix}^* , τ , \mathcal{SO}
 - 2: Set $\mathbf{h} = 0 \in \mathbb{R}^{|\mathcal{S}|}$, $\xi = \frac{1}{|\mathcal{S}|} \mathbf{1} \in \mathbb{R}^{|\mathcal{S}|}$, $\pi_i = \frac{1}{|\mathcal{A}|} \mathbf{1} \in \mathbb{R}^{|\mathcal{A}|}$ for all $i \in \mathcal{S}$
 - 3: Set $T = \tau^2 (t_{mix}^*)^2 |\mathcal{S}| |\mathcal{A}|$
 - 4: Set $\beta = \frac{1}{t_{mix}^*} \sqrt{\frac{\log(|\mathcal{S}| |\mathcal{A}|)}{2|\mathcal{S}| |\mathcal{A}| T}}$, $\alpha = |\mathcal{S}| t_{mix}^* \sqrt{\frac{\log(|\mathcal{S}| |\mathcal{A}|)}{2|\mathcal{A}| T}}$, $M = 4t_{mix}^* + 1$
 - 5: **for** $t = 1, 2, 3, \dots, T$ **do**
 - 6: Sample (i, a) with probability $\xi_i \pi_{i,a}$
 - 7: Sample j with probability $p_{ij}(a)$ from \mathcal{SO}
 - 8: $\Delta \leftarrow \beta \cdot \frac{h_j^t - h_i^t + r_{ij}(a) - M}{\xi_{i,a}^t \pi_{i,a}^t}$
 - 9: $h_i \leftarrow \min\{h_i + \alpha, 2t_{mix}^*\}$, $h_j \leftarrow \max\{h_j - \alpha, -2t_{mix}^*\}$,
 - 10: $\xi_i \leftarrow \xi_i + \pi_{i,a} (\exp\{\Delta\} - 1)$, $\xi \leftarrow \operatorname{argmin}_{\hat{\xi}} \left\{ D_{KL}(\hat{\xi} || \xi) \mid \mathbf{1}^\top \hat{\xi} = 1, \hat{\xi} \geq 0, \hat{\xi} \geq \frac{1}{\sqrt{\tau} |\mathcal{S}|} \mathbf{1} \right\}$
 - 11: $\pi_{i,a} \leftarrow \pi_{i,a} \cdot \exp\{\Delta\}$, $\pi_i \leftarrow \pi_i / \|\pi_i\|_1$
 - 12: $\pi^{t+1} \leftarrow \pi$
 - 13: $t \leftarrow t + 1$
 - 14: **end for**
 - 15: **Output:** $\hat{\pi} = \frac{1}{T} \sum_{t=1}^T \pi^t$
-

5 Sample Complexity and Run Time Analysis

In this section, we establish the sample complexity for the Primal-Dual π Learning method given by Algorithms 1. We also show that Algorithm 1 applies to the computation problem of MDP and gives a sublinear run-time algorithm.

5.1 Primal-Dual Convergence

Each iteration of Algorithm 1 performs a primal-dual update for the minimax problem (3). Our first result concerns the convergence of the primal-dual iteration.

Theorem 1 (Finite-Iteration Duality Gap). *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$ be an arbitrary MDP tuple satisfying Assumptions 1, 2. Then the sequence of iterates generated by Algorithm 1 satisfies*

$$\frac{1}{T} \sum_{t=1}^T \mathbf{E} \left[\sum_{a \in \mathcal{A}} (\mathbf{h}^* - P_a \mathbf{h}^* - \mathbf{r}_a)^\top \mu_a^t \right] + \bar{v}^* \leq \tilde{\mathcal{O}} \left(t_{mix}^* \sqrt{\frac{|\mathcal{S}| |\mathcal{A}|}{T}} \right),$$

where $\mu_{i,a}^t = \xi_i^t \pi_{i,a}^t$ for $i \in \mathcal{S}, a \in \mathcal{A}, t = 1, \dots, T$.

Theorem 1 establishes a finite-time error bound of a particular “duality gap.” It characterizes the level of violation of the linear complementarity condition. Our proof shares a similar spirit as that of Theorem 1 in [25]. Note that the analysis of [25] does not easily extend to the average-reward MDP and the π learning method. As a result, we have to develop a separate new convergence analysis. The complete proof is established through a series of lemmas, which we defer to Appendix.

5.2 Sample Complexity for Achieving ϵ -Optimal Policies

We have shown that the expected duality gap diminishes at a certain rate as Algorithm 1 iterates. It remains to analyze how many time steps are needed for the duality gap to become sufficiently small, and how a small duality gap would imply a near-optimal policy. We obtain the following result.

Lemma 2. For any policy π , its stationary distribution ν^π and average reward \bar{v}^π satisfies

$$\bar{v}^\pi = (\nu^\pi)^\top \sum_{a \in \mathcal{A}} \mathbf{diag}(\pi_a) ((P_a - I)\mathbf{h}^* + \mathbf{r}_a),$$

and

$$\bar{v}^* - \bar{v}^\pi = (\nu^\pi)^\top \sum_{a \in \mathcal{A}} \mathbf{diag}(\pi_a) (\bar{v}^* \cdot \mathbf{1} + (I - P_a)\mathbf{h}^* - \mathbf{r}_a).$$

Proof. Consider an arbitrary policy π . Let ν^π be the stationary distribution under policy π , so we have $(\nu^\pi)^\top P^\pi = (\nu^\pi)^\top$. Then we obtain the first result

$$\bar{v}^\pi = (\nu^\pi)^\top \mathbf{r}^\pi = (\nu^\pi)^\top ((P^\pi - I)\mathbf{h}^* + \mathbf{r}^\pi) = (\nu^\pi)^\top \sum_{a \in \mathcal{A}} \mathbf{diag}(\pi_a) ((P_a - I)\mathbf{h}^* + \mathbf{r}_a).$$

Using the fact that $(\nu^\pi)^\top \sum_{a \in \mathcal{A}} \mathbf{diag}(\pi_a) \mathbf{1} = 1$, we obtain the second result. \blacksquare

Now we are ready to show that the π learning method outputs an approximate policy whose average reward is close to the optimal average reward. Our second main result is as follows.

