

Reinforcement Learning for Joint Optimization of Multiple Rewards

Mridul Agarwal

Purdue University

West Lafayette IN 47907

AGARW180@PURDUE.EDU

Vaneet Aggarwal

Purdue University

West Lafayette IN 47907

VANEET@PURDUE.EDU

Editor: Animashree Anandkumar

Abstract

Finding optimal policies which maximize long term rewards of Markov Decision Processes requires the use of dynamic programming and backward induction to solve the Bellman optimality equation. However, many real-world problems require optimization of an objective that is non-linear in cumulative rewards for which dynamic programming cannot be applied directly. For example, in a resource allocation problem, one of the objectives is to maximize long-term fairness among the users. We notice that when an agent aim to optimize some function of the sum of rewards is considered, the problem loses its Markov nature. This paper addresses and formalizes the problem of optimizing a non-linear function of the long term average of rewards. We propose model-based and model-free algorithms to learn the policy, where the model-based policy is shown to achieve a regret of $\tilde{O}\left(LKDS\sqrt{\frac{A}{T}}\right)$ for K objectives combined with a concave L -Lipschitz function. Further, using the fairness in cellular base-station scheduling, and queueing system scheduling as examples, the proposed algorithm is shown to significantly outperform the conventional RL approaches.

1. Introduction

Many practical applications of sequential decision making often have multiple objectives. For example, a hydro-power project may have conflicting gains with respect to power generation and flood management (Castelletti et al., 2013). Similarly, a building climate controller can have conflicting objectives of saving energy and maximizing comfort of the residents of the building (Kwak et al., 2012). Video streaming applications also account for multiple objectives like stall duration and average video quality (Elgabli et al., 2018). Access of files from cloud storage aims to optimize the latency of file download and the cost to store the files (Xiang et al., 2015). Many applications also require to allocate resources fairly to multiple clients (Lan et al., 2010) which can be modelled as optimizing a function of the objectives of the individual clients. This paper aims to provide a novel formulation for decision making among multiple objectives using reinforcement learning approaches and to analyze the performance of the proposed algorithms.

We consider a setup where we want to optimize a possibly nonlinear joint objective function of long-term rewards of all the objectives (or, different objectives). As an example,

many problems in resource allocation for networking and computation resources use fairness among the long-term average rewards of the users as the metric of choice (Lan et al., 2010; Kwan et al., 2009; Bu et al., 2006; Li et al., 2018; Aggarwal et al., 2011; Margolies et al., 2016; Wang et al., 2014; Ibrahim et al., 2010), which is a non-linear metric. For fairness optimization, a controller wants to optimize a fairness objective among the different agents, e.g., proportional fairness, α -fairness, or improve the worst-case average reward of the users (Altman et al., 2008). In such situations, the overall joint objective function cannot be written as sum utility at each time instant. This prohibits the application of standard single-agent reinforcement learning based policies as the backward induction step update cannot be directly applied here. For example, if a process has 2 agents and $T > 1$ steps, and all the resource was allocated to the first agent till $T - 1$ steps. Then, at T^{th} step the resource should be allocated to the second agent to ensure fairness. This requires the need to track past allocation of all the resources and not just the current state of the system. We also note that the optimal policy cannot take a deterministic action in a state in general, and thus the optimal policy is not a deterministic policy in general. Consider a case where a scheduler needs to fairly allocate a resource between two users, and the system has only one state. A deterministic policy will allocate the resource to only one of the user, and hence is not optimal. We define a novel multi-agent formulation, making several practical assumptions, which optimizes the joint function of the average per-step rewards of the different objectives to alleviate the need for maintaining history.

SARSA and Q-Learning algorithms (Sutton and Barto, 2018b), and their deep neural network based DQN algorithm (Mnih et al., 2015) provide policies that depend only on the current state, hence are sub-optimal. Further, these algorithms learn a Q-value function which can be computed based on a dynamic programming approach, which is not valid in our work. Using evaluations on fair resource allocation and network routing problems, we verify that algorithms based on finding fixed point of Bellman equations do not perform well. This further motivates the need for novel RL based algorithms to optimize non-linear functions.

We further note that even though multi-agent reinforcement learning algorithms have been widely studied, (Tan, 1993; Shoham et al., 2003; Buşoniu et al., 2010; Ono and Fukumoto, 1996), there are no convergence proofs to the optimal joint objective function without the knowledge of the transition probability, to the best of our knowledge. This paper assumes no knowledge of the state transition probability of the objectives and aims to provide algorithms for the decision making of the different objectives. We provide two algorithms; The first is a model-based algorithm that learns the transition probability of the next state given the current state and action. The second algorithm is model-free, which uses policy gradients to find the optimal policy.

The proposed model-based algorithm uses posterior sampling with Dirichlet distribution. We show that the proposed algorithm converges to an optimal point when the joint objective function is Lipschitz continuous. Since the optimal policy is a stochastic policy, policy search space is not-finite. We show that the problem is convex under a certain class of functions and can be efficiently solved. In setups where the joint objective function is max-min, the setup reduces to a linear programming optimization problem. In addition, we show that the proposed algorithm achieves a regret bound sub-linear in the number of time-steps and number of objectives. To obtain the regret bound, we use a Bellman error based analysis

to analyze stochastic policies. The Bellman error quantifies the difference in rewards for deviating from the true MDP for one step and then following the true MDP thereafter. Using this analysis, our regret bound characterizes the gap between the optimal objective and the objective achieved by the algorithm in T time-steps. We show a regret bound of $\tilde{O}\left(KDS\sqrt{\frac{A}{T}}\right)$, where K, T denotes the number of objectives, and time steps, respectively.

The proposed model-free algorithm can be easily implemented using deep neural networks for any differentiable objective function. Further, we note that the reward functions of the different objectives can be very different, and can optimize different metrics for the objectives. As long as there is a joint objective function, the different objectives can make decisions to optimize this function and achieve the optimal decision at convergence.

The proposed framework works for any number of objectives, while is novel even for a single agent ($K = 1$). In this case, the agent wishes to optimize a non-linear concave function of the average reward. Since this function is not assumed to be monotone, optimizing the function is not equivalent to optimizing the average reward. For any general non-linear concave function, regret bound is analyzed for model-based case.

We also present evaluation results for both the algorithms for optimizing proportional fairness of multiple agents connecting to a cellular base station. We compare the obtained policies with existing asymptotically optimal algorithm (Blind Gradient Estimator or BGE) of optimizing proportional fairness for wireless networks (Margolies et al., 2016) and SARSA based RL solution proposed by (Perez et al., 2009). We developed a simulation environment for wireless network for multiple number of agents and states for each agent. The proposed algorithm significantly outperform the SARSA based solution, and it nearly achieves the performance of the asymptotically optimal BGE algorithm. We also considered α -fairness for an infinite state space to show the scalability of the proposed model-free algorithm. In this case, the domain-specific algorithm was not available, while we show that we outperform Deep Q-Network (DQN) based algorithm (Mnih et al., 2015). Finally, a queueing system is considered which models multiple roads merging into a single lane. The queue selection problem is modeled using the proposed framework and the proposed approach is shown to improve the fair latency reward metric among the queues significantly as compared to the DQN and the longest-queue-first policies.

Key contributions of our paper are:

- A structure for joint function optimization with multiple objectives based on average per-step rewards.
- Pareto Optimality guarantees when the joint objective is an element-wise monotone function.
- A model-based algorithm using posterior sampling with Dirichlet distribution, and its regret bounds.
- A model-free policy gradient algorithm which can be efficiently implemented using neural networks.
- Evaluation results and comparison with existing heuristics for optimizing fairness in cellular networks, and queueing systems.

The rest of the paper is organized as follows. Section 2 describes related works in the field of RL and MARL. Section 3 describes the problem formulation. Pareto optimality of the proposed framework is shown in Section 4. The proposed model based algorithm and model free algorithm are described in Sections 5 and 6, respectively. In Section 7, the proposed algorithms are evaluated for cellular scheduling problem. Section 8 concludes the paper with some future work directions.

2. Related Work

Reinforcement learning for single agent has been extensively studied in past (Sutton and Barto, 2018b). Dynamic Programming was used in many problems by finding cost to go at each stage (Puterman, 1994; Bertsekas, 1995). These models optimize linear additive utility and utilize the power of Backward Induction.

Following the success of Deep Q Networks (Mnih et al., 2015), many new algorithms have been developed for reinforcement learning (Schulman et al., 2015; Lillicrap et al., 2015; Wang et al., 2015; Schulman et al., 2017). These papers focus on single agent control, and provide a framework for implementing scalable algorithms. Sample efficient algorithms based on rate of convergence analysis have also been studied for model based RL algorithms (Agrawal and Jia, 2017; Osband et al., 2013), and for model free Q learning (Jin et al., 2018). However, sample efficient algorithms use tabular implementation instead of a deep learning based implementation.

Owing to high variance in the policies obtained by standard Markov Decision Processes and Reinforcement Learning formulations, various authors worked in reducing the *risk* in RL approaches (Garcia and Fernández, 2015). Even though the risk function (e.g., Conditional Value at Risk (CVaR)) is non-linear in the rewards, this function is not only a function of long-term average rewards of the single agent but also a function of the higher order moments of the rewards of the single agent. Thus, the proposed framework does not apply to the risk measures. However, for both the risk measure and general non-linear concave function of average rewards, optimal policy is non-Markovian.

Non-Markovian Decision Processes is a class of decision processes where either rewards, the next state transitions, or both do not only depends on the current state and actions but also on the history of states and actions leading towards the current state. One can augment the state space to include the history along with the current state and make the new process Markovian (Thiébaux et al., 2006). However, this increases the memory footprint of any Q-learning algorithm. (McCallum, 1995) considers only H states of history to construct an approximate MDP and then use Q-learning. (Li et al., 2006) provide guarantees on Q-learning for non-MDPs where an agent observes and work according to an abstract MDP instead of the ground MDP. The states of the abstract MDP are an abstraction of the states of the ground MDP. (Hutter, 2014) extend this setup to work with abstractions of histories. (Majeed and Hutter, 2018) consider a setup for History-based Decision Process (HDP). They provide convergence guarantees for Q-learning algorithm for a sub-class of HDP where for histories h and h' , $Q(h, a) = Q(h', a)$ if the last observed state is identical for both h and h' . They call this sub-class Q -value uniform Decision Process (QDP) and this subsumes the abstract MDPs. We note that our work is different from these as the Q -values constructed using joint objective is not independent of history.

In most applications such as financial markets, swarm robotics, wireless channel access, etc., there are multiple agents that make a decision (Bloembergen et al., 2015), and the decision of any agent can possibly affect the other agents. In early work on multi-agent reinforcement learning (MARL) for stochastic games (Littman, 1994), it was recognized that no agent works in a vacuum. In his seminal paper, Littman (Littman, 1994) focused on only two agents that had opposite and opposing goals. This means that they could use a single reward function which one tried to maximize and the other tried to minimize. The agent had to work with a competing agent and had to behave to maximize their reward in the worst possible case. In MARL, the agents select actions simultaneously at the current state and receive rewards at the next state. Different from the frameworks that solve for a Nash equilibrium in a stochastic game, the goal of a reinforcement learning algorithm is to learn equilibrium strategies through interaction with the environment (Tan, 1993; Shoham et al., 2003; Buşoniu et al., 2010; Ono and Fukumoto, 1996; Shalev-Shwartz et al., 2016).

(Roijers et al., 2013; Liu et al., 2014; Nguyen et al., 2020) considers the multi-objective Markov Decision Processes. Similar to our work, they consider function of expected cumulative rewards. However, they work with linear combination of the cumulative rewards whereas we consider a possibly non-linear function f . Further, based on the joint objective as a function of expected average rewards, we provide regret guarantees for our algorithm. For joint decision making, (Zhang and Shah, 2014, 2015) studied the problem of fairness with multiple agents and related the fairness to multi-objective MDP. They considered maximin fairness and used Linear Programming to obtain optimal policies. For general functions, linear programming based approach provided by (Zhang and Shah, 2014) will not directly work. This paper also optimizes joint action of agents using a centralized controller and propose a model based algorithm to obtain optimal policies. Based on our assumptions, maximin fairness becomes a special case of our formulation and optimal policies can be obtained using the proposed model based algorithm. We also propose a model free reinforcement learning algorithm that can be used to obtain optimal policies for any general differentiable functions of average per-step rewards of individual agents. Recently, (Jiang and Lu, 2019) considered the problem of maximizing fairness among multiple agents. However, they do not provide a convergence analysis for their algorithm. We attempt to close this gap in the understanding of the problem of maximizing a concave and Lipschitz function of multiple objectives with our work.

3. Problem Formulation

We consider an infinite horizon discounted Markov decision process (MDP) \mathcal{M} defined by the tuple $(\mathcal{S}, \mathcal{A}, P, r^1, r^2, \dots, r^K, \rho_0)$. \mathcal{S} denotes a finite set of state space of size S , and \mathcal{A} denotes a finite set of A actions. $P : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]^S$ denotes the probability transition distribution. K denotes the number of objectives and $[K] = \{1, 2, \dots, K\}$ is the set of K objectives. Let $r^k : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ be the bounded reward function for objective $k \in [K]$. Lastly, $\rho_0 : \mathcal{S} \rightarrow [0, 1]$ is the distribution of initial state. We motivate our choice of bounds on rewards from the fact that many problems in practice require explicit reward shaping. Hence, the controller or the learner is aware of the bounds on the rewards. We consider the bounds to be $[0, 1]$ for our case which is easy to satisfy by reward shaping.

