Thai Duong
School of Electrical Engineering and
Computer Science,
Oregon State University,
Corvallis, OR 97331
e-mail: duong@eecs.oregonstate.edu

Duong Nguyen-Huu
School of Electrical Engineering and
Computer Science,
Oregon State University,
Corvallis, OR 97331
e-mail: nguyendu@eecs.oregonstate.edu

Thinh Nguyen
School of Electrical Engineering and
Computer Science,
Oregon State University,
Corvallis, OR 97331
e-mail: thinhq@eecs.oregonstate.edu

Adiabatic Markov Decision Process: Convergence of Value Iteration Algorithm

Markov decision process (MDP) is a well-known framework for devising the optimal decision-making strategies under uncertainty. Typically, the decision maker assumes a stationary environment which is characterized by a time-invariant transition probability matrix. However, in many real-world scenarios, this assumption is not justified, thus the optimal strategy might not provide the expected performance. In this paper, we study the performance of the classic value iteration algorithm for solving an MDP problem under nonstationary environments. Specifically, the nonstationary environment is modeled as a sequence of time-variant transition probability matrices governed by an adiabatic evolution inspired from quantum mechanics. We characterize the performance of the value iteration algorithm subject to the rate of change of the underlying environment. The performance is measured in terms of the convergence rate to the optimal average reward. We show two examples of queuing systems that make use of our analysis framework. [DOI: 10.1115/1.4032875]

1 Introduction

1.1 MDP. The theory of MDP aims to study optimal decision-making processes under uncertainty. It is widely used in economics, engineering, operation research, and artificial intelligence. In an MDP setting, there is a controller which interacts with its environment by taking actions based on its observations at every discrete-time step. Each action by the controller induces a change in the environment. Typically, the environment is described by a finite set of states. An action will move the environment from the current state to some other states with certain probabilities. Associated with each action in each state is a reward given to the controller. The goal of the controller is to maximize the expected cumulative reward or average reward over some finite or infinite number of time steps by making sequential decisions based on its current observations.

It is not difficult to find many applications of the MDP framework. A classic application of MDP is the warehouse example in operation research. In this setup, a company's business is to buy and sell a number of merchandizes. To operate smoothly, it uses a warehouse to store the merchandizes, which allows shipments to the buyers promptly. Everyday, it has to make the decision on how many and which items it should buy and store in its warehouse subject to the uncertainty of the market demands. Buying too many items would incur high storage costs while buying too little would run the risk of not having the items ready for shipping and thus reducing profits. The MDP framework enables the company to decide on the optimal action, i.e., how many and which items it should buy on a given day in order to maximize the expected cumulative reward, i.e., its profits over a month, a year, or indefinitely. Naturally, the optimal action should base on the environmental states, i.e., the current status of different items in the stock and the current market demands.

A solution of an MDP problem is an optimal policy. A policy/decision rule is a mapping from the states to the action. The optimal policy would produce the maximum expected cumulative reward. For the infinite-horizon MDP models, to be discussed subsequently, there are a number of classic algorithms for finding such optimal policies. These algorithms include value iteration [1], policy iteration [2], and linear programing [3,4], which all are based on the Bellman equations [1,5]. Also, all assume a stationary policy, i.e., a policy that does not change with time. This assumption is justified as it is well-known that for a stationary environment, there exists an optimal policy that is also stationary. Fundamentally, the MDP framework relies on the assumption that a given policy will induce a stationary dynamics on the states. Moreover, the state changes are characterized by a time-invariant transition probability matrix P.

1.2 Nonstationary Environment. For many real-world scenarios, this assumption is not justified, thus the optimal policy might not provide the expected performance. In this paper, we study the performance of the classic value iteration algorithm for solving an MDP problem under nonstationary environments. Specifically, the nonstationary environment is modeled as a sequence of time-variant transition probability matrices governed by an adiabatic evolution inspired from quantum mechanics [6–9]. Formally, the transition probability matrix P_i^d at time step i induced by decision rule d is determined by

$$P_i^d = \Phi(i)P_0^d + (1 - \Phi(i))P_f^d, \quad \forall d \tag{1}$$

where P_0^d and P_f^d are the transition probability matrices induced by the decision rule d at time step 0 and ∞, and $\Phi(.)$ characterizes the rate of change of the system with $\Phi(0) = 1$ and $\Phi(\infty) = 0$. The decision rule will be formalized in Sec. 2.

The above transition model can be applied in two interesting scenarios. In the first scenario, P_i^d models the actual dynamics of the underlying nonstationary environment under the decision rule d at step i. In other words, the environment is initially characterized by P_0^d, then its dynamic is characterized by $\Phi(\cdot)$ and converges to P_f^d according to $\Phi(\cdot)$. In the second scenario, the environment is assumed to be stationary and is characterized by P_f^d. However, the estimation of the environmental parameters is initially inaccurate and thought to be P_0^d. Thus, the actions/decisions are made based on inaccurate knowledge of the environment. Over time, the estimations of the environmental parameters become increasingly more accurate, i.e., P_i^d getting closer to P_f^d. Therefore, using an MDP algorithm such as value iteration will eventually produce the optimal solution due to

increasing accuracy. Over time, the decisions are closer to the optimal ones and eventually converge to the optimal policy. However, we characterize the performance of the value iteration algorithm subject to the rate of change in the environment as characterized by $\Phi(\cdot)$. The performance is measured in terms of the convergence rate to the optimal average reward. We present two queuing system examples illustrating the two scenarios above that make use of our analysis framework.

1.3 Related Work. A learning scheme in varying environment is presented by Szita et al. in Ref. [10]. This work introduces the concept of ε–MDP where at each step, the transition matrix is perturbed by small noise bounded by ε around a base MDP. It shows that the value iteration algorithm for ε–MDP converges to a suboptimal value which is δ-close to the optimal value, where $\delta \propto \varepsilon$. In addition, the paper shows that Q-learning can return a suboptimal decision rule/policy in varying environments. On the other hand, in this paper we consider the change in transition matrix (environment) in adiabatic settings, which is characterized by a nonincreasing function Φ. Given the function Φ, we show that the value iteration algorithm converges to the optimal value of the final MDP (base MDP), derive bounds on the convergence rate of the value iteration algorithm, and therefore can predict its stopping time.

The adiabatic evolution was first studied in quantum mechanics by Born and Fock [6]. They provided a first important adiabatic theorem for unitary matrices. Basically, the theorem shows that in the evolution from the initial Hamiltonian to the final one, a system converges to the ground state of the final Hamiltonian after a long time. Recently, another type of adiabatic theorem in Markov chain was presented for linear evolution by Kovchegov [8] and for general evolution by Bradford and Kovchegov [9]. In Ref. [8], the adiabatic time of a time homogeneous Markov chain with linear evolution, which is the chain's mixing time, is shown as an order of the square of the mixing time of the final transition matrix. Generally, Bradford et al. showed that the stable adiabatic time, which is a general concept of adiabatic time, is an order of the maximum mixing time to the power of four [11]. The idea of applying adiabatic evolution to partially observable MDP is mentioned in Ref. [12].

For queuing application, an adiabatic approach to analysis of adaptive queuing policies was proposed in Ref. [13]. It shows that for continuous-time queuing systems with arrival rate estimation, the queue converges to the final distribution after a long time with high probability. It also evaluates the adiabatic time for the queuing system. In our work [14], an MDP in nonstationary environments is modeled as a sequence of time-variant transition matrices governed by adiabatic evolution. The convergence rate of the value iteration algorithm is partly evaluated.

1.4 Our Contribution. In this paper, we study the MDP in nonstationary environment modeled as an adiabatic evolution. Specifically, the convergence of value iteration algorithm is evaluated by an upper bound on the distance between the actual average reward in value iteration and the optimal average reward. As a result, we can derive the necessary time for the value iteration algorithm to get ε-close to the optimal one. Application to an $M/M/1/K$ continuous-time queue with estimated arrival rate and a discrete-time queue with changing arrival rate is shown to verify the results.

Section 2 provides some background on the theory of MDP and value iteration, which are necessary for the development of our results. In Sec. 3, we formulate the problem in terms of the distance from the average reward to the optimal one using value iteration algorithm under changing environment. The theoretical results on the convergence rate of adiabatic value iteration algorithm are also presented in this section. Section 4 formulates adiabatic MDP for two examples of queuing systems based on the

theoretical results. Finally, some conclusions and future work on noisy stochastic matrices are provided in Sec. 5.

2 Mathematical Preliminaries

In this section, we present some definitions, notations, and some propositions for stationary environments.

2.1 MDP

2.1.1 Definitions. A typical discrete-time MDP represents a dynamic system and is specified by a finite set of states S, representing the possible states of the system, a set of control actions A_s for each state $s \in S$, a transition probability matrix $P^{|S| \times |S|}$, and a reward function r. The transition probability specifies the dynamics of the system whose each entry $P_{ij} \triangleq P(s_{t+1} = j | s_t = i, a_t = a)$ represents the conditional probability of the system moving to state $s_{t+1} = j$ in the next time step after taking an action a in the current state $s_t = i$. The dynamics are Markovian in the sense that the probability of the next state j depends only on the current state i and the action a and not on any previous history. The reward function $r(s, a)$ assigns a real number to the state s and the action a, so that $r(s, a)$ represents the immediate reward of being in state s and taking action a. A policy $\pi = \{d_1, d_2, \ldots\}$ is a sequence of decision rules. Each decision rule d_t is a mapping from states to actions at each time step: $d_t : S \to A$ and induces a corresponding transition probability matrix, where $A = \cup_{s \in S} A_s$. The policy π is called stationary if its actions depend only on the state s, independent of time, i.e., $\pi = \{d, d, d, \ldots\}$. A stationary policy induces a time-invariant transition probability matrix. Every policy π is associated with a value function $V^\pi(s)$ such that $V^\pi(s)$ gives the expected cumulative reward achieved by π when starting in state s.

Let Π be the policy space. It can be MD (Markovian and deterministic), HD (history dependent and deterministic), MR (Markovian and randomized), or HR (history dependent and randomized) [5]. Clearly, $\Pi^{\text{MD}} \subset \Pi^{\text{MR}} \subset \Pi^{\text{HR}}$ and $\Pi^{\text{MD}} \subset \Pi^{\text{HD}} \subset \Pi^{\text{HR}}$. Moreover, for each $\pi \in \Pi^{\text{HD}}, s \in S$, there exists a $\pi' \in \Pi^{\text{MD}}$ such that their transition matrices are the same [15]. Therefore, we only consider Markovian deterministic policy.

DEFINITION 1 (unichain MDP). *An MDP is unichain if the transition matrix corresponding to every deterministic stationary policy consists of a single recurrent class plus a possibly empty set of transient states.*

The solution to an MDP problem is an optimal policy π^* that maximizes the expected cumulative reward over some finite or infinite number of time steps. The former and latter are termed finite-horizon MDP and infinite-horizon MDP, respectively. An infinite-horizon model has two typical forms of reward functions: the discounted and the average reward functions. The discounted reward function is defined as

$$V_{\text{dis}}^\pi(s) = E_s^\pi \left\{ \sum_{t=1}^{\infty} \alpha^t r_t(s_t, a_t) \right\} \tag{2}$$

where $0 < \alpha < 1$ denotes a given discount factor that provides convergence of $V^\pi(s)$, but also carries the notion of discounting the future reward, i.e., putting less emphasis on the rewards in the far future than those in the near future. The average reward function is defined as

$$V_{\text{ave}}^\pi(s) = \lim_{N \to \infty} \frac{v_N}{N} \tag{3}$$

where $v_N^\pi(s) = E_s^\pi [\sum_{t=1}^{N} r(s_t, a_t)]$.

