

Exact finite approximations of average-cost countable Markov Decision Processes

Arie Leizarowitz

Faculty of Mathematics
Technion, Haifa 32000
Israel

Adam Shwartz

Faculty of Electrical Engineering
Technion, Haifa 32000
Israel

March 2007

Abstract

For a countable-state Markov decision process we introduce an embedding which produces a finite-state Markov decision process. The finite-state embedded process has the same optimal cost, and moreover, it has the same dynamics as the original process when restricting to the approximating set. The embedded process can be used as an approximation which, being finite, is more convenient for computation and implementation.

Keywords: Markov Decision Processes, countable state space, finite approximations, average cost criterion.

AMS subject classification: 60G70, 60J10, 62L20.

1 Introduction

In this paper we develop a tool that is useful in studying countable state Markov Decision Processes (MDPs). A Markov Decision Process is a controlled dynamical system with probabilistic transitions, that are influenced by the control actions (for precise definitions see §§ 1.1). We consider discrete-time MDPs with a discrete *state space* X which is either *finite* or *countably infinite*, to which we will refer in the sequel as *countable*. The cost under consideration is the long-time average cost.

Countable MDPs are obviously more difficult to study, analytically and numerically, than finite state MDPs. Several approaches were developed to deal with this issue. The first approach is to reduce the state space by clustering together “equivalent” states: see e.g. [GDG] and references therein. This approach provides a smaller state space and exact relations, but requires a very special structure of the MDP in order for the derived model to have a finite state space. Namely, equivalent states must have the same transition probability into and out of the state, under any action, and the same immediate cost. This special structure seldom exists in applications.

The second approach deals with the approximation of countable MDPs by finite state MDPs using a truncation of the state space. Existing results show that as the size of the approximating MDP increases, its cost function and, under some conditions its optimal policies approach those of the original, countable MDP. See, e.g. [C, A1, A2] and references therein. This approach is applicable in greater generality, but typically provides approximations without an error estimate—thus the results are “asymptotic” in nature.

We propose a different approach, with the advantage that the optimal cost of the approximating, finite MDP agrees with that of the countable MDP. Moreover, restricted to the approximating set the optimal policies agree as well. Thus the term “exact approximations.” We establish this approach for general, compact action spaces.

We conclude this section with a precise statement of the problem. In §2 we introduce the main idea—the finite embedding, and prove its existence. In §3 we show that the embedding possesses the desired properties. We discuss some extensions in §4. Section 5 develops some applications.

1.1 Problem formulation

Consider a process with state space $X \subset \{0, 1, 2, 3, \dots\}$. When the system occupies state $i \in X$, then the controller can influence its behavior by choosing an *action* a from the compact *action set* A_i , $i \in X$, which is a subset of the *action space* A . Choosing an action $a \in A_i$ has a twofold effect:

- (i) A *running cost* $c(i, a)$ is incurred,
- (ii) The system transits from state i to j according to the *transition probability* $P(j|i, a)$.

Thus an MDP is defined in terms of a quadruplet

$$\mathcal{M} = (X, \{A_i\}, c(i, a), P(j|i, a)).$$

The state and action at time $k \geq 0$ are denoted x_k and a_k respectively, so that the system's behavior on the infinite time interval is described in terms of the stochastic process $\{(x_k, a_k)\}_{k=0}^{\infty}$.

Admissible policies. A policy π is a law which is used to choose the actions $a_k \in A_{x_k}$. It is admissible if its choice at time k depends only on the history $(x_0, a_0, \dots, x_{k-1}, a_{k-1}, x_k)$ of the system up to time k . A policy can be either deterministic or randomized, so that the choice of a_k may be made according to a probability measure on A_{x_k} . A (possibly randomized) policy which depends only on x_k is called a “stationary Markov policy”. Such policies generate a state process $\{x_k\}_{k=0}^{\infty}$ which is a Markov chain with stationary transition probabilities.

The cost criterion. An admissible policy π generates the stochastic processes $\{x_k\}_{k=0}^{\infty}$ and $\{a_k\}_{k=0}^{\infty}$, and the expected cost flow

$$C_N(i, \pi) = E_i^\pi \sum_{k=0}^{N-1} c(x_k, a_k), \quad N \geq 1. \quad (1.1)$$

The expectation E_i^π in (1.1) is with respect to the probability measure P_i^π induced by π on the set of sequences $\{(x_k, a_k)\}_{k=0}^{\infty}$ with $x_0 = i$. We address the optimal control problem of minimizing the functional

$$\pi \mapsto J_i(\pi) = \liminf_{N \rightarrow \infty} \frac{1}{N} C_N(i, \pi), \quad x_0 = i,$$

over all admissible policies, and call $J_i(\pi)$ the *expected long-run average cost*. It is convenient to introduce the notation $g^*(\mathcal{M})$ for the optimal cost: this notation makes explicit the model \mathcal{M} under consideration. An *optimal policy* realizes the minimal long-run average cost, and an ϵ -optimal policy realizes it up to ϵ .

