

A Strongly Polynomial Algorithm for Controlled Queues

Alexander Zadorojnyi, Guy Even

School of Electrical Engineering, Tel-Aviv University, Tel-Aviv 69978, Israel
{sasha@eng.tau.ac.il, <http://www.eng.tau.ac.il/~sasha/>; guy@eng.tau.ac.il, <http://www.eng.tau.ac.il/~guy/>}

Adam Schwartz

Department of Electrical Engineering, Technion, Haifa 32000, Israel,
adam@ee.technion.ac.il, <http://www.ee.technion.ac.il/~adam/>

We consider the problem of computing optimal policies of finite-state finite-action Markov decision processes (MDPs). A reduction to a continuum of constrained MDPs (CMDPs) is presented such that the optimal policies for these CMDPs constitute a path in a graph defined over the deterministic policies. This path contains, in particular, an optimal policy of the original MDP. We present an algorithm based on this new approach that finds this path, and thus an optimal policy. In the general case, this path might be exponentially long in the number of states and actions. We prove that the length of this path is polynomial if the MDP satisfies a coupling property. Thus we obtain a strongly polynomial algorithm for MDPs that satisfies the coupling property. We prove that discrete time versions of controlled $M/M/1$ queues induce MDPs that satisfy the coupling property. The only previously known polynomial algorithm for controlled $M/M/1$ queues in the expected average cost model is based on linear programming (and is not known to be strongly polynomial). Our algorithm works both for the discounted and expected average cost models, and the running time does not depend on the discount factor.

Key words: Markov decision process; constrained Markov decision process; controlled queues; linear programming; $M/M/1$ queue; optimization

MSC2000 subject classification: Primary: 90C40, 68Q25; secondary: 90C05

ORMS subject classification: Primary: Optimal control (Markov: finite state), queues (algorithms, birth death, optimization); secondary: programming (linear)

History: Received July 16, 2008; revised April 19, 2009. Published online in *Articles in Advance* October 7, 2009.

1. Introduction. The problem of designing a strongly polynomial algorithm for finding an optimal policy in a Markov decision process (MDP) has been a long-standing open problem (Blondel and Tsitsiklis [3]). The parameters of an MDP are as follows: n —the number of states, k —the number of actions, and B —the length of the input in bits. In the discounted cost model, there is an additional parameter $\beta < 1$ called the discount factor. Recently, Ye [24] presented a strongly polynomial combinatorial algorithm for the discounted cost model. This algorithm is based on a predictor-corrector interior-point algorithm.

The well-known algorithms for solving MDPs are: value iteration, policy iteration, and linear programming (d'Epenoux [5], Kallenberg [8], Littman et al. [11], Puterman [16]). The running times of the value iteration and policy iteration algorithms in the discounted cost model are polynomial in n , k , B , and $1/(1-\beta)$ (Littman et al. [11], Tseng [22], Ye [24]). The dependence on $1/(1-\beta)$ implies that the algorithm is not strongly polynomial (e.g., when $\beta = 1 - 2^{-n}$). The only nontrivial upper bound on the number of iterations of the policy iteration algorithm (for two actions) that does not depend on the discount factor is $O(2^n/n)$ (Mansour and Singh [13]).

In the expected average cost model, the only polynomial algorithm is based on a reduction discovered nearly 50 years ago to linear programming (de Ghellinck [4], Derman [6], Manne [12]). Linear programming is not known to have strongly polynomial algorithms (Schrijver [18]). Hence the problem of developing a strongly polynomial algorithm for MDPs remains open in the expected average cost model.

1.1. Contribution. We introduce a new approach for solving MDPs in the discounted cost model and expected average cost model. The approach is based on adding an artificial constraint with parameter α to obtain a continuum of constrained MDPs, denoted by $\text{CMDP}(\alpha)$. We consider the whole range of values for α , so that it also includes the value that an optimal policy of the MDP attains. Our approach is based on a new structural lemma that proves that the set of optimal policies of $\text{CMDP}(\alpha)$ (for all values of α) constitutes a path in a graph over the deterministic policies. We present an algorithm that finds all the deterministic policies along the path. The optimal policy of the MDP is simply the min-cost policy along this path. We cannot rule out the possibility that this path may be exponentially long, and hence the running time of this algorithm might be exponential.