We use a stochastic policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ which returns the probability of selecting action $a \in \mathcal{A}$ for any given state $s \in \mathcal{S}$. Following policy π on the MDP, the agent observes a sequence of random variables $\{S_t, A_t\}_t$ where S_t denotes the state of the agent at time t and A_t denotes the action taken by the agent at time t . The expected discounted long term reward and expected per step reward of the objective k are given by $J_\pi^{P,k}$ and $\lambda_\pi^{P,k}$, respectively, when the joint policy π is followed. Formally, for discount factor $\gamma \in (0, 1)$, $J_\pi^{P,k}$ and $\lambda_\pi^{P,k}$ are defined as

$$J_\pi^{P,k} = \mathbb{E}_{S_0, A_0, S_1, A_1, \dots} \left[\lim_{\tau \rightarrow \infty} \sum_{t=0}^{\tau} \gamma^t r^k(S_t, A_t) \right] \quad (1)$$

$$S_0 \sim \rho_0(S_0), A_t \sim \pi(A_t|S_t), S_{t+1} \sim P(\cdot|S_t, A_t)$$

$$\lambda_\pi^{P,k} = \mathbb{E}_{S_0, A_0, S_1, A_1, \dots} \left[\lim_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau} r^k(S_t, A_t) \right] \quad (2)$$

$$= \lim_{\gamma \rightarrow 1} (1 - \gamma) J_\pi^{P,k} \quad (3)$$

Equation (3) follows from the Laurent series expansion of J_π^k (Puterman, 1994). For brevity, in the rest of the paper $\mathbb{E}_{S_t, A_t, S_{t+1}; t \geq 0}[\cdot]$ will be denoted as $\mathbb{E}_{\rho, \pi, P}[\cdot]$, where $S_0 \sim \rho_0$, $A_t \sim \pi(\cdot|S_t)$, $S_{t+1} \sim P(\cdot|S_t, A_t)$. The expected per step reward satisfies the following Bellman equation

$$h_\pi^{P,k}(s) + \lambda_\pi^{P,k} = \mathbb{E}_{a \sim \pi} \left[r^k(s, a) \right] + \mathbb{E}_{a \sim \pi} \left[\sum_{s'} P(s'|s, a) h_\pi^{P,k}(s') \right] \quad (4)$$

where $h_\pi^k(s)$ is the bias of policy π for state s . We also define the discounted value function $V_{\gamma, \pi}^{P,k}(s)$ and Q-value functions $Q_{\gamma, \pi}^{P,k}(s, a)$ as follows:

$$V_{\gamma, \pi}^{P,k}(s) = \mathbb{E}_{A_0, S_1, A_1, \dots} \left[\lim_{\tau \rightarrow \infty} \sum_{t=0}^{\tau} \gamma^t r^k(S_t, A_t) | S_0 = s \right] \quad (5)$$

$$\begin{aligned} Q_{\gamma, \pi}^{P,k}(s, a) &= \mathbb{E}_{S_1, A_1, \dots} \left[\lim_{\tau \rightarrow \infty} \sum_{t=0}^{\tau} \gamma^t r^k(S_t, A_t) | S_0 = s, A_0 = a \right] \\ &= r^k(s, a) + \mathbb{E}_{a \sim \pi} \left[\sum_{s'} P(s'|s, a) V_{\gamma, \pi}^{P,k}(s') \right] \end{aligned} \quad (6)$$

Further, the bias $h_\pi^{P,k}(s)$ and the value function $V_{\gamma, \pi}^{P,k}$ are related as

$$h_\pi^{P,k}(s_1) - h_\pi^{P,k}(s_2) = \lim_{\gamma \rightarrow 1} \left(V_{\gamma, \pi}^{P,k}(s_1) - V_{\gamma, \pi}^{P,k}(s_2) \right), \text{ where } s_1, s_2 \in \mathcal{S}. \quad (7)$$

For notation simplicity, we may drop the superscript P when discussing about variables for the true MDP.

Note that each policy induces a Markov Chain on the states \mathcal{S} with transition probabilities $P_{\pi, s}(s') = \sum_{a \in \mathcal{A}} \pi(a|s) P(s'|s, a)$. After defining a policy, we can now define the diameter of the MDP \mathcal{M} as:

Definition 1 (Diameter) Consider the Markov Chain induced by the policy π on the MDP \mathcal{M} . Let $T(s'|\mathcal{M}, \pi, s)$ be a random variable that denotes the first time step when this Markov Chain enters state s' starting from state s . Then, the diameter of the MDP \mathcal{M} is defined as:

$$D(\mathcal{M}) = \max_{\pi} \max_{s' \neq s} \mathbb{E} [T(s'|\mathcal{M}, \pi, s)] \quad (8)$$

Further, starting from an arbitrary initial state distribution, the state distribution may take a while to converge to the steady state distribution. For any policy π , let $P_{\pi,s}^t = (P_{\pi,s})^t$ be the t -step probability distribution of the states when policy π is applied to MDP with transition probabilities P starting from state s .

We are now ready to state our first assumption on the MDP. We assume that the Markov Decision Process is ergodic. This implies that: **1.** for any policy all states, $s \in \mathcal{S}$, communicate with each other; **2.** for any policy, the process converges to the steady state distribution exponentially fast. Formally, we have

Assumption 1 The Markov Decision Process, \mathcal{M} , is ergodic. Then, we have,

1. The diameter, D , of \mathcal{M} is finite.
2. For any policy π , for some $C > 0$ and $0 \leq \rho < 1$, we have

$$\|P_{\pi,s}^t - d_{\pi}\|_{TV} \leq C\rho^t \quad (9)$$

where d_{π} is the steady state distribution induced by policy π on the MDP.

The agent aim to collaboratively optimize the function $f : \mathbb{R}^K \rightarrow \mathbb{R}$, which is defined over the long-term rewards of the individual objectives. We make certain practical assumptions on this joint objective function f , which are listed as follows:

Assumption 2 The objective function f is jointly concave. Hence for any arbitrary distribution \mathcal{D} , the following holds.

$$f(\mathbb{E}_{\mathbf{x} \sim \mathcal{D}}[\mathbf{x}]) \geq \mathbb{E}_{\mathbf{x} \sim \mathcal{D}}[f(\mathbf{x})]; \mathbf{x} \in \mathbb{R}^K \quad (10)$$

The objective function f represents the utility obtained from the expected per step reward for each objective. These utility functions are often concave to reduce variance for a risk averse decision maker (Pratt, 1964). To model this concave utility function, we assume the above form of Jensen's inequality.

Assumption 3 The function f is assumed to be a L - Lipschitz function, or

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq L \|\mathbf{x} - \mathbf{y}\|_1; \mathbf{x}, \mathbf{y} \in \mathbb{R}^K \quad (11)$$

Assumption 3 ensures that for a small change in long run rewards for any objective does not cause unbounded changes in the objective.

Based on Assumption 2, we maximize the function of expected sum of rewards for each objective. Further to keep the formulation independent of time horizon or γ , we maximize the function over expected per-step rewards of each objective. Hence, our goal is to find the optimal policy as the solution for the following optimization problem.

$$\pi^* = \arg \max_{\pi} f(\lambda_{\pi}^1, \dots, \lambda_{\pi}^K) \quad (12)$$

If $f(\cdot)$ is also monotone, we note that the optimal policy in (12) can be shown to be Pareto optimal. The detailed proof will be presented later in Section 4.

Any online algorithm \mathcal{A} starting with no prior knowledge will require to obtain estimates of transition probabilities P and obtain rewards $r_k, \forall k \in [K]$ for each state action pair. Initially, when algorithm \mathcal{A} does not have good estimates of the model, it accumulates a regret for not working as per optimal policy. We define a time dependent regret R_T to achieve an optimal solution defined as the difference between the optimal value of the function and the value of the function at time T , or

$$R_T = \mathbb{E}_{S_t, A_t} \left[\left| f(\lambda_{\pi^*}^1, \dots, \lambda_{\pi^*}^K) - f\left(\frac{1}{T} \sum_{t=0}^T r^1(S_t, A_t), \dots, \frac{1}{T} \sum_{t=0}^T r^K(S_t, A_t)\right) \right| \right] \quad (13)$$

The regret defined in Equation (13) is the expected deviation between the value of the function obtained from the expected rewards of the optimal policy and the value of the function obtained from the observed rewards from a trajectory. Following the work of (Roijers et al., 2013), we note that the outer expectation comes for running the decision process for a different set of users or running a separate and independent instance for the same set of users. Since the realization can be different from the expected rewards, the function values can still be different even when following the optimal policy.

We note that we do not require $f(\cdot)$ to be monotone. Thus, even for a single objective, optimizing $f(\mathbb{E}[\sum_t r_t])$ is not equivalent to optimizing $\mathbb{E}[\sum_t r_t]$. Hence, the proposed framework can be used to optimize functions of cumulative or long term average reward for single objective as well.

In the following sections, we first show that the joint-objective function of average rewards allows us to obtain Pareto-optimal policies with an additional assumption of monotonicity. Then, we will present a model-based algorithm to obtain this policy π^* , and regret accumulated by the algorithm. We will present a model-free algorithm in Section 6 which can be efficiently implemented using Deep Neural Networks.

4. Obtaining Pareto-Optimal Policies

Many multi-objective or multi-agent formulations require the policies to be Pareto-Optimal (Roijers et al., 2013; Sener and Koltun, 2018; Van Moffaert and Nowé, 2014). The conflicting rewards of various agents may not allow us to attain simultaneous optimal average rewards for any agent with any joint policy. Hence, an envy-free Pareto optimal policy is desired. We now provide an additional assumption on the joint objective function, and show that the optimal policy satisfying Equation (12) is Pareto optimal.

Assumption 4 *If f is an element-wise monotonically strictly increasing function. Or, $\forall k \in [K]$, the function satisfies,*

$$x^k > y^k \implies f(x^1, \dots, x^k, \dots, x^K) > f(x^1, \dots, y^k, \dots, x^K) \quad (14)$$

Element wise increasing property motivates the agents to be strategic as by increasing its per-step average reward, agent can increase the joint objective. Based on Equation (14), we notice that the solution for Equation (12) is Pareto optimal.

Definition 2 A policy π^* is said to be **Pareto optimal** if and only if there is exists no other policy π such that the average per-step reward is at least as high for all agents, and strictly higher for at least one agent. In other words,

$$\forall k \in [K], \lambda_{\pi^*}^k \geq \lambda_{\pi}^k \text{ and } \exists k, \lambda_{\pi^*}^k > \lambda_{\pi}^k \quad (15)$$

Theorem 3 Solution of Equation (12), or the optimal policy π^* is Pareto Optimal.

Proof We will prove the result using contradiction. Let π^* be the solution of Equation (12) and not be Pareto optimal. Then there exists some policy π for which the following equation holds,

$$\forall k \in [K], \lambda_{\pi}^k \geq \lambda_{\pi^*}^k \text{ and } \exists k, \lambda_{\pi}^k > \lambda_{\pi^*}^k \quad (16)$$

From element-wise monotone increasing property in Equation (14), we obtain

$$f(\lambda_{\pi^*}^1, \dots, \lambda_{\pi}^k, \dots, \lambda_{\pi^*}^K) > f(\lambda_{\pi^*}^1, \dots, \lambda_{\pi^*}^k, \dots, \lambda_{\pi^*}^K) \quad (17)$$

$$= \arg \max_{\pi'} f(\lambda_{\pi'}^1, \dots, \lambda_{\pi'}^K) \quad (18)$$

This is a contradiction. Hence, π^* is a Pareto optimal solution. ■

This result shows that algorithms presented in this paper can be used to optimally allocate resources among multiple agents using average per step allocations.

5. Model-based Algorithm

RL problems typically optimize the cumulative rewards, which is a linear function of rewards at each time step because of the addition operation. This allows the Bellman Optimality Equation to require only the knowledge of the current state to select the best action to optimize future rewards (Puterman, 1994). However, since our controller is optimizing a joint non-linear function of the long-term rewards from multiple sources, Bellman Optimality Equations cannot be written as a function of the current state exclusively. Our goal is to find the optimal policy as solution of Equation (12). Using average per-step reward and infinite horizon allows us to use Markov policies. An intuition into why this works is there is always infinite time available to optimize the joint objective f .

The individual long-term average-reward for each agent is still linearly additive ($\frac{1}{\tau} \sum_{t=0}^{\tau} r^k(S_t, A_t)$). For infinite horizon optimization problems (or $\tau \rightarrow \infty$), we can use steady state distribution of the state to obtain expected cumulative rewards. For all $k \in [K]$, we use

$$\lambda_{\pi}^k = \sum_{s \in S} \sum_{a \in A} r^k(s, a) d_{\pi}(s, a) \quad (19)$$

where $d_{\pi}(s, a)$ is the steady state joint distribution of the state and actions under policy π . Equation (19) suggests that we can transform the optimization problem in terms of optimal

policy to optimal steady-state distribution. Thus, we have the joint optimization problem in the following form which uses steady state distributions

$$d^* = \arg \max_d f \left(\sum_{s \in \mathcal{S}, a \in \mathcal{A}} r^1(s, a)d(s, a), \dots, \sum_{s \in \mathcal{S}, a \in \mathcal{A}} r^K(s, a)d(s, a) \right) \quad (20)$$

with the following set of constraints,

$$\sum_{a \in \mathcal{A}} d(s', a) = \sum_{s \in \mathcal{S}, a \in \mathcal{A}} P(s'|s, a)d(s, a) \quad \forall s' \in \mathcal{S} \quad (21)$$

$$\sum_{s \in \mathcal{S}, a \in \mathcal{A}} d(s, a) = 1 \quad (22)$$

$$d(s, a) \geq 0 \quad \forall s \in \mathcal{S}, a \in \mathcal{A} \quad (23)$$

Constraint (21) denotes the transition structure for the underlying Markov Process. Constraint (22), and constraint (23) ensures that the solution is a valid probability distribution. Since $f(\dots)$ is jointly concave, arguments in Equation (20) are linear, and the constraints in Equation (21,22,23) are linear, this is a convex optimization problem. Since convex optimization problems can be solved in polynomial time (Bubeck et al., 2015), we can use standard approaches to solve Equation (20). After solving the optimization problem, we find the optimal policy from the obtained steady state distribution $d^*(s, a)$ as,

$$\pi^*(a|s) = \frac{Pr(a, s)}{Pr(s)} = \frac{d^*(a, s)}{\sum_{a \in \mathcal{A}} d^*(s, a)} \quad (24)$$

The proposed model-based algorithm estimates the transition probabilities by interacting with the environment. We need the steady state distribution d_π to exist for any policy π . We note that when the priors of the transition probabilities $P(\cdot|s, a)$ are a Dirichlet distribution for each state and action pair, such a steady state distribution exists. Proposition 4 formalizes the result of the existence of a steady state distribution when the transition probability is sampled from a Dirichlet distribution.

Proposition 4 *For MDP \mathcal{M} with state space \mathcal{S} and action space \mathcal{A} , let the transition probabilities P come from a Dirichlet distribution. Then, any policy π for \mathcal{M} will have a steady state distribution \hat{d}_π given as*

$$\hat{d}_\pi(s') = \sum_{s \in \mathcal{S}} \hat{d}_\pi(s) \left(\sum_{a \in \mathcal{A}} \pi(a|s) P(s'|s, a) \right) \quad \forall s' \in \mathcal{S}. \quad (25)$$

Proof The transition probabilities $P(s'|s, a)$ follow Dirichlet distribution, and hence they are strictly positive. Further, as the policy $\pi(a|s)$ is a probability distribution on actions conditioned on state, $\pi(a|s) \geq 0$, $\sum_a \pi(a|s) = 1$. So, there is a non zero transition probability to reach from state $s \in \mathcal{S}$ to state $s' \in \mathcal{S}$.