We need to exclude one case in which it is not possible to approximate a countable MDP by a finite one.

Definition 1.1 *Let σ be a stationary Markov policy of an MDP \mathcal{M} , generating the state-action process $\{(x_k, a_k)\}$. We say that σ is a drifting policy if for every finite set $F \subset X$ and every initial condition $i \in F$,*

$$P_i^\sigma(x_k \in F) \rightarrow 0 \text{ as } k \rightarrow \infty. \quad (1.2)$$

We say that an MDP \mathcal{M} is a drifting MDP if there exists a constant $\delta > 0$ such that any stationary Markov policy of \mathcal{M} , say σ , that satisfies

$$J(\sigma) < g^*(\mathcal{M}) + \delta$$

is drifting. An MDP is called non-drifting if it is not a drifting MDP.

Obviously it is not possible to approximate a drifting MDP by a finite MDP while preserving both cost structure, optimality and dynamics. Moreover, if the MDP is drifting so that the condition in Definition 1.1 is violated, then for a certain $\delta > 0$, for every $0 < \nu < \delta$, every stationary Markov policy with average cost smaller than $g^* + \nu$, eventually leaves every finite set F with probability 1. But this yields the existence of a Markov policy (not necessarily stationary) with average cost g^* which is drifting, namely eventually leaves every finite set with probability 1. Assuming that this does not happen implies that the requirement holds. In this paper we will therefore consider only non-drifting MDPs.

2 The embedding

We wish to associate with the given MDP \mathcal{M} a finite state MDP \mathcal{M}_0 ,

$$\mathcal{M}_0 = (X_0, A_0, Q_0(j|i, a), c_0(i, a))$$

with cost flow $C_N^0(i, \pi)$ defined as in (1.1), in such a manner that the two MDPs will share common optimality properties and, in some sense, will share

transition probabilities. To this end we introduce the notion of embedding, whose exact definition is presented below. In §3 we prove that this definition implies the desired exact approximation property.

We will next define the embedding notion, which will be followed by a construction of an embedding \mathcal{M}_0 . Denote by $\{(x_k, a_k)\}_{k=0}^\infty$ the generic state-action process of \mathcal{M} and by $\{(\xi_k, \alpha_k)\}_{k=0}^\infty$ the generic state-action process of \mathcal{M}_0 . Given a finite subset $Z \subset X$ define

$$\eta = \inf\{k \geq 0 : x_k \notin Z\} \quad \tau = \inf\{k > \eta : x_k \in Z\} . \quad (2.1)$$

Thus if the initial state is in Z then η is the first exit time and τ is the first return time, while if the initial state is not in Z then $\eta = 0$ and again τ is the first return time. For a stopping time $\nu > 0$ define

$$C_\nu(i, \pi) = E_i^\pi \sum_{k=0}^{\nu-1} c(x_k, a_k)$$

and note that if $\nu = N$, a deterministic integer, this definition agrees with (1.1). (If, however, ν is a random variable then $C_\nu(i, \pi) \neq C_N(i, \pi) \mid_{N=\nu}$, the latter being a random variable.) Given a finite subset $Z_0 \subset X_0$, define η_0 , τ_0 and C_ν^0 analogously.

Definition 2.1 *We say that \mathcal{M}_0 is embedded in \mathcal{M} if there exist subsets $Z_0 \subset X_0$ and $Z \subset X$, and a one-to-one mapping $e : Z_0 \mapsto Z$ from Z_0 onto Z such that, for any stationary Markov policy σ of \mathcal{M} under which Z contains at least one recurrent state, the following holds.*

There exists a stationary Markov policy σ_0 of \mathcal{M}_0 , such that if the processes start with initial states $x_0 \in Z$ and $\xi_0 = e^{-1}(x_0) \in Z_0$ respectively, then x_τ and ξ_{τ_0} are identically distributed (under the probability measures $P_{x_0}^\sigma$ and $P_{\xi_0}^{\sigma_0}$ respectively) and, if the term on the right is finite,

$$C_\tau(x_0, \sigma) = C_{\tau_0}^0(\xi_0, \sigma_0).$$

Embedding means that if we restrict attention only to the states Z in X and to the corresponding states Z_0 in X_0 , then the performance of any stationary Markov policy σ on \mathcal{M} can be imitated by the performance of some stationary Markov policy σ_0 on \mathcal{M}_0 . This imitation can be achieved when considering finite subsets F of the state space X on which σ generates a nontrivial dynamics, namely the states in F are not all transient under σ .