Theorem 2 (Sample Complexity of Single-Run π Learning (Algorithm 1)). Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$ be an arbitrary MDP tuple satisfying Assumptions 1, 2, let $\epsilon > 0$. Then by letting Algorithm 1 run for the following number of iterations/samples

$$T = \Omega \left((\tau \cdot t_{mix}^*)^2 \cdot \frac{|\mathcal{S}| |\mathcal{A}|}{\epsilon^2} \right)$$

it outputs an approximate policy $\hat{\pi}$ such that $\bar{v}^{\hat{\pi}} \geq \bar{v}^* - \epsilon$ with probability at least $2/3$.

Proof. Consider the policy given by $\hat{\pi} = \frac{1}{T} \sum_{t=1}^T \pi^t$. Note that $\frac{1}{\sqrt{T}|\mathcal{S}|} \mathbf{1} \leq \nu^{\hat{\pi}} \leq \frac{\sqrt{T}}{|\mathcal{S}|} \mathbf{1}$ (by Assumption 1) and $\frac{1}{\sqrt{T}|\mathcal{S}|} \mathbf{1} \leq \xi^t$ (since $\mu^t \in \mathcal{H}$). Then we have

$$\nu^{\hat{\pi}} \leq \frac{\sqrt{T}}{|\mathcal{S}|} \mathbf{1} = \tau \cdot \frac{1}{\sqrt{T}|\mathcal{S}|} \mathbf{1} \leq \tau \xi^t.$$

According to Lemma 2, we have

$$\begin{aligned} \bar{v}^* - \bar{v}^{\hat{\pi}} &= (\nu^{\hat{\pi}})^\top \sum_{a \in \mathcal{A}} \mathbf{diag}(\hat{\pi}_a) (\bar{v}^* \cdot \mathbf{1} + (I - P_a)\mathbf{h}^* - \mathbf{r}_a) \\ &= \frac{1}{T} \sum_{t=1}^T (\nu^{\hat{\pi}})^\top \sum_{a \in \mathcal{A}} \mathbf{diag}(\pi_a^t) (\bar{v}^* \cdot \mathbf{1} + (I - P_a)\mathbf{h}^* - \mathbf{r}_a) \\ &\leq \tau \cdot \frac{1}{T} \sum_{t=1}^T (\xi^t)^\top \sum_{a \in \mathcal{A}} \mathbf{diag}(\pi_a^t) (\bar{v}^* \cdot \mathbf{1} + (I - P_a)\mathbf{h}^* - \mathbf{r}_a) \\ &= \tau \cdot \left(\frac{1}{T} \sum_{t=1}^T \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t (\mathbf{h}^* - P_a \mathbf{h}^* - \mathbf{r}_a)_i + \bar{v} \right), \end{aligned}$$

where the inequality uses the fact $\nu^{\hat{\pi}} \leq \tau \xi^t$ for all t (due to the dual constraint \mathcal{H}) and the primal feasibility $(\bar{v}^* \cdot \mathbf{1} + (I - P_a)\mathbf{h}^* - \mathbf{r}_a) \geq 0$ for all $a \in \mathcal{A}$. We use the Markov inequality and obtain that

$$\bar{v}^* - \bar{v}^{\hat{\pi}} \leq \frac{3\tau}{2} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{E} \left[\sum_{a \in \mathcal{A}} (\mathbf{h}^* - P_a \mathbf{h}^* - \mathbf{r}_a)^\top \mu_a^t \right] + \bar{v}^* \right)$$

with probability at least $2/3$. Now if we pick $T = \Omega(\tau^2 (t_{mix}^*)^2 \frac{|\mathcal{S}| |\mathcal{A}|}{\epsilon^2})$ and apply Theorem 1, we obtain that $\bar{v}^{\hat{\pi}} \geq \bar{v}^* - \epsilon$ with probability at least $2/3$. \blacksquare

5.3 Boosting The Success Probability to $1 - \delta$

Our next aim is to achieve an ϵ -optimal policy with probability that is arbitrarily close to 1. To do this, we need to run Algorithm 1 for sufficiently many trials and pick the best outcome. This requires us be able to evaluate multiple candidate policies and select the best one out of many. In the next lemma, we show that it is possible to approximately evaluate any policy π within ϵ -precision using $\tilde{\mathcal{O}}(\frac{t_{mix}^*}{\epsilon^2})$ samples.

