

Online Reinforcement Learning of Optimal Threshold Policies for Markov Decision Processes

Arghyadip Roy, *Member, IEEE*, Vivek Borkar, *Fellow, IEEE*, Abhay Karandikar, *Member, IEEE* and Prasanna Chaporkar, *Member, IEEE*

Abstract—To overcome the *curse of dimensionality and modeling* of Dynamic Programming (DP) methods to solve Markov Decision Process (MDP) problems, Reinforcement Learning (RL) methods are adopted in practice. Contrary to traditional RL algorithms which do not consider the structural properties of the optimal policy, we propose a structure-aware learning algorithm to exploit the ordered multi-threshold structure of the optimal policy, if any. We prove the asymptotic convergence of the proposed algorithm to the optimal policy. Due to the reduction in the policy space, the proposed algorithm provides remarkable improvements in storage and computational complexities over classical RL algorithms. Simulation results establish that the proposed algorithm converges faster than other RL algorithms.

Index Terms—Markov Decision Process, Stochastic Approximation Algorithms, Reinforcement Learning, Stochastic Control, Online Learning of Threshold Policies.

I. INTRODUCTION

Markov Decision Process (MDP) [2] is a framework which is widely used for the optimization of stochastic systems (e.g., queue [3], inventory and production management) to make optimal temporal decisions. Dynamic Programming (DP) methods [2] to compute the *optimal policy* suffer from the *curse of dimensionality* [4, Chapter 4.1], [5] in many practical applications as they are computationally inconvenient due to extremely high dimension of the iterates. Furthermore, they suffer from the *curse of modeling* since the knowledge of the underlying transition probabilities (which often depend on the statistics of unknown system parameters) required by DP methods, may not be available beforehand. RL techniques [6] address the curse of modeling by learning the optimal policy iteratively. RL algorithms do not require any prior knowledge regarding the transition probabilities of the underlying model. RL being sampling based, updates only one component at a time, reducing per iterate computation at the expense of speed. Examples are: Q-learning [7] iteratively evaluates the Q-function of every state-action pair using a combination of *exploration* and *exploitation*. In [8]–[10], upper confidence bound based exploration improves the convergence speed over classical ϵ -greedy exploration. PDS learning algorithm [4], [11] removes the requirement of action exploration and results in faster convergence than Q-learning. Virtual Experience (VE) learning algorithm in [12] updates multiple PDSs at a time. Faster convergence is achieved at the cost of increased computational complexity.

However, popular RL techniques [4], [6], [7], [11], [13], [14] do not exploit the known structural properties of the optimal policy

This paper is a substantially expanded and revised version of the work in [1].

Arghyadip Roy was with Department of Electrical Engineering, Indian Institute of Technology Bombay, Mumbai, 400076, India when the work was done and is currently with Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, 61820, USA. e-mail: arghyad4@illinois.edu. Vivek Borkar, Abhay Karandikar and Prasanna Chaporkar are with the Department of Electrical Engineering, Indian Institute of Technology Bombay. e-mail: {borkar, karandi, chaporkar}@ee.iitb.ac.in. Abhay Karandikar is currently Director, Indian Institute of Technology Kanpur (on leave from IIT Bombay), Kanpur, 208016, India. e-mail: karandi@iitk.ac.in.

if any, and consider the set of all policies as the policy search space. However, in operations research and communications literature, various structural properties of the optimal policy including threshold structure, transience of certain states and index rules [15], [16] are often established using monotonicity, convexity/concavity and sub-modularity/super-modularity properties of value functions of states. If one can exploit these structural properties to reduce the search space while learning, then faster convergence can be achieved with a reduction in the computational complexity. A few works in the literature [17]–[20] focus on the exploitation of the structural properties [16] while learning the optimal policy. Q-learning based approaches in [17], [19], in every iteration, project the value functions to guarantee the monotonicity in system state. Although improved convergence speed is obtained, the per-iteration computational complexity does not improve over Q-learning. In [18], a scheme based on piecewise linear approximation of the value function, is proposed. However, as the approximation becomes better, the complexity increases. In [21], a Stochastic Approximation (SA) [22] approach based on simultaneous perturbation is proposed to compute the optimal thresholds. It uses a combination of ideas from renewal theory and Monte Carlo simulation. The low complexity Q-learning algorithm proposed in [23] exploits the threshold nature of the optimal policy. The proposed algorithm estimates the optimal policy for a subset of sets, and then policy interpolation is performed for the unvisited states. The performance of policy interpolation is further improved by a policy refinement. However, the proposed algorithm provides a near-optimal policy after a finite number of visits to a set of state-action pairs.

We consider a scenario where the optimal policy has a multi-threshold structure, and the thresholds for different events are ordered. Therefore, learning the optimal policy is equivalent to learning the value of threshold for each event. Motivated by this, we propose a Structure-Aware Learning for Multiple Thresholds (SALMUT) algorithm which considers only the set of ordered threshold policies. We consider a two timescale approach where the value functions of the states and threshold parameters are updated in faster and slower timescale, respectively. Threshold parameters are updated based on the gradients of the average reward w.r.t the threshold. The proposed scheme results in reductions in storage and computational complexities (by amortizing the computation, and hence the complexity of the original problem, over several iterations) compared to traditional RL schemes. We prove that the proposed algorithm converges to the optimal policy asymptotically. Simulation results exhibit that SALMUT converges faster than classical RL schemes due to reduction in policy search space. The techniques in this paper can be employed to learn the optimal policy in problems involving optimality of threshold policies [15], [24]–[27]. We illustrate this using an example of a finite buffer multi-server queue with multiple customer classes (similar to [3], [28]).

The proposed SALMUT algorithm can be adopted for single threshold case in [1], without any modification. In [29], a structure-aware learning algorithm learns a single parameterized threshold where the thresholds for different parameter values are independent of each other. However, in this paper, the thresholds have to

satisfy certain ordering constraints and hence, are not independent. Therefore, the threshold update scheme in the slower timescale and corresponding convergence behavior differs significantly from [29]. To the best of our knowledge, contrary to other works [17]–[20], we for the first time consider the threshold vector as a parameter while learning to reduce a non-linear iteration (involving maximization over a set of actions) into a linear iteration for a quasi-static value of threshold and thereby, achieve a significant reduction in the per-iteration computational complexity.

II. SYSTEM MODEL & PROBLEM FORMULATION

Consider a continuous time MDP problem where we aim to obtain the optimal control policy in a multi-event scenario. Let the system state in the state space ($\mathcal{S} \times \mathcal{I}$, say) be (s, i) where $s \in \{0, 1, \dots, W\}$, and $i \in \{0, 1, \dots, N\}$ denotes the event type. Since we have a finite-state regular Markov chain [30], it is sufficient to observe the system state only at the decision epochs [31]. Let the action space \mathcal{A} consists of two actions, viz., A_1 , and A_2 . The transition probability from state (s, i) to (s', i') under action a ($p((s, i), (s', i'), a)$, say) can be factored into two parts, viz., the deterministic and the probabilistic transitions due to the chosen action ($p(s, s', a)$, say) and the next event, respectively. Let the mean transition rate from state (s, i) be denoted by $v(s)$. Let the arrival times of events in state s be independent exponentially distributed with means $\lambda_i(s)$ respectively, so that the next event is i with probability $\frac{\lambda_i(s)}{\sum_{j=0}^N \lambda_j(s)}$.