Now, note that all the entries of the transition probability matrix are strictly positive. And, hence the Markov Chain induced over the MDP \mathcal{M} by any policy π is 1) irreducible,

as it is possible to reach any state from any other state, and 2) aperiodic, as it is possible to reach any state in a single time step from any other state. Together, we get the existence of the steady-state distribution (Lawler, 2018). \blacksquare

To complete the setup for our algorithm, we make few more assumptions stated below.

Assumption 5 *The transition probabilities $P(\cdot|s, a)$ of the Markov Decision Process have a Dirichlet prior for all state action pairs (s, a) .*

Since we assume that transition probabilities of the MDP \mathcal{M} follow Dirichlet distributions, all policies on \mathcal{M} have a steady-state distribution.

Algorithm 1 Model-Based Joint Decision Making Algorithm

```

1: procedure MODEL BASED ONLINE( $\mathcal{S}, \mathcal{A}, [K], f, r$ )
2:   Initialize  $N(s, a, s') = 1 \ \forall (s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ 
3:   Initialize  $\nu_0(s, a) = 0 \ \forall (s, a) \in \mathcal{S} \times \mathcal{A}$ 
4:   Initialize  $\pi_1(a|s) = \frac{1}{|\mathcal{A}|} \ \forall (a, s) \in \mathcal{A} \times \mathcal{S}$ 
5:   Initialize  $e = 1$ 
6:   for time steps  $t = 1, 2, \dots$  do
7:     Observe current state  $S_t$ 
8:     Sample action to play  $A_t \sim \pi_e(\cdot|S_t)$ 
9:     Play  $A_t$ , obtain reward  $r_t \in [0, 1]^K$  and observe next state  $S_{t+1}$ 
10:    Update  $N(S_t, A_t, S_{t+1}) \leftarrow N(S_t, A_t, S_{t+1}) + 1, \nu_e(S_t, A_t) \leftarrow \nu_e(S_t, A_t) + 1$ 
11:    if  $\nu_e(S_t, A_t) \geq \max\{1, \sum_{e' < e} \nu_{e'}(S_t, A_t)\}$  then
12:      Update epoch counter  $e \leftarrow e + 1$ 
13:      Initialize  $\nu_e(s, a) = 0 \ \forall (s, a) \in \mathcal{S} \times \mathcal{A}$ 
14:       $P_e(s'|s, a) \sim \text{Dir}(N(s, a, \cdot)) \ \forall (s, a)$ 
15:      Solve steady state distribution  $d(s, a)$  as the solution of the optimization
      problem in Equations (20-23)
16:      Obtain optimal policy  $\pi_e$  as
          
$$\pi_e(a|s) = \frac{d(s, a)}{\sum_{a \in \mathcal{A}} d(s, a)}$$

17:    end if
18:  end for
19: end procedure

```

5.1 Algorithm Description

Algorithm 1 describes the overall procedure that estimates the transition probabilities and the reward functions. The algorithm takes as input the state space \mathcal{S} , action space \mathcal{A} , set of agents $[K]$, the reward structure r , and the objective function f . It initializes the next state visit count for each state-action pair $N(s, a, s')$ by one for Dirichlet sampling. For initial exploration, the policy uses a uniform distribution over all the actions. The algorithm

proceeds in epochs e with t_e denoting the time step t at which epoch e starts. Also, $\nu_e(s, a)$ stores the number of times a state-action pair is visited in epoch e . We assume that the controller is optimizing for infinite horizon and thus there is no stopping condition for epoch loop in Line 6. For each time index $t \in [t_e, t_{e+1})$ in epoch e , the controller observes the state, samples and plays the action according to $\pi_e(\cdot|s)$, and observes the rewards for each agent and next state. It then updates the state visit count for the observed state and played action pair. We break the current epoch e if the total number of visitations of any state action pair in current epoch exceeds total visitations in previous epoch. This epoch breaking strategy ensures that each state-action triggers epoch switches exponentially slower. This increasing epoch duration bounds the policy switching cost (Jaksch et al., 2010). After breaking the epoch, we sample transition probabilities from the updated posterior and find a new optimal policy for the sampled MDP by solving the optimization framework described in Equation(20).

5.2 Regret

We now prove the regret bounds of Algorithm 1 in the form of the following theorem. We first give the high level ideas used in obtaining the bounds on regret. The Lipschitz property of the objective function f allows to write the regret as sum of regrets for individual per-step average reward λ_π^k . Now, we can divide the regret into regret incurred in each epoch e . Then, we use the posterior sampling lemma (Lemma 1 from (Osband et al., 2013)) to obtain the equivalence between the value of the function f for the optimal policy of the true MDP \mathcal{M} and the value of the function for the optimal value of the sampled MDP $\widehat{\mathcal{M}}$ conditioned on the observed state evolution. Now, for any epoch e , note that for each reward k , the agent plays policy π_e optimized for the sampled MDP $\widehat{\mathcal{M}}$ on the true MDP \mathcal{M} . We break the difference between the two terms, the per-step average reward $\lambda_{\pi_e}^{P_e, k}$ of the policy π_e for the sampled MDP with transition probability P_e and the rewards obtained into two terms. The first term is the difference $\lambda_{\pi_e}^{P_e, k}$ and $\lambda_{\pi_e}^{P, k}$. We will denote our immediate discussion on bounding this term. The second deviation is the difference between $\lambda_{\pi_e}^{P, k}$ and the observed rewards r_t^k for $t \in [t_e, t_{e+1})$, which we can bound using Azuma-Hoeffding's inequality.

To compute the regret incurred by the optimal policy π_e for the sampled MDP on the true MDP \mathcal{M} , we use Bellman error. For some policy π , we define Bellman error $B_\pi^{\tilde{P}, k}(s, a)$ for the infinite horizon MDPs as the difference between the cumulative expected rewards obtained for deviating from the system model with transition \tilde{P} for one step by taking action a in state s and then following policy π . We have:

$$\begin{aligned} B_{\gamma, \pi}^{\tilde{P}, k}(s, a) &= \left(Q_{\gamma, \pi}^{\tilde{P}, k}(s, a) - r(s, a) - \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) V_{\gamma, \pi}^{\tilde{P}, k}(s, a) \right) \\ B_\pi^{\tilde{P}, k}(s, a) &= \lim_{\gamma \rightarrow 1} B_{\gamma, \pi}^{\tilde{P}, k}(s, a) \end{aligned} \tag{26}$$

We relate the Bellman error defined in Equation (26) to the gap between the expected per step reward $\lambda_\pi^{\tilde{P}, k}$ for running policy π on an MDP with transition probability \tilde{P} and the expected per step reward $\lambda_\pi^{P, k}$ for running policy π on the true MDP in the following lemma:

Lemma 5 *The difference of long-term average rewards for running the policy π_e on the MDP, $\lambda_\pi^{\tilde{P}, k}$, and the average long-term average rewards for running the policy π on the true*

MDP, $\lambda_\pi^{\tilde{P},k}$, is the long-term average Bellman error as

$$\lambda_\pi^{\tilde{P},k} - \lambda_\pi^{P,k} = \sum_{s,a} d_\pi(s,a) B_\pi^{\tilde{P},k}(s,a) = \mathbb{E}_{(s,a) \sim \pi, P} \left[B_\pi^{\tilde{P},k}(s,a) \right] \quad (27)$$

where d_π is the occupancy measure generated by policy π on the true MDP and $\mathbb{E}_{(s,a) \sim \pi, P}[\cdot]$ denotes expectation over s, a sampled when running policy π on MDP with transition probability P .

Proof [Proof Sketch] We start by writing $Q_{\gamma, \pi}^{\tilde{P},k}$ in terms of the Bellman error. Now, subtracting $V_{\gamma, \pi}^{P,k}$ from $V_{\gamma, \pi}^{\tilde{P},k}$ and using the fact that $\lambda_{\pi_e}^{P,k} = \lim_{\gamma \rightarrow 1} (1 - \gamma) V_{\gamma, \pi}^{P,k}$ and $\lambda_{\pi_e}^{\tilde{P},k} = \lim_{\gamma \rightarrow 1} (1 - \gamma) V_{\gamma, \pi_e}^{\tilde{P},k}$ we obtain the required result. A complete proof is provided in Appendix B.1. \blacksquare

After relating the gap between the long-term average rewards of policy π_e on the two MDPs, we now want to bound the sum of Bellman error over an epoch. For this, we first bound the Bellman error for a particular state action pair s, a in the form of following lemma. We have,

Lemma 6 For an MDP with rewards $r^k(s, a)$ and transition probability $\tilde{P}(s'|s, a)$ such that $\|\tilde{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 \leq \epsilon_{s,a}$, the Bellman error $B_{\pi_e}^{\tilde{P},k}(s, a)$ for state-action pair s, a is upper bounded as

$$B_\pi^{\tilde{P},k}(s, a) \leq \left\| \tilde{P}(\cdot|s, a) - P(\cdot|s, a) \right\|_1 \|h_\pi^{\tilde{P},k}(\cdot)\|_\infty \leq \tilde{D} \min \{ \epsilon_{s,a}, 2 \}, \quad (28)$$

where \tilde{D} is the diameter of the MDP with transition probability \tilde{P} .

Proof [Proof Sketch] We start by noting that the Bellman error essentially bounds the impact of the difference in value obtained because of the difference in transition probability to the immediate next state. A complete proof is provided in Appendix B.2. \blacksquare

A major part of the regret analysis is how well the learned model estimates the true system model. For that, bound the deviation of the estimates of the estimated transition probabilities of the Markov Decision Processes \mathcal{M} . For that we use ℓ_1 deviation bounds from (Weissman et al., 2003). Consider, the following event,

$$\mathcal{E}_t = \left\{ \left\| \bar{P}_t(\cdot|s, a) - P(\cdot|s, a) \right\|_1 \leq \sqrt{\frac{14S \log(2AT)}{\max\{1, n_t(s, a)\}}} \forall (s, a) \in \mathcal{S} \times \mathcal{A} \right\} \quad (29)$$

where $n_t(s, a) = \sum_{t'=1}^t \mathbf{1}_{\{S_{t'}=s, a_{t'}=a\}}$ and $\bar{P}(s'|s, a) = \left(\sum_{t'=1}^t \mathbf{1}_{\{S_{t'}=s, a_{t'}=a, S_{t'+1}=s'\}} \right) / n_t(s, a)$ of the empirical estimate of the transition probabilities. Then we have, the following lemma:

Lemma 7 The probability that the event \mathcal{E}_t fails to occur is upper bounded by $\frac{1}{20t^6}$.

We also want to ensure that the behavior of the sampled MDP from the posterior distribution is identical to the behavior of the true MDP conditioned on the observed state, action evolution. For this, we use the following lemma.

Lemma 8 (Posterior Sampling Lemma 1 (Osband et al., 2013)) For any $\sigma(H_t)$ -measurable function g , if P follows distribution ϕ , then for transition probabilities P_e sampled from $\phi(\cdot|H_t)$ we have,

$$\mathbb{E}[g(P)|\sigma(H_t)] = \mathbb{E}[g(P_e)|\sigma(H_t)] \quad (30)$$

Lastly, we have the following lemma to bound the number of policy switches to ensure that after every policy switch, the stochastic process of the state evolution gets sufficient steps to converge towards stationary distribution.

Lemma 9 [(Jaksch et al., 2010)[Proposition 18]] The total number of epochs E for the Algorithm 1 with $\nu_e(s, a) \leq \sum_{e'}^{e-1} \max\{1, \nu_{e'}(s, a)\}$ for any s, a , up to step $T \geq SA$ is upper bounded as

$$E \leq 1 + 2AS + AS \log_2 \left(\frac{T}{SA} \right) \quad (31)$$

Using Lemma 5, Lemma 6, Lemma 8 and Lemma 7, we can now bound the regret of Algorithm 1 in the form of following theorem.

Theorem 10 The regret R_T of Algorithm 1 for MDP with Dirichlet priors and diameter D is bounded.

$$R_T \leq \tilde{O} \left(LKDS \sqrt{\frac{A}{T}} + \frac{1}{T} \frac{SACD}{1-\rho} \right) \quad (32)$$

Proof We use the Lipschitz continuity of the function to break the scalarized objective into long-term average reward regrets of individual objectives. Using Lipschitz continuity, the total regret becomes the sum of individual regrets.

$$\mathbb{E}[R_T] = \mathbb{E} \left[\left| f \left(\dots, \lambda_{\pi^*}^{P,k}, \dots \right) - f \left(\dots, \frac{1}{T} \sum_{t=0}^T r^k(S_t, A_t), \dots \right) \right| \right] \quad (33)$$

$$= \mathbb{E} \left[\frac{L}{T} \sum_{k=1}^K \left| T \lambda_{\pi^*}^{P,k} - \sum_{t=0}^T r^k(S_t, A_t) \right| \right] \leq \frac{LK}{T} \max_{k \in [K]} \mathbb{E} \left[\left| T \lambda_{\pi^*}^{P,k} - \sum_{t=0}^T r^k(S_t, A_t) \right| \right]. \quad (34)$$

We can divide regret for any objective $k \in [K]$ over T time steps into regret accumulated over episodes as:

$$\mathbb{E} \left[\left| T \lambda_{\pi^*}^{P,k} - \sum_{t=1}^T r^k(S_t, A_t) \right| \right] = \mathbb{E} \left[\left| \sum_{e=1}^E \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi^*}^{P,k} - \sum_{t=1}^T r^k(S_t, A_t) \right) \right| \right] \quad (35)$$

Now, we note that the regret in each episode is conditioned on filtration H_{t_e} , and is $\sigma(H_{t_e})$ -measurable. Hence, we can use (Osband et al., 2013, Posterior Sampling Lemma) and (Osband et al., 2013, Regret Equivalence Theorem) to obtain the equivalence between the per-step average reward of the optimal policy for the true MDP and the per-step average reward of the optimal policy of the sampled MDP. We have:

$$\mathbb{E} \left[\left| \sum_{e=1}^E \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi^*}^{P,k} - \sum_{t=1}^T r^k(S_t, A_t) \right) \right| \right] \leq \mathbb{E} \left[\sum_{e=1}^E \left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi^*}^{P,k} - \sum_{t=1}^T r^k(S_t, A_t) \right) \right| \right] \quad (36)$$

$$= \mathbb{E} \left[\sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi^*}^{P,k} - \sum_{t=1}^T r^k(S_t, A_t) \right) \right| \middle| H_{t_e} \right] \right] \quad (37)$$

$$= \mathbb{E} \left[\sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_e,k} - \sum_{t=1}^T r^k(S_t, A_t) \right) \right| \middle| H_{t_e} \right] \right] \quad (38)$$

$$= \mathbb{E} \left[\sum_{e=1}^E \left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_e,k} - \sum_{t=1}^T r^k(S_t, A_t) \right) \right| \right] \quad (39)$$

We now consider two cases. The first case, (a), is where the estimated system model or the sampled system model are not close to the true system model. The second case, (b), is where both the estimated system model and the sampled system model are close the true system model. The total regret can be bounded by using the law of total expectation. Also, to reduce notational clutter, we calculate the regret conditioned on the number of episodes E or $\mathbb{E}[R(T)|E]$. We can then remove the dependency on number of episodes by considering the largest possible number of episodes from Lemma 9.