Lemma 3 (Approximate Policy Evaluation). *There exists an algorithm that outputs an approximate value \bar{Y} such that $\bar{\mathbf{v}}^\pi - \epsilon \leq \bar{Y} \leq \bar{\mathbf{v}}^\pi + \epsilon$ with probability at least $1 - \delta$ in $\tilde{\mathcal{O}}(\frac{t_{mix}^*}{\epsilon^2} \log(\frac{1}{\delta}))$ time steps.*

Proof. Consider the algorithm that generates a sequence of L consecutive state transitions according to the \mathcal{SO} and outputs the empirical mean reward, which we denote by \bar{Y} . Note that \bar{Y} is the empirical mean of L Markov random variables in $[0, 1]$. We apply the McDiarmid inequality for Markov chains to the L -step empirical reward \bar{Y} and obtain

$$\mathbf{P}(|\bar{Y} - \bar{v}^\pi| \geq \epsilon) \leq 2 \exp\left(-\frac{L\epsilon^2}{t_{mix}^*}\right)$$

When $L \geq \frac{t_{mix}^*}{\epsilon^2} \log(\frac{1}{\delta})$, we have $|\bar{Y} - \bar{v}^\pi| < \epsilon$ with probability at least $1 - \delta$. ■

Now we prove that by repeatedly running Algorithm 1 and using approximate policy evaluation, one can compute a near-optimal policy with probability arbitrarily close to 1. The main arguments are (1) the best policy out of multiple trials must be close-to-optimal with high probability; (2) the policy evaluation is nearly accurate with high probability, therefore the output policy (which performs the best in policy evaluation) is also close-to-optimal. Our main result is as follows.

Theorem 3 (Overall Sample Complexity). *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$ be an arbitrary MDP tuple satisfying Assumptions 1, 2 and let $\epsilon > 0$ and $\delta \in (0, 1)$ be arbitrary values. Then there exists an algorithm that draws the following number of state transitions*

$$T = \Omega\left((\tau \cdot t_{mix}^*)^2 \cdot \frac{|\mathcal{S}||\mathcal{A}|}{\epsilon^2} \log \frac{1}{\delta}\right)$$

and outputs an approximate policy $\hat{\pi}$ such that $\bar{v}^{\hat{\pi}} \geq \bar{v}^ - \epsilon$ with probability at least $1 - \delta$.*

Proof. We describe an approach that runs Algorithm 1 for multiple times in order to achieve an ϵ -optimal policy with probability $1 - \delta$:

1. We first run Algorithm 1 for K independent trials with precision parameter $\frac{\epsilon}{3}$, and we denote the output policies by $\pi^{(1)}, \dots, \pi^{(K)}$. The total running time is $K \cdot N_{\frac{\epsilon}{3}}$, where $N_{\frac{\epsilon}{3}}$ is the number of samples needed by Algorithm 1. According to Theorem 1, each trial generates an $\epsilon/3$ -optimal policy with probability at least $2/3$.
2. For each output policy $\pi^{(k)}$, we conduct approximate value evaluation for L time steps and obtain an approximate evaluation $\bar{Y}^{(k)}$ with precision level $\frac{\epsilon}{3}$ and fail probability $\frac{\delta}{2K}$. According to Lemma 3, we have

$$\bar{Y}^{(k)} - \bar{\mathbf{v}}^{\pi^{(k)}} \in \left[-\frac{\epsilon}{3}, \frac{\epsilon}{3}\right],$$

with probability at least $1 - \frac{\delta}{2K}$, and this step takes $K \cdot L = K \cdot \tilde{\mathcal{O}}(\frac{t_{mix}^*}{\epsilon^2} \log(\frac{K}{\delta}))$ time steps.

3. Output $\hat{\pi} = \pi^{(k^*)}$ such that $k^* = \operatorname{argmax}_{k=1, \dots, K} \bar{Y}^{(k)}$.

The number of samples required by the above procedure is $\tilde{\mathcal{O}}(N_{\frac{\epsilon}{3}} \log \frac{1}{\delta} + L \log \frac{1}{\delta})$. The space complexity is $\mathcal{O}(|\mathcal{S}||\mathcal{A}|)$.

Now we verify that $\hat{\pi}$ is indeed near-optimal with probability at least $1 - \delta$, as long as K is chosen appropriately. Let $\mathcal{K} = \left\{k \in [K] \mid \bar{\mathbf{v}}^{\pi^{(k)}} \geq \bar{\mathbf{v}}^* - \frac{\epsilon}{3}\right\}$, which can be interpreted as the set of successful trails of Algorithm 1. Consider the event where $\mathcal{K} \neq \emptyset$ and all policy evaluation errors belong to the small interval $[-\frac{\epsilon}{3}, \frac{\epsilon}{3}]$. In this case, we have $\bar{\mathbf{v}}^{\pi^{(k)}} - \frac{\epsilon}{3} \leq \bar{Y}^{(k)} \leq \bar{\mathbf{v}}^{\pi^{(k)}} + \frac{\epsilon}{3}$ for all k and $\mathbf{v}^{\pi^{(k)}} \geq \bar{\mathbf{v}}^* - \frac{1}{3}\epsilon$ if $k \in \mathcal{K}$. As long as \mathcal{K} is nonempty, the output policy which has the largest value of $\bar{Y}^{(k)}$ must satisfy $\bar{Y}^{(k)} \geq \bar{\mathbf{v}}^* - \frac{2}{3}\epsilon$. Since the policy evaluation error is bounded by $\frac{\epsilon}{3}$, it follows that this policy must be ϵ -optimal. We use the union bound to obtain

$$\begin{aligned} \mathbf{P}\left(\bar{\mathbf{v}}^{\hat{\pi}} < \bar{\mathbf{v}}^* - \epsilon\right) &\leq \mathbf{P}\left(\{\mathcal{K} = \emptyset\} \cup \left\{\exists k : \bar{Y}^{(k)} - \bar{\mathbf{v}}^{\pi^{(k)}} \notin \left[-\frac{\epsilon}{3}, \frac{\epsilon}{3}\right]\right\}\right) \\ &\leq \mathbf{P}(\mathcal{K} = \emptyset) + \mathbf{P}\left(\exists k : \bar{Y}^{(k)} - \bar{\mathbf{v}}^{\pi^{(k)}} \notin \left[-\frac{\epsilon}{3}, \frac{\epsilon}{3}\right]\right) \\ &\leq \prod_{k=1}^K \mathbf{P}\left(\bar{\mathbf{v}}^{\pi^{(k)}} < \bar{\mathbf{v}}^* - \frac{\epsilon}{3}\right) \\ &\quad + \sum_{k=1}^K \mathbf{P}\left(\bar{Y}^{(k)} - \bar{\mathbf{v}}^{\pi^{(k)}} \notin \left[-\frac{\epsilon}{3}, \frac{\epsilon}{3}\right]\right) \\ &\leq (1/3)^K + K \cdot \frac{\delta}{2K}. \end{aligned}$$