We now establish that for certain finite sets $Z \subset X$ there exists a finite state $\mathcal{M}_0 = (X_0, A_0, P_0, c_0)$ and an embedding $e(\cdot)$ of \mathcal{M}_0 in \mathcal{M} . This embedding will be useful and significant for stationary Markov policies σ which induce a nontrivial dynamics on Z . The embedding is such that if $e(\cdot)$ is defined on Z_0 , then $Z = e(Z_0)$ and

$$\#(X_0) = 2\#(Z_0).$$

In §3 we will employ the embedding result to establish existence and characterize optimal policies for certain countable state MDPs. As we shall see there, we need the embedding to be such that, in addition to the cost flow, the expected return times agree, that is, $E\tau = E\tau_0$.

2.1 The existence of embedding

Theorem 2.2 *Let \mathcal{M} be a countable state MDP and let $Z \subset X$ be a finite set. Then there exists a finite state MDP \mathcal{M}_0 with state space X_0 and an embedding $e : Z_0 \mapsto Z$ of \mathcal{M}_0 in \mathcal{M} such that*

$$\#(X_0) = 2\#(Z_0).$$

The result has non-trivial content provided that $Z \subset X$ contains a recurrent state of some stationary Markov policy of \mathcal{M} .

Proof. We will define an MDP \mathcal{M}_0 , and will then establish that it has the properties asserted in the theorem. Denote

$$Z = \{z_1, \dots, z_n\},$$

let Z_0 be a finite set

$$Z_0 = \{s_1, \dots, s_n\} \tag{2.2}$$

and let $e : Z_0 \mapsto Z$ be defined by

$$e(s_i) = z_i, \quad i = 1, 2, \dots, n. \tag{2.3}$$

With each state s_i in Z_0 we associate a state ω_i , and we then define

$$X_0 = \{s_1, \dots, s_n\} \cup \{\omega_1, \dots, \omega_n\}. \tag{2.4}$$

We have to specify \mathcal{M}_0 as a quadruplet

$$(X_0, (A_0)_s, c_0(s, a), P_0(s'|s, a))$$

and define explicitly the terms A_0 , c_0 and P_0 . We define $A_0(s_i)$ ($= (A_0)_{s_i}$) by

$$A_0(s_i) = A(e(s_i)) \text{ for every } 1 \leq i \leq n. \quad (2.5)$$

The definition of the action sets $A_0(\omega_i)$ will be given below.

We next define the transition probabilities $P_0(s|s_i, a)$, $s \in X_0$, $a \in A_0(s_i)$. First, for $s_i \in Z_0$

$$P_0(s|s_i, a) = \begin{cases} P(e(s_j)|e(s_i), a) & \text{if } s = s_j \in Z_0 \\ 1 - \sum_{j=1}^n P_0(s_j|s_i, a) & \text{if } s = \omega_i \\ 0 & \text{if } s = w_j, j \neq i. \end{cases} \quad (2.6)$$

The corresponding cost is defined by

$$c_0(s_i, a) = c(e(s_i), a) \quad (2.7)$$

for $a \in A_0(s_i) = A(e(s_i))$.

We now fix $1 \leq i \leq n$ and define the action sets and the transition probabilities for the state $\omega_i \in X_0 \setminus Z_0$. Let σ be a stationary Markov policy of \mathcal{M} such that Z contains at least one recurrent state. Recall the definitions (2.1) and let

$$q_j(\sigma) = P^\sigma\{x_\tau = z_j \mid x_{\tau-1} = z_i\}. \quad (2.8)$$

This is the probability that the process $\{x_k\}$ will first enter Z through state z_j , conditioned on having left Z from z_i while employing the action $a \in A(z_i)$ specified by σ . *The action for state ω_i* induced by σ is the collection of $n+1$ nonnegative numbers

$$\alpha(\sigma) = (\lambda q_1(\sigma), \dots, \lambda q_n(\sigma), c(\sigma)) \quad (2.9)$$

where the constant λ satisfies $0 < \lambda \leq 1$. These two parameters λ and $c(\sigma)$ which define the action $\alpha(\sigma)$ will be specified below. The quantities $q_1(\sigma), \dots, q_n(\sigma)$ are the probabilities associated in (2.8) with the fixed state z_i and the stationary Markov policy σ . We define the action set of ω_i to be the set

$$A_0(\omega_i) = \bigcup_{\sigma} \alpha(\sigma), \quad (2.10)$$

where the union is over all the stationary Markov policies under which Z contains a recurrent state. Of course, two different policies σ_1 and σ_2 may give rise to the same action, that is $\alpha(\sigma_1) = \alpha(\sigma_2)$.