For discounted reward function with discount factor $0 < \alpha < 1$, Puterman [5] showed that the convergence of value iteration for discounted reward function is controlled by function α^n, from which we can easily find the convergence rate. Therefore,

in this paper, we only consider the average reward function criteria presented in Appendix A.

2.1.2 Span Seminorm. The span seminorm $sp(v)$ of a vector v is used to evaluate the convergence rate of value iteration for MDP using average reward criterion [5] (see Appendix B for definition and properties of span seminorm).

PROPOSITION 1. *Let $v \in V$ and $d \in D$, where D is the set of decision rule [5], then*

$$sp(P^d v) \leq \delta_d sp(v)$$

where delta coefficient $\delta_d = 1 - \min_{s,u \in S \times S} \sum_{j \in S} \min\{P^d(j|s), P^d(j|u)\}$, and P^d is the transition matrix corresponding to the decision rule d.

Furthermore, $0 \leq \delta_d \leq 1$, and there exists a $v \in V$ such that $sp(P^d v) = \delta_d sp(v)$.

2.1.3 Value Iteration. The value iteration algorithm is an iterative algorithm for finding an ε-optimal policy for the infinite-horizon MDP. More precisely, given an ε, the value iteration algorithm guarantees to produce a reward value within an ε of the optimal value. The key to the value iteration algorithm is that each step of the algorithm can be viewed as applying a contracting operator L on v. Running the algorithm iteratively, or equivalently, applying the operator L repeatedly, will guarantee that v will converge to the optimal value based on Bellman equation. Specifically, for a unichain MDP, at each iteration n, we have: $v_{n+1} = Lv_n$, where L is defined as $Lv = \max_{d \in D}\{r^d + P^d v\}$, and r^d and P^d denote the reward and the transition probability matrix induced by the decision rule d. The pseudocode for the value iteration algorithm with average reward objective is shown below.

DEFINITION 2 (the value iteration). *The algorithm for the value iteration with average reward criteria is shown below [5]:*

(1) *Choose any initial reward vector v_0, for a given $\varepsilon > 0$. Let $n = 0$.*
(2) *For each $s \in S$, we have: $v_{n+1}(s) = \max_{a \in A}\{r(s,a) + \sum_{j \in S} p(j|s,a)v_n(j)\}$.*
(3) *Increasing n until $sp(v_{n+1} - v_n) < \varepsilon$, then choose: $d_\varepsilon \in \text{argmax}\{r(s,a) + \sum_{j \in S} p(j|s,a)v_n(j)\}$*

where $sp(v)$ is the span seminorm of vector v.

We note that the ε-optimal policy approaches to an optimal policy as ε reduces to zero when the number of iterations becomes infinity.

DEFINITION 3 (gamma coefficient). *The gamma coefficient is defined as follows [5]:*

$$\gamma = \max_{s \in S, a \in A_s, s' \in S, a' \in A_{s'}} \left[1 - \sum_{j \in S} \min\{p(j|s,a), p(j|s',a')\} \right]$$

Easily, we can see that the gamma coefficient is an upper bound of the delta coefficient δ_d for all decision rule d. Therefore, from Proposition 1, we have the following proposition:

PROPOSITION 2. *Let $v \in V$ and $d \in D$, where D is the set of decision rule, then*

$$sp(P^d v) \leq \gamma sp(v)$$

PROPOSITION 3 (convergence of value iteration). *For unichain MDPs [5], we have*

$$sp(v^{n+2} - v^{n+1}) \leq \gamma sp(v^{n+1} - v^n)$$

where $\gamma = \max_{s \in S, a \in A_s, s' \in S, a' \in A_{s'}}[1 - \sum_{j \in S}\min\{p(j|s,a), p(j|s',a')\}]$. Then, if $\gamma < 1$, the value iteration algorithm will stop after a finite step n.

3 Adiabatic MDP

This section formalizes the adiabatic evolution and provides the necessary background for performance study of value iteration algorithm in nonstationary setting.

3.1 Preliminaries on Adiabatic-Time Evolution. We introduce adiabatic-time framework to model nonstationary environments. The initial adiabatic setting was first described in quantum mechanics by Born and Fock [6]. Recently, the adiabatic setting for both discrete-time and continuous-time Markov chain was proposed by Kovchegov [8], which is shown below:

$$P_i = \Phi(i)P_0 + (1 - \Phi(i))P_f$$
$$= P_f + \Phi(i)(P_0 - P_f) \tag{4}$$

where P_0 and P_f are the initial transition matrix and the final transition matrix, respectively. P_i is the transition matrix at time i. $\Phi(\cdot)$ is a function to characterize the adiabatic evolution such that $\Phi(i) : [0, +\infty) \to [0,1]$, $\Phi(0) = 1$, and $\lim_{i \to \infty}\Phi(i) = 0$.

Because of the properties of $\Phi(\cdot)$, we have: $\lim_{i \to \infty}P_i = P_f$. For each application, we have different $\Phi(\cdot)$ function. In this paper, we suppose that Φ is given.

The following propositions and corollaries are needed for the development of convergence of value iteration with adiabatic setting.

PROPOSITION 4 (the bound on gamma coefficients). *Given $P_i = \Phi(i)P_0 + (1 - \Phi(i))P_f, \Phi(i) \in [0,1]$*

$$\gamma_i \leq \Phi(i)\gamma_0 + (1 - \Phi(i))\gamma_f$$

Proof. See Appendix C for details.

COROLLARY 1. *Given $P_i = \Phi(i)P_0 + (1 - \Phi(i))P_f$ with nonincreasing function $\Phi(i) > 0, \Phi(i) \in [0,1]$ for any $n_0 \leq i$*

$$\gamma_i \leq \max(\gamma_{n_0}, \gamma_f) \tag{5}$$

Proof. See Appendix C for details.

Note that one way to ensure that $0 < \gamma < 1$ is that for each decision rule d, there exists a column of the corresponding P having all positive entries.

3.2 Problem Formulation

3.2.1 Adiabatic MDP Model. Typically, the value iteration algorithm is used to find an ε-optimal stationary policy in an offline manner using a number of iterations, assuming a stationary environment such that every stationary policy π induces a time-invariant transition probability matrix. The resulting policy is then used in an online manner with the assumption that the environment is stationary and governed by the time-invariant transition probability matrices in the value iteration algorithm. In this paper, we study an adiabatic MDP setting, in which we assume that the "environment" is no longer stationary. Instead, it might change at every iteration of the value iteration algorithm, resulting in a sequence of time-variant transition probability matrices under a stationary policy. The precise meaning of the environment will be clear shortly.

Instead of running the value iteration algorithm offline to find an ε-optimal policy, we apply the decision rule found after each iteration immediately and repeatedly in an online manner. Our goal is to determine how good the reward is for such a scheme. The analysis of such a setting is useful in the rapidly changing environments where decisions must be made quickly. Unlike the classic MDP setting where for each decision rule d, there is a time-invariant transition probability matrix P^d, in our setting, for a fixed decision rule d, there is sequence of time-variant transition probability matrices

$$P_i^d = \Phi(i)P_0^d + (1 - \Phi(i))P_f^d$$
$$= P_f^d + \Phi(i)(P_0^d - P_f^d), \quad \forall d \tag{6}$$

where Φ is a function such that

$$\Phi(i) : [0, +\infty) \rightarrow [0, 1], \Phi(0) = 1, \lim_{i \rightarrow \infty} \Phi(i) = 0 \qquad (7)$$

$\Phi(\cdot)$ characterizes the change in the environment at the iteration i of the value iteration algorithm, and P_i^d is the induced transition probability matrix due to the decision rule found at the iteration i of the value iteration algorithm. A slowly changing $\Phi(i)$ implies a slow change in the environment. We note that the notion of optimal reward is not well defined if the environment fluctuates arbitrarily. Thus, in the model above, we assume that the environment will approach to a final stationary environment characterized by P_f^d to ensure a well-defined reward. This can be seen as $\lim_{i \rightarrow \infty} P_i^d = P_f^d$. In this paper, we also assume that the function $\Phi(\cdot)$ is the same for all decision rules d as seen in the example below.

A simple example. Consider a discrete-time queue with size $K = 2$. Let p and q be the probabilities that a packet arrives and departs the queue, respectively. The state-space is $S = \{0, 1, 2\}$. The action is the value of q. For this queue, the transition matrix is

$$P = \begin{bmatrix} 1 - p(1-q) & p(1-q) & 0 \\ q(1-p) & pq + (1-p)(1-q) & p(1-q) \\ 0 & q(1-p) & 1 - q(1-p) \end{bmatrix}$$

If we let p change over time following the rule: $p_i = \phi(i) p_0 + (1 - \phi(i)) p_f$, where p_0 and p_f are the initial and final arrival probabilities, respectively. Since each entry in the transition matrix is a linear function of p, for any decision rule d which is a mapping from state-space S to a set of values of q, the transition matrix is changing following the rule $P_i^d = \Phi(i) P_0^d + (1 - \Phi(i)) P_f^d$, where $\Phi(\cdot) = \phi(\cdot)$. Note that the function $\Phi(\cdot)$ is the same for all decision rules d.

We now articulate a bit more on the meaning of the environment. Note that the induced P_i^d depends on both actions and environments. Therefore, a change in the environment implies possible changes in the underlying environments, or the set of actions, or the combination of both over time. For example, let us consider a queuing system in which the controller attempts to send packets (actions) at some varying rates based on the number of packets in the queue in order to maximize a given reward. In one scenario, we assume the traffic arrival rate at the queue increases steadily from a initial rate of λ_0 to a final rate of λ_f. As a result, P_i^d varies as the underlying environment changes over time. In another scenario, the arrival rate of the packets remains the same, however, the controller has inaccurate estimation of the arrival rate initially due to few observations. Consequently, it makes the decision on what rate it should send based on an inaccurate arrival rates, and P_i^d characterizes the change based on its decision rule d at iteration i. However, over time with more observations, its estimation of the arrival rate becomes more accurate. Therefore, its decision rule approaches the optimal one for which the state dynamic is characterized by P_f^d. We will discuss these two examples in more detail in Sec. 4.

3.2.2 Convergence Rate of Adiabatic MDP. It is important to emphasize again that the environment will be asymptotically stationary corresponding to an induced transition probability matrix P_f^d for any decision rule d. In addition, there exists an optimal decision rule d^* corresponding to a transition probability matrix $P_f^{d^*}$, which can be obtained when running the value iteration algorithm under a stationary environment [5]. Importantly, running the value iteration algorithm in an adiabatic setting for a sufficiently large number of steps would produce the same optimal decision rule d_f^* as that of the classic value iteration algorithm and also the same average reward g^* (see Appendix A)

$$g^* = P_{d^*}^\infty r_f^{d^*} = \pi_{P_f}^{d^*} r_f^{d^*} e$$

where by Cesaro mean, $P_{d^*}^\infty = \lim_{N \rightarrow \infty} 1/N \sum_{n=1}^{N} (P_f^{d^*})^{n-1} =$

$$\lim_{n \rightarrow \infty} (P_f^{d^*})^n = \begin{bmatrix} \pi_{P_f}^{d^*} \\ \pi_{P_f}^{d^*} \\ \pi_{P_f}^{d^*} \\ \cdots \\ \pi_{P_f}^{d^*} \end{bmatrix}, \pi_{P_f}^{d^*} \text{ is the stationary distribution for } P_f^{d^*},$$

and $e = [11 \ldots 1]^{\mathrm{T}}$.