Now, $v(s) = \sum_{i=0}^N \lambda_i(s)$. Hence,

$$p((s, i), (s', i'), a) = p(s, s', a) \frac{\lambda_{i'}(s')}{v(s')}.$$

Let the non-negative reward rate ($r((s, i), a)$, say) obtained by choosing action A_2 for event i be R_i , where $R_i > R_j$ for $i < j$. Therefore, $r((s, i), a) = R_i \mathbb{1}_{\{a=A_2\}}$. Let the non-negative cost rate in state (s, i) be $h(s)$ (independent of i).

Let \mathcal{Q} be the set of stationary policies. Since the zero state is reachable from any state with positive probability, the underlying Markov chain is unichain and hence, a unique stationary distribution exists. Let the infinite horizon average reward (independent of the initial state) under policy $Q \in \mathcal{Q}$ be ρ_Q . We aim to maximize $\rho_Q = \lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}_Q[R(t)]$, where $R(t)$ is the total reward till time t . The DP equation describing the necessary condition for optimality in a semi-Markov decision process ($\forall \{(s, i), (s', i')\} \in \mathcal{S} \times \mathcal{I}$) is

$$\begin{aligned} \bar{V}(s, i) = & \max_{a \in \mathcal{A}} [r((s, i), a) + \sum_{s', i'} p((s, i), (s', i'), a) \bar{V}(s', i') \\ & - \rho \bar{\beta}((s, i), a)] - h(s), \end{aligned}$$

where $\bar{V}(s, i)$, ρ and $\bar{\beta}((s, i), a)$ denote the value function of state (s, i) , the optimal average reward and the mean transition time from state (s, i) upon choosing action a , respectively. We rewrite the DP equation after substituting the values of $r((s, i), a)$ and transition probabilities as

$$\begin{aligned} \bar{V}(s, i) = & \max_{a \in \mathcal{A}} [R_i \mathbb{1}_{\{a=A_2\}} + \sum_{s', i'} p(s, s', a) \frac{\lambda_{i'}(s')}{v(s')} \bar{V}(s', i') \\ & - \rho \bar{\beta}((s, i), a)] - h(s). \end{aligned}$$

We define $V(s) := \sum_{i=0}^N \frac{\lambda_i(s)}{v(s)} \bar{V}(s, i)$. Therefore, the following relations hold.

$$\begin{aligned} \bar{V}(s, i) = & \max_{a \in \mathcal{A}} [R_i \mathbb{1}_{\{a=A_2\}} + \sum_{s'} p(s, s', a) V(s') \\ & - \rho \bar{\beta}((s, i), a)] - h(s), \end{aligned}$$

and (since $\frac{\lambda_i(s)}{v(s)}$ is independent of a)

$$\begin{aligned} V(s) = & \max_{a \in \mathcal{A}} \left[\sum_i \frac{\lambda_i(s)}{v(s)} R_i \mathbb{1}_{\{a=A_2\}} + \sum_{s'} p(s, s', a) V(s') \right. \\ & \left. - \rho \bar{\beta}(s, a) \right] - h(s), \end{aligned} \quad (1)$$

where $\bar{\beta}(s, a) = \sum_i \frac{\lambda_i(s)}{v(s)} \bar{\beta}((s, i), a)$. Using (1), instead of considering $(s, i) \in \mathcal{S} \times \mathcal{I}$, we can consider the system state as $s \in \mathcal{S}$ with value function $V(s)$ and transition probability $p(s, s', a)$, and the analysis remains unaffected. However, in this model, the reward rate is the weighted average of original reward rates. The sojourn times being exponentially distributed, following [2], we obtain

$$\begin{aligned} V(s) = & \max_{a \in \mathcal{A}} \left[\sum_i \frac{\lambda_i(s)}{v(s)} R_i \mathbb{1}_{\{a=A_2\}} + \sum_{s'} p(s, s', a) V(s') \right] \\ & - h(s) - \rho, \end{aligned} \quad (2)$$

where $p(s, s, a) = 1 - \sum_{s' \neq s} p(s, s', a)$. (2) is the DP equation for an equivalent discrete time MDP having controlled transition probabilities $p(s, s', a)$ which is used throughout the rest of the paper. This problem can be solved using Relative Value Iteration Algorithm (RVIA) as follows.

$$\begin{aligned} V_{n+1}(s) = & \max_{a \in \mathcal{A}} \left[\sum_i \frac{\lambda_i(s)}{v(s)} R_i \mathbb{1}_{\{a=A_2\}} + \sum_{s'} p(s, s', a) V_n(s') \right] \\ & - V_n(s^*) - h(s), \end{aligned} \quad (3)$$

where $s^* \in \mathcal{S}$ is a fixed state and $V_n(s)$ is the estimate of value function of state s at n^{th} iteration.

III. EXPLOITATION OF STRUCTURAL PROPERTIES IN RL

We assume that the optimal policy is of threshold-type where it is optimal to choose action A_2 only upto a threshold $\tau(i) \in \mathcal{S}$ which is a non-increasing function of i . In this section, we propose an RL algorithm which exploits the knowledge regarding the existence of a threshold-based optimal policy.

Given that the optimal policy is threshold in nature where the optimal action changes from A_2 to A_1 at $\tau(i)$ for i^{th} event, the knowledge of $\tau(0), \dots, \tau(N)$ uniquely characterizes the optimal policy. However, computation of these threshold parameters requires the knowledge of the event probabilities in state s (governed by $\lambda_i(s)$). When the $\lambda_i(\cdot)$ s are unknown, then we can learn these ordered thresholds instead of learning the optimal policy from the set of all policies including the non-threshold policies. We devise an iterative update rule for a threshold vector of dimensionality $(N+1)$ so that the threshold vector iterate converges to the optimal threshold vector.

We consider the set of threshold policies where the thresholds for different events are ordered ($\tau(i) \geq \tau(j)$ for $i < j$) and represent them as policies parameterized by the threshold vector $\tau = [\tau(0), \tau(1), \dots, \tau(N)]^T$ where $\tau(0) \geq \tau(1) \geq \dots \geq \tau(N)$. In this context, we redefine the notations associated with the MDP to reflect their dependence on τ . We intend to compute the gradient of the average reward w.r.t τ and improve the policy by updating τ in the direction of the gradient. Let the transition probability from state s to state s' corresponding to the threshold vector τ be $P_{ss'}(\tau)$. Hence, $P_{ss'}(\tau) = P(X_{n+1} = s' | X_n = s, \tau)$. Let the value function of state s , the average reward of the Markov chain and the stationary probability of state s parameterized by τ be denoted by $V(s, \tau)$, $\sigma(\tau)$ and $\pi(s, \tau)$, respectively.

The optimal policy can be computed using (3) if we know the state transition probabilities and $\lambda_i(s)$ s. When these parameters are unknown, theory of SA [22] enables us to replace the expectation

operation in (3) by averaging over time and still converge to the optimal policy. Let $a(n)$ be a positive step-size sequence satisfying $\sum_{n=1}^{\infty} a(n) = \infty$; $\sum_{n=1}^{\infty} (a(n))^2 < \infty$. Let $b(n)$ be another step-size sequence which apart from the above properties satisfies $\lim_{n \rightarrow \infty} \frac{b(n)}{a(n)} = 0$. We update the value function of the system state (based on the type of event) at any given iteration and keep the value functions of other states unchanged. Let S_n be the state of the system at n^{th} iteration. Let $\gamma(s, n) = \sum_{x=1}^n \mathbb{1}_{\{S_x=s\}}$. For a fixed threshold vector τ (i.e., a fixed policy), max operator in (2) goes away, and the resulting system becomes a linear system. Therefore we have

$$V(s) = \sum_i \frac{\lambda_i(s)}{v(s)} R_i \mathbb{1}_{\{l(s, \tau(i))=A_2\}} + \sum_{s'} p(s, s', a) V(s') - h(s) - \rho,$$

where $l(s, \tau(i)) = A_2$ if $s < \tau(i)$, $l(s, \tau(i)) = A_1$ else.