We now define the transition probabilities from ω_i . If $\alpha \in A_0(\omega)$, then

$$\alpha = (\alpha_1, \dots, \alpha_n, \alpha_{n+1}) = (\lambda q_1, \dots, \lambda q_n, c) \quad (2.11)$$

for some constant $0 < \lambda \leq 1$, and it follow that

$$\sum_{j=1}^n \alpha_j = \lambda.$$

We then define $P_0(s|\omega_i, \alpha)$ by

$$P_0(s|\omega_i, \alpha) = \begin{cases} \alpha_j & \text{if } s = s_j \in Z_0 \\ 1 - \lambda & \text{if } s = \omega_i \\ 0 & \text{if } s \notin Z_0 \cup \{\omega_i\}. \end{cases} \quad (2.12)$$

Thus for every choice of $0 < \lambda \leq 1$, the conditional probability to enter Z_0 through s_j , given that the process did enter Z_0 , is q_j , independent of λ . However, the value of λ determines the expected time that would elapse until entrance, and we choose λ in such manner that this expected time turns out to be equal to the corresponding time for \mathcal{M} . Namely, if σ_0 is the policy that uses $\alpha(\sigma)$ in state ω_i , then λ is chosen such that

$$1 + E_{\omega_i}^{\sigma_0} \tau_0 = E_{z_i}^{\sigma}[\tau \mid \eta = 1] \quad (2.13)$$

is satisfied.

Remark 2.3 *Note that if z_i is not accessible from the state in Z which is recurrent under σ then we do not need to specify actions for ω_i , while if it is accessible then necessarily it is recurrent, so that the expectation on the right-hand side of (2.13) is finite. An appropriate value of λ can clearly be chosen since for $\lambda = 1$ the left-hand side equals 1, while as $\lambda \rightarrow 0$ this expression diverges.*

We note that if it is possible to leave Z from z_i in one step then

$$E_{z_i}^{\sigma}[\tau \mid \eta = 1] = 1 + \left[\sum_{x_j \notin Z} P(x_j|z_i, \sigma(z_i)) \right]^{-1} \sum_{x_j \notin Z} P(x_j|z_i, \sigma(z_i)) E_{x_j}^{\sigma} \tau.$$

If, however, this is not possible, than we can define the action in $A_0(\omega_i)$ that corresponds to σ in an arbitrary manner. Finally note that the normalizing constant (in square bracket) above is just $P_0(\omega_i|s_i, \alpha(\sigma))$.

We next consider the cost associated with the action α , namely $c_0(\omega_i, \alpha)$, which is chosen to be such that

$$C_\tau(z_i, \sigma) = C_{\tau_0}^0(e^{-1}(z_i), \sigma_0) \quad (2.14)$$

holds for all $z_i \in Z$. We distinguish between two cases: If under σ the process $\{x_k\}$ starting at z_i does not leave Z , then equality holds in (2.14) by definition. If the process does leave Z , then

$$C_\tau(z_i, \sigma) = E_{z_i}^\sigma \sum_{k=0}^{\tau-1} c(x_k, a_k)$$

can be expressed in the form

$$c(z_i, \sigma(z_i)) + \sum_{x_j \notin Z} P(x_j | z_i, \sigma(z_i)) C_\tau(x_j, \sigma) + \sum_{z_j \in Z} P(z_j | z_i, \sigma(z_i)) C_\tau(z_j, \sigma).$$

We recall that $c_0(\omega_i, \alpha)$ is actually the $n+1$ component of $\alpha(\sigma)$ in (2.9), denoted $c(\sigma)$. It follows that setting

$$c_0(\omega_i, \alpha) = \sum_{x_j \notin Z} P(x_j | z_i, \sigma(z_i)) C_\tau(x_j, \sigma) [E_{\omega_i}^{\sigma_0} \tau_0]^{-1} \quad (2.15)$$

ensures the desired equality (2.14). Thus to each σ there corresponds an action $\alpha = \alpha(\sigma)$, and a cost

$$c_0(\omega_i, \alpha) = \alpha_{n+1} = c(\sigma)$$

which in view of the explicit expression (2.15), indeed depends on the state and action alone.