We overcome the problem of a long path by introducing a coupling property. We prove that, if the coupling property holds and if a specific artificial constraint is chosen, then the length of the path is polynomial (i.e., $n \cdot k$). Hence the algorithm becomes strongly polynomial. We prove that the coupling property is satisfied in discrete versions of controlled birth-death processes such as single server controlled $M/M/1$ queues. Such controlled

birth-death processes are among the most studied examples of MDPs (Yadin and Naor [23], Altman [1], Kitaev and Rykov [9], Tijms [21]).

When the coupling property holds, the running time of the algorithm is $O(n^4 \cdot k^2)$. This running time holds both in the discounted cost model and the expected average cost model. This compares with the running time of Ye's [24] algorithm, which is $O(n^4 \cdot k^4 \cdot \log(nk/(1-\beta)))$. Thus, in addition to coping with the expected average cost model, we reduce the running time in the discounted cost model.

1.2. Organization. In §2, we briefly overview definitions related to MDPs and CMDPs. In §§3 and 4, we present two properties: uniqueness and coupling. We prove that uniqueness can be obtained by randomly perturbing the cost vector. We prove that the coupling property holds in discrete time controlled $M/M/1$ queues. In §5, we study the structure of optimal policies of $\text{CMDP}(\alpha)$, for all values of α . Lemma 5.4 proves that these optimal policies are a path in a graph over the deterministic policies. In §6, we present a new algorithm for computing an optimal policy of an MDP. In §7, we present a strongly polynomial algorithm that works under the assumption that the coupling property holds. We conclude with a discussion of the assumptions that the MDP is irreducible and satisfies the uniqueness property.

2. Background. In this section, we briefly overview the topics of MDPs, CMDPs, and their linear programming formulations. See Altman [1], Puterman [16], Ross [17], and Tijms [20, 21] for more material on these topics.

2.1. Definition of MDP and CMDP. An MDP is a four-tuple $\langle X, U, P, c \rangle$, where $X = \{0, \dots, n-1\}$ is a finite set of *states*, $U = \{0, \dots, k-1\}$ is a finite set of *actions*, $P: X^2 \times U \rightarrow [0, 1]$ is a *transition probability function*, and $c: X \times U \rightarrow \mathbb{R}$ is a *cost function*. The probability of the transition from state x to state y when the action u is chosen is specified by the function P and denoted by $P(y | x, u)$. The cost associated with selecting the action u when in state x equals $c(x, u)$. We often refer to the cost function as a vector $c \in \mathbb{R}^{nk}$.

An MDP is a generalization of a Markov chain, where in a Markov chain there is only one possible action in each state. For simplicity, we assume that the initial state is fixed and we denote it by x_0 . In fact, Assumption 2.1 implies that, in the expected average cost model, the initial state does not affect the optimal policy. In the discounted cost model, one could assume any initial probability distribution over the states.

Time is discrete, and in each time unit t , let x_t denote the random variable that equals the state at time t . Similarly, let u_t denote the random variable that equals the action selected at time t . The sequence of states $\{x_t\}_{t=1}^\infty$ defines an infinite-random walk over the set of states X .

A (stationary) policy¹ is a function $\pi: X \times U \rightarrow [0, 1]$, such that $\sum_{u \in U} \pi(x, u) = 1$ for every $x \in X$. A policy controls the action selected in each state as follows: the probability of selecting action u in state x equals $\pi(x, u)$. If for a state x and an action u , the policy π satisfies $\pi(x, u) = 1$, then we say that π is *deterministic* in state x . In this case, we abuse notation and write $\pi(x) = u$. If there exists an action u such that $0 < \pi(x, u) < 1$, then we say that π is *randomized* in state x . A *deterministic* policy is a policy that is deterministic in all states.