The update of value function of state s (corresponding to the i^{th} event) is done using the following scheme:

$$\begin{aligned} V_{n+1}(s, \tau) &= (1 - a(\gamma(s, n))) V_n(s, \tau) + a(\gamma(s, n)) [-h(s) + \\ &\quad R_i \mathbb{1}_{\{l(s, \tau(i))=A_2\}} + V_n(s', \tau) - V_n(s^*, \tau)], \quad (4) \\ V_{n+1}(\tilde{s}, \tau) &= V_n(\tilde{s}, \tau), \forall \tilde{s} \neq s, \end{aligned}$$

where $V_n(s, \tau)$ denotes the value function of state s at n^{th} iteration provided the threshold vector is τ . This is known as the primal RVIA which is performed in the faster timescale. The scheme (4) works for a fixed value of threshold vector. To obtain the optimal value of τ , the threshold vector needs to be iterated in a slower timescale $b(\cdot)$. The idea is to learn the optimal threshold vector by computing $\nabla \sigma(\tau)$ and update the value of threshold in the direction of the gradient. This scheme is similar to stochastic gradient routine:

$$\tau_{n+1} = \tau_n + b(n) \nabla \sigma(\tau_n), \quad (5)$$

where τ_n is the threshold vector at n^{th} iteration. Conditions on step sizes ensure that the value function and threshold vector iterates are updated in different timescales. From the slower timescale, the value functions seem to be quasi-equilibrated, whereas from the faster timescale, the threshold vector appears to be quasi-static (known as the ‘‘leader-follower’’ behavior).

Given a threshold vector τ , it is assumed that the transition from state s for i^{th} event is driven by the rule $P_1(s'|s)$ if $s < \tau(i)$ and by the rule $P_0(s'|s)$ otherwise. Under rule $P_1(s'|s)$, the system moves from state s to state $s' = s + 1$ following action A_2 if $s < \tau(i)$. On the other hand, rule $P_0(s'|s)$ dictates that the system remains in state s following action A_1 . For a fixed τ , (4) is updated using the above rule. Let $g_s(\tau(i)) = R_i \mathbb{1}_{\{l(s, \tau(i))=A_2\}} - h(s)$. The following assumption is made on $P_{ss'}(\tau)$ and $g_s(\tau)$ to embed the discrete parameter τ into a continuous domain later.

Assumption 1. $P_{ss'}(\tau)$ and $g_s(\tau)$ are bounded and twice differentiable functions of τ . It has bounded first and second derivatives.

For a given event, the threshold policy chooses the rule $P_1(\cdot|s)$ upto a threshold and follows the rule $P_0(\cdot|s)$, thereafter. Therefore the threshold policy is defined at discrete points and does not satisfy Assumption 1 as the derivative is undefined. To address this issue, we propose an approximation (\approx interpolation to continuous domain) of the threshold policy, which resembles a step function, so that the derivative exists at every point. This results in a randomized policy which in state s , chooses policies $P_0(s'|s)$ and $P_1(s'|s)$ with probabilities $f(s, \tau)$ and $1 - f(s, \tau)$, respectively. In other words,

$$P_{ss'}(\tau) \approx P_0(s'|s) f(s, \tau) + P_1(s'|s) (1 - f(s, \tau)). \quad (6)$$

Intuitively, $f(s, \tau)$ should allocate similar probabilities to $P_0(s'|s)$ and $P_1(s'|s)$ near the threshold. As we move away towards the

left (right) direction, the probability of choosing $P_0(s'|s)$ ($P_1(s'|s)$) should decrease. The following function is chosen as a convenient approximation as it is continuously differentiable and the derivative is non-zero at every point.

$$f(s, \tau(i)) = \frac{e^{(s-\tau(i)-0.5)}}{1 + e^{(s-\tau(i)-0.5)}}. \quad (7)$$

Similar to (6), we approximate $g_s(\tau)$ as

$$g_s(\tau(i)) \approx -f(s, \tau(i)) h(s) + (1 - f(s, \tau(i))) (R_i - h(s)).$$

Remark 1. Note that although the state space is discrete, individual threshold vector component iterates may take values in the continuous domain. However, only an ordinal comparison dictates which action needs to be chosen in the current state.

Remark 2. Instead of the sigmoid function in (7), the function $f(s, \tau(i)) = \mathbb{1}_{\{s \geq \tau(i)+1\}} + (s - \tau(i)) \mathbb{1}_{\{s < \tau(i) < s+1\}}$ which uses approximation only when $s < \tau(i) < s + 1$, could have been chosen. This function suffers lesser approximation error than that of (7). However, it may lead to slow convergence since the derivative of the function and hence the gradient becomes zero outside $s < \tau(i) < s + 1$. Although the derivative of sigmoid function decays exponentially fast too, we observe in simulations that the convergence behavior with (7) is better.

Under Assumption 1, the following proposition [32, Proposition 1] provides a closed form expression for the gradient of $\sigma(\tau)$. The proposition stated next is extended to policy gradient theorem in [33].

Proposition 1.

$$\nabla \sigma(\tau) = \sum_{s \in \mathcal{S}} \pi(s, \tau) (\nabla g_s(\tau) + \sum_{s' \in \mathcal{S}} \nabla P_{ss'}(\tau) V(s', \tau)). \quad (8)$$

Remark 3. The first term on the right of (8) was dropped earlier, the reason given being that ‘‘the reward is independent of the parameter’’. This is incorrect because the indicator function multiplying the reward is not so. Nevertheless, it turns out (see Fig. 1a and 1b) that the performance is hardly affected by the omission of this term. This is presumably because this term is significant only in a small neighborhood of the threshold. This mistake is present in our earlier works [1], [29] and similar remarks apply there.¹

Based on the proposed approximation (7), we set to devise an online update rule for the threshold vector in the slower timescale $b(\cdot)$, following (5). We evaluate $\nabla P_{ss'}(\tau)$ as a representative of $\nabla \sigma(\tau)$ since the stationary probabilities in (8) can be replaced by averaging over time. Using (6), we get

$$\nabla P_{ss'}(\tau) = (P_0(s'|s) - P_1(s'|s)) \nabla f(s, \tau). \quad (9)$$

We incorporate a multiplying factor of $\frac{1}{2}$ in the right hand side of (9) since multiplication by a constant term does not alter the scheme. The physical significance of this operation is that at every iteration, transitions following rules $P_1(\cdot|s)$ and $P_0(\cdot|s)$ are adopted with equal probabilities. $\nabla f(s, \tau)$ depends on the system state and threshold vector at any given iteration. Similar online rule can be devised for $\nabla g_s(\tau)$ using an identical procedure.

Based on this, when i^{th} event occur, the online update rule for the i^{th} component of the threshold vector is as follows.

$$\begin{aligned} \tau_{n+1}(i) &= \Omega_i [\tau_n(i) + b(n) \nabla f(s, \tau_n(i)) ((-1)^{\beta_n} \hat{h}_{\beta_n}(s, i) + \\ &\quad (-1)^{\alpha_n} V_n(\tilde{s}, \tau_n(i)))] \end{aligned}$$

¹We thank Prof. Aditya Mahajan, McGill University for pointing out this error.

where α_n and β_n are i.i.d. random variables, each of which can take values 0 and 1 with equal probabilities. If $\alpha_n = 1$, then the transition is governed by the rule $P_1(\cdot|\cdot)$, else by $P_0(\cdot|\cdot)$. In other words, the next state is \hat{s} with probability $\alpha_n P_1(\hat{s}|s) + (1 - \alpha_n) P_0(\hat{s}|s)$. Similarly, $\hat{h}_\beta(s, i) = -\beta h(s) + (1 - \beta)(R_i - h(s))$ where the obtained reward is $-h(s)$ and $R_i - h(s)$ with equal probabilities.