The definition of \mathcal{M}_0 is thus complete, and it follows from the definition that \mathcal{M}_0 indeed has the properties asserted in the theorem. \square

3 Existence of optimal policies

In order to use the embedding result it is needed that some optimal policy will have a recurrent state within some finite set. The following is a simple condition under which there exists a finite subset Z as required in the theorem presented in the previous section. Suppose that we have an estimate

$$g^*(\mathcal{M}) < \gamma, \quad (3.1)$$

for some γ , and moreover, for some ordering of the states $\{x_j\}$ the following holds:

$$\liminf_{j \rightarrow \infty} \{\min\{c(x_j, a) : a \in A(x_j)\}\} > \gamma. \quad (3.2)$$

It clearly follows from (3.1) and (3.2) that if $J(\sigma) < \gamma$, then some finite set Z contains a recurrent state of σ . Such a set is, e.g.,

$$Z = \{x : \min\{c(x, a) : a \in A(x)\} < \gamma\}. \quad (3.3)$$

An estimate as in (3.1) does not require a computation of the optimal policy, but can be obtained by restricting to a special type of policies.

Fix some state z and denote $s = e^{-1}(z)$. Define $\nu = \inf\{k > 0 : x_k = z\}$, and let ν_0 be defined analogously.

Theorem 3.1 *Fix a stationary Markov policy σ such that $z \in Z$ is recurrent under σ , and let \mathcal{M}_0 be an embedding as above. Let σ_0 be the stationary Markov policy of \mathcal{M}_0 associated with σ . If*

$$C_\nu(z, \sigma) = E_z^\sigma \sum_{k=0}^{\nu-1} c(x_k, a_k)$$

is well defined then

$$\lim_{N \rightarrow \infty} \frac{1}{N} C_N(z, \sigma) = \lim_{N \rightarrow \infty} \frac{1}{N} C_N^0(e^{-1}(z), \sigma_0).$$

Remark 3.2 *Conditions under which the average cost does not depend on the initial state are standard, and therefore we shall not elaborate on this point.*

The proof applies without change when the condition that $C_\nu(z, \sigma)$ is well defined is replaced by the condition that the immediate costs $c(x, a)$ are all nonnegative.

Proof: We denote $s = e^{-1}(z)$, and by construction s is recurrent under σ_0 . We note that s_j is accessible from s under σ_0 if and only if $z_j = e(s_j)$ is accessible from z under σ . Therefore we may assume that all states in Z (resp. Z_0) communicate, and ignore transient states. It is also convenient to ignore (or remove from Z) states from which z can be reached only by leaving Z .

Let the random times ν and ν_0 be as in the sentence that precedes Theorem 3.1. Under the recurrence assumption the limit of the average cost flow exists. It follows from Theorem 17.2.1 of [MT] that under σ the process pair $\{x_k, a_k\}$ possesses an invariant probability measure which we denote by π . Moreover, almost surely under P_z^σ we have

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} c(x_k, a_k) = \frac{E_z^\sigma \sum_{k=0}^{\nu-1} c(x_k, a_k)}{E_z^\sigma \nu} = E_\pi c(x, a),$$

and analogously for \mathcal{M}_0 :

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} c_0(\xi_k, \alpha_k) = \frac{E_s^{\sigma_0} \sum_{k=0}^{\nu_0-1} c_0(\xi_k, \alpha_k)}{E_s^{\sigma_0} \nu_0} = E_{\pi_0} c_0(\xi, \alpha).$$

In view of our assumptions concerning recurrence and existence of cycle costs, this implies that

$$\lim_{N \rightarrow \infty} \frac{1}{N} C_N(z, \sigma) = \frac{E_z^\sigma \sum_{k=0}^{\nu-1} c(x_k, a_k)}{E_z^\sigma \nu}$$

and similarly for C^0 . It therefore suffices to establish that

$$\frac{E_z^\sigma \sum_{k=0}^{\nu-1} c(x_k, a_k)}{E_z^\sigma \nu} = \frac{E_s^{\sigma_0} \sum_{k=0}^{\nu_0-1} c_0(\xi_k, \alpha_k)}{E_s^{\sigma_0} \nu_0}. \quad (3.4)$$

We note that the numerators in (3.4) are the cycle costs corresponding to z and s , assumed to be well defined. We first deal with the numerator in the left-hand side of (3.4). Let I_A denote the indicator of the set A , that is

$$I_A(x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{if } x \notin A. \end{cases}$$

Recalling the definitions of η and τ we have

$$\begin{aligned} & E_z^\sigma \sum_{k=0}^{\nu-1} c(x_k, a_k) \\ &= E_z^\sigma \sum_{k=0}^{\min(\nu, \eta)-1} c(x_k, a_k) + E_z^\sigma I_{\{\eta < \nu\}} \left(\sum_{k=\eta}^{\tau-1} c(x_k, a_k) + \sum_{k=\tau}^{\nu-1} c(x_k, a_k) \right). \end{aligned}$$