DEFINITION 2.1. A policy π is *strictly 1-randomized* if: (i) It is deterministic in all states but one state. (ii) Let x denote the state in which π is not deterministic. Then, the set $\{u: \pi(x, u) > 0\}$ contains only two actions.

The goal is to find a policy that minimizes the cost $C(\pi)$ defined below. We consider two cost models: (1) discounted cost and (2) expected average cost.

2.2. Discounted cost model. In the discounted cost model, the parameter $\beta \in (0, 1)$ specifies the rate in which future costs are reduced. Let $P^\pi(x_t = x, u_t = u)$ denote the probability of the event $x_t = x$ and $u_t = u$ when the initial state equals x_0 and the (randomized) policy is π . The expected cost $E_t^\pi[c(x_t, u_t)]$ equals

$$E_t^\pi[c(x_t, u_t)] = \sum_{x \in X, u \in U} c(x, u) \cdot P^\pi(x_t = x, u_t = u).$$

The infinite-horizon discounted expected cost $C(\pi)$ is defined by

$$C(\pi) \triangleq (1 - \beta) \cdot \sum_{t=0}^{\infty} \beta^t \cdot E_t^\pi[c(x_t, u_t)]. \tag{1}$$

¹ By the general theory of MDPs and CMDPs (Puterman [16], Altman [1]), under our conditions, there exists an optimal stationary policy. Therefore we restrict our attention to such policies.

2.3. Expected average cost model. In the expected average cost model, the cost $C(\pi)$ is defined by

$$C(\pi) \triangleq \lim_{T \rightarrow \infty} \left(\frac{\sum_{t=0}^{T-1} E_t^\pi [c(x_t, u_t)]}{T} \right). \quad (2)$$

It can be shown that this limit exists for every stationary policy (Puterman [16]).

2.4. Definition of CMDP. A constrained MDP is an MDP with an additional input consisting of a cost function $d: X \times U \rightarrow \mathbb{R}$ and a parameter α . The cost $D(\pi)$ of π is defined similarly to $C(\pi)$ in both models based on $E_t^\pi [d(x_t, u_t)] = \sum_{x \in X, u \in U} d(x, u) \cdot P^\pi(x_t = x, u_t = u)$. The additional input defines the constraint $D(\pi) = \alpha$ that a feasible policy must satisfy. The optimization problem in $\text{CMDP}(\alpha)$ is to find a policy π that satisfies the constraint $D(\pi) = \alpha$ and minimizes $C(\pi)$.

2.5. Occupation measures. Every policy π induces a probability measure over the state-action pairs. We call this probability measure the *occupation measure* corresponding to π and denote it by ρ_π . The definition of ρ_π depends on the cost model.

In the discounted cost model, $\rho(x, u) \triangleq (1 - \beta) \cdot \sum_{t=0}^{\infty} \beta^t \cdot P^\pi(x_t = x, u_t = u)$. In the expected average cost model, $\rho(x, u) \triangleq \lim_{T \rightarrow \infty} ((\sum_{t < T} P^\pi(x_t = x, u_t = u)) / T)$.

Given an occupation measure $\rho(x, u)$ over $X \times U$, the policy π^ρ induced by ρ is defined by $\pi^\rho(x, u) \triangleq \rho(x, u) / \sum_{u'} \rho(x, u')$. (Note that if $\sum_{u'} \rho(x, u') = 0$, then one may define $\pi^\rho(x, u)$ arbitrarily as long as $\sum_u \pi^\rho(x, u) = 1$.)

We refer to an occupation measure ρ as deterministic (resp. strictly 1-randomized) if $\rho = \rho_\pi$ for a deterministic (resp. strictly 1-randomized) policy π .

2.6. Irreducibility assumption.

DEFINITION 2.2 (IRREDUCIBILITY). An MDP is *irreducible* if every deterministic policy π induces an irreducible Markov chain.

Throughout this paper, we assume the following.

ASSUMPTION 2.1. We assume that the MDP is irreducible.