The projection operator Ω_i ensures that $\tau(i)$ iterates remain bounded in a specific interval, as specified later. Recall that, we have assumed that the threshold for i^{th} event is a non-increasing function of i . Therefore, the i^{th} component of the threshold vector iterates should always be less than or equal to the $(i - 1)^{\text{th}}$ component. The first component of τ is considered to be a free variable which can choose any value in $[0, W]$. The projection operator $\Omega_i, \{i > 0\}$ ensures that $\tau_n(i)$ remains bounded in $[0, \tau_n(i - 1)]$. To be precise,

$$\begin{aligned} \Omega_0 : x &\mapsto 0 \vee (x \wedge W) \in [0, W], \\ \Omega_i : x &\mapsto 0 \vee (x \wedge \tau(i - 1)) \in [0, \tau(i - 1)]. \quad \forall i > 0. \end{aligned}$$

The framework of SA enables us to obtain the effective drift in (9) by performing averaging. Therefore the online RL scheme where the value functions and the threshold vector are updated in the faster and the slower timescale, respectively, is as follows. We suppress the parametric dependence of V on τ .

$$\begin{aligned} V_{n+1}(s) &= (1 - a(\gamma(s, n)))V_n(s) + a(\gamma(s, n))[-h(s) + \\ &\quad R_i \mathbb{1}_{\{I(s, \tau(i))=A_2\}} + V_n(s') - V_n(s^*)], \quad (10) \\ V_{n+1}(\tilde{s}) &= V_n(\tilde{s}), \forall \tilde{s} \neq s, \end{aligned}$$

and

$$\begin{aligned} \tau_{n+1}(i) &= \Omega_i[\tau_n(i) + b(n)\nabla f(s, \tau_n(i))((-1)^{\beta_n} \hat{h}_{\beta_n}(s, i) + \\ &\quad (-1)^{\alpha_n} V_n(\hat{s}))], \\ \tau_{n+1}(i') &= \tau_n(i'), \quad \forall i' < i, \\ \tau_{n+1}(i') &= \Omega_i[\tau_n(i')], \quad \forall i' > i, \end{aligned} \quad (11)$$

where for current state s , transition to next state s' in (10) refers to a single run of a simulated chain as commonly seen in RL, and \hat{s} in (11) is determined separately following the probability distribution $\alpha_n P_1(\hat{s}|s) + (1 - \alpha_n) P_0(\hat{s}|s)$. The immediate reward is determined using $-\beta_n h(s) + (1 - \beta_n)(R_i - h(s))$. The physical significance of (11) is that when i^{th} type of event occurs, then the i^{th} component of τ is updated. However, since the components are provably ordered, we need to update the i^{th} components too where $i' > i$. We no longer need to update the components for $i' < i$ since the order is already preserved while updating the i^{th} component. Also, in (10), the reward function is taken to be $R_i \mathbb{1}_{\{I(s, \tau(i))=A_2\}}$ when i^{th} event occurs. The expectation operation in (3) is mimicked by the averaging over time implicit in an SA scheme. Note that contrary to [29], where due to the independence among the threshold parameters, only one threshold is updated at a time, in this paper, multiple threshold parameters may need to get updated to capture the ordering constraints.

Remark 4. *Instead of the two-timescale approach adopted in this paper, a multi-timescale approach where each individual threshold is updated in a separate timescale, may be chosen. However, since the updates of thresholds are coupled only through the ordering constraints, they can be updated in the same timescale. Moreover, in practice, a multi-timescale scheme may not work well since the fastest (slowest) timescale may be too fast (slow), leading to higher fluctuations and/or very slow speed.*

Theorem 1. *Update rules (10) and (11) converge to the optimal policy almost surely (a.s.).*

Proof: Proof is given in Appendix I.

Remark 5. *Unlike the value function iterates, the threshold vector iterates do not require local clocks of individual elements for convergence. However, all components of the threshold vector are required to get updated comparably often. In other words, their frequencies of update should be bounded away from zero which generally holds true for stochastic gradient approaches [22, Chapter 7].*

Based on the foregoing analysis, we describe the resulting SALMUT algorithm in Algorithm 1. On a decision epoch, action a is chosen based on the current value of threshold vector. Based on the event, value function of current state s is updated then using (10) in the faster timescale. The threshold vector is also updated following (11) in the slower timescale. Note that the value function is updated one component at a time. However, multiple components of the threshold vector may need to be updated in a single iteration. The scheme resembles actor-critic method [13] with policy gradient for the actor part (see (11)) and post-decision framework for the critic (see (10)).

Algorithm 1 Two-timescale SALMUT algorithm

- 1: Initialize $n \leftarrow 1, V(s) \leftarrow 0, \forall s \in \mathcal{S}$ and $\tau \leftarrow \bar{0}$.
 - 2: **while** TRUE **do**
 - 3: **if** i^{th} event occurs **then**
 - 4: Choose action a based on current value of $\tau(i)$.
 - 5: **end if**
 - 6: Update value function of state s using (10).
 - 7: Update threshold τ using (11).
 - 8: Update $s \leftarrow s'$ and $n \leftarrow n + 1$.
 - 9: **end while**
-

Remark 6. *Even if there does not exist an optimal threshold policy for a given MDP problem, the techniques in this paper can be applied to learn the best threshold policy (locally at least) asymptotically. Threshold policies are easy to implement and often provide comparable performances to that of the optimal policy, with a significantly lower storage complexity.*

IV. COMPLEXITY ANALYSIS

In this section, we compare the storage and computational complexities of SALMUT algorithm with those of existing learning schemes and summarize in Table I. Q-learning and PDS learning

TABLE I: Complexities of RL algorithms.

Algorithm	Computational complexity	Storage complexity
Q-learning [6], [7]	$O(\mathcal{A})$	$O(\mathcal{S} \times \mathcal{A})$
Monotone Q-learning [17], [19]	$O(\mathcal{A})$	$O(\mathcal{S} \times \mathcal{A})$
MH-Q-learning	$O(\log \mathcal{A})$	$O(\mathcal{S} \times \mathcal{A})$
OQL [9]	$O(\log \mathcal{A})$	$O(\mathcal{S} \times \mathcal{A})$
EEQL [10]	$O(\log \mathcal{A})$	$O(\mathcal{S} \times \mathcal{A})$
PDS learning [4], [11]	$O(\mathcal{A})$	$O(\mathcal{S})$
MH-PDS learning	$O(\log \mathcal{A})$	$O(\mathcal{S})$
VE learning [12]	$O(\mathcal{V} \times \mathcal{A})$	$O(\mathcal{S})$
Grid learning [20]	$O(\mathcal{W} \times \mathcal{A})$	$O(\mathcal{S})$
Adaptive appx. learning [18]	$O(q(\delta) \mathcal{A})$	$O(\mathcal{S})$
SALMUT	$O(1)$	$O(\mathcal{S})$

need to store the value function of every state-action pair and every PDS, respectively. While updating the value function, both choose the best one after evaluating $|\mathcal{A}|$ functions. Since only one state-action pair is updated at a time, the computational complexity associated with remaining operations is constant. Thus, the storage complexities of Q-learning and PDS learning are $O(|\mathcal{S}| \times |\mathcal{A}|)$ and $O(|\mathcal{S}|)$, respectively. The per-iteration computational complexity of