However, the convergence rates to this final reward g^* are quite different for the classic MDP and adiabatic MDP settings. One would expect that the rate in the former setting would be faster since the environment does not change, thus the value iteration algorithm can learn faster than that of the latter. Therefore, our primary focus of this paper is to characterize the convergence rate of the value iteration algorithm in an adiabatic setting given the dynamics specified by $\Phi(\cdot)$. Specifically, we want to find an integer N such that $\forall n > N$

$$E = \left\| \frac{v_n}{n} - g^* \right\|_\infty \leq \varepsilon \qquad (8)$$

Note that the index n is both the step index in value iteration algorithm and the time step. In other words, we run the value iteration algorithm while the environment is changing. The reason for this online manner is that it might take a long time to wait until the environment is stable then run the classic value iteration algorithm or the environment keeps changing.

The difference between the classic value iteration and the value iteration in adiabatic setting is illustrated in Figs. 1 and 2. We can see that, compared to the classic value iteration algorithm, the set of transition matrix P, where we are looking for the optimal one, is changing at every time step, i.e., also the step in value iteration algorithm. In the value iteration in adiabatic setting, we apply the same operator on different set of matrices at each step. Denote L_i as the operator L applied on the set of matrices at step i.

Finding N is not trivial. Therefore, we will provide a lower bound on N which depends on $\Phi(\cdot)$ as well as the set of all possible matrices P_f^d.

3.3 Main Results

THEOREM 1 (main result 1). *Consider a unichain adiabatic-time MDP with S and A_s both finite, $|r(s, a)|$ bounded by a number M. Suppose $0 < \gamma = \max(\gamma_f, \gamma_{n_0}) < 1$, then*

$$E \leq \frac{\|v_{n_0} - n_0 g^*\|_\infty}{n} + \frac{1}{n} \sum_{i=n_0}^{n-1} \left[\left(\prod_{k=n_0}^{i-1} \gamma_k \right) sp(v_{n_0+1} - v_{n_0}) \right.$$
$$\left. + \gamma M \left(e_i + 2 \sum_{j=n_0+1}^{i} sp(v_j) |\Delta \Phi_{j-1}| \left(\prod_{k=j}^{i-1} \gamma_k \right) \right) \right]$$
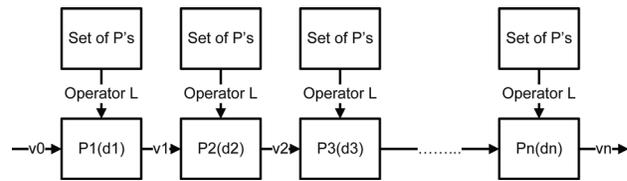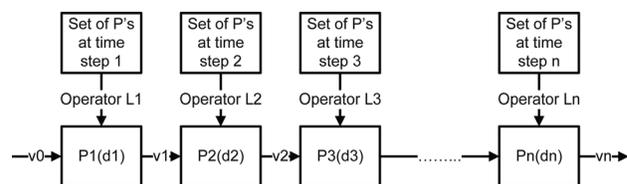

Fig. 1   The classic value iteration

Fig. 2   The value iteration in an adiabatic setting

for any n_0, where $\prod_{k=u}^{v}(\cdot) = 1$ if $v < u$, $e_i = \sum_{j=i+1}^{\infty}(sp(v_j)|\Delta\Phi_{j-1}|)$, $2Y = \max_{d\in D}|||P_0 - P_f||_{\infty}$.

Proof. From Eq. (6), for any $d \in D$ we have

$$|P_i^d - P_{i-1}^d| \leq |\Phi(i) - \Phi(i-1)||P_0^d - P_f^d|$$
$$\leq |\Phi(i) - \Phi(i-1)||P_0^d - P_f^d| \quad (9)$$
$$\leq 2Y|\Delta\Phi_{i-1}|, \text{ where fixed } 2Y = \max_{d\in D}|P_0 - P_f|$$

Consider $E = ||(v_n/n) - g^*||_{\infty}$

$$E = \left\|\frac{v_n}{n} - g^*\right\|_{\infty}$$
$$= \left\|\frac{v_n - ng^*}{n}\right\|_{\infty}$$
$$= \left\|\sum_{i=n_0}^{n-1}\left(\frac{v_{i+1} - v_i - g^*}{n}\right) + \frac{v_{n_0} - n_0 g^*}{n}\right\|_{\infty}$$
$$\leq \sum_{i=n_0}^{n-1}\frac{||v_{i+1} - v_i - g^*||_{\infty}}{n} + \frac{||v_{n_0} - n_0 g^*||_{\infty}}{n}$$

First, we bound $||v_{i+1} - v_i - g^*||_{\infty}$. Let $x_i = \text{argmax}_{s\in S}(v_{i+1} - v_i) = \text{argmax}_{s\in S}(L_i v_i - L_{i-1}v_{i-1})$ and $y_i = \text{argmin}_{s\in S}(v_{i+1} - v_i) = \text{argmin}_{s\in S}(L_i v_i - L_{i-1}v_{i-1})$. Let d_i be the optimal decision rule corresponding to the operator L_i. Let d_{i-1} be the optimal decision rule corresponding to the operator L_{i-1}. Let L^d be the operator that we apply the decision rule d: $L^d v = r^d + P^d v$. Since $L_{i-1}v_{i-1}(x_i) \geq L_{i-1}^{d_i}v_{i-1}(x_i)$

$$L_i v_i(x_i) - L_{i-1}v_{i-1}(x_i) \leq L_i^{d_i}v_i(x_i) - L_{i-1}^{d_i}v_{i-1}(x_i)$$
$$\leq (r_{d_i}(x_i) + P_i^{d_i}v_i(x_i)) - (r_{d_i}(x_i) + P_{i-1}^{d_i}v_{i-1}(x_i))$$
$$\leq P_i^{d_i}v_i(x_i) - P_{i-1}^{d_i}v_{i-1}(x_i)$$
$$\leq (P_i^{d_i} - P_{i-1}^{d_i})v_i(x_i) + P_{i-1}^{d_i}(v_i - v_{i-1})(x_i)$$

Let $\alpha_i = \text{argmax}_{s\in S}(v_i(s))$, $\beta_i = \text{argmin}_{s\in S}(v_i(s))$, $\Delta P_i^{d_i} = P_i^{d_i} - P_{i-1}^{d_i}$. We have: $a = \Delta P_i^{d_i}v_i(x_i) = (\Delta P_i^{d_i}(x_i, \cdot))v_i = (\Delta P_i^{d_i}(x_i, \cdot))(v_i - v_i(\beta_i)e)$ where $e = [11...1]^T$ since $P_i^{d_i}$ and $P_{i-1}^{d_i}$ are the transition matrices, then $\sum_{s\in S}(\Delta P_i^{d_i}(x_i, s)) = 0$. Then,

$$a = \sum_{s\in S}\Delta P_i^{d_i}(x_i, s)\left(v_i(s) - v_i(\beta_i)\right)$$
$$\leq \sum_{\Delta P_i^{d_i}(x_i, s))\geq 0}\Delta P_i^{d_i}(x_i, s)\left(v_i(s) - v_i(\beta_i)\right)$$
$$\leq \left(v_i(\alpha_i) - v_i(\beta_i)\right)\sum_{\Delta P_i^{d_i}(x_i, s))\geq 0}\Delta P_i^{d_i}(x_i, s)$$
$$\leq \left(v_i(\alpha_i) - v_i(\beta_i)\right)\frac{||P_i^{d_i} - P_{i-1}^{d_i}||_{\infty}}{2}$$
$$\leq Y|\Delta\Phi_{i-1}|sp(v_i) \text{ since Eq. (9).}$$

Then, we have

$$L_i v_i(x_i) - L_{i-1}v_{i-1}(x_i) \leq YMsp(v_i)|\Delta\Phi_{i-1}| + P_{i-1}^{d_i}(v_i - v_{i-1})(x_i) \quad (10)$$

Similarly,

$$L_i v_i(y_i) - L_{i-1}v_{i-1}(y_i) \geq -YMsp(v_i)|\Delta\Phi_{i-1}| + P_{i-1}^{d_{i-1}}(v_i - v_{i-1})(y_i) \quad (11)$$

Consider $b = P_{i-1}^{d_i}(v_i(x_i) - v_{i-1}(x_i)) = P_{i-1}^{d_i}(v_i - v_{i-1})(x_i) \leq v_i(x_{i-1}) - v_{i-1}(x_{i-1})$ since $x_{i-1} = \text{argmax}_{s\in S}(v_i - v_{i-1})$. Therefore, $L_i v_i(x_i) - L_{i-1}v_{i-1}(x_i) \leq YMsp(v_i)|\Delta\Phi_{i-1}| + v_i(x_{i-1}) - v_{i-1}(x_{i-1})$. Similarly, $L_i v_i(y_i) - L_{i-1}v_{i-1}(y_i) \geq -YMsp(v_i)|\Delta\Phi_{i-1}| + v_i(y_{i-1}) - v_{i-1}(y_{i-1})$.

Now, by keep on expanding, we have: $v_{t+1}(x_t) - v_t(x_t) = L_t v_t(x_t) - L_{t-1}v_{t-1}(x_t) \leq YM\sum_{j=i+1}^{t}(sp(v_j)|\Delta\Phi_{j-1}|) + (v_{i+1}(x_i) - v_i(x_i))$. Let $t \to \infty$, when $t = \infty$, we know exactly P_f and run value iteration algorithm for it, we will have a reward received at one time step $\lim_{t\to\infty}(v_{t+1}(x_t) - v_t(x_t)) = g^*$ which is the optimal average reward [5].

Therefore, $g^* \leq YM\sum_{j=i+1}^{\infty}(sp(v_j)|\Delta\Phi_{j-1}|) + (v_{i+1}(x_i) - v_i(x_i))$ for any i.

Similarly, we have: $g^* \geq -YM\sum_{j=i+1}^{\infty}(sp(v_j)|\Delta\Phi_{j-1}|) + (v_{i+1}(y_i) - v_i(y_i))$ for any i.