We start by characterizing how close are the estimated transition probability and the sampled transition probability are to the true transition probability. We use ℓ_1 distance metric for this. For all s, a , we construct the set of probability distributions $P'(\cdot|s, a)$,

$$\mathcal{P}_t = \left\{ P' : \|\bar{P}_t(\cdot|s, a) - P'(\cdot|s, a)\|_1 \leq \sqrt{\frac{14S \log(2AT)}{\max\{1, n_t(s, a)\}}} \right\} \quad (40)$$

where $n_t = \sum_{t'} \mathbf{1}_{\{S_{t'}=s, a_{t'}=a\}}$. Using the construction of the set \mathcal{P}_t , we can now define the events \mathcal{E}_t , and $\hat{\mathcal{E}}_t$ as:

$$\mathcal{E}_t = \{P \in \mathcal{P}_t\}, \text{ and } \hat{\mathcal{E}}_t = \{P_e \in \mathcal{P}_t\} \quad (41)$$

Further, note that \mathcal{P}_t is $\sigma(H_t)$, $H_t = \{s_1, a_1, \dots, s_t, a_t\}$ measurable and hence from Lemma 8 we have $\mathbb{P}(\hat{\mathcal{E}}_t) = \mathbb{P}(P_e \in \mathcal{P}_t) = \mathbb{P}(P \in \mathcal{P}_t) = \mathbb{P}(\mathcal{E}_t)$.

We first bound the regret for the case where the system model is not well estimated or the sampled system from the posterior distribution is far from the true system model. This is equivalent to considering the event in Equation (40) does not occur or the complementary events $\mathcal{E}_t^c, \hat{\mathcal{E}}_t^c$. We already bounded the probability of this event in Lemma 7 using result from (Weissman et al., 2003). In particular, we have:

$$\sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_{e,k}} - r^k(S_t, A_t) \right) \right| \middle| \mathcal{E}_{t_e}^c \cup \widehat{\mathcal{E}}_{t_e}^c \right] \leq \sum_{e=1}^E \mathbb{E} \left[\sum_{t=t_e}^{t_{e+1}-1} 1 \middle| \mathcal{E}_{t_e}^c \cup \widehat{\mathcal{E}}_{t_e}^c \right] \quad (42)$$

$$\leq \sum_{e=1}^E \sum_{s,a} \nu_k(s, a) \mathbb{P}(\mathcal{E}_{t_e}^c) \quad (43)$$

$$\leq \sum_{e=1}^E t_e \mathbb{P}(\mathcal{E}_{t_e}^c \cup \widehat{\mathcal{E}}_{t_e}^c) \quad (44)$$

$$\leq \sum_{t=1}^T t \left(\mathbb{P}(\mathcal{E}_t^c) + \mathbb{P}(\widehat{\mathcal{E}}_{t_e}^c) \right) \quad (45)$$

$$= 2 \sum_{t=1}^T t \mathbb{P}(\mathcal{E}_t^c) \quad (46)$$

$$\leq 2 \sum_{t=1}^{T^{1/4}} t \mathbb{P}(\mathcal{E}_t^c) + 2 \sum_{t=T^{1/4}+1}^T t \mathbb{P}(\mathcal{E}_t^c) \quad (47)$$

$$\leq 2 \sum_{t=1}^{T^{1/4}} t \cdot 1 + 2 \sum_{t=T^{1/4}+1}^T t \frac{1}{t^6} \quad (48)$$

$$\leq 2\sqrt{T} + 2 \int_{t=T^{1/4}}^{\infty} \frac{1}{t^5} \quad (49)$$

$$\leq 2\sqrt{T} + 2 \frac{1}{4T} \leq 4\sqrt{T} \quad (50)$$

where Equation (42) follows from the fact that rewards are bounded by 1 and the modulus operator is not required on the positive sum. Equation (45) follows from the fact that $\sum_{s,a} \nu_k(s, a) \leq \sum_{s,a} N_k(s, a) = t_k$. Further, Equation (48) follows from Lemma 7. This completes the first case (a).

For the second case, we now break the regret into two terms as follows:

$$\begin{aligned} & \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_{e,k}} - r^k(S_t, A_t) \right) \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] \\ &= \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_{e,k}} - \lambda_{\pi_e}^{P,k} + \lambda_{\pi_e}^{P,k} - r^k(S_t, A_t) \right) \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] \end{aligned} \quad (51)$$

$$\begin{aligned} & \leq \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_{e,k}} - \lambda_{\pi_e}^{P,k} \right) \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] + \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P,k} - r^k(S_t, A_t) \right) \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] \\ &= R_1^k(T) + R_2^k(T). \end{aligned} \quad (52)$$

The first term, $R_1^k(T)$, denotes the gap of running the optimal policy for the sampled policy on the true MDP in an epoch e . We bound this term with the Bellman error defined

in Equation (26). The second term, $R_2^k(T)$, denotes the regret incurred from the deviation of the observed rewards and the expected per step rewards.

We now focus on the R_1^k term. We begin with using Lemma 5 to replace $\lambda_{\pi_e}^{P_e,k} - \lambda_{\pi_e}^{P,k}$. We have the following series of inequalities.

$$R_1^k(T) = \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} (\lambda_{\pi_e}^{P_e,k} - \lambda_{\pi_e}^{P,k}) \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] \quad (53)$$

$$\leq \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} (\lambda_{\pi_e}^{P_e^k} - \lambda_{\pi_e}^{P,k}) \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] \quad (54)$$

$$= \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \mathbb{E}_{(s,a) \sim \pi_e, P} [B_{\pi_e}^{P_e^k, k}(s, a)] \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] \quad (55)$$

where P_e^k is the transition probability for which π_e maximizes $\lambda_{\pi}^{P',k}$ for $P' \in \mathcal{P}_{t_e}$. From Lemma 12 in Appendix, the bias-span $\|h_{\pi_e}^{P_e^k, k}\|_{\infty}$ is upper bounded by D . Further, since $P' \in \mathcal{P}_{t_e}$ we have $\|P_e^k(\cdot|s, a) - P(\cdot|s, a)\|_1 \leq \sqrt{\frac{14S \log(2AT)}{1Vn_{t_e}(s,a)}}$ for all s, a . Hence, we have $B_{\pi_e}^{P_e^k, k} \leq \min\{2D, D\sqrt{\frac{14S \log(2AT)}{1Vn_{t_e}(s,a)}}\}$.

We now need to bound the expected value in Equation (55). Note that conditioned on filtration H_{t_e} the two expectations $\mathbb{E}_{s,a \sim \pi_e, P}[\cdot]$ and $\mathbb{E}_{s,a \sim \pi_e, P}[\cdot|H_{t_e-1}]$ are not equal as the former is the expected value of the long-term state distribution and the later is the long-term state distribution condition on initial state s_{t_e-1} . We now use Assumption 1 to obtain the following set of inequalities.

$$\begin{aligned} \mathbb{E}_{(s,a) \sim \pi_e, P} [B_{\pi_e}^{P_e^k, k}(s, a)] &= \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} [B_{\pi_e}^{P_e^k, k}(S_t, A_t) | H_{t_e-1}] \\ &\quad + \left(\mathbb{E}_{(s,a) \sim \pi_e, P} [B_{\pi_e}^{P_e^k, k}(s, a)] - \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} [B_{\pi_e}^{P_e^k, k}(S_t, A_t) | H_{t_e-1}] \right) \end{aligned} \quad (56)$$

$$\begin{aligned} &\leq \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} [B_{\pi_e}^{P_e^k, k}(S_t, A_t) | H_{t_e-1}] \\ &\quad + 2D \left(\left| \sum_{a,s} (\pi_e(a|s) d_{\pi_e}(s) - \pi_e(a|s) P_{\pi, S_{t_e-1}}^{t-t_e+1}(s)) \right| \right) \end{aligned} \quad (57)$$

$$\begin{aligned} &= \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} [B_{\pi_e}^{P_e^k, k}(S_t, A_t) | H_{t_e-1}] \\ &\quad + 2D \left(\left| \sum_s \left((d_{\pi_e}(s) - P_{\pi, S_{t_e-1}}^{t-t_e+1}(s)) \left(\sum_a \pi_e(a|s) \right) \right) \right| \right) \end{aligned} \quad (58)$$

$$\begin{aligned} &= \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} [B_{\pi_e}^{P_e^k, k}(S_t, A_t) | H_{t_e-1}] \\ &\quad + 2D \left(2 \|d_{\pi_e}(s) - P_{\pi, S_{t_e-1}}^{t-t_e+1}(s)\|_{TV} \right) \end{aligned} \quad (59)$$

$$\leq \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} [B_{\pi_e}^{P_e^k, k}(S_t, A_t) | H_{t_e-1}] + 4CD\rho^{t-t_e} \quad (60)$$

where Equation (59) comes from Assumption 1 for running policy π_e starting from state s_{t_e-1} for $t - t_e + 1$ steps and from Lemma 6. Equation (60) follows from the fact that

$\sum_a \pi(a|s) = 1$ and the fact that ℓ_1 -distance of probability distribution is twice the total variation distance (Levin and Peres, 2017, Proposition 4.2).

Using Equation (60) with Equation (55) we get,

$$R_1^k(T) \leq \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \mathbb{E}_{(s,a) \sim \pi_e, P} [B_{\pi_e}^{P,k}(s, a)] \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] \quad (61)$$

$$\leq \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} [B_{\pi_e}^{P,k}(S_t, A_t) | H_{t_e-1}] + 4CD\rho^{t-t_e} \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] \quad (62)$$

$$= \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} [B_{\pi_e}^{P,k}(S_t, A_t) | H_{t_e-1}] \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] + \sum_{e=1}^E \frac{4CD}{1-\rho} \quad (63)$$

$$= \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} [B_{\pi_e}^{P,k}(S_t, A_t) | H_{t_e-1}] \right| \middle| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right] + \frac{4ECD}{1-\rho} \quad (64)$$

where Equation (63) follows from summation of series ρ^{t-t_e} for $t \rightarrow \infty$.

We can now construct a Martingale sequence to bound the summation in Equation (64). We construct a Martingale sequence as

$$X_t^e = \mathbb{E}_{(S_t, A_t) \sim \pi_e, P} \left[\sum_{t=t_e}^{t_{e+1}-1} B_{\pi_e}^{P,k}(S_t, A_t) | H_{t-1} \right]; t_e \leq t < t_{e+1} \quad (65)$$

such that $|X_t^e - X_{t-1}^e| \leq 4D$ for all t, e . We can now use Azuma-Hoeffding's inequality to bound $X_{t_e}^e$ as:

$$\left| \left(X_{t_e}^e - \sum_{t=t_e}^{t_{e+1}-1} B_{\pi_e}^{P,k}(S_t, A_t) \right) \right| \leq 4D\sqrt{(t_{e+1} - t_e) \log(2/T)} \quad (66)$$

with probability at least $1 - 1/T$. This eventually gives an upper bound on X_t^e as:

$$|X_t^e| - \left| \sum_{t=t_e}^{t_{e+1}-1} B_{\pi_e}^{P,k}(S_t, A_t) \right| \leq \left| \left(X_t^e - \sum_{t=t_e}^{t_{e+1}-1} B_{\pi_e}^{P,k}(S_t, A_t) \right) \right| \quad (67)$$

$$\leq 4D\sqrt{(t_{e+1} - t_e) \log(2/T)} \quad (68)$$

$$\implies |X_t^e| \leq \left| \sum_{t=t_e}^{t_{e+1}-1} B_{\pi_e}^{P,k}(S_t, A_t) \right| + 4D\sqrt{(t_{e+1} - t_e) \log(2/T)} \quad (69)$$

Hence, for $n_{t_e}(s, a) = \sum_{t'=1}^{t_e-1} \mathbf{1}_{\{S_{t'}=s, a_{t'}=a\}} = \sum_{e'=1}^{e-1} \sum_{t'=t_{e'}}^{t_{e'+1}-1} \mathbf{1}_{\{S_{t'}=s, a_{t'}=a\}} = \sum_{e'=1}^{e-1} \nu_{e'}(s, a)$, we have

$$R_1(T) \leq \sum_{e=1}^E \mathbb{E} \left[\left| D \sum_{t=t_e}^{t_{e+1}-1} B_{\pi_e}^{P^k, k}(S_t, A_t) \right| + \left| 4D\sqrt{(t_{e+1} - t_e) \log(2T)} \right| \left| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right| + \frac{4ECD}{1-\rho} \right] \quad (70)$$

$$\leq \sum_{e=1}^E \mathbb{E} \left[\left| D \sum_{s,a} \nu_e(s, a) B_{\pi_e}^{P^k, k}(s, a) \right| + \left| 4D\sqrt{(t_{e+1} - t_e) \log(2T)} \right| \left| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right| + \frac{4ECD}{1-\rho} \right] \quad (71)$$

$$\leq \sum_{e=1}^E \mathbb{E} \left[\left| D \sum_{s,a} \nu_e \sqrt{\frac{14S \log(2AT)}{n_{t_e}(s, a)}} \right| + \left| 4D\sqrt{(t_{e+1} - t_e) \log(2T)} \right| \left| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right| + \frac{4ECD}{1-\rho} \right] \quad (72)$$

$$= \sum_{e=1}^E \mathbb{E} \left[\left| D \sum_{s,a} \nu_e \sqrt{\frac{14S \log(2AT)}{n_{t_e}(s, a)}} + 4D\sqrt{(t_{e+1} - t_e) \log(2T)} \right| \left| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right| + \frac{4ECD}{1-\rho} \right] \quad (73)$$

$$= D\sqrt{14S \log(2AT)} \mathbb{E} \left[\sum_{e=1}^E \sum_{s,a} \nu_e \frac{1}{\sqrt{n_{t_e}(s, a)}} \left| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right| \right] + \mathbb{E} \left[\sum_{e=1}^E 4D\sqrt{(t_{e+1} - t_e) \log(2T)} \left| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right| \right] + \frac{4ECD}{1-\rho} \quad (74)$$

$$\leq D\sqrt{14S \log(2AT)} \sum_{s,a} (\sqrt{2} + 1) \sqrt{N(s, a)} + 4D\sqrt{ET \log(2T)} + \frac{4ECD}{1-\rho} \quad (75)$$

$$\leq D(\sqrt{2} + 1) \sqrt{14S \log(2AT)} \sqrt{SAT} + 4D\sqrt{ET \log(2T)} + \frac{4ECD}{1-\rho} \quad (76)$$

where Equation (70) comes from replacing X_t^e of Equation (64) by Equation (69). Equation (73) follows from the fact that both terms in modulus operator in Equation (72) are non-negative. In Equation (75) the first term follows from (Jaksch et al., 2010)[Lemma 19] and second term follows from Cauchy-Schwarz inequality. Equation (76) follows from Cauchy-Schwarz inequality.