Q-learning and PDS learning is $O(|\mathcal{S}|)$. Since at each iteration, the monotone Q-learning algorithm [17], [19] projects the policy obtained using Q-learning within the set of monotone policies, the complexities are identical to those of Q-learning. VE learning [12] updates multiple PDSs at a time. Therefore, the computational complexity contains an additional term $|\mathcal{V}|$ which signifies the cardinality of the VE tuple. Similarly, grid learning [20] and adaptive approximation learning [18] are associated with additional factors $|\mathcal{W}|$ and $q(\delta)$, respectively. $|\mathcal{W}|$ and $q(\delta)$ depend on the depth of a quadtree used for value function approximation and the approximation error threshold δ , respectively. Note that computational complexities of Q-learning and PDS learning can be reduced to $O(\log |\mathcal{A}|)$ using a max-heap implementation (MH-Q-learning and MH-PDS learning in Table I) where the complexities of obtaining the best action and updating a value are $O(1)$ and $O(\log |\mathcal{A}|)$, respectively. The computational complexity of a max-heap implementation of Optimistic Q-learning (OQL) [9] and Exploration Enhanced Q-learning (EEQL) [10] is $O(\log |\mathcal{A}|)$. The storage complexity of OQL and EEQL is $O(|\mathcal{S}| \times |\mathcal{A}|)$. It is not clear whether the knowledge of structural properties can be encoded in Q-learning easily since the range of the threshold is large, viz., the entire state space. The storage complexity of SALMUT algorithm is $O(|\mathcal{S}|)$ as we need to store the value functions of states. SALMUT may require to update all components of the threshold vector at a time (See (11)). Furthermore, the update of value function involves the computation of a single function based on the current threshold (See (10)). Therefore, the per-iteration computational complexity is $O(1)$. Thus SALMUT provides significant improvements in storage and per-iteration computational complexities compared to other schemes. Note that the computational complexity of SALMUT does not depend on $|\mathcal{A}|$ and depends only on the number of events.

V. ILLUSTRATIVE EXAMPLE

We consider a queuing system with m identical servers, N customer classes and a finite buffer of size B . We investigate an optimal admission control problem. It is assumed that the arrival of class- i customers is a Poisson process with mean λ_i . Let the service time be exponentially distributed with mean $\frac{1}{\mu}$, irrespective of the customer class. Note that presence of m identical servers does not change the nature of the model. Similar analysis holds for the single server case also. The state of the system is (s, i) where $s \in \{0, 1, \dots, m+B\}$ denotes the total number of customers in the system and $i \in \{0, 1, \dots, N\}$ denotes the class type. Arrivals and departures of customers are taken as decision epochs. We take $i=0$ as the departure event and $i=1, \dots, N$ as an arrival of a 1, \dots , N^{th} class of customer. For example, states $(2, 0)$ and $(2, 1)$ correspond to a departure and an arrival of class 1 customer while there are 2 users in the system, respectively. \mathcal{A} consists of three actions, viz., blocking of an arriving user (A_1), admission of an arriving user (A_2) and continue/do nothing (A_0 , say) for departures. When $s = m+B$, then the only feasible action for an arrival (i.e., $i > 0$) is A_1 . We have, $p(s, s', A_2) = \mathbb{1}_{\{s'=s+1\}}$, $p(s, s', A_1) = \mathbb{1}_{\{s'=s\}}$ and $p(s, s', A_0) = \mathbb{1}_{\{s'=(s-1)^+\}}$. where $x^+ = \max\{x, 0\}$ and $v(s) = \sum_{i=1}^N \lambda_i + \min\{s, m\}\mu$. Let $\lambda_0(s) = \min\{s, m\}\mu$ and $\lambda_i(s) = \lambda_i, \forall i \neq 0$. Note that $\lambda_1(s), \dots, \lambda_N(s)$ do not depend on s . Furthermore, $r((s, i), a) = R_i \mathbb{1}_{\{a=A_2, i>0\}}$. $R_0 = 0$ corresponds to a departure event. In state (s, i) , a non-negative cost rate of $h(s)$ is incurred. $h(s)$ and $h(s+1) - h(s)$ are increasing in s (convex increasing in the discrete domain).

One application of this model is a make-to-stock production system that produces m items with N demand classes [34] and buffer size B . Satisfaction of a demand (requesting a single unit of the product) from class- i gives rise to reward rate R_i . The production time is

exponentially distributed with mean $\frac{1}{\mu}$. The inventory holding cost rate is $h(s)$.

Now, we derive that there exists a threshold based optimal policy which admits a class- i customer only upto a threshold $\tau(i)$ which is a non-increasing function of i . We prove these properties using the following lemma. The proof follows from [35, Theorem 3.1]. Detailed proof is given in [36, Lemma 1].

Lemma 1. $V(s+1) - V(s)$ is decreasing in s .

Theorem 2. The optimal policy is of threshold-type where it is optimal to admit class- i customers only upto a threshold $\tau(i) \in \mathcal{S}$ which is a non-increasing function of i .

Proof: For class- i customers, if A_1 is optimal in state s , then $R_i + V(s+1) \leq V(s)$ (Using (2)). From Lemma 1, $V(s+1) - V(s)$ is decreasing in s . This proves the existence of a threshold $\tau(i)$ for class- i customers. Since $R_i > R_j$ for $i < j$, $R_i + V(s+1) \leq V(s)$ implies $R_j + V(s+1) \leq V(s)$. Therefore, $\tau(i)$ is a non-increasing function of i . \square

The subsequent lemmas establish the unimodality of the average reward with respect to τ . Hence, Theorem 1 holds. Proofs are presented in Appendix II.

Lemma 2. $v_n(s+1) - v_n(s)$ is decreasing in n .

Lemma 3. $\sigma(\tau)$ is unimodal in τ .

For this particular problem, the improvement in computational complexity offered by SALMUT may not be significant since $|\mathcal{A}| = 3$. However, in general, for a large $|\mathcal{A}|$, the improvement in computational complexity may be remarkable.

VI. SIMULATION RESULTS

In this section, we compare the convergence speed of SALMUT algorithm with traditional RL algorithms. We simulate the finite buffer multi-server system with two customer classes. We take $\lambda_1 = \lambda_2 = 1$, s^{-1} , $m = B = 5$, $R_1 = 20$, $R_2 = 10$, $h(s) = 0.1s^2$, $a(n) = \frac{1}{(\lfloor \frac{n}{100} \rfloor + 2)^{0.6}}$ and $b(n) = \frac{10}{n}$. We exclude initial 10 burn-in period values of the iterates.