By choosing $K = \log(2/\delta) \geq \log_{1/3}(\frac{\delta}{2})$, we obtain $\mathbf{P}(\bar{\mathbf{v}}^{\hat{\pi}} < \bar{\mathbf{v}}^* - \epsilon) \leq \delta$. Then the output policy $\hat{\pi}$ is ϵ -optimal with probability at least $1 - \delta$. \blacksquare

The π learning method is not only useful in the setting of reinforcement learning. It also applies to the computational problem of approximating the optimal policy when the MDP model is explicitly given. We obtain the following sublinear run-time complexity for numerically solving the AMDP.

Theorem 4 (Sublinear Run Time for Ergodic MDP). *Let the $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r})$ be an MDP tuple that is specified in data structures that enable sampling state transitions in $\tilde{\mathcal{O}}(1)$ time. Then there exists an algorithm that takes \mathcal{M} as the input and outputs an approximate policy $\hat{\pi}$ such that $\bar{v}^{\hat{\pi}} \geq \bar{v}^* - \epsilon$ with probability at least $1 - \delta$ in run time*

$$\Omega\left((\tau \cdot t_{mix}^*)^2 \cdot \frac{|\mathcal{S}||\mathcal{A}|}{\epsilon^2} \log \frac{1}{\delta}\right).$$

Proof. Note that each iteration of Algorithm 1 draws one sample from the \mathcal{SO} and makes updates in $\tilde{\mathcal{O}}(1)$ time. Then the result of Theorem 4 follows straightforwardly from the sample complexity result of Theorem 3. \blacksquare

Theorem 4 suggests that one can approximately solve the AMDP problem in sublinear time. In particular, finding an approximately optimal policy does not even require reading most of the input entries. Remarkably, the sample complexity and run-time complexity of the π learning method happen to be equivalent to each other. Such an equivalence holds because the π learning method uses each new sample transition in a most computationally efficient way - making only a few coordinate updates to the value and policy vectors. Note that for general reinforcement learning methods, the sample complexity and run-time complexity are typically far from equal to each other.

6 Summary

We have developed a primal-dual π learning method for solving the undiscounted ergodic Markov decision problems by sampling state-to-state transitions. The method directly updates the value and policy estimates as new state transitions are observed. This method is model-free and can be implemented efficiently in $\mathcal{O}(|\mathcal{S}||\mathcal{A}|)$ space. We show that it achieves a sample complexity $\mathcal{O}((\tau \cdot t_{mix}^*)^2 \frac{|\mathcal{S}||\mathcal{A}|}{\epsilon^2})$ for ergodic average-reward Markov decision process, where τ is a parameter characterizing the range of stationary distributions and t_{mix}^* is an upper bound of mixing times across all policies.

The π learning method can be applied to approximating the optimal policy when the MDP is fully specified. When state transitions can be sampled in $\mathcal{O}(1)$ time, the π learning method can be used as a randomized algorithm and computes an ϵ -optimal policy in run time $\mathcal{O}((\tau \cdot t_{mix}^*)^2 \frac{|\mathcal{S}||\mathcal{A}|}{\epsilon^2})$, which is sublinear with respect to the input size. An open question is to investigate the roles of τ, t_{mix}^* in the complexity and potentially improve the complexity's dependence on these parameters.

References

- [1] Mohammad Gheshlaghi Azar, Rémi Munos, and Bert Kappen. On the sample complexity of reinforcement learning with a generative model. *arXiv preprint arXiv:1206.6461*, 2012.
- [2] Richard Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.
- [3] Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena Scientific, Belmont, MA, 1995.
- [4] Dimitri P Bertsekas. *Abstract dynamic programming*. Athena Scientific, Belmont, MA, 2013.
- [5] Dimitri P Bertsekas and John N Tsitsiklis. Neuro-dynamic programming: an overview. In *Proceedings of the 34th IEEE Conference on Decision and Control*, volume 1, pages 560–564. IEEE, 1995.
- [6] Yichen Chen and Mengdi Wang. Stochastic primal-dual methods and sample complexity of reinforcement learning. *arXiv preprint arXiv:1612.02516*, 2016.
- [7] Yichen Chen and Mengdi Wang. Lower bound on the computational complexity of discounted markov decision problems. *arXiv preprint arXiv:1705.07312*, 2017.
- [8] Kenneth L Clarkson, Elad Hazan, and David P Woodruff. Sublinear optimization for machine learning. *Journal of the ACM (JACM)*, 59(5):23, 2012.
- [9] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [10] Guy De Ghellinck. Les problemes de decisions sequentielles. *Cahiers du Centre d'Etudes de Recherche Opérationnelle*, 2(2):161–179, 1960.
- [11] F d'Epenoux. A probabilistic production and inventory problem. *Management Science*, 10(1):98–108, 1963.
- [12] Ronald A. Howard. *Dynamic programming and Markov processes*. The MIT press, Cambridge, MA, 1960.