We bound the $R_2^k(T)$ term similarly as:

$$R_2^k(T) = \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P^k, k} - r^k(S_t, A_t) \right) \right| \left| \mathcal{E}_{t_e} \cap \widehat{\mathcal{E}}_{t_e} \right| \right] \quad (77)$$

$$\leq 2\sqrt{ET \log(2T)} + \frac{4EC}{1-\rho} \quad (78)$$

with probability $1/T$.

Now, combining the regret using the law of total expectation, we have:

$$\begin{aligned} \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_e,k} - r^k(S_t, A_t) \right) \right| \right] &= \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_e,k} - r^k(S_t, A_t) \right) \right| \middle| \mathcal{E}_t^c \cup \widehat{\mathcal{E}}_t^c \right] \mathbb{P} \left(\mathcal{E}_t^c \cup \widehat{\mathcal{E}}_t^c \right) \\ &\quad + \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_e,k} - r^k(S_t, A_t) \right) \right| \middle| \mathcal{E}_t \cap \widehat{\mathcal{E}}_t \right] \mathbb{P} \left(\mathcal{E}_t \cap \widehat{\mathcal{E}}_t \right) \end{aligned} \quad (79)$$

$$\begin{aligned} &\leq \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_e,k} - r^k(S_t, A_t) \right) \right| \middle| \mathcal{E}_t^c \cup \widehat{\mathcal{E}}_t^c \right] \\ &\quad + \sum_{e=1}^E \mathbb{E} \left[\left| \sum_{t=t_e}^{t_{e+1}-1} \left(\lambda_{\pi_e}^{P_e,k} - r^k(S_t, A_t) \right) \right| \middle| \mathcal{E}_t \cap \widehat{\mathcal{E}}_t \right] \end{aligned} \quad (80)$$

$$\begin{aligned} &\leq 4\sqrt{T} + (\sqrt{2} + 1)DS\sqrt{14AT \log(2AT)} \\ &\quad + 4D\sqrt{ET \log(2T)} + 2\sqrt{ET \log(2T)} + \frac{4EC(D+1)}{1-\rho} \end{aligned} \quad (81)$$

where Equation (80) follows from the fact that probability for any event can be at most 1.

Combining everything gives the regret:

$$R(T) = \frac{LK}{T} \left(4\sqrt{T} + (\sqrt{2} + 1)DS\sqrt{14AT \log(2AT)} + (4D + 2)\sqrt{ET \log(2T)} + \frac{4EC(D+1)}{1-\rho} \right) \quad (82)$$

and completes the proof from using the upper bound of E from Lemma 9. ■

6. Model Free Algorithm

In the previous section, we developed a model based tabular algorithm for joint function optimization. However, as the state space, action space, or number of agents increase the tabular algorithm becomes infeasible to implement. In this section, we consider a policy gradient based algorithm which can be efficiently implemented using (deep) neural networks thus alleviating the requirement of a tabular solution for large MDPs.

For the model free policy gradient algorithm, we will use finite time horizon MDP, or $T < \infty$ in our MDP \mathcal{M} . This is a practical scenario where communication networks optimize fairness among users for finite duration (Margolies et al., 2016). We now describe a model free construction to obtain the optimal policy. We use a neural network parameterized by θ . The objective thus becomes to find optimal parameters θ^* , which maximizes,

$$\arg \max_{\theta} f \left((1-\gamma)J_{\pi_{\theta}}^1, \dots, (1-\gamma)J_{\pi_{\theta}}^K \right). \quad (83)$$

For model-free algorithm, we assume that the function is differentiable. In case, the function is not differentiable, sub-gradients of f can be used to optimize the objective (Nesterov, 2003). Gradient estimation for Equation (83) can be obtained using chain rule:

$$\nabla_{\theta} f = \sum_{k \in [K]} (1 - \gamma) \frac{\partial f}{\partial (1 - \gamma) J_{\pi}^k} \nabla_{\theta} J_{\pi}^k \quad (84)$$

For all k , J_{π}^k can be replaced with averaged cumulative rewards over N trajectories for the policy at i^{th} step, where a trajectory τ is defined as the tuple of observations, or $\tau = (s_0, a_0, r_0^1, \dots, r_0^K, s_1, a_1, r_1^1, \dots, r_1^K, \dots)$. Further, $\nabla_{\theta} J_{\pi}^k$ can be estimated using REINFORCE algorithm proposed in (Williams, 1992; Sutton et al., 2000) for all k , and is given as

$$\widehat{\nabla}_{\theta} J_{\pi}^k = \frac{1}{N} \sum_{j=1}^N \sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(A_{t,j} | S_{t,j}) \sum_{t'=t}^T r^k(S_{t',j}, A_{t',j}) \quad \forall k \in [K]. \quad (85)$$

Further, J_{π}^k is estimated as

$$\hat{J}_{\pi}^k = \frac{1}{N} \sum_{t=0}^T r^k(S_{t,j}, A_{t,j}) \quad \forall k \in [K]. \quad (86)$$

For a learning rate η , parameter update step to optimize the parameters becomes

$$\theta_{i+i} = \theta_i + \eta (1 - \gamma) \sum_{k \in [K]} \frac{\partial f}{\partial (1 - \gamma) J_{\pi}^k} \Big|_{J_{\pi}^k = \hat{J}_{\pi}^k} \widehat{\nabla}_{\theta} J_{\pi}^k \quad (87)$$

For example, consider alpha-fairness utility defined as

$$f((1 - \gamma) J_{\pi}^1, \dots, (1 - \gamma) J_{\pi}^K) = \sum_{k=1}^K \frac{1}{1 - \alpha} \left((J_{\pi}^k)^{1 - \alpha} - 1 \right) \quad (88)$$

Then the corresponding gradient estimate can be obtained as

$$\widehat{\nabla}_{\theta} f = \sum_{k \in [K]} \frac{\sum_{j=1}^N \sum_{t=0}^T \nabla_{\theta} \log \pi(A_{t,j} | S_{t,j}) \sum_{\tau=t}^T \gamma^{\tau} r_k(S_{\tau,j}, A_{\tau,j})}{(1 - \gamma)^{\alpha - 1} \left(\sum_{j=1}^N \sum_{t=0}^T \gamma^t r_k(S_{t,j}, A_{t,j}) \right)^{\alpha}}. \quad (89)$$

The proposed Model Free Policy Gradient algorithm for joint function optimization is described in Algorithm 2. The algorithm takes as input the parameters $\mathcal{S}, \mathcal{A}, [K], T, \gamma, f$ of MDP \mathcal{M} , number of sample trajectories N , and learning rate η as input. The policy neural network is initialized with weights θ randomly. It then collects N sample trajectories using the policy with current weights in Line 4. In Line 5, the gradient is calculated using Equations (84), (85), and (86) on the N trajectories. In optimization step of Line 6, the weights are updated using gradient ascent.

Algorithm 2 Model Free Joint Policy Gradient

```

1: procedure JOINT POLICY GRADIENT ( $\mathcal{S}, \mathcal{A}, [K], T, \gamma, f, N, \eta$ )
2:   Initialize  $\pi_{\theta_0}(a, s)$  ▷ Initialize the neural network with random weights  $\theta$ 
3:   for  $i = 0, 1, \dots$ , until convergence criteria do
4:     Collect  $N$  trajectories using policy  $\pi_{\theta_i}$ 
5:     Estimate gradient using Equation (84), (85), (86)
6:     Perform Gradient Ascent using Equation (87)
7:   end for
8:   Return  $\pi_{\theta}$ 
9: end procedure

```

7. Evaluations

In this section, we consider two systems. The first is the cellular scheduling, where multiple users connect to the base station. The second is a multiple-queue system which models multiple roads merging into a single lane. In both these systems, the proposed algorithms are compared with some baselines including the linear metric adaptation of reward at each time.

7.1 Cellular fairness maximization

The fairness maximization has been at the heart of many other resource allocation problems such as cloud resource management, manufacturing optimization, etc. (Perez et al., 2009; Zhang and Shah, 2015). The problem of maximizing wireless network fairness has been extensively studied in the past (Margolies et al., 2016; Kwan et al., 2009; Bu et al., 2006; Li et al., 2018). With increasing number of devices that need to access wireless network and ever upgrading network architectures, this problem still remains of practical interest. We consider two problem setup for fairness maximization, one with finite state space, and other with infinite state space. For finite state space, we evaluate both the proposed model-based algorithm (Algorithm 1) and the proposed model-free algorithm (Algorithm 2) while for infinite state space, we evaluate only model-free algorithm as tabular model-based algorithm cannot be implemented for this case.

7.1.1 PROBLEM SETUP

We consider fairness metric of the form of generalized α -fairness proposed in (Altman et al., 2008). For the rest of the paper, we will call this metric as α -fairness rather than generalized α -fairness. The problem of maximizing finite horizon α -fairness for multiple agents attached to a base station is defined as

$$C_{\alpha}(T) = \max_{\{a_{k,t}\}_{K \times T}} \sum_{k=1}^K \frac{1}{1-\alpha} \left(\left(\frac{1}{T} \sum_{t=1}^T a_{k,t} r_{k,t} \right)^{(1-\alpha)} - 1 \right) \quad (90)$$

$$\text{s.t. } \sum_{k=1}^K a_{k,t} = 1 \forall t \in \{1, 2, \dots, T\} \quad (91)$$

$$a_{i,j} \in 0, 1 \quad (92)$$

where, $a_{k,t} = 1$ if the agent k obtains the network resource at time t , and 0 otherwise. This implies, at each time t , the scheduler gives all the resources to only one of the attached user. Further, $r_{k,t}$ denotes the rate at which agent k can transmit at time t if allocated network resource. Cellular networks typically use Proportional Fair (PF) utility which is a special case of the above metric for $\alpha \rightarrow 1$ (Holma and Toskala, 2005), and is defined as:

$$C_1(T) = \max_{\{a_{k,t}\}_{K \times T}} \sum_{k=1}^K \log \left(\frac{1}{T} \sum_{t=1}^T a_{k,t} r_{k,t} \right) \quad (93)$$

We note that $r_{k,t}$ is only known causally limiting the use of offline optimization techniques and making the use of learning-based strategies for the problem important. We evaluate our algorithm for $\alpha = 2$ fairness, and proportional fairness for $T = 1000$.

7.1.2 PROPORTIONAL FAIRNESS

We let the number of users attached to the network, K , belong to the set $\{2, 4, 6\}$. The state space of each agent comes from its channel conditions. We assume that the channel for a agent can only be in two conditions $\{good, bad\}$, where the good and bad conditions for each agent could be different. The action at each time is a one-hot vector with the entry corresponding to the agent receiving the resources set to one. This gives $|\mathcal{S}| = 2^K$ (corresponding to the joint channel state of all agents), and $|\mathcal{A}| = K$ (K actions correspond to the agent that is selected in a time slot). Based on the channel state of agent, the scheduling decision determines the agent that must be picked in the time-slot. Rate $r_{k,t}$, for agent k at time t , is dependent on the state of the agent $s_{k,t}$ and is mentioned in Table 1. Each agent remains in the same state with probability of 0.8, and moves to a different state $s \sim U(\mathcal{S})$ with probability 0.2. The state transition model becomes,

$$\forall k, s_{k,t+1} = \begin{cases} s_{k,t}, & w.p. 0.8 \\ s \sim U(\{good, bad\}) & w.p. 0.2 \end{cases} \quad (94)$$

Agent state	$r_{1,t}$	$r_{2,t}$	$r_{3,t}$	$r_{4,t}$	$r_{5,t}$	$r_{6,t}$
<i>good</i>	1.50	2.25	1.25	1.50	1.75	1.25
<i>bad</i>	0.768	1.00	0.384	1.12	0.384	1.12

Table 1: Agent rate $r_{k,t}$ (in Mbps) based on agent state $s_{k,t}$. Rate values are practically observable data rates over a wireless network such as 4G-LTE.

We compare our model-based and model-free algorithms with practically implemented algorithm of Blind Gradient Estimation (Margolies et al., 2016; Bu et al., 2006) in network schedulers and SARSA based algorithm devised by (Perez et al., 2009). We first describe the algorithms used in evaluations for maximizing proportional fairness in finite state systems.

- **Blind Gradient Estimation Algorithm (BGE):** This heuristic allocates the resources based on the previously allocated resources to the agents. Starting from $t = 1$, this policy allocates resource to agent k_t^* at time t , where

$$k_t^* = \arg \max_{k \in [K]} \frac{r_{k,t}}{\sum_{t'=0}^{t-1} \alpha_{k,t'} r_{k,t'}}; \alpha_{k,t'} = \begin{cases} 1, & k = k_{t'}^* \\ 0, & k \neq k_{t'}^* \end{cases} \quad (95)$$

BGE is used as de facto standard for scheduling in cellular systems (Holma and Toskala, 2005), and has been shown to be asymptotically optimal for the proportional fairness metric (Viswanath et al., 2002).