In practical cases, when the average reward of the system does not change much over a suitable window, we may conclude that stopping condition is met as the obtained policy is close to the optimal policy with a high probability. The choice of window size $\sum_{k=p}^{p+n} a(k)$ over n is to eliminate the effect of diminishing step size affecting the convergence behavior. We choose a window size of 50 and stop when the ratio of maximum and minimum average rewards exceeds 0.95. Fig. 1a (1b) reveals that practical convergences for Q-learning, PDS learning and SALMUT algorithms are achieved in 1180 (1180), 580 (1180) and 426 (580) iterations, respectively. Since unlike PDS learning, Q-learning is associated with exploration mechanism, PDS learning converges faster than Q-learning. However, SALMUT converges faster than both Q-learning and PDS learning algorithms since it operates on a smaller policy space (threshold policies only). Due to efficient exploration, the performances of OQL [9] and EEQL [10] are better than that of Q-learning and slightly worse than that of SALMUT. Since for a given sample path, the sequence of actions are same, the performance of MH-Q-learning (MH-PDS learning) is identical to that of Q-learning (PDS learning). The performance of adaptive approximation learning algorithm [18] is identical to that of PDS learning as batch update of PDSs is not possible. Therefore, we do not show their performances in the plot. The convergence behavior of Renewal Monte Carlo (RMC) [21] algorithm (with discount factor=0.999) is slightly worse than that of SALMUT. Since the sampled Q-learning algorithm [23] provides a

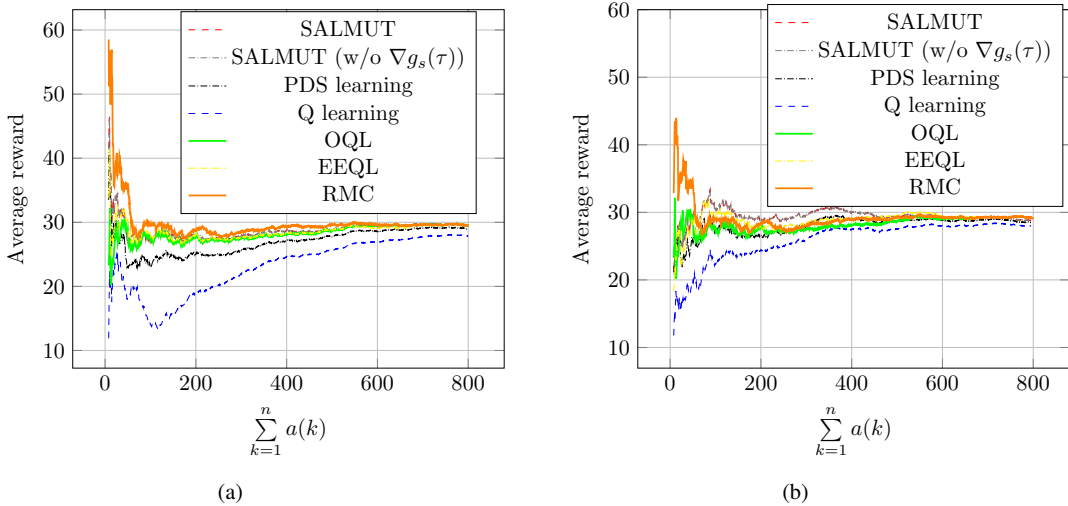


Fig. 1: Plot of average reward vs. $(\sum_{k=1}^n a(k))$ for different algorithms ((a) $\mu = 4s^{-1}$, (b) $\mu = 2s^{-1}$).

near-optimal policy after a finite number of visits to a set of state-action pairs, the convergence behavior may not be comparable.

VII. CONCLUSIONS & FUTURE DIRECTIONS

In this paper, we have proposed an RL algorithm which exploits the ordered multi-threshold nature of the optimal policy. The convergence of the proposed algorithm to the globally optimal threshold vector is established. The proposed scheme provides improvements in storage and computational complexities over traditional RL algorithms. Simulation results establish the improvement in convergence behavior with respect to state-of-the-art RL schemes. In future, this can be extended to develop RL algorithms for Constrained MDP problems by updating the Lagrange Multiplier (LM) in a slower timescale [22] than that of value functions. The LM and the threshold parameter can be updated in the same slower timescale without requiring a third timescale as they are independent of each other. Another possible future direction is to develop RL algorithms for restless bandits such as [37] since threshold policies often translate into index-based policies.

APPENDIX I PROOF OF THEOREM 1

We adopt the approach of viewing SA algorithms as a noisy discretization of a limiting Ordinary Differential Equation (ODE), similar to [1], [29]. Step size parameters are viewed as discrete time steps. Standard assumptions on step sizes ensure that the errors due to noise and discretization are negligible asymptotically. Therefore, the iterates closely follow the trajectory of the ODE and ensure a.s. convergence to the globally asymptotically stable equilibrium. Using the two timescale approach [22], we consider (10) for a fixed τ . Let $M_1 : \mathcal{R}^{|\mathcal{S}|} \rightarrow \mathcal{R}^{|\mathcal{S}|}$ be the following map

$$M_1(x) = \sum_{s'} P_{ss'}(\tau) \left[\sum_{i=0}^N \frac{\lambda_i(s)}{v(s)} R_i \mathbb{1}_{\{a=A_2\}} - h(s) + x(s') \right] - x(s^*), x \in \mathcal{R}^{|\mathcal{S}|}. \quad (12)$$

Note that the knowledge of $P_{ss'}(\tau)$ and $\lambda_i(s)$ are not required for the algorithm and is only required for analysis. For a fixed τ , (10) tracks the limiting ODE $\dot{V}(t) = M_1(V(t)) - V(t)$. $V(t)$ converges to the fixed point of $M_1(\cdot)$ (determined using $M_1(V) = V$) [38] which is the asymptotically stable equilibrium of the ODE, as $t \rightarrow \infty$. Analogous methodologies are adopted in [38], [39]. Note that

the approximation described in (7) does not impact the convergence argument since the fixed point of $M_1(\cdot)$ remains the same. Next we establish that the value function and threshold vector iterates are bounded.

Lemma 4. *The threshold vector and value function iterates are bounded a.s.*

Proof: Let $M_0 : \mathcal{R}^{|\mathcal{S}|} \rightarrow \mathcal{R}^{|\mathcal{S}|}$ be the following map

$$M_0(x) = \sum_{s'} P_{ss'}(\tau) x(s') - x(s^*), x \in \mathcal{R}^{|\mathcal{S}|}. \quad (13)$$

Clearly, if the reward and cost functions are zero, (12) is same as (13). Also, $\lim_{b \rightarrow \infty} \frac{M_1(bV)}{b} = M_0(V)$. The globally asymptotically stable equilibrium of the ODE $\dot{V}(t) = M_0(V(t)) - V(t)$ which is a scaled limit of ODE $\dot{V}(t) = M_1(V(t)) - V(t)$, is the origin. Boundedness of value functions and threshold vector iterates follow from [40] and (11), respectively. \square

Lemma 5. $V_n - V^{\tau_n} \rightarrow 0$, where V^{τ_n} is the value function of states for $\tau = \tau_n$ a.s.

Proof: Since the threshold vector iterates are updated in a slower timescale, value function iterates in the faster timescale treat the threshold vector iterates as fixed. Therefore, iterations for τ are $\tau_{n+1} = \tau_n + \gamma(n)$, where $\gamma(n) = O(b(n)) = o(a(n))$. Therefore, the limiting ODEs for value function and threshold vector iterates are $\dot{V}(t) = M_1(V(t)) - V(t)$ and $\dot{\tau}(t) = 0$, respectively. It is sufficient to consider the ODE $\dot{V}(t) = M_1(V(t)) - V(t)$ alone for a fixed τ because $\dot{\tau}(t) = 0$. The rest of the proof is analogous to that of [13]. \square

For the time being assume that $\sigma(\tau)$ is unimodal in τ . This is proved later. The lemmas presented next establish that threshold vector iterates τ_n converge to the optimal threshold vector τ^* and hence, (V_n, τ_n) converges to the optimal pair (V, τ^*) .