Denote $e_i = \sum_{j=i+1}^{\infty}(sp(v_j)|\Delta\Phi_{j-1}|)$, which represents the total error from the time step i to ∞. This error comes from the fact that we use the inaccurate matrix at each time step instead of the true one. Now, for any $s \in S$

$$-YMe_i - (v_{i+1}(x_i) - v_i(x_i)) + (v_{i+1}(y_i) - v_i(y_i)) \leq (v_{i+1}(s) - v_i(s)) - g^* \leq YMe_i + (v_{i+1}(x_i) - v_i(x_i)) - (v_{i+1}(y_i) - v_i(y_i))$$

Hence, if we use ∞-norm, we have

$$||v_{i+1} - v_i - g^*||_{\infty} \leq YMe_i + (v_{i+1}(x_i) - v_i(x_i)) - (v_{i+1}(y_i) - v_i(y_i))$$
$$\leq sp(v_{i+1} - v_i) + YMe_i$$

Now, we bound $sp(v_{i+1} - v_i)$. Since Eqs. (10) and (11), we have

$$sp(v_{i+1} - v_i) = (L_i v_i(x_i) - L_{i-1}v_{i-1}(x_i)) - (L_i v_i(y_i) - L_{i-1}v_{i-1}(y_i))$$
$$\leq 2YMsp(v_i)|\Delta\Phi_{i-1}| + P_{i-1}^{d_i}(v_i - v_{i-1})(x_i)$$
$$- P_{i-1}^{d_{i-1}}(v_i - v_{i-1})(y_i)$$
$$\leq 2YMsp(v_i)|\Delta\Phi_{i-1}| + \max_{s\in S}P_{i-1}^{d_i}(v_i - v_{i-1})(s)$$
$$- \min_{s\in S}P_{i-1}^{d_{i-1}}(v_i - v_{i-1})(s)$$
$$\leq 2YMsp(v_i)|\Delta\Phi_{i-1}| + sp([P_{i-1}^{d_i}/P_{i-1}^{d_{i-1}}](v_i - v_{i-1}))$$
$$\leq 2YMsp(v_i)|\Delta\Phi_{i-1}| + \gamma_{i-1}sp(v_i - v_{i-1})$$
$$\text{(Proposition 2)} \quad (12)$$

where $[P1/P2]$ denotes the stacked matrix in which the rows of $P1$ follow the rows of $P2$. Based on Definition 3, the gamma coefficient of the set of stacked matrices at time step $i-1$ is at most γ_{i-1}. Equation (12) is similar to Proposition 3 except there is an error $2YMsp(v_i)|\Delta\Phi_{i-1}|$.

Since $0 < \gamma_i \leq \gamma = \max(\gamma_{n_0}, \gamma_f) < 1, \forall i \geq n_0$ (Corollary 1), we have

$$sp(v_{i+1} - v_i) \leq \left(\prod_{k=n_0}^{i-1}\gamma_k\right)sp(v_{n_0+1} - v_{n_0})$$
$$+ 2YM\left(\sum_{j=n_0+1}^{i}sp(v_j)\right)|\Delta\Phi_{j-1}|\left(\prod_{k=j}^{i-1}\gamma_k\right) \text{ for all } i \geq n_0 \quad (13)$$

Then,

$$\sum_{i=n_0}^{n-1}\frac{||v_{i+1} - v_i - g^*||_{\infty}}{n} \leq \frac{1}{n}\sum_{i=n_0}^{n-1}\left[sp(v_{i+1} - v_i) + YMe_i\right]$$
$$\sum_{i=n_0}^{n-1}\frac{||v_{i+1} - v_i - g^*||_{\infty}}{n} \leq \frac{1}{n}\sum_{i=n_0}^{n-1}\left[\left(\prod_{k=n_0}^{i-1}\gamma_k\right)sp(v_{n_0+1} - v_{n_0})\right.$$
$$+ YM\left(e_i + 2\left(\sum_{j=n_0+1}^{i}sp(v_j)|\Delta\Phi_{j-1}|\right.\right.$$
$$\left.\left.\left.\times\left(\prod_{k=j}^{i-1}\gamma_k\right)\right)\right)\right]$$

Therefore, we have a upper bound of A which is as follows:

$$E \leq \frac{||v_{n_0} - n_0 g^*||_\infty}{n} + \frac{1}{n} \sum_{i=n_0}^{n-1} \left[\left(\prod_{k=n_0}^{i-1} \gamma_k \right) sp(v_{n_0+1} - v_{n_0}) \right.$$
$$\left. + YM \left(e_i + 2 \sum_{j=n_0+1}^{i} sp(v_j)|\Delta\Phi_{j-1}| \left(\prod_{k=j}^{i-1} \gamma_k \right) \right) \right] \quad (14)$$

Comments on Theorem 1. From Theorem 1, if we want to get an integer N so that for $n \geq N$, $E < \varepsilon$, we have to find out an n_0 so that we can bound the last term in the right-hand side of Eq. (14). If we can run value iteration algorithm until time steps n_0, we have some more information about the value vector v. Therefore, we can get a good bound as shown in Theorem 2. Whereas, if we are not able to run value iteration algorithm for some steps, we have to find a general bound on all the terms in the right-hand side of Eq. (14) as described in Theorem 3. As a consequence, we obtain a looser bound on the left-hand side. We will see this in Sec. 4.

Note that, in Theorem 1, we consider general function Φ while in Theorems 2 and 3, we consider positive nonincreasing function Φ.

THEOREM 2 (main result 2). *Consider a unichain adiabatic-time MDP with S and A_s both finite, $|r(s,a)|$ bounded by a number M. Suppose $0 < \gamma = \max(\gamma_f, \gamma_{n_0}) < 1$ and $\Phi(i)$ is a positive nonincreasing function on $[n_0, +\infty)$, then for*

$$n \geq \max \left(n_0, \frac{2}{\varepsilon} \left(||v_{n_0}||_\infty + \frac{sp(v_{n_0+1} - v_{n_0})}{1 - \gamma} + 2YM \frac{e_{n_0}}{1 - \gamma} \right) \right) \quad (15)$$

we guarantee: $E = ||(v_n)/n) - g^*||_\infty < \varepsilon$ where $e_i = \sum_{j=i+1}^{\infty} (sp(v_j) |\Delta\Phi_{j-1}|), 2Y = \max_{d \in D} ||P_0 - P_f||_\infty$, n_0 is the smallest integer satisfying $(2YM\Phi(n_0)/1 - \gamma) < 1$ and $(sp(v_{n_0}) + sp(v_{n_0+1} - v_{n_0})/1 - \gamma)\Phi(n_0)/(1 - 2YM\Phi(n_0)/1 - \gamma) \leq \varepsilon/2YM$.

Proof. By applying Theorem 1 with n_0, we have

$$E \leq \frac{||v_{n_0} - n_0 g^*||_\infty}{n} + \frac{1}{n} \sum_{i=n_0}^{n-1} \left[\left(\prod_{k=n_0}^{i-1} \gamma_k \right) sp(v_{n_0+1} - v_{n_0}) \right.$$
$$\left. + YM \left(e_i + 2 \sum_{j=n_0+1}^{i} sp(v_j)|\Delta\Phi_{j-1}| \left(\prod_{k=j}^{i-1} \gamma_k \right) \right) \right) \right]$$
$$\leq \frac{||v_{n_0} - n_0 g^*||_\infty}{n} + \frac{sp(v_{n_0+1} - v_{n_0})}{n(1 - \gamma)}$$
$$+ \frac{1}{n} \sum_{i=n_0}^{n-1} YM \left(e_i + 2 \sum_{j=n_0+1}^{i} \left(\gamma^{i-j} sp(v_j)|\Delta\Phi_{j-1}| \right) \right) \right)$$

since $0 < \gamma_i \leq \gamma = \max(\gamma_{n_0}, \gamma_f) < 1, \forall i \geq n_0$ (Corollary 1).

Given n_0, $||v_{n_0}||_\infty$ and $sp(v_{n_0+1} - v_{n_0})$ are fixed. Let $y_i = \sum_{j=n_0+1}^{i} (\gamma^{i-j} sp(v_j)|\Delta\Phi_{j-1}|))$. We have

$$y_{n_0} = 0$$
$$y_{n_0+1} = sp(v_{n_0+1})\Delta\Phi_{n_0}$$
$$y_{n_0+2} = \gamma y_{n_0+1} + sp(v_{n_0+2})\Delta\Phi_{n_0+1}$$
$$y_{n_0+3} = \gamma y_{n_0+2} + sp(v_{n_0+3})\Delta\Phi_{n_0+2}$$
$$\dots$$
$$y_n = \gamma y_{n-1} + sp(v_n)\Delta\Phi_{i-1}$$
$$\dots$$

Then, $(1 - \gamma) \sum_{i=n_0}^{\infty} y_i = \sum_{j=n_0+1}^{\infty} (sp(v_j)|\Delta\Phi_{j-1}|)$ or $\sum_{i=n_0}^{\infty} y_i = e_{n_0}/1 - \gamma$.

Now, we find conditions on n_0 so that $e_{n_0} = \sum_{j=n_0+1}^{\infty} (sp(v_j)|\Delta\Phi_{j-1}|) \leq \varepsilon/2YM$. In order to do that we

need to bound $sp(v_j)$ because we do not know $sp(v_j)$ ahead of time when $j > n_0$. By the triangle property of seminorm, we get

$$sp(v_{j+1}) \leq sp(v_j) + sp(v_{j+1} - v_j)$$
$$\leq sp(v_{j-1}) + sp(v_j - v_{j-1}) + sp(v_{j+1} - v_j)$$
$$\leq sp(v_{n_0}) + \sum_{i=n_0}^{j} sp(v_{i+1} - v_i) \quad (16)$$

From Eq. (13), we have

$$sp(v_{i+1} - v_i) \leq 2YM \left(\sum_{j=n_0+1}^{i} sp(v_j)|\Delta\Phi_{j-1}|\gamma^{i-j} \right)$$
$$+ \gamma^{i-n_0} sp(v_{n_0+1} - v_{n_0})$$
$$\leq 2YMy_i + \gamma^{i-n_0} sp(v_{n_0+1} - v_{n_0})$$

Then, for all $j \geq n_0$

$$sp(v_{j+1}) \leq sp(v_{n_0}) + 2YM \sum_{i=n_0}^{j} y_i + sp(v_{n_0+1} - v_{n_0}) \sum_{i=n_0}^{j} \gamma^{i-n_0}$$
$$\leq sp(v_{n_0}) + 2YM \sum_{i=n_0}^{\infty} y_i + sp(v_{n_0+1} - v_{n_0}) \sum_{i=n_0}^{\infty} \gamma^{i-n_0}$$
$$\leq sp(v_{n_0}) + 2YM \frac{e_{n_0}}{1 - \gamma} + \frac{sp(v_{n_0+1} - v_{n_0})}{1 - \gamma}$$

By plugging back into e_{n_0}, we get

$$e_{n_0} = \sum_{j=n_0+1}^{\infty} (sp(v_j)|\Delta\Phi_{j-1}|)$$
$$\leq \left(sp(v_{n_0}) + 2YM \frac{e_{n_0}}{1 - \gamma} + \frac{sp(v_{n_0+1} - v_{n_0})}{1 - \gamma} \right) \times \sum_{j=n_0+1}^{\infty} |\Delta\Phi_{j-1}|$$
$$\leq \left(sp(v_{n_0}) + 2YM \frac{e_{n_0}}{1 - \gamma} + \frac{sp(v_{n_0+1} - v_{n_0})}{1 - \gamma} \right) \Phi(n_0)$$

since $\Phi(i)$ is nonincreasing, $i \geq n_0$, then $|\Delta\Phi_{j-1}| = \Phi(j-1) - \Phi(j)$. In other words,

$$e_{n_0} \left(1 - \frac{2YM\Phi(n_0)}{1 - \gamma} \right) \leq \left(sp(v_{n_0}) + \frac{sp(v_{n_0+1} - v_{n_0})}{1 - \gamma} \right) \Phi(n_0)$$