- [13] Anatoli Juditsky, Arkadi Nemirovski, Claire Tauvel, et al. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 1(1):17–58, 2011.
- [14] Sham M Kakade. *On the sample complexity of reinforcement learning*. PhD thesis, University of London London, England, 2003.
- [15] Michael Kearns, Yishay Mansour, and Andrew Y Ng. A sparse sampling algorithm for near-optimal planning in large markov decision processes. *Machine Learning*, 49(2-3):193–208, 2002.
- [16] Michael J Kearns and Satinder P Singh. Finite-sample convergence rates for q-learning and indirect algorithms. In *Advances in neural information processing systems*, pages 996–1002, 1999.
- [17] Tor Lattimore and Marcus Hutter. PAC bounds for discounted MDPs. In *International Conference on Algorithmic Learning Theory*, pages 320–334. Springer, 2012.
- [18] Michael L Littman, Thomas L Dean, and Leslie Pack Kaelbling. On the complexity of solving Markov decision problems. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pages 394–402. Morgan Kaufmann Publishers Inc., 1995.
- [19] Yishay Mansour and Satinder Singh. On the complexity of policy iteration. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pages 401–408. Morgan Kaufmann Publishers Inc., 1999.
- [20] Arkadi Nemirovski and Reuven Y Rubinstein. An efficient stochastic approximation algorithm for stochastic saddle point problems. In *Modeling Uncertainty*, pages 156–184. Springer, 2005.
- [21] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [22] Bruno Scherrer. Improved and generalized upper bounds on the complexity of policy iteration. In *Advances in Neural Information Processing Systems*, pages 386–394, 2013.
- [23] A. Sidford, M. Wang, C. Wu, and Y. Ye. Variance reduced value iteration and faster algorithms for solving markov decision processes. In *SODA*, 2017.
- [24] Paul Tseng. Solving h-horizon, stationary markov decision problems in time proportional to $\log(h)$. *Operations Research Letters*, 9(5):287–297, 1990.
- [25] Mengdi Wang. Randomized linear programming solves the discounted Markov decision problem in nearly-linear running time. *arXiv preprint arXiv:1704.01869*, 2017.
- [26] Mengdi Wang and Yichen Chen. An online primal-dual method for discounted Markov decision processes. In *IEEE Conference of Decisions and Control*, 2016.
- [27] Yinyu Ye. A new complexity result on solving the Markov decision problem. *Mathematics of Operations Research*, 30(3):733–749, 2005.
- [28] Yinyu Ye. The simplex and policy-iteration methods are strongly polynomial for the Markov decision problem with a fixed discount rate. *Mathematics of Operations Research*, 36(4):593–603, 2011.

A Proof of Theorem 1: Duality Gap Analysis

In this section, we analyze the convergence of Algorithm 1. In what follows, we denote by \mathcal{F}_t the collection of random variables that are revealed up to the end of the t -th iteration. For two probability distributions p, q over a finite set X , we denote by $D_{KL}(p||q) = \sum_{x \in X} p(x) \log \frac{p(x)}{q(x)}$ the Kullback-Leibler divergence. We assume that Assumptions 1 and 2 hold throughout the analysis. We let

$$\mathbf{r}_a = \sum_{j \in \mathcal{S}} p_{ij}(a) r_{ij}(a), \quad \mathbf{r}^\pi = \sum_{a \in \mathcal{A}} \pi_i(a) \sum_{j \in \mathcal{S}} p_{ij}(a) r_{ij}(a), \quad \mathbf{r}^* = \sum_{a \in \mathcal{A}} \pi_i^*(a) \sum_{j \in \mathcal{S}} p_{ij}(a) r_{ij}(a).$$

We denote by $P^* = P^{\pi^*}$ the transition matrix under the optimal policy. Note that

$$P^\pi = \sum_{a \in \mathcal{A}} \mathbf{diag}(\pi_a) P_a, \quad P^* = \sum_{a \in \mathcal{A}} \mathbf{diag}(\pi_a^*) P_a,$$

where $\mathbf{diag}(\pi_a)$ is the diagonal matrix with $\pi_{1,a}, \pi_{2,a}, \dots, \pi_{|\mathcal{S}|,a}$ along its diagonal.

In addition, the updates on ξ^t and π^t can be equivalently written as updates on μ^t , given by

$$\begin{aligned} \mu_{i,a}^{t+1/2} &= \frac{\mu_{i,a}^t \cdot \exp(\Delta_{i,a}^{t+1})}{\sum_{i',a'} \mu_{i',a'}^t \cdot \exp(\Delta_{i',a'}^{t+1})}, \quad \forall i \in \mathcal{S}, a \in \mathcal{A}, \\ \mu^{t+1} &= \operatorname{argmin}_{\mu \in \mathcal{U}} D_{KL}(\mu || \mu^{t+1/2}), \end{aligned} \quad (6)$$

One can verify that $\mu^t \in \mathcal{U}$ and $\mathbf{h}^t \in \mathcal{H}$ for all t with probability 1.

Lemma 4. *The iterates generated by Algorithm 1 satisfy*

$$\mathbf{E} [D_{KL}(\mu^* || \mu^{t+1}) | \mathcal{F}_t] - D_{KL}(\mu^* || \mu^t) \leq \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} (\mu_{i,a}^t - \mu_{i,a}^*) \mathbf{E} [\Delta_{i,a}^{t+1} | \mathcal{F}_t] + \frac{1}{2} \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \mathbf{E} \left[\left(\Delta_{i,a}^{t+1} \right)^2 | \mathcal{F}_t \right], \quad (7)$$

for all t , with probability 1.