Lemma 6. *The threshold vector iterates $\tau_n \rightarrow \tau^*$ a.s.*

Proof: The limiting ODE for (11) is the gradient ascent scheme $\dot{\tau} = \nabla \sigma(\tau)$. Note that the gradient points inwards at $\tau(\cdot) = 0$ and $\tau(\cdot) = W$. Since $\sigma(\tau)$ is unimodal in τ , there does not exist any local maximum except τ^* which is the global maximum. This concludes the proof of the lemma. \square

Remark 7. *If the unimodality of the average reward with respect to*

the threshold vector does not hold, then convergence to only a local maximum can be guaranteed.

APPENDIX II

A. Proof of Lemma 2:

Proof: Proof methodologies are similar to [1].

$$v_{n+1}(s) = \sum_{i=1}^N \lambda_i \max\{R_i + v_n(s+1), v_n(s)\} + \min\{s, m\} \mu v_n((s-1)^+) + (1-v(s))v_n(s) - h(s), \quad (14)$$

and $\hat{v}_{i,n+1}(s) = \max\{v_n(s), R_i + v_n(s+1)\}$. We know, $Dv_n(s) = v_n(s+1) - v_n(s)$. We use induction to prove that $Dv_n(s)$ is decreasing in n . If $n=0$, $v_0(s) = 0$ and $Dv_0(s) = 0$. Using (14), it is easy to see that $Dv_1(s) < Dv_0(s)$ as $h(s+1) > h(s)$. Assuming that the claim holds for any n , i.e., $Dv_{n+1}(s) < Dv_n(s)$, we need to prove that $Dv_{n+2}(s) < Dv_{n+1}(s)$. To analyze the second and third terms in (14), we consider two cases.

(a) $s \geq m$: $m\mu Dv_{n+1}((s-1)^+) + (1 - \sum_{i=1}^N \lambda_i - m\mu)Dv_{n+1}(s)$ is less than $m\mu Dv_n((s-1)^+) + (1 - \sum_{i=1}^N \lambda_i - m\mu)Dv_n(s)$.

(b) $s < m$: $s\mu Dv_{n+1}((s-1)^+) + (1 - v(s+1))Dv_{n+1}(s)$ is less than $s\mu Dv_n((s-1)^+) + (1 - v(s+1))Dv_n(s)$.

We proceed to prove that $D\hat{v}_{i,n+2}(s) \leq D\hat{v}_{i,n+1}(s)$. Let at $(n+2)^{\text{th}}$ iteration, maximizing actions for the admission of class- i customers in states s and $(s+1)$ be denoted by $a_{i,1} \in \{A_1, A_2\}$ and $a_{i,2} \in \{A_1, A_2\}$, respectively. Let $b_{i,1}, b_{i,2} \in \{A_1, A_2\}$ be the maximizing actions in states s and $(s+1)$, respectively, at $(n+1)^{\text{th}}$ iteration. It is not possible to have $a_{i,2} = A_2$ and $b_{i,1} = A_1$. If $b_{i,1} = A_1$, then $Dv_n(s) \leq -R_i$. Therefore, we must have $Dv_n(s+1) < -R_i$ (Using Lemma 1). If $a_{i,2} = A_2$, then $Dv_{n+1}(s+1) \geq -R_i$ which contradicts the inductive assumption. Therefore, we consider the remaining cases. Note that if the inequality holds for any $a_{i,1}$ and $b_{i,2}$ for given $a_{i,2}$ and $b_{i,1}$, then the maximizing actions will satisfy the inequality too.

(a) $a_{i,2} = b_{i,1} = A_1$: We choose $a_{i,1} = b_{i,2} = A_1$ to get $D\hat{v}_{i,n+2}(s) - D\hat{v}_{i,n+1}(s) = D\hat{v}_{i,n+1}(s) - D\hat{v}_{i,n}(s) \leq 0$.

(b) $a_{i,2} = b_{i,1} = A_2$: Proof is similar to the preceding case by choosing $a_{i,1} = b_{i,2} = A_2$.

(c) $a_{i,2} = A_1, b_{i,1} = A_2$: Choose $a_{i,1} = A_2$ and $b_{i,2} = A_1$. $D\hat{v}_{i,n+2}(s) - D\hat{v}_{i,n+1}(s) = v_{n+1}(s+1) - R_i - v_{n+1}(s+1) - v_n(s+1) + R_i + v_n(s+1) = 0$.

Thus, $D\hat{v}_{i,n+2}(s) \leq D\hat{v}_{i,n+1}(s)$. Since this holds for every i and $h(s)$ is independent of n , this concludes the proof. \square

B. Proof of Lemma 3:

Proof: Proof idea is similar to that of [29]. We prove this lemma for i^{th} component of the threshold vector (viz., $\tau(i)$). If the optimal action for the admission of class- i customers in state s is A_1 , then $V(s+1) - V(s) \leq -R_i$. Since VIA converges to the optimal threshold vector τ^* , $\exists N_0 > 0$ such that $\forall n \geq N_0$, $v_n(s+1) - v_n(s) \leq -R_i, \forall s \geq \tau^*(i)$ and $v_n(s+1) - v_n(s) \geq -R_i, \forall s < \tau^*(i)$. Let $U_{i,n}, n \geq 1$ be the optimal threshold for class- i customers at n^{th} iteration of VIA. Hence, $U_{i,n} = \min\{s \in \mathbb{N}_0 : v_n(s+1) - v_n(s) \leq -R_i\}$. If no values of s satisfies the inequality, then $U_{i,n} = m + B$. Since $v_n(s+1) - v_n(s)$ is decreasing in n (Lemma 2), $U_{i,n}$ is monotonically decreasing in n , and $\lim_{n \rightarrow \infty} U_{i,n} = \tau^*(i)$.

Consider a re-designed problem where for a given threshold vector τ' such that $\tau^*(i) \leq \tau'(i) \leq m + B$, action A_1 is not allowed in any state $s < \tau'(i)$. Note that Lemma 2 holds for this re-designed

problem also. Let $n_{\tau'(i)}$ be the first iteration of VIA when the threshold reduces to $\tau'(i)$. The value function iterates for the original and re-designed problem are same for $n \leq n_{\tau'(i)}$ because in the original problem also A_1 is never chosen as the optimal action in states $s < \tau'(i)$ at these iterations. Hence, $n_{\tau'(i)}$ must be finite and the inequality $v_n(\tau'(i)+1) - v_n(\tau'(i)) \leq -R_i$ is true for both the problems after $n_{\tau'(i)}$ iterations. Using Lemma 2, this inequality holds $\forall n \geq \tau'(i)$. Therefore, in the re-designed problem, $U_{i,n}$ converges to $\tau'(i)$. Thus, the threshold policy with $\tau'(i)$ is superior than that with $\tau'(i)+1$. Since this holds for arbitrary choice of $\tau'(i)$, average reward monotonically decreases with $\tau'(i)$, $\forall \tau'(i) > \tau^*(i)$.

If we have $\sigma(\tau) \geq \sigma(\tau + e_i)$ (where $e_i \in \mathbb{R}^N$ is a vector with all zeros except the i^{th} element being '1'), then we must have $\tau(i) \geq \tau^*(i)$. Therefore, $\sigma(\tau + e_i) \geq \sigma(\tau + 2e_i)$. Hence, the average reward is unimodal in $\tau(i)$. Since the proof holds for any i , this concludes the proof of the lemma. \square

ACKNOWLEDGMENT

Works of Arghyadip Roy, Abhay Karandikar and Prasanna Chaporkar are supported by the Ministry of Electronics and Information Technology (MeitY), Government of India as part of "5G Research and Building Next Gen Solutions for Indian Market" project. Work of Vivek Borkar is supported by the CEFIPRA grant for "Machine Learning for Network Analytics" and a J. C. Bose Fellowship.