By the construction of the embedding,

$$E_z^\sigma \sum_{k=0}^{\min(\nu, \eta)-1} c(x_k, a_k) = E_s^{\sigma_0} \sum_{k=0}^{\min(\nu_0, \eta_0)-1} c(\xi_k, \alpha_k)$$

since, while x_k is in Z , both transition probabilities and immediate costs agree. Also

$$E_z^\sigma I_{\{\eta < \nu\}} \sum_{k=\eta}^{\tau-1} c(x_k, a_k) = E_s^{\sigma_0} I_{\{\eta_0 < \nu_0\}} \sum_{k=\eta_0}^{\tau_0-1} c(\xi_k, \alpha_k)$$

by the definition of the costs $c_0(\omega_i, \alpha)$. Finally, using the Markov property,

$$E_z^\sigma I_{\{\eta < \nu\}} \sum_{k=\tau}^{\nu-1} c(x_k, a_k) = \sum_{z_j \in Z} P_z^\sigma(\eta < \nu, x_\tau = z_j) C_\nu(z_j, \sigma). \quad (3.5)$$

Now write

$$P_z^\sigma(\eta < \nu, x_\tau = z_j) = \sum_{z_i \in Z} P_z^\sigma(\eta < \nu, x_\tau = z_j \mid x_{\eta-1} = z_i) P_z^\sigma(x_{\eta-1} = z_i).$$

Recalling that ν is the return time to state z , we express the first probability on the right-hand side as

$$P_z^\sigma(\eta < \nu, x_\tau = z_j \mid x_{\eta-1} = z_i) = P_z^\sigma(x_t \neq z, 1 \leq t < \eta, x_\tau = z_j \mid x_{\eta-1} = z_i).$$

We now observe that the right hand side describes the conditional probability of two events: one before the conditioning, one after. Since this is a Markov process we have conditional independence and so

$$P_z^\sigma(\eta < \nu, x_\tau = z_j \mid x_{\eta-1} = z_i) = P_z^\sigma(\eta < \nu \mid x_{\eta-1} = z_i) \cdot P_z^\sigma(x_\tau = z_j \mid x_{\eta-1} = z_i).$$

It follows from (3.5) and the above computation that

$$\begin{aligned} & E_z^\sigma I_{\{\eta < \nu\}} \sum_{k=\tau}^{\nu-1} c(x_k, a_k) \\ &= \sum_{z_j, z_i \in Z} P_z^\sigma(\eta < \nu \mid x_{\eta-1} = z_i) \cdot P_z^\sigma(x_\tau = z_j \mid x_{\eta-1} = z_i) \cdot P_z^\sigma(x_{\eta-1} = z_i) \cdot C_\nu(z_j, \sigma). \end{aligned} \quad (3.6)$$

Similarly to (3.5) we have the following expression for σ_0 :

$$E_s^\sigma I_{\{\eta_0 < \nu_0\}} \sum_{k=\tau_0}^{\nu_0-1} c(\xi_k, \alpha_k) = \sum_{s_j \in Z_0} P_s^{\sigma_0}(\xi_{\tau_0} = s_j) C_{\nu_0}^0(s_j, \sigma_0). \quad (3.7)$$

We now repeat the discussion that appears in the text between equations (3.5) and (3.6) for the embedded process. Since all the probabilities and conditional probabilities in (3.6) agree with the corresponding quantities of the embedded process, in view of (3.5) and (3.7) it remains to establish that

$$C_\nu(z_j, \sigma) = C_{\nu_0}^0(s_j, \sigma_0). \quad (3.8)$$

We proceed as before to consider two cases. We compute $C_\nu(z_j, \sigma)$ as follows:

$$\begin{aligned} C_\nu(z_j, \sigma) &= E_{z_j}^\sigma \sum_{k=0}^{\nu-1} c(x_k, a_k) \\ &= E_{z_j}^\sigma \sum_{k=0}^{\min(\nu, \eta)-1} c(x_k, a_k) + E_{z_j}^\sigma \left\{ I_{\{\eta < \nu\}} \sum_{k=\eta}^{\nu-1} c(x_k, a_k) \right\}. \end{aligned} \quad (3.9)$$

The first term once again agrees with the embedded chain. Conditioning on the exit point we can write the second term as

$$E_{z_j}^\sigma \left\{ I_{\{\eta < \nu\}} \sum_{k=\eta}^{\nu-1} c(x_k, a_k) \right\} = \sum_{z_i \in Z} P_{z_j}^\sigma(\eta < \nu, x_{\eta-1} = z_i) C_\nu(z_i, \sigma), \quad (3.10)$$

and similarly

$$E_{s_j}^{\sigma_0} \left\{ I_{\{\eta_0 < \nu_0\}} \sum_{k=\eta_0}^{\nu_0-1} c_0(\xi_k, \alpha_k) \right\} = \sum_{s_i \in Z_0} P_{s_j}^{\sigma_0}(\eta_0 < \nu_0, \xi_{\eta_0-1} = s_i) C_{\nu_0}^0(s_i, \sigma_0). \quad (3.11)$$