Since $\lim_{i \to \infty} \Phi(i) = 0$, then there exists n_0 so that

$$\frac{2YM\Phi(n_0)}{1 - \gamma} < 1 \quad (17)$$

Then, we only need to run the value iteration until the following condition is satisfied:

$$e_{n_0} \leq \frac{\left(sp(v_{n_0}) + \frac{sp(v_{n_0+1} - v_{n_0})}{1 - \gamma} \right) \Phi(n_0)}{\left(1 - \frac{2YM\Phi(n_0)}{1 - \gamma} \right)} \leq \frac{\varepsilon}{2YM} \quad (18)$$

Now, for $n \geq n_0$

$$E \leq \frac{||v_{n_0} - n_0 g^*||_\infty}{n} + \frac{1}{n} \frac{sp(v_{n_0+1} - v_{n_0})}{1 - \gamma} + \frac{YM \sum_{i=n_0}^{n-1} e_i}{n}$$
$$+ \frac{1}{n} 2YM \frac{e_{n_0}}{1 - \gamma}$$
$$\leq \frac{1}{n} \left(||v_{n_0} - n_0 g^*||_\infty + \frac{sp(v_{n_0+1} - v_{n_0})}{1 - \gamma} + 2YM \frac{e_{n_0}}{1 - \gamma} \right)$$
$$+ \frac{YM(n - n_0) \frac{\varepsilon}{2YM}}{n}$$
$$\leq \frac{1}{n} \left(||v_{n_0} - n_0 g^*||_\infty + \frac{sp(v_{n_0+1} - v_{n_0})}{1 - \gamma} + 2YM \frac{e_{n_0}}{1 - \gamma} \right) + \frac{\varepsilon}{2}$$

Let $1/n(||v_{n_0} - n_0 g^*||_\infty + (sp(v_{n_0+1} - v_{n_0})/1 - \gamma) + 2YM(e_{n_0}/1 - \gamma)) \le \varepsilon/2$, or $n \ge 2/\varepsilon(||v_{n_0} - n_0 g^*||_\infty + (sp(v_{n_0+1} - v_{n_0})/1 - \gamma) + 2YM(e_{n_0}/1 - \gamma))$. Then, for all $n \ge \max(n_0, 2/\varepsilon(||v_{n_0} - n_0 g^*||_\infty + (sp(v_{n_0+1} - v_{n_0})/1 - \gamma) + 2YM(e_{n_0}/1 - \gamma))$, we guarantee $E \le \varepsilon$.

Comments on Theorem 2. From Theorem 2, if we run value iteration algorithm until step n_0 satisfying the conditions in the theorem, then we can predict that number of steps necessary for value iteration algorithm to ensure that $E < \varepsilon$. If running value iteration algorithm to find n_0 is not satisfactory then in that case, we provide an alternative on the upper bound without n_0 as shown in Theorem 3. However, this bound is looser then the bound in Theorem 2.

THEOREM 3 (main result 3). *Consider a unichain adiabatic-time MDP with S and A_s both finite, $|r(s,a)|$ bounded by a number M. Suppose $0 < \gamma = \max(\gamma_f, \gamma_{n_0})$, $\gamma' = \max(\gamma_f, \gamma_0) < 1$, and $\Phi(i)$ is a positive nonincreasing function on $[n_0, +\infty)$, then for*

$$n \ge \frac{2}{\varepsilon}\left(2n_0 M + ||v_0||_\infty + \frac{\varepsilon}{1-\gamma}\right.$$
$$\left. + \frac{M + (1+\gamma)\left[\frac{M(1-(\gamma')^{n_0})}{1-\gamma'} + (\gamma')^{n_0}(sp(v_0))\right]}{1-\gamma}\right) \quad (19)$$

we guarantee: $E = ||(v_n/n) - g^*||_\infty < \varepsilon$ *where* $2Y = \max_{d \in D} ||P_0 - P_f||_\infty$, n_0 *is the smallest integer satisfying*

$$\left[\frac{M}{1-\gamma'} + \gamma'^{(n_0)} sp(v_0)\right]\Phi(n_0) \le \frac{\varepsilon}{2YM} \quad (20)$$

Proof. By applying Theorem 1 with n_0, we have

$$E \le \frac{||v_{n_0} - n_0 g^*||_\infty}{n} + \frac{1}{n}\sum_{i=n_0}^{n-1}\left[\left(\prod_{k=n_0}^{i-1}\gamma_k\right)sp(v_{n_0+1} - v_{n_0})\right.$$
$$\left. + YM\left(e_i + 2\sum_{j=n_0+1}^{i} sp(v_j)|\Delta\Phi_{j-1}|\left(\prod_{k=j}^{i-1}\gamma_k\right)\right)\right]$$
$$\le \frac{||v_{n_0} - n_0 g^*||_\infty}{n} + \frac{1}{n}\frac{sp(v_{n_0+1} - v_{n_0})}{1-\gamma}$$
$$+ \frac{1}{n}\sum_{i=n_0}^{n-1}YM\left(e_i + 2\sum_{j=n_0+1}^{i}\left(\gamma^{i-j}sp(v_j)|\Delta\Phi_{j-1}|\right)\right)$$

We have

$$||v_{n_0} - n_0 g^*||_\infty \le ||r_{d_{n_0-1}} + P_i^{d_{n_0-1}} v_{n_0-1} - n_0 g^*||_\infty \le ||(r_{d_{n_0-1}} - g^*)||_\infty$$
$$+ ||P_i^{d_{n_0-1}}(v_{n_0-1} - (n_0-1)g^*)||_\infty$$
$$\le ||(r_{d_{n_0-1}} - g^*)||_\infty + ||(v_{n_0-1} - (n_0-1)g^*)||_\infty$$
$$\le \sum_{k=1}^{n_0}||r_{d_{k-1}} - g^*||_\infty + ||v_0||_\infty \le 2n_0 M + ||v_0||_\infty$$

since $0 < |r_{d_{k-1}}|, g^* < M$
and

$$sp(v_{n_0+1} - v_{n_0}) = sp(r_{d_{n_0}} + P_{d_{n_0}} v_{n_0} - v_{n_0})$$
$$\le sp(r_{d_{n_0}}) + sp(P_{d_{n_0}} v_{n_0}) + sp(v_{n_0})$$
$$\le M + (1+\gamma)sp(v_{n_0})$$

Consider

$$sp(v_i) = sp\left(r^{d_i} + P_{i-1}^{d_{i-1}} v_{i-1}\right)$$
$$\le \left[sp(r^{d_i}) + sp\left(P_{i-1}^{d_{i-1}} v_{i-1}\right)\right]$$
$$\le [M + \gamma_{i-1} sp(v_{i-1})] \quad \text{(Proposition 2)}$$
$$\le M\left[1 + \sum_{j=1}^{i-1}\prod_{k=1}^{j}\gamma_{i-k}\right] + \prod_{k=1}^{i}\gamma_{i-k}(sp(v_0))$$
$$\le \frac{M(1 - (\gamma')^i)}{1-\gamma'} + (\gamma')^i(sp(v_0)), \text{ where } \gamma' = \max(\gamma_0, \gamma_f) \quad (21)$$

Then,

$$sp(v_{n_0+1} - v_{n_0}) \le M + (1+\gamma)sp(v_{n_0})$$
$$\le M + (1+\gamma)\left[\frac{M(1-(\gamma')^{n_0})}{1-\gamma'} + (\gamma')^{n_0}(sp(v_0))\right] \quad (22)$$

Let $y_i = \sum_{j=n_0+1}^{i}(\gamma^{i-j} sp(v_j)|\Delta\Phi_{j-1}|))$. Then $\sum_{i=n_0}^{\infty} y_i = e_{n_0}/1 - \gamma$ as shown in Theorem 2 proof. Now, we find conditions on n_0 so that $e_{n_0} = \sum_{j=n_0+1}^{\infty}(sp(v_j)|\Delta\Phi_{j-1}|) \le (\varepsilon/2YM)$.
Since

$$sp(v_j) \le M(1-(\gamma')^j)/1 - \gamma' + (\gamma')^j sp(v_0)$$
$$\le (M/1-\gamma') + \gamma'^{(n_0)} sp(v_0)$$

for all $j > n_0$

$$e_{n_0} \le \sum_{j=n_0+1}^{\infty}(sp(v_j)|\Delta\Phi_{j-1}|)\left[\frac{M}{1-\gamma'} + (\gamma')^{n_0} sp(v_0)\right]\sum_{j=n_0+1}^{\infty}(|\Delta\Phi_{j-1}|)$$
$$\le \left[\frac{M}{1-\gamma'} + \gamma'^{(n_0)} sp(v_0)\right]\Phi(n_0) \quad (23)$$

since $\Phi(i)$ is nonincreasing, $i \ge n_0$, then $|\Delta\Phi_{j-1}| = \Phi(j-1) - \Phi(j)$. Easily, we can see $(M/(1-\gamma')) + \gamma'^{(n_0)} sp(v_0)$ is bounded.

Therefore, there exists n_0 so that $e_{n_0} \le [M/1-\gamma' + \gamma'^{(n_0)} sp(v_0)]\Phi(n_0) \le \varepsilon/2YM$. Then for all $i \ge n_0$, $e_i \le e_{n_0} \le \varepsilon/2YM$. Now, for $n \ge n_0$

$$E \le \frac{||v_{n_0} - n_0 g^*||_\infty}{n} + \frac{1}{n}\frac{sp(v_{n_0+1} - v_{n_0})}{1-\gamma} + \frac{YM\sum_{i=n_0}^{n-1}e_i}{n} + \frac{1}{n}2YM\frac{e_{n_0}}{1-\gamma}$$
$$\le \frac{1}{n}\left(2n_0 M + ||v_0||_\infty + \frac{M + (1+\gamma)\left[\frac{M(1-(\gamma')^{n_0})}{1-\gamma'} + (\gamma')^{n_0}(sp(v_0))\right]}{1-\gamma} + \frac{\varepsilon}{1-\gamma}\right) + \frac{YM(n-n_0)\frac{\varepsilon}{2YM}}{n}$$
$$\le \frac{1}{n}\left(2n_0 M + ||v_0||_\infty + \frac{M + (1+\gamma)\left[\frac{M(1-(\gamma')^{n_0})}{1-\gamma'} + (\gamma')^{n_0}(sp(v_0))\right]}{1-\gamma} + \frac{\varepsilon}{1-\gamma}\right) + \frac{\varepsilon}{2}$$

Let $1/n((M + (1+\gamma)[M(1-(\gamma')^{n_0})/1-\gamma'+(\gamma')^{n_0}(sp(v_0))]/1-\gamma) + 2n_0 M + ||v_0||_\infty + \varepsilon/1-\gamma) \leq (\varepsilon/2)$, or

$$n \geq 2/\varepsilon(2n_0 M + ||v_0||_\infty + (\varepsilon/1-\gamma) + M + (1+\gamma)[M(1-(\gamma')^{n_0})/ 1-\gamma' + (\gamma')^{n_0}(sp(v_0))]/1-\gamma).$$

Then, $E \leq \varepsilon$.

Comments on Theorem 3. From Theorem 3, we can predict the number of steps necessary for the value iteration algorithm to ensure that $E < \varepsilon$ without running value iteration algorithm for a while. This is useful for some applications that we would like to have a prediction in advance although the prediction might not be good. Note that Eqs. (21) and (22) together show the existence of n_0 in Theorem 2.