2.7. Linear programming formulation of CMDPs. In this section, we formulate MDP and $\text{CMDP}(\alpha)$ as linear programs (LPs). We denote the LP corresponding to MDP (resp. $\text{CMDP}(\alpha)$) by LP (resp. $\text{LP}(\alpha)$). The LP is of the form $\min\{c^T \cdot \rho \mid A\rho = b, \rho \geq 0\}$. The $\text{LP}(\alpha)$ is of the form $\min\{c^T \cdot \rho \mid A \cdot \rho = b, d^T \cdot \rho = \alpha, \rho \geq 0\}$ (the transpose of a row vector v is denoted by v^T). The matrix A and the vector b in the LPs depend on the number of states and actions, as well as on the values of the transition probabilities and cost.

Given an MDP $\langle X, U, P, c \rangle$, where $X = \{0, 1, \dots, n-1\}$ and $U = \{0, 1, \dots, k-1\}$, we represent the cost function c as a column vector in \mathbb{R}^{nk} indexed by pairs in $X \times U$; namely, $c_{x,u} = c(x, u)$. We begin with the LP formulation in the discounted cost model.

2.8. Discounted cost model. We define the matrix A as follows. For each action $u \in U$, let $P(u)$ denote the $n \times n$ square matrix whose entries are defined by $P(u)_{y,x} \triangleq P(y \mid x, u)$. The matrix A is an $n \times (nk)$ matrix obtained by concatenating the square matrices $I - \beta P(u)$; namely, $A = [I - \beta P(0), \dots, I - \beta P(k-1)]$. The column vector $b \in \mathbb{R}^n$ is defined by $b = (1 - \beta, 0, \dots, 0)^T$, where the zeroth coordinate corresponds to the initial state.

The occupation measure is the variable of the LP and is represented by the column vector $\rho \in \mathbb{R}^{nk}$ indexed by pairs in $X \times U$. For a state-action pair (x, u) , the component $\rho_{x,u}$ denotes the value of the occupation measure $\rho(x, u)$.

2.9. Expected average cost model. In the expected average cost model, the matrix A is an $(n+1) \times (nk)$ matrix obtained by adding a row $\vec{1}$ consisting of ones to the concatenation of the matrices $I - P(u)$. The vector b is a unit vector, where the coordinate of the one corresponds to the row $\vec{1}$ in A . Note that the constraint $\vec{1} \cdot \rho = 1$ implies that $\rho(x, u)$ is a probability distribution.

The following theorem was proved for various cost models in de Ghellinck [4], d'Epenoux [5], Derman [6], and Manne [12]. A more recent textbook proof appears in Altman [1, Theorem 3.3].

THEOREM 2.1 (EQUIVALENCE OF $\text{CMDP}(\alpha)$ AND $\text{LP}(\alpha)$). $\text{CMDP}(\alpha)$ is feasible if and only if $\text{LP}(\alpha)$ is feasible. Moreover, if ρ^* is an optimal solution of $\text{LP}(\alpha)$, then π^{ρ^*} is an optimal policy of $\text{CMDP}(\alpha)$.

3. The uniqueness property. Consider an MDP and a fixed cost function $d(\cdot)$.

DEFINITION 3.1 (UNIQUENESS). An MDP satisfies the *uniqueness* property if the following holds for every $\alpha \in \mathbb{R}$: If π^* is deterministic and optimal for CMDP(α) and if $\pi \neq \pi^*$ is any stationary policy, then $D(\pi) \neq \alpha$ or $C(\pi) > C(\pi^*)$.

Uniqueness has the following geometric interpretation. Consider the polytope generated by all the deterministic occupation measure (i.e., the feasible solutions of LP). Intersect this polytope with a hyperplane $d^T \cdot \rho = \alpha$ to obtain the feasible solution of LP(α). If this intersection has an optimal solution that is a deterministic occupation measure, then this optimal solution is unique.

The following proposition follows from the fact that every basic feasible solution (BFS) of LP(α) is deterministic or strictly 1-randomized (Theorem 5.1).