By construction

$$P_{z_j}^\sigma(\eta < \nu, x_{\eta-1} = z_i) = P_{s_j}^{\sigma_0}(\eta_0 < \nu_0, \xi_{\eta_0-1} = s_i),$$

and since z is accessible from z_j , $P_{z_j}^\sigma(\eta < \nu) < 1$. Iterating equations (3.9), (3.10) and (3.11) we see that the costs for both models agree, up to a last term that goes to zero geometrically fast with the number of iterations. It follows that the numerators of both sides in (3.4) agree, and the proof for the denominators is similar. Thus (3.4) is established, and the proof of the theorem is complete. \square

Theorem 3.3 *Let \mathcal{M} be a Markov Decision Process, and suppose that the state z is recurrent under a stationary Markov optimal policy σ , and that the cycle cost is finite. Then for any embedding such that $z \in Z$, the optimal cost of \mathcal{M}_0 agrees with that of \mathcal{M} . Moreover, \mathcal{M}_0 has an optimal policy σ_0 that agrees with σ on corresponding states of Z and Z_0 .*

The theorem assumes explicitly that there exists a stationary Markov optimal policy. This holds for most applications: for conditions see for example [P] and references therein.

4 Extensions

First, note that the requirement that the cycle cost is finite holds whenever the immediate costs are bounded, since we assume recurrence. Moreover, as noted above, this requirement is not needed when the immediate costs are all of the same sign (positive or negative).

The following result is immediate, but nonetheless useful.

Theorem 4.1 *Fix some i . Suppose the stationary policies σ and σ' are such that the associated actions $\alpha(\sigma)$ and $\alpha(\sigma')$ have costs starting at ω_i that satisfy*

$$\alpha_{n+1}(\sigma) = c(\sigma) > c(\sigma') = \alpha_{n+1}(\sigma'),$$

while

$$\alpha_i(\sigma) = \alpha_i(\sigma') \text{ for every } i = 1, \dots, n.$$

Then the action $\alpha(\sigma)$ may be eliminated from $A_0(\omega_i)$.

This is quite clear from the definition, and in fact this follows from standard results of action elimination in MDPs [L].

Next note that, even if the excursion costs C_τ are difficult to calculate, any approximation of C_τ and of the mean excursion times leads to a non-exact, approximate embedding, in the sense that optimal costs are not equal anymore. However, it is easy to see that the approximation is continuous in the sense that as the approximations of C_τ and $E\tau$ improve, the costs (including optimal costs) of the embedded model approach those of the original MDP.

We now outline the extension to constrained MDPs, where a detailed description of the model may be found in [A2]. In addition to the usual four

components of an MDP we define a collection of immediate cost functions $\{d_k(x, a), k = 1, \dots, K\}$. Define $J_i^k(\pi)$ in the same way that $J_i(\pi)$ is defined, but with d_k replacing the immediate cost c . The constrained optimization problem is to minimize the functional $J_i(\pi)$, subject to the constraints

$$J_i^k(\pi) \leq V_k$$

for some prescribed constants $V_k, 1 \leq k \leq K$.

Standard approximations of constrained problems are more difficult to handle and establish than unconstrained approximations of optimization problem. The reason for this is that when we require the approximate model to satisfy the hard constraints

$$J_i^{0,k}(\pi) \leq V_k, k = 1, 2, \dots, K,$$

then clearly we may lose continuity, in the sense that even if the original problem is feasible (that is, there exist policies satisfying the constraints), an approximation of the required type may not be feasible [A2]. However, using our exact approximation, this difficulty does not arise.

The embedding results hold for this model, with the following minor modification. Recall the definitions (2.11) and (2.15) of the action in \mathcal{M}_0 . Define $d^k(\sigma)$ as in (2.15) and define $\alpha(\sigma)$ by

$$\alpha = (\alpha_1, \dots, \alpha_n, \alpha_{n+1}, \alpha_{n+2}, \dots, \alpha_{n+K+1}) = (\lambda q_1, \dots, \lambda q_n, c, d^1, \dots, d^K). \quad (4.1)$$

Then the same arguments show that for the embedded chain, all costs agree with those of the original model, so that we may approximate the countable chain by a finite chain.