- **DQN Algorithm:** This algorithm based on SARSA (Sutton and Barto, 2018a) and DQN Mnih et al. (2015). The reward at each time τ is the fairness of the system at time τ , or $f_\tau = C_\alpha(t)$ The DQN neural network consists of two fully connected hidden layers with 100 units each with ReLU activation and one output layer with linear activation. We use $\gamma = 0.99$, $\epsilon = 0.05$, and Adam optimizer with learning rate $\eta = 0.01$ to optimize the network. The batch size is 64 and the network is trained for 1500 episodes.
- **Vanilla Policy-Gradient Algorithm:** This algorithm is based on the REINFORCE policy gradient algorithm (Williams, 1992). Similar to the SARSA algorithm, the reward at each time τ is the fairness of the system at time τ , or $f_\tau = C_\alpha(t)$ We use $\gamma = 0.99$, and learning rate $\eta = 1 \times 10^{-4}$.
- **Proposed Model Based Algorithm:** We describe the algorithm for infinite horizon, so we maximize the policy for infinite horizon proportional fairness problem by discounting the rewards as

$$\lim_{T \rightarrow \infty} \sum_{k=1}^K \log \left(\frac{1}{T} \sum_{t=1}^T \gamma^t \alpha_{k,t} r_{k,t} \right) \tag{96}$$

The learned policy is evaluated on finite horizon environment of $T = 1000$. We keep $\gamma = 0.99$ for implementation of Algorithm 1. We use a fixed episode length of 100 and update policy after every $\tau = 100$ steps. In Algorithm 1, the convex optimization is solved using CVXPY Diamond and Boyd (2016).

- **Proposed Model Free Algorithm:** Since $\log(\cdot)$ is differentiable, the gradient in Equation (84) is evaluated using Equation (89) at $\alpha = 1$. The neural network consists of a single hidden layer with 200 neurons, each having ReLU activation function. The output layer uses softmax activation. The value of other hyperparameters are $\gamma = 0.99$, $\eta = 1 \times 10^{-3}$, and batch size $N = 100$. The algorithm source codes for the proposed algorithms have been provided at (Agarwal and Aggarwal, 2019).

Proportional Fairness Simulation Results: We trained the SARSA algorithm and the model based Algorithm 1 for 5000 time steps for each value of K . To train model free Algorithm 2 and the Vanilla Policy Gradient algorithm, we used 1000 batches where each batch contains 36 trajectories of length 1000 time steps. Note that the Blind Gradient Estimation algorithm doesn't need training as it selects the agent based on observed rewards. For all the algorithms we performed a grid search to find the hyperparameters.

We show the performance of policies implemented by each of the algorithm. Each policy is run 50 times and median and inter-quartile range is shown in Figure 1 for each policy. The policy performance for $K = 2$, $K = 4$, and $K = 6$ is shown in Figure 1(a), Figure 1(c), and Figure 1(e), respectively.

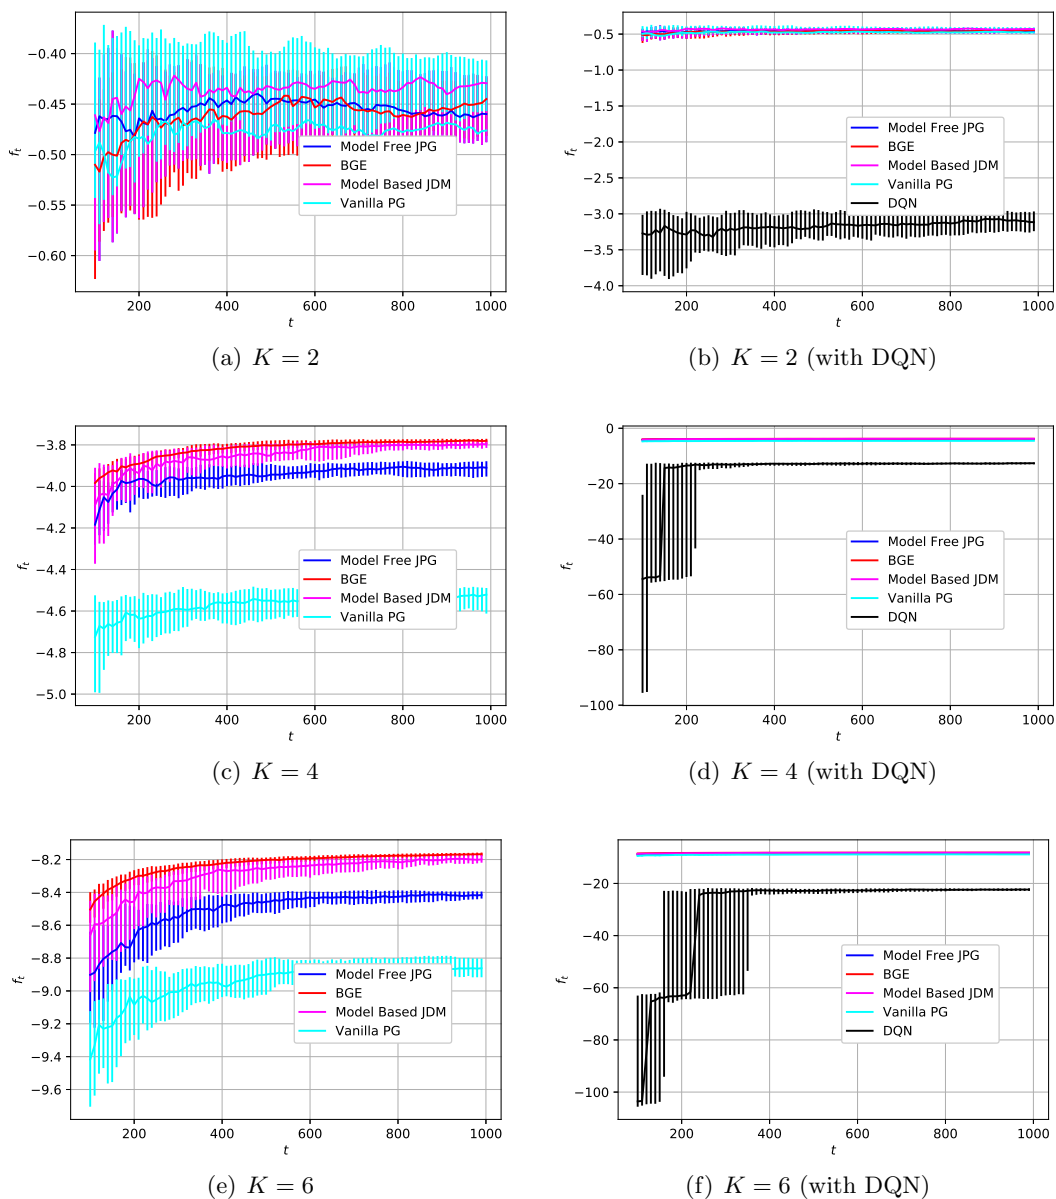


Figure 1: Proportional Fairness for Cellular Scheduling, f_t v/s t (Best viewed in color)

We note that the performance of model-based algorithm (Algorithm 1) and that of the model-free algorithm (Algorithm 2) are close. For $K = 2$, the model-based algorithm outperforms the BGE algorithm. For $K \in \{4, 6\}$, the gap between the model based algorithm and BGE algorithm is because of the finite time horizon. The proposed framework assumes an infinite horizon framework, but the algorithm is trained for finite time horizon. The regret of $\tilde{O}\left(LDKS\sqrt{\frac{A}{T}}\right)$ also guarantees that the proposed algorithm converges towards optimal policies for large T .

From Figure 1(b), Figure 1(d), and Figure 1(f), we note that the DQN algorithm performs much worse than expected. The reason for this is that the joint objective function of fairness is non-linear and is not properly modelled with standard RL formulation. Also, from Figure 1(a), we note that for $K = 2$, policy gradient algorithm which uses fairness till time t can still learn a good policy, but the performance is still not at par with the proposed framework. This is because, using the value of the joint objective as reward works as a linear approximation of the joint reward function. Hence, if the approximation is worse, the policy gradient algorithm with joint objective as reward will not be able to optimize the true reward function. Note that as K increases, the approximation error $\sum_k \log(\lambda_k) - \lambda_k$ increase. In the next experiment, we will demonstrate that if the approximation error is too large, the policy gradient algorithm can perform even worse than the DQN algorithm.

7.1.3 $\alpha = 2$ FAIRNESS

We now evaluate our algorithm with the metric of α -fairness, where no optimal baseline is known. We also consider a large state-space to show the scalability of the proposed model-free approach. We consider a Gauss-Markov channel model (Ariyakhajorn et al., 2006) for modeling the channel state to the different users, and let the number of users be $K = 8$. Under Gauss-Markov Model, channel state of each user k varies as,

$$X_{k,t} = \sqrt{1 - \beta^2}X_{k,t-1} + \beta\epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, 1). \quad (97)$$

We assume $X_{k,0} \sim \mathcal{N}(0, 1)$ for each $k \in [K]$. The rate for each user i at time t and in channel state $X_{i,t}$ is given as,

$$r_{k,t} = P_k |X_{k,t}|^2, \quad (98)$$

where P_k is multiplicative constant for the k^{th} user, which indicates the average signal-to-noise ratio to the user. We let $P_k = k^{-0.2}$.

Since the state space is infinite, we only evaluate the model free algorithm. The gradient update equation is defined in Equation (89) with $\alpha = 2$. The neural network consists of a single hidden layer with 200 neurons, each having ReLU activation function. We use stochastic gradient ascent with learning rate $\eta = 1 \times 10^{-3}$ to train the network. The value of other hyperparameters are $\gamma = 0.99$, and batch size $N = 36$. The network is trained for 1000 epochs.

Also, since no optimal baseline is known, we compare the model free algorithm with the Deep Q-Network (DQN) algorithm (Mnih et al., 2015) and the Policy Gradient algorithm Williams (1992). Reward for both DQN algorithm and the Policy Gradient algorithm at time τ is taken as the value $C_2(\tau)$. For DQN, the neural network consists of two fully connected

hidden layers with 100 units each with ReLU activation and one output layer with linear activation. We use Adam optimizer with learning rate 0.01 to optimize the DQN network. The batch size is 64 and the network is trained for 1500 episodes. For the Policy Gradient algorithm, we choose a single hidden layer of 200 neurons, with learning rate 1×10^{-5} . Similar to the implementation of the Algorithm 2, we select batch size $N = 36$ and train the network for 1000 epochs.

α -fairness Simulations Results The results for α -fairness are provided in Figure 2. Each policy is run 50 times and median and we show the inter-quartile range for each policy. As a baseline, we also consider a strategy that chooses a user in each time uniformly at random, this strategy is denoted as “uniform random” (Figure 2(a)) We note that the DQN algorithm and the policy gradient algorithm are not able to outperform the uniform random policy while the proposed model free policy outperforms the uniform random policy. Further Figure 2(b) shows the detrimental effect of using incorrect gradients. The linear approximation has larger error for $\alpha = 2$ joint objective than compared to the proportional fairness joint objective. The standard policy gradient with fairness as rewards now performs as worse as the DQN algorithm.

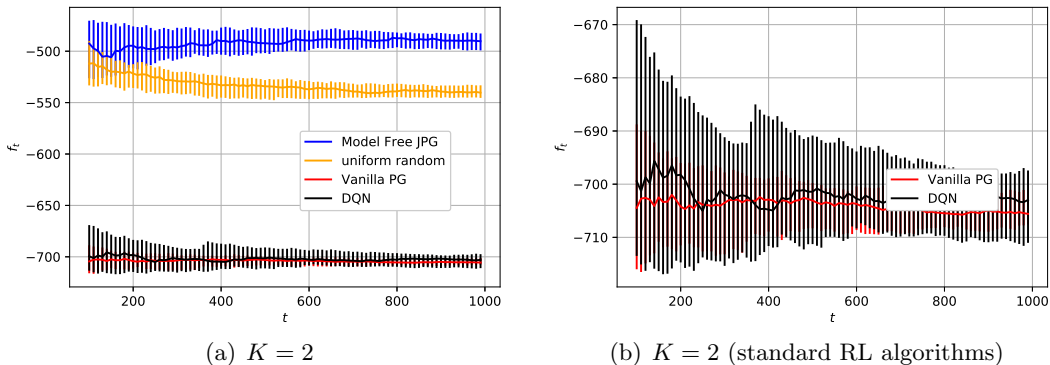


Figure 2: Alpha Fairness for Cellular Scheduling, f_t v/s t (Best viewed in color)

7.2 Multiple Queues Latency Optimization

We consider a problem, where multiple roads merge into a single lane, which is controlled by a digital sign in front of each road to indicate which road’s vehicle proceeds next. This problem can be modeled as having K queues with independent arrival patterns, where the arrival at queue $k \in \{1, \dots, K\}$ is Bernoulli with an arrival rate of λ_k . At each time, the user at the head of the selected one of out these K queues is served. The problem is to determine which queue user is served at each time. Such problems also occur in processor scheduling systems, where multiple users send their requests to the server and the server decides which user’s computation to do next (Haldar and Subramanian, 1991).

In such a system, latency is of key concern to the users. The authors of (Zhang et al., 2019) demonstrated that the effect of latency on the Quality of Experience (QoE) to the

end user is a ‘‘sigmoid-like’’ function. Thus, we define the latency of a user w as

$$QoE(w) = \frac{1 - e^{-3}}{1 + e^{(w-3)}}. \quad (99)$$

Note that in Equation (99), $QoE(w)$ remains close to 1 for small wait times ($w \leq 1$), and close to 0 for high wait times ($w \geq 10$). We let the service distribution of the queue be deterministic, where each user takes one unit of time for service. We model the problem as a non-linear multi-agent system. The different queues are the agents. The state is the queue lengths of the different queues. The action at each time is to determine which of the non-empty queue is selected. The reward of each agent at time t , $r_{k,t}$ is zero if queue k is not selected at time t , and if the QoE of the latency of the user served, if queue k is selected at time t , where the latency of a user w is the time spent by a user w in the system (from entering the queue to being served). We again consider α -fairness for $\alpha = 2$, and proportional fairness for comparisons.

7.2.1 $\alpha = 2$ FAIRNESS

We consider $K = 8$ queues in our system. We let the arrival rate λ_i in each queue be as given in Table 2. We also assume that each queue has bounded capacity of 100, and the user is dropped if the queue is full. The overall reward function among different agents is chosen as α -fairness, where $\alpha = 2$. The objective can be written as

$$C_2(T) = \max_{\{a_{k,t}\}_{K \times T}} \sum_{k=1}^K -1 \left(\frac{1}{T} \sum_{t=1}^T a_{k,t} r_{k,t} \right)^{-1} \quad (100)$$

λ_1	λ_2	λ_3	λ_4	λ_5	λ_6	λ_7	λ_8
0.2	0.1	0.05	0.25	0.15	0.21	0.01	0.3

Table 2: Arrival rates λ_k (in number of packets per step) for $\alpha = 2$ fairness

Since the number of states is large, we only evaluate the model-free algorithm with $T = 1000$. The gradient update equation for policy gradient algorithm as given in Equation (89) is used for $\alpha = 2$. Stochastic Gradient Ascent with learning rate $\eta = 5 \times 10^{-3}$ is used to train the network. The value of discount factor γ is set to 0.99 and the batch size N is kept as 30.