According to Ref. [16], the delta coefficient of a matrix is an upper bound of the second largest eigenvalue modulus (SLEM) λ^*. Thus, the gamma coefficient is also an upper bound of λ^*. Then, the term $1/1-\gamma$ is a upper bound of the relaxation time $t_{\rm rel} = 1/1 - \lambda^*$ of P_f^d for all $d \in D$. Moreover, the relaxation time is proportional to the mixing time of P_f^d or the convergence rate of the corresponding Markov chain [17]. Therefore, in any of the theorems above, the convergence rate of the adiabatic-time MDP is proportional to the convergence rate of a Markov chain with $P_f^{d_f}$. This is reasonable to the intuition that when n approaches infinity, the environment stops changing and the value iteration process becomes a Markov chain with the decision rule d_f^*.

4 Application of Adiabatic MDP: Queuing Systems

In this section, we apply the previous result to a continuous queuing system with estimated λ.

4.1 Queuing Theory Basics.
Queuing systems are used in many applications in engineering and management, such as transferring packets over a network, waiting line of customers, etc. [18–20]. A queuing system typically includes a queue and one or many servers as seen in Fig. 3. It is characterized by the arrival rate λ (of packets or customers), the service rate μ (of packets or customers), the number of servers, and the system's size.

For continuous-time queuing system, the arrival process to the queue is modeled by a Poisson process with the arrival rate λ. The arrival rate λ could be known or estimated by taking the average number of arrival packets per unit time. Similarly, the service or departure process is also modeled by a Poisson process with the service rate μ. The arrival sequence and departure sequence are independent or Markovian. In this paper, we consider queue with only one server. According to Kendall's notation, this queuing system is called $M|M|1|K$, where the first M stands for Markovian arrival process, the second M stands for Markovian departure process, 1 is the number of server, and K is the size of the system including 1 in the server and $K-1$ in the queue. This queuing system can be modeled as a continuous-time Markov chain with transition matrix in time interval t: $P^t = e^{Qt}$, where Q is the generator matrix [21,22]

$$Q = \begin{pmatrix} -\lambda & \lambda & 0 & 0 & \dots \\ \mu & -(\mu+\lambda) & \lambda & 0 & \dots \\ \ddots & \ddots & \ddots & \ddots & \ddots \\ \dots & 0 & \mu & -(\mu+\lambda) & \lambda \\ \dots & 0 & 0 & \mu & -\mu \end{pmatrix}$$

For the discrete-time queuing system, time is divided into small discrete-time slots. Suppose the time slot is small enough that at most one packet can come in. Then, the arrival process can be modeled as a Bernoulli process with parameter p defined as the probability of one packet arrived in a time slot. This probability could be changed over time due to the network's environment or the source. Similarly, the departure process is also modeled as a Bernoulli process with parameter q which is the probability that one packet is served and departs the queue. In this paper, we also consider the discrete-time queuing system with only one server and system's size K. This queuing can be modeled as a Markov chain with tridiagonal transition matrix as follows assuming that if there is no packet in the queue, a packet coming in will be served immediately:

$$P = \begin{bmatrix} 1-p(1-q) & p(1-q) & 0 & \dots & 0 \\ q(1-p) & pq+(1-p)(1-q) & p(1-q) & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & q(1-p) & pq+(1-p)(1-q) & p(1-q) \\ 0 & \dots & \dots & q(1-p) & 1-q(1-p) \end{bmatrix}$$

A more general analysis of discrete-time queuing system could be found in Ref. [23].

4.2 Application to Continuous-Time Queuing System With Arrival Rate Estimation

4.2.1 MDP Formulation for Continuous-Time Queuing System.
Consider an $M|M|1|K$ queue with fixed arrival rate λ which is unknown. We estimate λ at time $i\Delta t$ denoted as $\hat{\lambda}_i$ and

Fig. 3 Queuing systems

decide packet departure rate, $\mu_i = f(\hat{\lambda}_i)$ based on this estimate as follows:

$$\hat{\lambda}_i = \frac{1}{i\Delta t}\sum_{k=1}^{i} X_k$$

$$\mu_i = f(\hat{\lambda}_i) = (1+\beta_i)\hat{\lambda}_i$$

where $X_k \sim \text{Poisson}(\lambda\Delta t)$ is the number of packets in the kth slot of duration Δt, and $\beta_i > 0$ is an action we choose from a set of constant numbers. The goal is to find the departure rate that maximizes the reward. Intuitively, high departure rate incurs high costs (high power consumed) and low departure rate may lead to overflow [24]. We define the state s as the number of packets in the queue; therefore, $s \in S = \{0, 1, 2, \dots, K-1, K\}$. Let d be the decision rule which is a mapping which maps each state to a value

of β_i or $\beta_i(s) = d_i(s)$. The generator matrix in time interval $(i\Delta t, (i+1)\Delta t]$ is shown as follows:

$$Q_i = \begin{pmatrix} -\lambda_i & \lambda_i & 0 & 0 & \dots \\ \mu_i & -(\mu_i + \lambda_i) & \lambda_i & 0 & \dots \\ \ddots & \ddots & \ddots & \ddots & \ddots \\ \dots & 0 & \mu_i & -(\mu_i + \lambda_i) & \lambda \\ \dots & 0 & 0 & \mu_i & -\mu_i \end{pmatrix} \quad (24)$$

Note that at each time step, we have different Q_i since the arrival rate is estimated over time. The corresponding transition probability matrix $P(i\Delta t, (i+1)\Delta t)$

$$P_i = P(i\Delta t, (i+1)\Delta t) = e^{Q_i \Delta t} \quad (25)$$

where

$$Q_i = \hat{\lambda}_i \begin{pmatrix} -1 & 1 & 0 & 0 & \dots \\ 1+\beta_i(1) & -(2+\beta_i(1)) & 1 & 0 & \dots \\ \ddots & \ddots & \ddots & \ddots & \ddots \\ \dots & 0 & 1+\beta_i(K-1) & -(2+\beta_i(K-1)) & 1 \\ \dots & 0 & 0 & 1+\beta_i(K) & -(1+\beta_i(K)) \end{pmatrix}$$

Suppose $\beta \in [0,1]$. If we only care about the delay of packets in the queue and the cost of high serving rate, we can set and normalize the reward function: $r(s,\beta) = -M(s - (K)\beta)^2 / \max_{s,\beta}(s - K\beta)^2$ or we can define the cost function: $c(s,\beta) = M(s - K\beta)^2 / \max_{s,\beta}(s - K\beta)^2$. In this reward function, for a value of s, the cost is high if we use either too high departure rate (high cost) or too low departure rate (overflow cost). Since $s \in [0, K]$, we have to scale $\beta \in [0, 1]$ by K.

4.2.2 Simulation Results. Let $Z_i = \sum_{k=1}^{i} X_k \sim \text{Poisson}(i\lambda\Delta t), E[Z_i] = \text{VAR}[Z_i] = i\lambda\Delta t$. By applying Chernoff's bound, we get

$$P(Z_i > (1+a_i)i\lambda\Delta t) < e^{(-i\lambda\Delta t)} \frac{(ei\lambda\Delta t)^{(1+a_i)i\lambda\Delta t}}{((1+a_i)i\lambda\Delta t)^{(1+a_i)i\lambda\Delta t}}$$

Since $e^{(-i\lambda\Delta t)}(ei\lambda\Delta t)^{(1+a_i)i\lambda\Delta t}/((1+a_i)i\lambda\Delta t)^{(1+a_i)i\lambda\Delta t}$ is a decreasing function and approaches 0, when $a_i > 0$. Therefore, given a small α, we can find $a_i = f(i\lambda\Delta t)$ which is a decreasing function of i so that $e^{(-i\lambda\Delta t)}(ei\lambda\Delta t)^{(1+a_i)i\lambda\Delta t}/((1+a_i)i\lambda\Delta t)^{(1+a_i)i\lambda\Delta t} < \alpha$ and $\lim_{i\to\infty} a_i = 0$. Then, $P((Z_i/i\Delta_t) > (1+a_i)\lambda) < \alpha$ or

$$P(\hat{\lambda}_i > (1+a_i)\lambda) < \alpha$$

Therefore, with probability α, we have $\hat{\lambda}_i > (1+a_i)\lambda$. Similarly, with probability α, we have $\hat{\lambda}_i < (1-b_i)\lambda$ where $b_i > 0$ can be calculated using the Chernoff's bound. Both a_i and b_i tend to 0 when i reaches infinity. Overall, with probability $1 - 2\alpha$, $(1-b_i)\lambda \leq \hat{\lambda}_i \leq (1+a_i)\lambda$. This is illustrated in Fig. 4.

To use a decreasing function $\Phi(\cdot)$ for queue, we consider the cases where $\hat{\lambda}_i$ follows the upper and lower bounds on $\hat{\lambda}_i$ above to model the dynamic of P_i in Eq. (25), i.e., $\hat{\lambda}_i = (1+a_i)\lambda$ and $\hat{\lambda}_i = (1-b_i)\lambda$, respectively. These are two worst cases for $\hat{\lambda}_i$ with probability $1 - 2\alpha$. We know that $P_i = P_f + \Phi(i)(P_0 - P_f)$, hence

$$|P_i - P_f|_\infty = \Phi(i)|P_0 - P_f|_\infty$$

or $\Phi(i) = |P_i - P_f|_\infty / |P_0 - P_f|_\infty$ can be computed in term of a_i or b_i, respectively.

Since a_i and b_i decrease to 0 when i tends to ∞, $\Phi(i)$ decreases to 0 from $\Phi(0) = 1$. Therefore, $\Phi(i)$ is a decreasing function satisfying our conditions on $\Phi(i)$. We now show the bound on the average reward provided by Theorems 1 and 3 applied to the following scenario.