REFERENCES

- [1] A. Roy, V. Borkar, A. Karandikar, and P. Chaporkar, "A structure-aware online learning algorithm for Markov decision processes," in *ACM EAI VALUETOOLS*, 2019, pp. 71–78.
- [2] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [3] E. B. Çil, E. L. Örmeci, and F. Karaesmen, "Effects of system parameters on the optimal policy structure in a class of queueing control problems," *Queueing Systems*, vol. 61, no. 4, pp. 273–304, 2009.
- [4] W. B. Powell, *Approximate Dynamic Programming: Solving the curses of dimensionality*. John Wiley & Sons, 2007.
- [5] R. Bellman, "Dynamic programming," *Princeton University Press*, 1957.
- [6] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press Cambridge, 1998.
- [7] C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
- [8] C. Jin, Z. Allen-Zhu, S. Bubeck, and M. I. Jordan, "Is Q-learning provably efficient?" in *Advances in Neural Information Processing Systems*, 2018, pp. 4863–4873.
- [9] C.-Y. Wei, M. Jafarnia-Jahromi, H. Luo, H. Sharma, and R. Jain, "Model-free reinforcement learning in infinite-horizon average-reward Markov decision processes," *arXiv preprint arXiv:1910.07072*, 2019.
- [10] M. Jafarnia-Jahromi, C.-Y. Wei, R. Jain, and H. Luo, "A model-free learning algorithm for infinite-horizon average-reward MDPs with near-optimal regret," *arXiv preprint arXiv:2006.04354*, 2020.
- [11] N. Salodkar, A. Borkar, A. Karandikar, and V. S. Borkar, "An on-line learning algorithm for energy efficient delay constrained scheduling over a fading channel," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 4, pp. 732–742, 2008.
- [12] N. Mastrorade and M. van der Schaar, "Joint physical-layer and system-level power management for delay-sensitive wireless communications," *IEEE Transactions on Mobile Computing*, vol. 12, no. 4, pp. 694–709, 2012.
- [13] V. S. Borkar, "An actor-critic algorithm for constrained Markov decision processes," *Systems & control letters*, vol. 54, no. 3, pp. 207–213, 2005.
- [14] D. P. Bertsekas, *Dynamic programming and optimal control*. Athena scientific Belmont, MA, 1995, vol. 1, no. 2.
- [15] M. Agarwal, V. S. Borkar, and A. Karandikar, "Structural properties of optimal transmission policies over a randomly varying channel," *IEEE Transactions on Automatic Control*, vol. 53, no. 6, pp. 1476–1491, 2008.
- [16] J. E. Smith and K. F. McCardle, "Structural properties of stochastic dynamic programs," *Operations Research*, vol. 50, no. 5, pp. 796–809, 2002.

- [17] S. Kunnunkal and H. Topaloglu, "Exploiting the structural properties of the underlying Markov decision problem in the Q-learning algorithm," *INFORMS Journal on Computing*, vol. 20, no. 2, pp. 288–301, 2008.
- [18] F. Fu and M. van der Schaar, "Structure-aware stochastic control for transmission scheduling," *IEEE Transactions on Vehicular Technology*, vol. 61, no. 9, pp. 3931–3945, 2012.
- [19] M. H. Ngo and V. Krishnamurthy, "Monotonicity of constrained optimal transmission policies in correlated fading channels with ARQ," *IEEE Transactions on Signal Processing*, vol. 58, no. 1, pp. 438–451, 2010.
- [20] N. Sharma, N. Mastrorarde, and J. Chakareski, "Accelerated structure-aware reinforcement learning for delay-sensitive energy harvesting wireless sensors," *arXiv preprint arXiv:1807.08315*, 2018.
- [21] J. Chakravorty and A. Mahajan, "Remote estimation over a packet-drop channel with markovian state," *IEEE Transactions on Automatic Control*, vol. 65, no. 5, pp. 2016–2031, 2019.
- [22] V. S. Borkar, *Stochastic approximation: A dynamical systems viewpoint*. Cambridge University Press, 2008.
- [23] L. Liu and U. Mitra, "On sampled reinforcement learning in wireless networks: Exploitation of policy structures," *IEEE Transactions on Communications*, vol. 68, no. 5, pp. 2823–2837, 2020.
- [24] A. Sinha and P. Chaporkar, "Optimal power allocation for a renewable energy source," in *IEEE NCC*, 2012, pp. 1–5.
- [25] G. Koole, "Structural results for the control of queueing systems using event-based dynamic programming," *Queueing systems*, vol. 30, no. 3-4, pp. 323–339, 1998.
- [26] G. A. Brouns and J. Van Der Wal, "Optimal threshold policies in a two-class preemptive priority queue with admission and termination control," *Queueing Systems*, vol. 54, no. 1, pp. 21–33, 2006.
- [27] M. H. Ngo and V. Krishnamurthy, "Optimality of threshold policies for transmission scheduling in correlated fading channels," *IEEE Transactions on Communications*, vol. 57, no. 8, 2009.
- [28] E. B. Çil, E. L. Örmeci, and F. Karaesmen, "Structural results on a batch acceptance problem for capacitated queues," *Mathematical Methods of Operations Research*, vol. 66, no. 2, pp. 263–274, 2007.
- [29] A. Roy, V. Borkar, P. Chaporkar, and A. Karandikar, "Low complexity online radio access technology selection algorithm in LTE-WiFi hetnet," *IEEE Transactions on Mobile Computing*, vol. 19, no. 2, pp. 376–389, 2019.
- [30] R. W. Wolff, *Stochastic modeling and the theory of queues*. Pearson College Division, 1989.
- [31] A. Kumar, "Discrete event stochastic processes," *Lecture Notes for Engineering Curriculum*, 2012. [Online]. Available: <https://ece.iisc.ac.in/~anurag/books/anurag/spqt.pdf>
- [32] P. Marbach and J. N. Tsitsiklis, "Simulation-based optimization of Markov reward processes," *IEEE Transactions on Automatic Control*, vol. 46, no. 2, pp. 191–209, 2001.
- [33] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *Advances in neural information processing systems*, 2000, pp. 1057–1063.
- [34] A. Y. Ha, "Inventory rationing in a make-to-stock production system with several demand classes and lost sales," *Management Science*, vol. 43, no. 8, pp. 1093–1103, 1997.
- [35] G. Koole, *Monotonicity in Markov reward and decision chains: Theory and applications*. Now Publishers Inc, 2007, vol. 1.
- [36] A. Roy, V. Borkar, A. Karandikar, and P. Chaporkar, "Online reinforcement learning of optimal threshold policies for Markov decision processes," *arXiv preprint arXiv:1912.10325*, 2019.
- [37] V. S. Borkar and K. Chadha, "A reinforcement learning algorithm for restless bandits," in *IEEE ICC*, 2018, pp. 89–94.
- [38] V. R. Konda and V. S. Borkar, "Actor-critic-type learning algorithms for Markov decision processes," *SIAM Journal on control and Optimization*, vol. 38, no. 1, pp. 94–123, 1999.
- [39] J. Abounadi, D. Bertsekas, and V. S. Borkar, "Learning algorithms for Markov decision processes with average cost," *SIAM Journal on Control and Optimization*, vol. 40, no. 3, pp. 681–698, 2001.
- [40] V. S. Borkar and S. P. Meyn, "The ODE method for convergence of stochastic approximation and reinforcement learning," *SIAM Journal on Control and Optimization*, vol. 38, no. 2, pp. 447–469, 2000.