Proof. By using the relation (6), we have

$$\begin{aligned} D_{KL}(\mu^* || \mu^{t+1/2}) - D_{KL}(\mu^* || \mu^t) &= \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^* \log \frac{\mu_{i,a}^*}{\mu_{i,a}^{t+1/2}} - \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^* \log \frac{\mu_{i,a}^*}{\mu_{i,a}^t} \\ &= \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^* \log \frac{\mu_{i,a}^t}{\mu_{i,a}^{t+1/2}} \\ &= \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^* \log \frac{Z}{\exp(\Delta_{i,a}^{t+1})} \\ &= \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^* \log(Z) - \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^* \Delta_{i,a}^{t+1} \\ &= \log Z - \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^* \Delta_{i,a}^{t+1}, \end{aligned} \quad (8)$$

where $Z = \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \exp(\Delta_{i,a}^{t+1})$. According to (5), we have $h_j^t - h_i^t + r_{ij}(a) - M \leq 2t_{mix}^* - (-2t_{mix}^*) + 1 - 4t_{mix}^* - 1 \leq 0$ because $h_i^t \in [-2t_{mix}^*, 2t_{mix}^*]$, $r_{ij}(a) \in [0, 1]$ and $M = 4t_{mix}^* + 1$. It

follows that $\Delta_{i,a}^{t+1} \leq 0$ for all $i \in \mathcal{S}, a \in \mathcal{A}$ with probability 1. Then we derive

$$\begin{aligned} \log Z &= \log \left(\sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \exp(\Delta_{i,a}^{t+1}) \right) \leq \log \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \left(1 + \Delta_{i,a}^{t+1} + \frac{1}{2} \left(\Delta_{i,a}^{t+1} \right)^2 \right) \\ &= \log \left(1 + \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \Delta_{i,a}^{t+1} + \frac{1}{2} \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \left(\Delta_{i,a}^{t+1} \right)^2 \right) \quad (9) \\ &\leq \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \Delta_{i,a}^{t+1} + \frac{1}{2} \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \left(\Delta_{i,a}^{t+1} \right)^2, \end{aligned}$$

where the first inequality uses the fact $e^x \leq 1 + x + \frac{1}{2}x^2$ if $x \leq 0$ and the second inequality uses the fact $\log(1+x) \leq x$ for all x . We combine (8) and (9) and take conditional expectation $\mathbf{E}[\cdot | \mathcal{F}_t]$ on both sides, then we obtain

$$\begin{aligned} &\mathbf{E} \left[D_{KL}(\mu^* || \mu^{t+1/2}) | \mathcal{F}_t \right] - D_{KL}(\mu^* || \mu^t) \\ &\leq \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} (\mu_{i,a}^t - \mu_{i,a}^*) \mathbf{E} \left[\Delta_{i,a}^{t+1} | \mathcal{F}_t \right] + \frac{1}{2} \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \mathbf{E} \left[\left(\Delta_{i,a}^{t+1} \right)^2 | \mathcal{F}_t \right], \quad (10) \end{aligned}$$

Finally, we note that $D_{KL}(\mu^* || \mu^{t+1}) \leq D_{KL}(\mu^* || \mu^{t+1/2})$ due to the information projection step (see [9] Theorem 11.6.1 on page 367) and that $\mu^* \in \mathcal{U}$. By combining the preceding two relations, we have obtained (7). \blacksquare

Lemma 5. *The iterates generated by Algorithm 1 satisfy*

$$\sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \mathbf{E} \left[\left(\Delta_{i,a}^{t+1} \right)^2 | \mathcal{F}_t \right] \leq 4|\mathcal{S}||\mathcal{A}|(4t_{mix}^* + 1)^2 \beta^2,$$

for all $t \geq 1$ with probability 1.

Proof. We have

$$\begin{aligned} \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \mathbf{E} \left[\left(\Delta_{i,a}^{t+1} \right)^2 | \mathcal{F}_t \right] &= \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t \cdot \xi_i^t \pi_{i,a}^t \cdot \sum_{j \in \mathcal{S}} p_{ij}(a) \left(\beta \cdot \frac{(h_j^t - h_i^t + r_{ij}(a) - M)}{\xi_i^t \pi_{i,a}^t} \right)^2 \\ &= \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{j \in \mathcal{S}} p_{ij}(a) \left(\beta \cdot (h_j^t - h_i^t + r_{ij}(a) - M) \right)^2 \\ &\leq \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{j \in \mathcal{S}} p_{ij}(a) \left(\beta \cdot 2 \cdot (4t_{mix}^* + 1) \right)^2 \\ &= 4|\mathcal{S}||\mathcal{A}| \beta^2 (4t_{mix}^* + 1)^2, \end{aligned}$$

where the inequality uses the fact that $h^t \in \mathcal{H}$. \blacksquare

Lemma 6. *The iterates generated by Algorithm 1 satisfy*

$$\mathbf{E} \left[D_{KL}(\mu^* || \mu^{t+1}) | \mathcal{F}_t \right] \leq D_{KL}(\mu^* || \mu^t) + \beta \sum_{a \in \mathcal{A}} (\mu_a^t - \mu_a^*)^\top \left((P_a - I) \mathbf{h}^t + \mathbf{r}_a \right) + 2|\mathcal{S}||\mathcal{A}|(4t_{mix}^* + 1)^2 \beta^2, \quad (11)$$

for all $t \geq 0$, with probability 1.