We compare the proposed algorithm with the DQN Algorithm for Q-learning implementation. We use fairness at time t or $C_2(t)$ as the reward for DQN network. The network consists of two fully connected hidden layers with 100 units each, ReLU activation function, and one output layer with linear activation. Adam optimizer with learning rate 0.01 is used to optimize the network. The batch size is 64 and the network is trained for 500 episodes.

We also compare the proposed algorithm with the *Longest Queue First* (LQF) Algorithm, which serves the longest queue of the system. This algorithm doesn’t require any learning, and has no hyperparameters.

We again run each policy for 50 times and median and we show the inter-quartile range for each policy. The results for fairness maximization for this queuing system are provided in

Figure (3). We note that the overall objective decreases for all the policies. This is because the queue length is increasing and each packet has to wait for longer time on an average to be served till the queue becomes steady. At the end of the episode, the proposed policy gradient algorithm outperforms both the DQN and the LQF policy. We note that the during the start of the episode, LQF is more fair because the queues are almost empty, and serving the longest queue would decrease the latency of the longest queue. However, serving the longest queue is not optimal in the steady state.




Figure 3: Alpha Fairness for the Queueing System, f_t v/s t (Best viewed in color)

7.2.2 PROPORTIONAL FAIRNESS

We also compare the performance of the policy learned by our algorithm with the policy learned by DQN algorithm for another synthetic example where we reduce the arrival rates in order to make the system less loaded. We let the arrival rate λ_i in each queue be as given in Table 3. We also assume that each queue has bounded capacity of 10, and the user is dropped if the queue is full. The overall reward function among different agents is chosen as weighted-proportional fairness. The objective can be written as

$$C(T) = \max_{\{a_{k,t}\}_{K \times T}} \sum_{k=1}^K w_k \log \left(\frac{1}{T} \sum_{t=1}^T a_{k,t} r_{k,t} \right) \quad (101)$$

The weights w_k are given in Table 4. The weights were generated from a Normal distribution with mean 1, and variance 0.1. After sampling, the weights were normalized to make $\sum_k w_k = 1$.

λ_1	λ_2	λ_3	λ_4	λ_5	λ_6	λ_7	λ_8
0.014	0.028	0.042	0.056	0.069	0.083	0.097	0.11

Table 3: Arrival rates λ_k (in number of packets per step) for proportional fairness

Again, we only evaluate the model-free algorithm with $T = 1000$. The gradient update equation for policy gradient algorithm as given in Equation (89) is used for $\alpha = 1$. Stochastic


w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
0.146	0.112	0.145	0.119	0.119	0.123	0.114	0.122

 Table 4: Weights w_k for weighted proportional fairness

Gradient Ascent with learning rate $\eta = 5 \times 10^{-1}$ is used to train the network. The value of discount factor γ is set to 0.999 and the batch size N is kept as 80.

We compare the proposed algorithm with the DQN Algorithm for Q-learning implementation. We use fairness at time t or $C_2(t)$ as the reward for DQN network. The network consists of two fully connected hidden layers with 100 units each, ReLU activation function, and one output layer with linear activation. Adam optimizer with learning rate 0.01 is used to optimize the network. The batch size is 64 and the network is trained for 5000 episodes.

The results for fairness maximization for this queuing system are provided in Figure 4. Similar to previous cases, we run each policy for 50 times and median and we show the inter-quartile range for each policy. We note that compared to previous case, the objective increases as the episode progress. This is because the arrival rates are low, and queue lengths are short. For the entire episode, the proposed policy gradient algorithm outperforms the DQN policy. This is because of incorrect Q-learning (and DQN) cannot capture the non-linear functions of rewards, which is weighted proportional fairness in this case.


 Figure 4: Weighted Proportional Fairness for the Queueing System, f_t v/s t (Best viewed in color)

8. Conclusion

This paper presents a novel average per step reward based formulation for optimizing joint objective function of long-term rewards of each objective for infinite horizon setup. In case of finite horizon, Markov policies may not be able to optimize the joint objective function, hence an average reward per step formulation is considered. A tabular model based algorithm which uses Dirichlet sampling to obtain regret bound of $\tilde{O}\left(LKDS\sqrt{\frac{A}{T}}\right)$ for K objectives scalarized using a L -Lipschitz concave function over a time horizon T is provided where S is

the number of states and D is the diameter of the underlying Markov Chain and A is the number of actions available to the centralized controller. Further, a model free algorithm which can be efficiently implemented using neural networks is also proposed. The proposed algorithms outperform standard heuristics by a significant margin for maximizing fairness in cellular scheduling problem, as well as for a multiple-queue queueing system.

Possible future works include modifying the framework to obtain actions from policies instead of probability values for infinite action space, and obtaining decentralized policies by introducing a message passing architecture.

References

- Mridul Agarwal and Vaneet Aggarwal. Source Code for Non-Linear Reinforcement Learning. <https://github.rcac.purdue.edu/Clan-labs/non-markov-RL>, 2019.
- Vaneet Aggarwal, Rittwik Jana, Jeffrey Pang, KK Ramakrishnan, and NK Shankaranarayanan. Characterizing fairness for 3g wireless networks. In *2011 18th IEEE Workshop on Local & Metropolitan Area Networks (LANMAN)*, pages 1–6. IEEE, 2011.
- Shipra Agrawal and Randy Jia. Optimistic posterior sampling for reinforcement learning: worst-case regret bounds. In *Advances in Neural Information Processing Systems*, pages 1184–1194, 2017.
- Eitan Altman, Konstantin Avrachenkov, and Andrey Garnaev. Generalized α -fair resource allocation in wireless networks. In *2008 47th IEEE Conference on Decision and Control*, pages 2414–2419. IEEE, 2008.
- Jinthana Ariyakhajorn, Pattana Wannawilai, and Chanboon Sathitwiriawong. A comparative study of random waypoint and gauss-markov mobility models in the performance evaluation of manet. In *2006 International Symposium on Communications and Information Technologies*, pages 894–899. IEEE, 2006.
- Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA, 1995.
- Daan Bloembergen, Karl Tuyls, Daniel Hennes, and Michael Kaisers. Evolutionary dynamics of multi-agent learning: A survey. *Journal of Artificial Intelligence Research*, 53:659–697, 2015.
- T Bu, L Li, and R Ramjee. Generalized proportional fair scheduling in third generation wireless data networks. In *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, pages 1–12. IEEE, 2006.
- Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Lucian Buşoniu, Robert Babuška, and Bart De Schutter. Multi-agent reinforcement learning: An overview. In *Innovations in multi-agent systems and applications-1*, pages 183–221. Springer, 2010.
- A Castelletti, Francesca Pianosi, and Marcello Restelli. A multiobjective reinforcement learning approach to water resources systems operation: Pareto frontier approximation in a single run. *Water Resources Research*, 49(6):3476–3486, 2013.
- Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- Anis Elgabli, Vaneet Aggarwal, Shuai Hao, Feng Qian, and Subhabrata Sen. Lbp: Robust rate adaptation algorithm for svc video streaming. *IEEE/ACM Transactions on Networking*, 26(4):1633–1645, 2018.

- Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- Sibsankar Haldar and DK Subramanian. Fairness in processor scheduling in time sharing systems. *ACM SIGOPS Operating Systems Review*, 25(1):4–18, 1991.
- Harri Holma and Antti Toskala. *WCDMA for UMTS.: Radio Access for Third Generation Mobile Communications*. John Wiley & Sons, 2005.
- Marcus Hutter. Extreme state aggregation beyond mdps. In *International Conference on Algorithmic Learning Theory*, pages 185–199. Springer, 2014.
- Shadi Ibrahim, Hai Jin, Lu Lu, Song Wu, Bingsheng He, and Li Qi. Leen: Locality/fairness-aware key partitioning for mapreduce in the cloud. In *2010 IEEE Second International Conference on Cloud Computing Technology and Science*, pages 17–24. IEEE, 2010.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(Apr):1563–1600, 2010.
- Jiechuan Jiang and Zongqing Lu. Learning fairness in multi-agent systems. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873, 2018.
- Jun-young Kwak, Pradeep Varakantham, Rajiv Maheswaran, Milind Tambe, Farrokh Jazizadeh, Geoffrey Kavulya, Laura Klein, Burcin Becerik-Gerber, Timothy Hayes, and Wendy Wood. Saves: A sustainable multiagent application to conserve building energy considering occupants. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 21–28, 2012.
- Raymond Kwan, Cyril Leung, and Jie Zhang. Proportional fair multiuser scheduling in lte. *IEEE Signal Processing Letters*, 16(6):461–464, 2009.
- Tian Lan, David Kao, Mung Chiang, and Ashutosh Sabharwal. *An axiomatic theory of fairness in network resource allocation*. IEEE, 2010.
- Gregory F Lawler. *Introduction to stochastic processes*. Chapman and Hall/CRC, 2018.
- David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- Lihong Li, Thomas J Walsh, and Michael L Littman. Towards a unified theory of state abstraction for mdps. *ISAIM*, 4:5, 2006.
- Xiaoshuai Li, Rajan Shankaran, Mehmet A Orgun, Gengfa Fang, and Yubin Xu. Resource allocation for underlay d2d communication with proportional fairness. *IEEE Transactions on Vehicular Technology*, 67(7):6244–6258, 2018.

- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994.
- Chunming Liu, Xin Xu, and Dewen Hu. Multiobjective reinforcement learning: A comprehensive overview. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45(3): 385–398, 2014.
- Sultan Javed Majeed and Marcus Hutter. On q-learning convergence for non-markov decision processes. In *IJCAI*, pages 2546–2552, 2018.
- Robert Margolies, Ashwin Sridharan, Vaneet Aggarwal, Rittwik Jana, NK Shankaranarayanan, Vinay A Vaishampayan, and Gil Zussman. Exploiting mobility in proportional fair cellular scheduling: Measurements and algorithms. *IEEE/ACM Transactions on Networking (TON)*, 24(1):355–367, 2016.
- R Andrew McCallum. Instance-based utile distinctions for reinforcement learning with hidden state. In *Machine Learning Proceedings 1995*, pages 387–395. Elsevier, 1995.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.
- Thanh Thi Nguyen, Ngoc Duy Nguyen, Peter Vamplew, Saeid Nahavandi, Richard Dazeley, and Chee Peng Lim. A multi-objective deep reinforcement learning framework. *Engineering Applications of Artificial Intelligence*, 96:103915, 2020.
- Norihiko Ono and Kenji Fukumoto. Multi-agent reinforcement learning: A modular approach. In *Second International Conference on Multiagent Systems*, pages 252–258, 1996.
- Ian Osband, Daniel Russo, and Benjamin Van Roy. (more) efficient reinforcement learning via posterior sampling. In *Advances in Neural Information Processing Systems*, pages 3003–3011, 2013.
- Julien Perez, Cécile Germain-Renaud, Balázs Kégl, and Charles Loomis. Responsive elastic computing. In *Proceedings of the 6th international conference industry session on Grids meets autonomic computing*, pages 55–64. ACM, 2009.
- John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1/2):122–136, 1964. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1913738>.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1994. ISBN 0471619779.

- Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. *J. Artif. Int. Res.*, 48(1):67–113, October 2013. ISSN 1076-9757.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. *arXiv preprint arXiv:1810.04650*, 2018.
- Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.
- Yoav Shoham, Rob Powers, and Trond Grenager. Multi-agent reinforcement learning: a critical survey. *Web manuscript*, 2003.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018a.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018b.
- Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
- Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pages 330–337, 1993.
- Sylvie Thiébaux, Charles Gretton, John Slaney, David Price, and Froduald Kabanza. Decision-theoretic planning with non-markovian rewards. *Journal of Artificial Intelligence Research*, 25:17–74, 2006.
- Kristof Van Moffaert and Ann Nowé. Multi-objective reinforcement learning using sets of pareto dominating policies. *The Journal of Machine Learning Research*, 15(1):3483–3512, 2014.
- P. Viswanath, D. N. C. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. *IEEE Transactions on Information Theory*, 48(6):1277–1294, June 2002. ISSN 0018-9448. doi: 10.1109/TIT.2002.1003822.
- Wei Wang, Baochun Li, and Ben Liang. Dominant resource fairness in cloud computing systems with heterogeneous servers. In *IEEE INFOCOM 2014-IEEE Conference on Computer Communications*, pages 583–591. IEEE, 2014.

- Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando De Freitas. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*, 2015.
- Tsachy Weissman, Erik Ordentlich, Gadiel Seroussi, Sergio Verdu, and Marcelo J Weinberger. Inequalities for the l1 deviation of the empirical distribution. 2003.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Yu Xiang, Tian Lan, Vaneet Aggarwal, and Yih-Farn R Chen. Joint latency and cost optimization for erasure-coded data center storage. *IEEE/ACM Transactions on Networking*, 24(4):2443–2457, 2015.
- Chongjie Zhang and Julie A Shah. Fairness in multi-agent sequential decision-making. In *Advances in Neural Information Processing Systems*, pages 2636–2644, 2014.
- Chongjie Zhang and Julie A Shah. On fairness in decision-making under uncertainty: Definitions, computation, and comparison. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Xu Zhang, Siddhartha Sen, Daniar Kurniawan, Haryadi Gunawi, and Junchen Jiang. E2e: Embracing user heterogeneity to improve quality of experience on the web. In *Proceedings of the ACM Special Interest Group on Data Communication, SIGCOMM '19*, pages 289–302, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-5956-6. doi: 10.1145/3341302.3342089. URL <http://doi.acm.org/10.1145/3341302.3342089>.

Appendix A. Proof of Auxiliary Lemmas

A.1 Bounding the bias span for any MDP for any policy

Lemma 11 (Bounded Span of MDP) *For an MDP \mathcal{M} with rewards $r^k(s, a)$ and transition probabilities P , for any stationary policy π with average reward λ_π^k , the difference of bias of any two states s , and s' , is upper bounded by the diameter of the MDP D as:*

$$h_\pi^{P,k}(s) - h_\pi^{P,k}(s') \leq D \quad \forall s, s' \in \mathcal{S}. \quad (102)$$

Proof Consider two states $s, s' \in \mathcal{S}$. Also, let τ be a random variable defined as:

$$\tau = \min\{t \geq 1 : S_t = s', S_1 = s\} \quad (103)$$

Then, for any policy π , we have the following Bellman operator

$$\begin{aligned} h_\pi^{P,k}(s) &= r_\pi^k(s) - \lambda_\pi^k + \langle P_\pi(\cdot|s), h_\pi^{P,k} \rangle \\ &= Th_\pi^{P,k}(s)(s) \end{aligned} \quad (104)$$

where $P_\pi(s'|s) = \sum_a \pi(a|s)P(s'|s, a)$ and $r_\pi^k(s) = \sum_a \pi(a|s)r^k(s, a)$.