PROPOSITION 3.1. *An MDP satisfies the uniqueness property if for every $\alpha \in \mathbb{R}$, every deterministic policy π^* , and every deterministic or strictly 1-randomized policy $\pi \neq \pi^*$, if π^* is optimal for CMDP(α), then π is not optimal for CMDP(α).*

Uniqueness is, in a sense, a generic property; that is, it holds for most values of the parameters. We show this by adding a small random perturbation $\varepsilon \in \mathbb{R}^{nk}$ to the cost vector c to obtain the perturbed cost vector $c_\varepsilon = c + \varepsilon$. Given any positive μ_1 and μ_2 , we choose the components of the vector ε randomly and independently, so that the cost differs from that of the original model by at most μ_2 , and the probability that uniqueness does not hold is at most μ_1 .

Let $C_\varepsilon(\pi)$ denote the cost of a policy π with respect to the perturbed cost vector c_ε . Define each coordinate ε_i of ε by $\varepsilon_i \stackrel{\Delta}{=} (r_i / (2^{p_1})) \cdot 2^{-p_2}$, where p_1, p_2 are positive integers and r_i is uniformly distributed over the set $\{0, \dots, 2^{p_1} - 1\}$. The following lemma proves that a random perturbation meets the requirements, while increasing the length of each component of the cost vector c by $O(n \cdot \log k + \log 1/(\mu_1 \cdot \mu_2))$ bits. This is done by choosing appropriate values for p_1, p_2 .

LEMMA 3.1. *If $p_1 \geq \log_2(k^{3n}/\mu_1)$ and $p_2 \geq \log_2(nk/\mu_2)$, then (1) the uniqueness property holds with probability at least $1 - \mu_1$, and (2) for every policy π , $|C(\pi) - C_\varepsilon(\pi)| \leq \mu_2$.*

PROOF. We prove part (1) as follows. Fix a realization of the vector ε , and suppose that c_ε does not obtain uniqueness for CMDP(α). This implies that there exists a deterministic policy π that is optimal with respect to the perturbed cost c_ε and is not unique. Let ρ_π denote the occupation measure corresponding to π . Since ρ_π is not the only optimal solution of LP(α) (with respect to the perturbed cost vector c_ε), there exists a BFS ρ that is also optimal (with respect to the same perturbed cost vector c_ε). Since both ρ_π and ρ are optimal, it follows that

$$c_\varepsilon \cdot \rho_\pi = c_\varepsilon \cdot \rho. \tag{3}$$

We conclude that the event that perturbation by ε fails implies the existence of an α and a pair $\rho_\pi \neq \rho$ of occupation measures that satisfy (1) $d \cdot \rho_\pi = d \cdot \rho = \alpha$, (2) ρ_π is induced by a deterministic policy π , (3) ρ is a BFS of LP(α), and (4) $c_\varepsilon \cdot \rho_\pi = c_\varepsilon \cdot \rho$. Since ε is random, the quantities c_ε , π , ρ_π , ρ , which depend on ε , are random as well. By the proof of Theorem 5.1, every BFS corresponds to a deterministic or strictly 1-randomized policy.

Let R_α denote the collection of all pairs (ρ_1, ρ_2) of BFS of LP(α) such that ρ_1 corresponds to a deterministic policy. By Theorem 5.1, ρ_2 corresponds to a deterministic or to a strictly 1-randomized policy. Note that R_α does not depend on ε and is not a random set. Let $R = \bigcup_\alpha R_\alpha$.

We claim that $|R| < k^{3n}$. There are k^n deterministic policies, thus we need to consider at most k^n values of α . For each α , there are at most $k^n + \binom{k^n}{2} < k^{2n}$ BFSs of LP(α). Indeed, a BFS is deterministic or strictly 1-randomized. We now bound the number of strictly 1-randomized basic feasible solutions of LP(α). Every strictly 1-randomized policy is a convex combination of two deterministic policies that disagree in a single state (there are less than $\binom{k^n}{2}$ such pairs). For each such pair of deterministic policies, at most one convex combination induces a BFS of LP(α). This follows from Proposition 5.6, because if every convex combination is optimal, then none is an extreme point of LP(α). Therefore the number of strictly 1-randomized BFSs is bounded by $\binom{k^n}{2}$, and $|R| < k^{3n}$ as claimed.