Proof. For arbitrary $i \in \mathcal{S}$ and $a \in \mathcal{A}$, we have

$$\frac{1}{\beta} \cdot \mathbf{E} \left[\Delta_{i,a}^{t+1} \mid \mathcal{F}_t \right] = \sum_{j \in \mathcal{S}} p_{ij}(a) h_j^t - h_i^t + \sum_{j \in \mathcal{S}} p_{ij}(a) r_{ij}(a) - M = (P_a \mathbf{h}^t - \mathbf{h}^t + \mathbf{r}_a)_i - M.$$

It follows that

$$\begin{aligned} \frac{1}{\beta} \cdot \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} (\mu_{i,a}^t - \mu_{i,a}^*) \mathbf{E} \left[\Delta_{i,a}^{t+1} \mid \mathcal{F}_t \right] &= \sum_{a \in \mathcal{A}} \sum_{i \in \mathcal{S}} (\mu_{i,a}^t - \mu_{i,a}^*) [(P_a \mathbf{h}^t - \mathbf{h}^t + \mathbf{r}_a)_i - M] \\ &= \sum_{a \in \mathcal{A}} (\mu_a^t - \mu_a^*)^\top ((P_a - I) \mathbf{h}^t + \mathbf{r}_a), \end{aligned}$$

where the second equality comes from the fact $\sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t = \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^* = 1$ (because $\mu^t \in \mathcal{U}$, $\mu^* \in \mathcal{U}$). We further apply Lemmas 2-5 and complete the proof. \blacksquare

Lemma 7. *The iterates generated by Algorithm 1 satisfy for all $t \geq 0$ with probability 1 that*

$$\mathbf{E} [\|\mathbf{h}^{t+1} - \mathbf{h}^*\|^2 \mid \mathcal{F}_t] \leq \|\mathbf{h}^t - \mathbf{h}^*\|^2 + 2\alpha(\mathbf{h}^t - \mathbf{h}^*)^\top \left(\sum_{a \in \mathcal{A}} (I - P_a)^\top \mu_a^t \right) + \mathcal{O}(\alpha^2). \quad (12)$$

Proof. According to the updates of Algorithm 1, we have

$$\mathbf{h}^{t+1} = \mathbf{Proj}_{\mathcal{H}} [\mathbf{h}^t + \mathbf{d}^t],$$

where $\mathbf{Proj}_{\mathcal{H}}$ denotes the Euclidean projection onto $\mathcal{H} = \{\mathbf{h} \mid \|\mathbf{h}\|_\infty \leq 2t_{mix}^*\}$. By using the nonexpansive property of $\Pi_{\mathcal{H}}$ and $\mathbf{h}^* \in \mathcal{H}$, we further obtain

$$\begin{aligned} \mathbf{E} [\|\mathbf{h}^{t+1} - \mathbf{h}^*\|^2 \mid \mathcal{F}_t] &= \mathbf{E} [\|\mathbf{Proj}_{\mathcal{H}}[\mathbf{h}^t + \mathbf{d}^t] - \mathbf{h}^*\|^2 \mid \mathcal{F}_t] \leq \mathbf{E} [\|\mathbf{h}^t + \mathbf{d}^t - \mathbf{h}^*\|^2 \mid \mathcal{F}_t] \\ &= \|\mathbf{h}^t - \mathbf{h}^*\|^2 + 2(\mathbf{h}^t - \mathbf{h}^*)^\top \mathbf{E} [\mathbf{d}^t \mid \mathcal{F}_t] + \mathbf{E} [\|\mathbf{d}^t\|^2 \mid \mathcal{F}_t], \end{aligned}$$

for all t with probability 1. We can verify that

$$\mathbf{E} [\mathbf{d}^t \mid \mathcal{F}_t] = \alpha(I - P^{\pi_t})^\top \xi^t = \alpha \sum_{a \in \mathcal{A}} (I - P_a)^\top \mathbf{diag}(\pi_a^t) \xi^t = \alpha \sum_{a \in \mathcal{A}} (I - P_a)^\top \mu_a^t.$$

and

$$\mathbf{E} [\|\mathbf{d}^t\|^2 \mid \mathcal{F}_t] = \mathcal{O}(\alpha^2).$$

Finally we combine all preceding inequalities and obtain (12). \blacksquare

Lemma 8. *We define for short that*

$$\mathcal{E}^t = D_{KL}(\mu^* \parallel \mu^t) + \frac{1}{2|\mathcal{S}|(t_{mix}^*)^2} \|\mathbf{h}^t - \mathbf{h}^*\|^2, \quad \mathcal{G}^t = \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu_{i,a}^t (\mathbf{h}^* - P_a \mathbf{h}^* - \mathbf{r}_a)_i + \bar{v}^*.$$

Let $\alpha = |\mathcal{S}|(t_{mix}^*)^2 \beta$. *The iterates generated by Algorithm 1 satisfy for all t with probability 1 that*

$$\mathbf{E} [\mathcal{E}^{t+1} \mid \mathcal{F}_t] \leq \mathcal{E}^t - \beta \mathcal{G}^t + \beta^2 \tilde{\mathcal{O}}(|\mathcal{S}| |\mathcal{A}| (t_{mix}^*)^2). \quad (13)$$

Proof. Let $\alpha = |\mathcal{S}|(t_{mix}^*)^2\beta$. We multiply (12) with $\frac{1}{2|\mathcal{S}|(t_{mix}^*)^2}$ and takes its sum with (11), obtaining

$$\begin{aligned} \mathbf{E} [\mathcal{E}^{t+1} | \mathcal{F}_t] &\leq \mathcal{E}^t + \beta^2 \tilde{\mathcal{O}}(|\mathcal{S}||\mathcal{A}|(t_{mix}^*)^2) \\ &\quad + \beta \left(\sum_{a \in \mathcal{A}} (\mu_a^t - \mu_a^*)^\top ((P_a - I)\mathbf{h}^t + \mathbf{r}_a) + (\mathbf{h}^t - \mathbf{h}^*)^\top \left(\sum_{a \in \mathcal{A}} (I - P_a)^\top \mu_a^t \right) \right). \end{aligned}$$