We also define another operator,

$$\bar{T}h(s) = \begin{cases} \min_{s,a} r^k(s, a) - \lambda_\pi^k + \langle P_\pi(\cdot|s), h \rangle, & s \neq s' \\ h_\pi^{P,k}(s'), & s = s' \end{cases} \quad (105)$$

Note that $(T - \bar{T})h_\pi^{P,k}(s) = r_\pi^k(s, a) - \min_{s,a} r^k(s, a) \geq 0$, for all s . Hence, we have $\bar{T}h(s) \leq Th(s) = h_\pi^{P,k}(s)$, for all s . Further, for any two vectors u, v , where all the elements of u are not smaller than w we have $\bar{T}u \geq \bar{T}w$. Hence, we have $\bar{T}^n h_\pi^{P,k}(s) \leq h_\pi^{P,k}(s)$ for all s . Unrolling the recurrence, we have

$$h_\pi^{P,k}(s) \geq \bar{T}^n h_\pi^{P,k}(s) = \mathbb{E} \left[-(\lambda_\pi^k - \min_{s,a} r^k(s, a))(n \wedge \tau) + h_\pi^{P,k}(S_{n \wedge \tau}) \right] \quad (106)$$

For $\lim n \rightarrow \infty$, we have $h_\pi^{P,k}(s) \geq h_\pi^{P,k}(s') - D$, completing the proof. \blacksquare

A.2 Bounding the bias span for MDP with transition probabilities P_e^k

Lemma 12 (Bounded Span of optimal MDP) *For a MDP with rewards $r^k(s, a)$ and transition probabilities $P_e^k \in \mathcal{P}_{t_e}$, for policy π_e , the difference of bias of any two states s , and s' , is upper bounded by the diameter of the true MDP D as:*

$$h_{\pi_e}^{P_e^k,k}(s) - h_{\pi_e}^{P_e^k,k}(s') \leq D \quad \forall s, s' \in \mathcal{S}. \quad (107)$$

Proof Note that $\lambda_{\pi_e}^{P_e^k,k} \geq \lambda_{\pi_e}^{P',k}$ for all $P' \in \mathcal{P}_{t_e}$. Now, consider the following Bellman equation:

$$\begin{aligned} h_{\pi_e}^{P_e^k,k}(s) &= r_{\pi_e}^k(s, a) - \lambda_{\pi_e}^{P_e^k,k} + \langle P_{\pi_e, e}^k(\cdot|s), h_{\pi_e}^{P_e^k,k} \rangle \\ &= Th_{\pi_e}^{P_e^k,k}(s) \end{aligned} \quad (108)$$

where $r_{\pi_e}(s) = \sum_a \pi_e(a|s)r^k(s, a)$ and $P_{\pi_e, e}^k(s'|s) = \sum_a \pi(a|s)P_e^k(s'|s, a)$.

Consider two states $s, s' \in \mathcal{S}$. Also, let τ be a random variable defined as:

$$\tau = \min\{t \geq 1 : S_t = s', S_1 = s\} \quad (109)$$

We also define another operator,

$$\bar{T}h(s) = \begin{cases} \min_{s,a} r^k(s, a) - \lambda_{\pi_e}^{P_e^k, k} + \langle P_{\pi_e}(\cdot|s), h \rangle, & s \neq s' \\ h_{\pi_e}^{P_e^k, k}(s'), & s = s' \end{cases} \quad (110)$$

where $P_{\pi_e}(\cdot|s) = \sum_a \pi_e(a|s)P(s'|s, a)$.

Now, for any $s \in \mathcal{S}$, note that

$$h(s) = Th(s) \quad (111)$$

$$= \max_{P' \in \mathcal{P}_{t_e}} \left(r_{\pi_e}^k(s) - \lambda_{\pi_e}^{P_e^k, k} + \langle P', h \rangle \right) \quad (112)$$

$$\geq r_{\pi_e}^k(s) - \lambda_{\pi_e}^{P_e^k, k} + \langle P_{\pi_e}, h \rangle \quad (113)$$

$$\geq \min_{s,a} r^k(s, a) - \lambda_{\pi_e}^{P_e^k, k} + \langle P_{\pi_e}, h \rangle \quad (114)$$

$$= \bar{T}h(s) \quad (115)$$

Further, for any two vectors u, v , where all the elements of u are not smaller than w we have $\bar{T}u \geq \bar{T}w$. Hence, we have $\bar{T}^n h_{\pi}^{P_e^k, k}(s) \leq h_{\pi}^{P_e^k, k}(s)$ for all s . Unrolling the recurrence, we have

$$h_{\pi}^{P_e^k, k}(s) \geq \bar{T}^n h_{\pi}^{P_e^k, k}(s) = \mathbb{E} \left[-(\lambda_{\pi}^k - \min_{s,a} r^k(s, a))(n \wedge \tau) + h_{\pi}^{P_e^k, k}(S_{n \wedge \tau}) \right] \quad (116)$$

For $\lim n \rightarrow \infty$, we have $h_{\pi}^{P_e^k, k}(s) \geq h_{\pi}^{P_e^k, k}(s') - D$, completing the proof. \blacksquare

Appendix B. Proof of Lemmas from main text

B.1 Proof of Lemma 5

Proof Note that for all $s \in \mathcal{S}$, we have:

$$V_{\gamma, \pi}^{\tilde{P}, k}(s) = \mathbb{E}_{a \sim \pi} \left[Q_{\gamma, \pi}^{\tilde{P}, k}(s, a) \right] \quad (117)$$

$$= \mathbb{E}_{a \sim \pi} \left[B_{\gamma, \pi}^{\tilde{P}, k}(s, a) + r(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) V_{\gamma, \pi}^{\tilde{P}, k}(s') \right] \quad (118)$$

where Equation (118) follows from the definition of the Bellman error for state action pair s, a .

Similarly, for the true MDP, we have,

$$V_{\gamma, \pi}^{P, k}(s) = \mathbb{E}_{a \sim \pi} \left[Q_{\gamma, \pi}^{P, k}(s, a) \right] \quad (119)$$

$$= \mathbb{E}_{a \sim \pi} \left[r(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) V_{\gamma, \pi}^{P, k}(s') \right] \quad (120)$$

Subtracting Equation (120) from Equation (118), we get:

$$V_{\gamma, \pi}^{\tilde{P}, k}(s) - V_{\gamma, \pi}^{P, k}(s) = \mathbb{E}_{a \sim \pi} \left[B_{\gamma, \pi}^{\tilde{P}, k}(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) \left(V_{\gamma, \pi}^{\tilde{P}, k} - V_{\gamma, \pi}^{P, k} \right) (s') \right] \quad (121)$$

$$= \mathbb{E}_{a \sim \pi} \left[B_{\gamma, \pi}^{\tilde{P}, k}(s, a) \right] + \gamma \sum_{s' \in \mathcal{S}} P_{\pi} \left(V_{\gamma, \pi}^{\tilde{P}, k} - V_{\gamma, \pi}^{P, k} \right) (s') \quad (122)$$

Using the vector format for the value functions, we have,

$$\bar{V}_{\gamma}^{\pi, \tilde{P}} - \bar{V}_{\gamma}^{\pi, P} = (I - \gamma P_{\pi})^{-1} \mathbb{E}_{a \sim \pi} \left[B_{\gamma, \pi}^{\tilde{P}, k}(s, a) \right] \quad (123)$$

Now, converting the value function to average per-step reward we have,

$$\lambda_{\pi}^{\tilde{P}, k} \mathbf{1}_{\mathcal{S}} - \lambda_{P, \pi}^k \mathbf{1}_{\mathcal{S}} = \lim_{\gamma \rightarrow 1} (1 - \gamma) \left(\bar{V}_{\gamma}^{\pi, \tilde{P}} - \bar{V}_{\gamma}^{\pi, P} \right) \quad (124)$$

$$= \lim_{\gamma \rightarrow 1} (1 - \gamma) (I - \gamma P_{\pi})^{-1} \mathbb{E}_{a \sim \pi} \left[B_{\gamma, \pi}^{\tilde{P}, k}(s, a) \right] \quad (125)$$

$$= \left(\sum_{s, a} d_{\pi}(s, a) B_{\pi}^{\tilde{P}, k}(s, a) \right) \mathbf{1}_{\mathcal{S}} \quad (126)$$

where the last equation follows from the definition of occupancy measures by Puterman (1994), and the existence of the limit $\lim_{\gamma \rightarrow 1} B_{\gamma, \pi}^{\tilde{P}, k}$ in Equation (125) from Equation (134). \blacksquare

B.2 Proof of Lemma 6

Proof Starting with the definition of Bellman error in Equation (26), we get

$$B_{\pi}^{\tilde{P},k}(s, a) = \lim_{\gamma \rightarrow 1} \left(Q_{\gamma, \pi}^{\tilde{P},k}(s, a) - \left(r(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) V_{\gamma, \pi}^{\tilde{P},k}(s') \right) \right) \quad (127)$$

$$= \lim_{\gamma \rightarrow 1} \left(\left(r(s, a) + \gamma \sum_{s' \in \mathcal{S}} \tilde{P}(s'|s, a) V_{\gamma, \pi}^{\tilde{P},k}(s') \right) - \left(r(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) V_{\gamma, \pi}^{\tilde{P},k}(s') \right) \right) \quad (128)$$

$$= \lim_{\gamma \rightarrow 1} \gamma \sum_{s' \in \mathcal{S}} \left(\tilde{P}(s'|s, a) - P(s'|s, a) \right) V_{\gamma, \pi}^{\tilde{P},k}(s') \quad (129)$$

$$= \lim_{\gamma \rightarrow 1} \gamma \left(\sum_{s' \in \mathcal{S}} \left(\tilde{P}(s'|s, a) - P(s'|s, a) \right) V_{\gamma, \pi}^{\tilde{P},k}(s') + V_{\gamma, \pi}^{\tilde{P},k}(s) - V_{\gamma, \pi}^{\tilde{P},k}(s) \right) \quad (130)$$

$$= \lim_{\gamma \rightarrow 1} \gamma \left(\sum_{s' \in \mathcal{S}} \left(\tilde{P}(s'|s, a) - P(s'|s, a) \right) V_{\gamma, \pi}^{\tilde{P},k}(s') - \sum_{s' \in \mathcal{S}} \tilde{P}(s'|s, a) V_{\gamma, \pi}^{\tilde{P},k}(s) + \sum_{s' \in \mathcal{S}} P(s'|s, a) V_{\gamma, \pi}^{\tilde{P},k}(s) \right) \quad (131)$$

$$= \lim_{\gamma \rightarrow 1} \gamma \left(\sum_{s' \in \mathcal{S}} \left(\tilde{P}(s'|s, a) - P(s'|s, a) \right) \left(V_{\gamma, \pi}^{\tilde{P},k}(s') - V_{\gamma, \pi}^{\tilde{P},k}(s) \right) \right) \quad (132)$$

$$= \left(\sum_{s' \in \mathcal{S}} \left(\tilde{P}(s'|s, a) - P(s'|s, a) \right) \lim_{\gamma \rightarrow 1} \gamma \left(V_{\gamma, \pi}^{\tilde{P},k}(s') - V_{\gamma, \pi}^{\tilde{P},k}(s) \right) \right) \quad (133)$$

$$= \left(\sum_{s' \in \mathcal{S}} \left(\tilde{P}(s'|s, a) - P(s'|s, a) \right) h_{\pi}^{\tilde{P},k}(s') \right) \quad (134)$$

$$\leq \left\| \tilde{P}(\cdot|s, a) - P(\cdot|s, a) \right\|_1 \|h_{\pi}^{\tilde{P},k}(\cdot)\|_{\infty} \quad (135)$$

$$\leq \epsilon_{s,a} \tilde{D} \quad (136)$$

where Equation (129) comes from the assumption that the rewards are known to the agent. Equation (133) follows from the fact that the difference between value function at two states is bounded. Equation (134) comes from the definition of bias term Puterman (1994) where h is the bias of the policy π when run on the sampled MDP. Equation (135) follows from Hölder's inequality. In Equation (136), $\|h(\cdot)\|_{\infty}$ is bounded by the diameter \tilde{D} of the sampled MDP (from Lemma 11). Also, the ℓ_1 norm of probability vector difference is bounded from the definition.

Additionally, note that the ℓ_1 norm in Equation (135) is bounded by 2. Thus the Bellman error is loose upper bounded by $2\tilde{D}$ for all state-action pairs. \blacksquare

B.3 Proof of Lemma 7

Proof From the result of Weissman et al. (2003), the ℓ_1 distance of a probability distribution over S events with n samples is bounded as:

$$\begin{aligned} \mathbb{P}\left(\|P(\cdot|s, a) - \hat{P}(\cdot|s, a)\|_1 \geq \epsilon\right) &\leq (2^S - 2) \exp\left(-\frac{n(s, a)\epsilon^2}{2}\right) \\ &\leq (2^S) \exp\left(-\frac{n(s, a)\epsilon^2}{2}\right) \end{aligned} \quad (137)$$

Thus, for $\epsilon = \sqrt{\frac{2}{n(s, a)} \log(2^S 20SAT^7)} \leq \sqrt{\frac{14S}{n(s, a)} \log(2At)} \leq \sqrt{\frac{14S}{n(s, a)} \log(2AT)}$, we have

$$\mathbb{P}\left(\|P(\cdot|s, a) - \hat{P}(\cdot|s, a)\|_1 \geq \sqrt{\frac{14S}{n(s, a)} \log(2At)}\right) \leq (2^S) \exp\left(-\frac{n(s, a)}{2} \frac{2}{n(s, a)} \log(2^S 20SAT^7)\right) \quad (138)$$

$$= 2^S \frac{1}{2^S 20SAT^7} \quad (139)$$

$$= \frac{1}{20AST^7} \quad (140)$$

We sum over the all the possible values of $n(s, a)$ till t time-step to bound the probability that the event \mathcal{E}_t does not occur as:

$$\sum_{n(s, a)=1}^t \frac{1}{20SAT^7} \leq \frac{1}{20SAT^6} \quad (141)$$

Finally, summing over all the s, a , we get

$$\mathbb{P}\left(\|P(\cdot|s, a) - \hat{P}(\cdot|s, a)\|_1 \geq \sqrt{\frac{14S}{n(s, a)} \log(2At)} \forall s, a\right) \leq \frac{1}{20t^6} \quad (142)$$

■