Simulation Scenario 1: Let $\varepsilon = 0.01, \lambda = 40, \Delta_t = 1, \alpha = 0.1, M = 1, K = 10, A = \{0.2, 0.4, 0.6, 0.8\}, v_0 = 0e$.
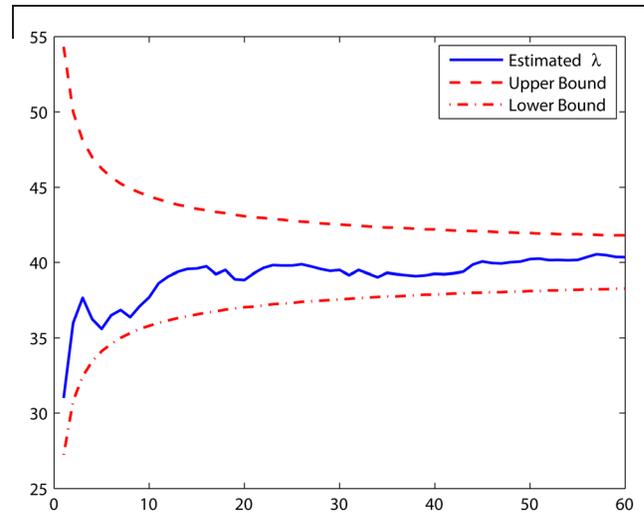

Fig. 4 An example of estimated $\hat{\lambda}_i$ and it bounds for actual $\lambda = 40$
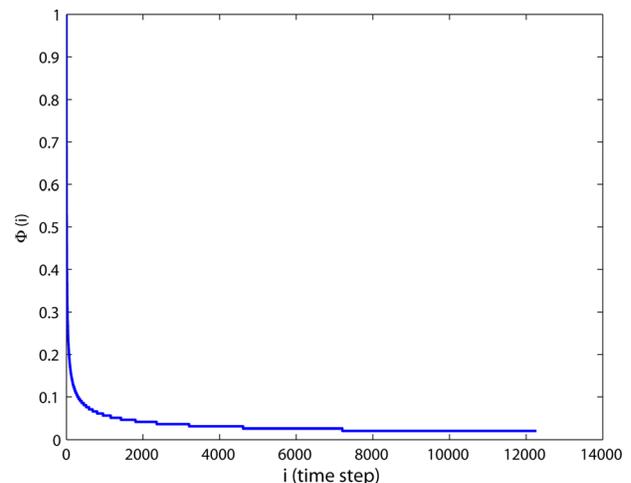

Fig. 5 The $\Phi(\cdot)$ function for simulation scenario 1

Now, we consider the upper bound as the estimated arrival rate $\hat{\lambda}_i = (1 + a_i\lambda)$. The function $\Phi(\cdot)$ is shown in Fig. 5. From Theorem 2, we obtained $n_0 = 1$ and $N = 29$. From Theorem 3, we obtained $n_0 = 58$ and $N = 23,870$. We can see that the value of N
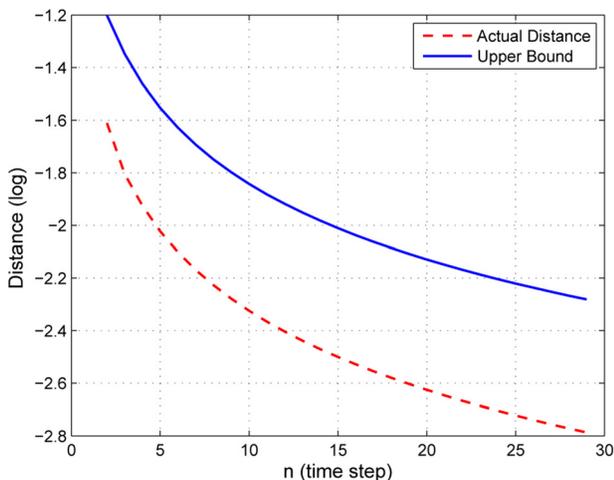
Fig. 6   The actual distance and its upper bound for $\hat{\lambda}_i = (1+a_i)\lambda$ from Theorem 2 (simulation scenario 1)
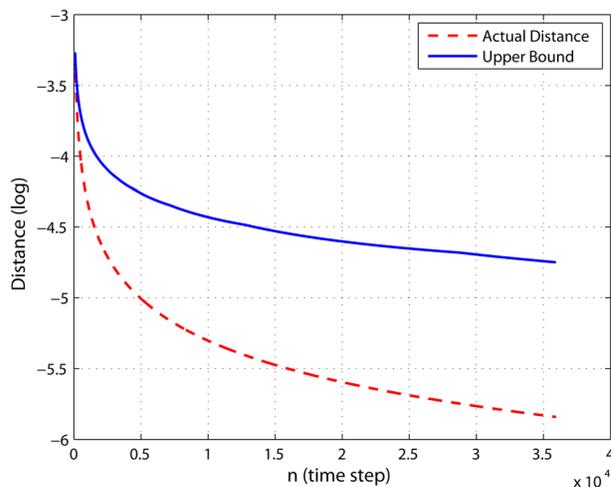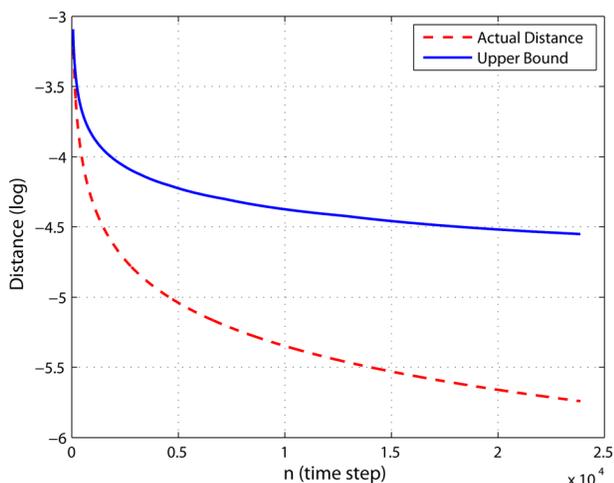

Fig. 7   The actual distance and its upper bound for $\hat{\lambda}_i = (1+a_i)\lambda$ from Theorem 3 (simulation scenario 1)

Fig. 8   The actual distance and its upper bound for $\hat{\lambda}_i = (1-b_i)\lambda$ from Theorem 2 (simulation scenario 1)
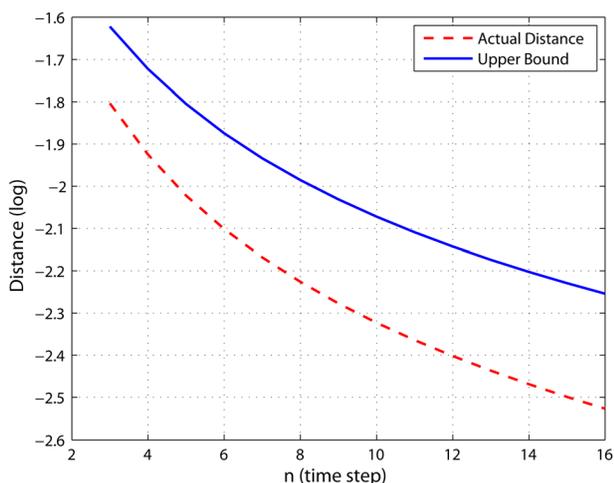

Fig. 9   The actual distance and its upper bound for $\hat{\lambda}_i = (1-b_i)\lambda$ from Theorem 3 (simulation scenario 1)

Table 1   The optimal decision rule for continuous-time queuing system example

State	0	1	2	3	4	5	6	7	8	9	10
Action	0.2	0.2	0.2	0.2	0.4	0.4	0.6	0.6	0.8	0.8	0.8

from Theorem 2 is much better than the value from Theorem 3 as expected.

Figures 6 and 7 show the actual distance to the optimal reward obtained from the value iteration algorithm and its upper bound by Theorems 2 and 3, respectively. As seen, they are well correlated with each other.

Similarly, for the lower bound on estimated arrival rate $\hat{\lambda}_i = (1-b_i)\lambda$, we have $n_0 = 2$ and $N = 16$ from Theorem 2 and $n_0 = 88$ and $N = 35,925$ from Theorem 3. As anticipated, the results from Theorem 2 are better than ones from Theorem 3.

The actual distance and its upper bound for each theorem are shown in Figs. 8 and 9. As seen, they are also well correlated with each other.

From both value iteration for nonstationary environment with adiabatic-time setting and value iteration for stationary environment, we have the optimal decision rule as shown in Table 1. The optimal average reward $g^* = -0.0235$. We can see that the optimal policy is reasonable since we need high departure rate for high queue occupancy and vice versa.

4.3   Application to Discrete-Time Queuing System

4.3.1   MDP Formulation. In this section, we show a simple example illustrating the application of our framework to the time-varying underlying environments. Specifically, we consider a discrete-time queuing system of size $K = 2$.
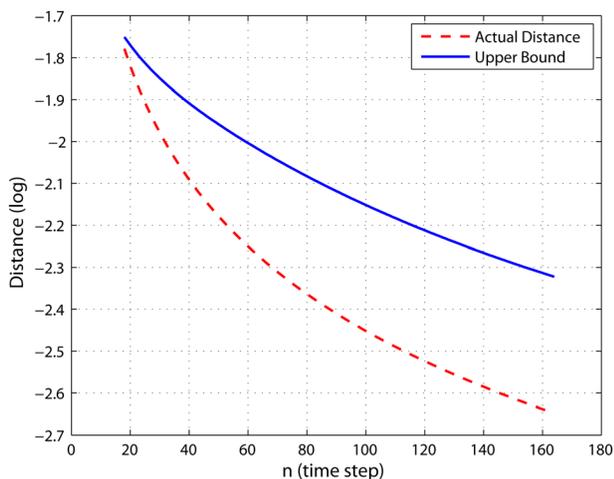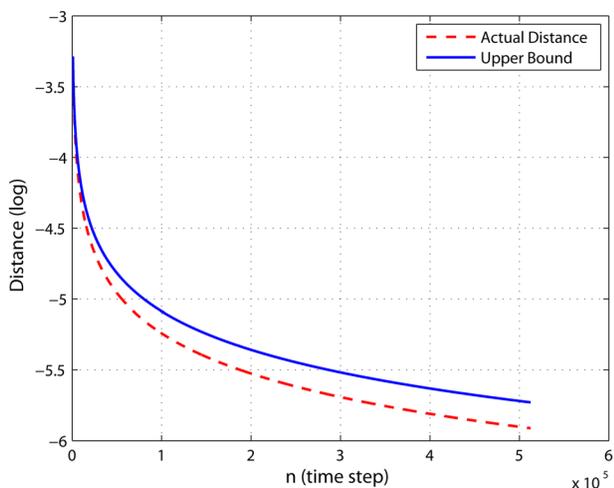
Assume at each time step, there are probabilities p and q that a packet will be arriving and departing the queue, respectively. In this case, the state-space $S = \{0, 1, 2\}$, the time-varying environment is described by changing the values of p over time, while the action is the value of q. The transition matrix has the following form:

$$P = \begin{bmatrix} 1-p(1-q) & p(1-q) & 0 \\ q(1-p) & pq+(1-p)(1-q) & p(1-q) \\ 0 & q(1-p) & 1-q(1-p) \end{bmatrix}$$

Because all the entries in the matrix P are linear functions of p, if we set $p_i = \phi(i)p_0 + (1-\phi(i))p_f$ to model the change in the

Table 2 The reward $r(s, a)$ for discrete-time queuing system

$r(s, a)$		A								
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
s	0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1
	1	0.1	0.2	0.3	0.4	0.5	0.4	0.3	0.2	0.1
	2	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9

Table 3 The optimal decision rule for discrete-time queuing system example

State	0	1	2
Action	0.1	0.5	0.9

From Theorem 2, we obtained $n_0 = 17$ and $N = 164$. From Theorem 2, we obtained $n_0 = 811$ and $N = 512, 509$. Once again, the results from Theorem 2 are better than ones from Theorem 3. Figures 10 and 11 show the simulated distance to the optimal reward obtained from the value iteration algorithm and its upper bound from Theorems 2 and 3. As seen, they are very correlated.

From both value iteration for nonstationary environment with adiabatic-time setting and value iteration for stationary environment, we have the optimal decision rule as shown in Table 3. The optimal average reward $g^* = 0.7185$. We can see that the optimal policy is reasonable since we need high departure rate for high queue occupancy and vice versa.

5 Conclusions and Future Work

We provide an analysis framework for studying the value iteration algorithm under the adiabatic setting. We provide theoretical bounds on the convergence rate of the value iteration algorithm with the average reward objective. Specifically, our work provides a lower bound on the number of time iterations in the value iteration algorithm needed in order to ensure that the resulting policy produces an average reward that is ε-close to the optimal average reward value. In the future, we can apply these results to noisy stochastic matrices with parameter estimation. This is more general than the continuous-time queuing systems we considered previously.