5 Examples

Consider the problem of controlling a single queue. New jobs arrive according to an i.i.d. sequence of Bernoulli random variables with mean λ , and join an infinite queue. The job at the head of the queue is served, and (independently) the probability of completion of service (representing the speed of service) is the control variable a . Assume that there exists some $I > 0$ and $\mu > \lambda > \delta$ so that

$$A(x) = [\delta, \mu] \text{ for } x < I \text{ and } A(x) = \{\mu\} \text{ for } x \geq I.$$

That is, service rate is controlled for small queue size, but maximal rate must be used if the queue is large. If we assume further that, for $x \geq I$ the immediate cost is sub-linear, that is, it satisfies $c(x, a) = c(x, \mu) \leq c \cdot x$ for some positive c , then all our assumptions hold: state 0 is recurrent under any policy and cycle costs are finite. If in addition $c(x, \mu) = c \cdot x$ for $x \geq I$ then the embedded model \mathcal{M}_0 can be computed explicitly (and easily). Thus, whatever the immediate cost $c(x, a)$, $x \leq I$ the control problem is reduced to a finite-state embedded problem. In fact, in this problem only one ω_i need be defined since the only way to exit the set $Z = \{0, 1, \dots, I\}$ is through state I .

This example can be extended as follows. Suppose we do not assume that the action space is restricted for $x \geq I$. Instead assume that

$$c(x, a) < \min\{c(y, a) : a \in A_y, y \geq I\}$$

for all $x < I$ and all a . It follows that μ is the optimal action at $x \geq I$, and we can apply Theorem 4.1 so that the previous conclusions apply.

In general, since this is a skip-free, one-dimensional problem, our results allows an easy decoupling of the behavior for $x < I$ from that for $x \geq I$. The situation is more complicated if the skip-free assumption is violated, namely either batch arrivals or batch service or both are allowed. However, as is clear from the proof of Theorem 3.3, we can write an implicit expression for the cost using the cost flows until the first hitting time of $\{0, 1, \dots, I - 1\}$, to obtain an explicit expression for the embedded model.

Consider now a multi-dimensional queueing problem. Jobs of type $1, \dots, K$ arrive according to a K -dimensional process $B(t)$ of i.i.d. vectors. The k th coordinate represents arrivals of customers of type k . Customers join infinite queues, one queue for each type. A single server chooses at each point in time which queue to serve, and serves the job at the head of the line. If job of type k is served, then the service will succeed with probability μ_k , and then the job will leave the queue.

If we impose the condition that some queue must be served as long as not all queues are empty, then the empty state will be recurrent under mild conditions. For example, it is sufficient to assume that the total number of arrivals at any unit time interval is bounded, and that the condition

$$\sum_{k=1}^K \frac{EB_k}{\mu_k} < 1 .$$

Let $Q(t)$ be the vector of queue-sizes at time t . This is the state of our MDP, and the control is the choice of queue to serve. This is easily seen to be an MDP, once immediate costs $c(q, a)$ are chosen. It is then natural to choose

$$Z = \left\{ Q : \sum_{k=1}^K Q_k(t) \geq Q_0 \right\}$$

and approximate the infinite model with a finite one.

Suppose for some Q_0 we have that if

$$\sum_{k=1}^K Q_k(t) \geq Q_0$$

then

$$c(x, a) = \sum_{k=1}^K c_k Q_k$$

for some positive coefficients $\{c_k\}$. Then the system simplifies considerably, and the computation of the hitting distributions and costs, required for our approximation, become feasible [BMM, W].

Acknowledgements:

This research was supported in part by the fund for promotion of research at the Technion, and in part by fund for promotion of sponsored research at the Technion.

Adam Shwartz holds the Julius M. and Bernice Naiman Chair in Engineering.

References

- [MT] S.P. Meyn, R.L. Tweedie, *Markov chains and stochastic stability*, Springer Verlag, New York, 1993.
- [A1] E. Altman, Denumerable constrained Markov decision processes and finite approximations *MOR* **19** 161–191, 1994.
- [A2] E. Altman, *Constrained Markov decision processes*, Chapman Hall/CRC, 1999.

- [BMM] J.S. Baras, D.-J. Ma and A.M. Makowski, K competing queues with geometric service requirements and linear costs: the μ -c rule is always optimal, *Systems and Control Letters* **6** 173–180, 1985.
- [C] R. Cavazos-Cadena, Finite state approximations for denumerable state discounted Markov decision processes, *J. Applied Math. and Optimization* **14** 27–47, 1986.
- [GDG] Givan, T. Dean and M. Greig, Equivalence notions and model minimization in markov decision processes, *Artificial Intelligence* **147** 163–223, 2003.
- [L] J.B. Lasserre, Detecting optimal and non-optimal actions in average cost Markov decision processes, *J. Appl. Prob.* 1995.
- [P] M. Puterman *Markov Decision processes*, Elsevier, 1990.
- [W] J. Walrand, *An introduction to queueing networks*. Prentice Hall, NJ 1988.