We have

$$\begin{aligned} &\sum_{a \in \mathcal{A}} (\mu_a^t - \mu_a^*)^\top ((P_a - I)\mathbf{h}^t + \mathbf{r}_a) + (\mathbf{h}^t - \mathbf{h}^*)^\top \left(\sum_{a \in \mathcal{A}} (I - P_a)^\top \mu_a^t \right) \\ &= \sum_{a \in \mathcal{A}} (\mu_a^t - \mu_a^*)^\top ((P_a - I)\mathbf{h}^t + \mathbf{r}_a) + (\mathbf{h}^t - \mathbf{h}^*)^\top \sum_{a \in \mathcal{A}} (I - P_a)^\top (\mu_a^t - \mu_a^*) \quad (\text{by the dual feasibility of } \mu^*) \\ &= \sum_{a \in \mathcal{A}} (\mu_a^t - \mu_a^*)^\top ((P_a - I)\mathbf{h}^* + \mathbf{r}_a) \\ &= \sum_{a \in \mathcal{A}} (\mu_a^t)^\top ((P_a - I)\mathbf{h}^* + \mathbf{r}_a) - \sum_{a \in \mathcal{A}} \bar{v}^* \cdot (\mu_a^*)^\top \mathbf{1} \quad (\text{by the linear complementarity condition for } \mathbf{h}^*, \mu^*) \\ &= \sum_{a \in \mathcal{A}} (\mu_a^t)^\top ((P_a - I)\mathbf{h}^* + \mathbf{r}_a) - \bar{v}^*, \end{aligned}$$

where the first equality uses the dual feasibility of μ^* of linear program (2):

$$\sum_{a \in \mathcal{A}} (I - P_a)^\top \mu_a^* = 0,$$

and the third equality uses the complementary condition of the linear programs (1)-(2):

$$\mu_{a,i}^* ((P_a - I)\mathbf{h}^* + \mathbf{r}_a - \bar{v}^* \cdot \mathbf{1})_i = 0, \quad \forall i \in \mathcal{S}, a \in \mathcal{A}.$$

Combining the preceding relations, we obtain (13). ■

Proof of Theorem 1. We claim that $\mathcal{E}^1 \leq \log(|\mathcal{S}||\mathcal{A}|) + 2$. To see this, we note that μ^1 is the uniform distribution (according to Step 2 of Algorithm 1) and $\mathbf{h}^0, \mathbf{h}^* \in \mathcal{H}$. Therefore we have $D_{KL}(\mu^* || \mu^1) \leq \log(|\mathcal{S}||\mathcal{A}|)$ and $\|\mathbf{h}^t - \mathbf{h}^*\|^2 \leq 4|\mathcal{S}|(t_{mix}^*)^2$ for all t . Then we have $\mathcal{E}^1 \leq D_{KL}(\mu^* || \mu^1) + \frac{1}{2|\mathcal{S}|(t_{mix}^*)^2} \|\mathbf{h}^1 - \mathbf{h}^*\|^2 \leq \log(|\mathcal{S}||\mathcal{A}|) + 2$.

We rearrange the terms of (13) and obtain

$$\mathcal{G}^t \leq \frac{1}{\beta} (\mathcal{E}^t - \mathbf{E} [\mathcal{E}^{t+1} | \mathcal{F}_t]) + \beta \tilde{\mathcal{O}}(|\mathcal{S}||\mathcal{A}|(t_{mix}^*)^2).$$

Summing over $t = 1, \dots, T$ and taking average, we have

$$\begin{aligned} \mathbf{E} \left[\sum_{t=1}^T \mathcal{G}^t \right] &\leq \frac{1}{\beta} \sum_{t=1}^T (\mathbf{E} [\mathcal{E}^t] - \mathbf{E} [\mathcal{E}^{t+1}]) + T\beta \tilde{\mathcal{O}}(|\mathcal{S}||\mathcal{A}|(t_{mix}^*)^2) \\ &= \frac{\mathbf{E} [\mathcal{E}^1] - \mathbf{E} [\mathcal{E}^T]}{\beta} + T\beta \tilde{\mathcal{O}}(|\mathcal{S}||\mathcal{A}|(t_{mix}^*)^2) \\ &\leq \frac{1}{\beta} (\log(|\mathcal{S}||\mathcal{A}|) + 2) + T\beta \tilde{\mathcal{O}}(|\mathcal{S}||\mathcal{A}|(t_{mix}^*)^2). \end{aligned}$$

where the inequality is based on the fact $\mathcal{E}^1 \leq \log(|\mathcal{S}||\mathcal{A}|) + 2$ and $\mathcal{E}^t \geq 0$. Therefore by taking

$$\beta = \frac{1}{t_{mix}^*} \sqrt{\frac{\log(|\mathcal{S}||\mathcal{A}|)}{2|\mathcal{S}||\mathcal{A}|T}}, \text{ we obtain } \mathbf{E} \left[\frac{1}{T} \sum_{t=1}^T \mathcal{G}^t \right] = \tilde{\mathcal{O}} \left(t_{mix}^* \sqrt{\frac{|\mathcal{S}||\mathcal{A}|}{T}} \right). \quad \blacksquare$$