Fig. 10 The actual distance and its upper bound from Theorem 2 (simulation scenario 2)

Fig. 11 The actual distance and its upper bound from Theorem 3 (simulation scenario 2)

Appendix A: Average Reward Criteria

In this section, we show some previous results on average reward criteria for stationary environment. Therefore, we have three following assumptions:

ASSUMPTION 1. *The reward $r(s, a)$ and the probability $p(j|s, a)$ are stationary.*

ASSUMPTION 2. *The reward $r(s, a)$ is bounded: $|r(s, a)| \leq M < \infty \quad \forall a \in A_s, s \in S$.*

ASSUMPTION 3. *The set of states S is finite.*

Note that the stationary probability $p(j|s, a)$ condition is only necessary for this section where the results for stationary environment are shown.

DEFINITION 4 (average reward). *Let $v_{N+1}^\pi(s) = E_s^\pi[\sum_{t=1}^N r(s_t, a_t)]$ be the total reward we get at the time step $N + 1$ starting from the state s under the policy π, where s_t and a_t are the state and action at time step t. Then, the average reward $g^\pi(s)$ is defined as follows:*

$$g^\pi(s) = \lim_{N \to \infty} \frac{1}{N} v_{N+1}^\pi(s) = \lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^N P_\pi^{n-1} r_{d_n}(s)$$

where d_n is the decision rule at the time step n under the policy π.

PROPOSITION 5. *Let S be countable, $\pi = d^\infty, P_d^\infty = \lim_{N \to \infty}(1/N) \sum_{n=1}^N (P^d)^{n-1}$ is stochastic [5], then $g^{d^\infty}(s)$ exists and*

$$g^{d^\infty}(s) = \lim_{N \to \infty} \frac{1}{N} v_{N+1}^{d^\infty}(s) = P_d^\infty r^d(s)$$

Obviously, for finite S, this proposition holds.

PROPOSITION 6. *Suppose P_d^∞ is stochastic. Then, if j and k are in the same closed recurrent class, $g(j) = g(k)$ [5]. Furthermore, if the chain is irreducible and aperiodic or has a single recurrent class and may have some transient states, $g(s)$ is a constant function.*

arrival rates, then: $P_i = \Phi(i)P_0 + (1 - \Phi(i))P_f$ where $\Phi(i) = \phi(i)$ for all i. Note that $\Phi(\cdot)$ satisfies the conditions for the adiabatic setting. For the reward function, when the number of packets s is small, we should not use high serving rate. Whereas, when the number of packets s is large, we have to use high serving rate to avoid overflow. Therefore, we choose the reward function as shown in Table 2.

4.3.2 Simulation Results. To examine the theoretical results, we run simulation using the following parameters:

Simulation Scenario 2: Let $\varepsilon = 0.01, p_0 = 0.4, p_f = 0.6, \phi(i) = 1/i + 1, v_0 = 0e, A = \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$.

Unichain MDP satisfies these conditions since it only has one single recurrent class plus possibly empty set of transient states.

Appendix B: Span-Seminorm and Its Properties

DEFINITION 5 (span seminorm). *The span seminorm of a vector* $v \in V$ *is defined as follows:*

$$sp(v) = \max_{s \in S} v(s) - \min_{s \in S} v(s)$$

where V is the vector space and s is the state.

This seminorm has the following properties:

(1) $sp(v) \geq 0$.
(2) $sp(u + v) \leq sp(u) + sp(v), \quad \forall u, v \in V$.
(3) $sp(kv) = |k| sp(v) \quad \forall k \in \mathbb{R}, v \in V$.
(4) $sp(v + ke) = sp(v) \quad \forall k \in \mathbb{R}$, where e is the constant vector $e = [11\ldots1]^{\mathrm{T}}$.
(5) $sp(v) = sp(-v)$.
(6) $sp(v) \leq 2||v||$.

Note that δ_d is called Hajnal measure or delta coefficient of P^d [16]. It is an upper bound on the SLEM $|\lambda_2|$ of P^d, and $\delta_d \leq 1 - \sum_{j \in S} \min_{s \in S} P^d(j|s)$.

Appendix C: Additional Details on Proof

PROPOSITION 4 (the bound on gamma coefficients). *Given* $P_i = \Phi(i)P_0 + (1 - \Phi(i))P_f, \Phi(i) \in [0, 1]$

$$\gamma_i \leq \Phi(i)\gamma_0 + (1 - \Phi(i))\gamma_f$$

Proof. Since $P_i = \Phi(i)P_0 + (1 - \Phi(i))P_f, \Phi(i) \in [0, 1]$, for all j, s, a, s', a', we have

$$\min\{p_i(j|s, a), p_i(j|s', a')\} \geq \Phi(i)\min\{p_0(j|s, a), p_0(j|s', a') \\ + (1 - \Phi(i))\min\{p_f(j|s, a), p_f(j|s', a')$$

Therefore,

$$\sum_{j \in S} \min\{p_i(j|s, a), p_i(j|s', a') \geq \Phi(i) \sum_{j \in S} \min\{p_0(j|s, a), p_0(j|s', a') \\ + (1 - \Phi(i)) \sum_{j \in S} \min\{p_f(j|s, a), \\ \times p_f(j|s', a')$$

or

$$1 - \sum_{j \in S} \min\{p_i(j|s, a), p_i(j|s', a') \leq \Phi(i)(1 - \sum_{j \in S} \min\{p_0(j|s, a), \\ \times p_0(j|s', a')) + (1 - \Phi(i)) \\ \times (1 - \sum_{j \in S} \min\{p_f(j|s, a), \\ \times p_f(j|s', a'))$$

In other words, $\gamma_i \leq \Phi(i)\gamma_0 + (1 - \Phi(i))\gamma_f$.

COROLLARY 1. *Given* $P_i = \Phi(i)P_0 + (1 - \Phi(i))P_f$ *with nonincreasing function* $\Phi(i) > 0, \Phi(i) \in [0, 1]$ *for any* $n_0 \leq i$

$$\gamma_i \leq \max(\gamma_{n_0}, \gamma_f) \tag{C1}$$

Proof. We have

$$\begin{aligned} P_i &= \Phi(i)P_0 + (1 - \Phi(i))P_f \\ &= P_f + \Phi(i)(P_0 - P_f) \\ &= P_f + \frac{\Phi(i)}{\Phi(n_0)}\left(\Phi(n_0)P_0 - \Phi(n_0)P_f\right) \\ &= P_f + \frac{\Phi(i)}{\Phi(n_0)}\left(\Phi(n_0)P_0 + \left(1 - \Phi(n_0)\right)P_f - P_f\right) \\ &= P_f + \frac{\Phi(i)}{\Phi(n_0)}\left(P_{n_0} - P_f\right) \end{aligned} \tag{C2}$$

Let $0 < \Phi'(i) = \Phi(i)/\Phi(n_0) \leq 1$, since $\Phi(\cdot)$ is a positive nonincreasing function. Then, $P_i = P_f + \Phi'(i)(P_{n_0} - P_f) = \Phi'(i)P_{n_0} + (1 - \Phi'(i))P_f$. By applying Proposition 4, we have

$$\gamma_i \leq \Phi(i)\gamma_{n_0} + (1 - \Phi(i))\gamma_f \leq \max(\gamma_{n_0}, \gamma_f)$$

References

[1] Bellman, R., 1957, *Dynamic Programming*, Princeton University Press, Princeton, NJ.
[2] Howard, R., 1960, *Dynamic Programming and Markov Processes*, MIT Press, Cambridge, MA.
[3] d'Epenoux, F., 1960, "Sur un probleme de production et de stockage dans laléatoire," Rev. Fr. Rech. Opér., **14**, pp. 3–16.
[4] Derman, C., 1970, *Finite State Markovian Decision Processes*, Academic Press, Orlando.
[5] Puterman, M. L., 1994, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, 1st ed., Wiley, New York.
[6] Born, M., and Fock, V., 1928, "Beweis des adiabatensatzes," Z. Phys., **51**, pp. 165–180.
[7] Messiah, A., 1962, *Quantum Mechanics*, 1st ed., Vol. 2, Wiley, New York.
[8] Kovchegov, Y., 2010, "A Note on Adiabatic Theorem for Markov Chains," Stat. Probab. Lett., **80**, pp. 186–190.
[9] Bradford, K., and Kovchegov, Y., 2011, "Adiabatic Times for Markov Chains and Applications," J. Stat. Phys., **143**(5), pp. 955–969.
[10] Szita, I., Takács, B., and Lörincz, A., 2002, "ε-MDPS: Learning in Varying Environments,," J. Mach. Learn. Res., **3**, pp. 145–174.
[11] Bradford, K., Kovchegov, Y., and Nguyen, T., 2016, "Stable Adiabatic Times for Markov Chains," Stochastics, **88**(4), pp. 567–585.
[12] Rosenwald, R., Meyer, D., and Schmitt, H., 2004, "Applications of Quantum Algorithms to Partially Observable Markov Decision Processes," 5th Asian Control Conference, Melbourne, Australia, June 20–23, Vol. 1, pp. 420–427.
[13] Zacharias, L., Nguyen, T., Kovchegov, Y., and Bradford, K., 2012, "Analysis of Adaptive Queuing Policies Via Adiabatic Approach," 2013 International Conference on Computing, Networking and Communications (ICNC), Network Algorithm and Performance Evaluation Symposium, San Diego, CA, Jan. 28–31, pp. 1053–1057.
[14] Duong, T., Nguyen-Huu, D., and Nguyen, T., 2013, "Adiabatic Markov Decision Process With Application to Queuing Systems," 47th Annual Conference on Information Sciences and Systems (CISS), Baltimore, MD, Mar. 20–22, pp. 1–6.
[15] Derman, C., and Strauch, R. E., 1966, "A Note on Memoryless Rules for Controlling Sequential Control Processes," Ann. Math. Stat., **37**(1), pp. 276–278.
[16] Seneta, E., 1981, *Non-Negative Matrices and Markov Chains*, Springer-Verlag, New York.
[17] Levin, A. D., Peres, Y., and Wilmer, E. L., 2008, *Markov Chains and Mixing Times*, American Mathematical Society, Providence, RI.
[18] Kleinrock, L., 1976, *Queuing Systems: Theory*, Vol. 1, Wiley, New York.
[19] Kleinrock, L., 1976, *Queuing Systems: Computer Applications*, Wiley, New York.
[20] Gautam, N., 2012, *Analysis of Queues: Methods and Applications (Operations Research Series)*, Taylor & Francis, Boca Raton.
[21] Lawler, G., 2006, *Introduction to Stochastic Processes* (Chapman and Hall/CRC Probability Series), Chapman & Hall/CRC, Boca Raton.
[22] Bremaud, P., 1999, *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues* (Texts in Applied Mathematics), Springer, Boca Raton.
[23] Briem, U., Theimer, T., and Kröner, H., 1991, "A General Discrete-Time Queuing Model: Analysis and Applications," International Teletraffic Congress, Vol. 13, pp. 13–19.
[24] Morrison, J. A., 1980, "Analysis of Some Overflow Problems With Queuing," Bell Syst. Tech. J., **59**(8), pp. 1427–1462.