Consider a pair $(\rho_1, \rho_2) \in R$. Without loss of generality, ρ_1 and ρ_2 disagree in the first coordinate. Let c_ε^1 denote the first coordinate of c_ε , and let c_ε^{-1} denote the vector c_ε with the first coordinate removed, so that $c_\varepsilon = (c_\varepsilon^1, c_\varepsilon^{-1})$. We use identical notation for any vector. The equation $c_\varepsilon \rho_1 = c_\varepsilon \rho_2$ implies that

$$c_\varepsilon^1 \rho_1^1 + c_\varepsilon^{-1} \cdot \rho_1^{-1} = c_\varepsilon^1 \rho_2^1 + c_\varepsilon^{-1} \cdot \rho_2^{-1}.$$

Now,

$$P(c_\varepsilon \cdot \rho_1 = c_\varepsilon \cdot \rho_2) = P(c_\varepsilon^{-1} \cdot (\rho_1^1 - \rho_2^1) = c_\varepsilon^{-1} \cdot (\rho_2^{-1} - \rho_1^{-1})) \leq 2^{-p_1}.$$

The last line follows from the fact that, given ρ_1, ρ_2 , and ε^{-1} , the event $c_\varepsilon^{-1} \cdot (\rho_1^1 - \rho_2^1) = c_\varepsilon^{-1} \cdot (\rho_2^{-1} - \rho_1^{-1})$ occurs for at most one value of ε^1 .

We now bound the probability that perturbation fails; namely, Equation (3) holds. Since $(\rho_\pi, \rho) \in R$,

$$P(c_\varepsilon \cdot \rho_\pi = c_\varepsilon \cdot \rho) \leq P(c_\varepsilon \cdot \rho_1 = c_\varepsilon \cdot \rho_2 \text{ for some } (\rho_1, \rho_2) \in R) \leq \sum_{(\rho_1, \rho_2) \in R} P(c_\varepsilon \cdot \rho_1 = c_\varepsilon \cdot \rho_2) \leq k^{3n} 2^{-p_1}.$$

We conclude that if $p_1 \geq \log_2(k^{3n}/\mu_1)$, then the probability of nonuniqueness is bounded by μ_1 .

Part (2) requires that the perturbation does not change the cost of the optimal policy by more than μ_2 . It suffices to show that, for every occupation measure ρ , $|(c_\varepsilon - c) \cdot \rho| \leq \mu_2$. Since ρ is an occupation measure, it follows that $|(c_\varepsilon - c) \cdot \rho| \leq \sum_i \varepsilon_i \leq n \cdot k \cdot 2^{-p_2}$. Hence, part (2) holds if $p_2 \geq \log_2(nk/\mu_2)$. \square

In the light of Lemma 3.1, we assume the following throughout this paper.

ASSUMPTION 3.1. *The MDP satisfies the uniqueness property.*

4. The coupling property.

DEFINITION 4.1. Two deterministic policies are *neighbors* if they disagree in a single state.

DEFINITION 4.2. Given a deterministic policy π and an action $j \neq \pi(i)$, the neighbor policy $\pi^{i,j}$ is defined by

$$\forall x \in X: \pi^{i,j}(x) \triangleq \begin{cases} j & \text{if } x = i, \\ \pi(x) & \text{otherwise.} \end{cases}$$

Thus two deterministic policies π and τ are *neighbors* if there exists a state i and an action j such that $\tau = \pi^{i,j}$.

Suppose that for every state i , there is a linear order over U . We denote the linear order over U corresponding to state i by \leq_i . In addition, we consider the natural linear order over the set of states $X = \{0, \dots, n - 1\}$.

The polynomial algorithm in §7 for finding an optimal policy depends on a property that we call the coupling property defined below.

DEFINITION 4.3 (COUPLING PROPERTY). The *coupling property* holds with respect to the linear orders $\{\leq_i\}_{i \in X}$ if for every deterministic policy π , every state i , and every action j ,

$$\pi(i) \leq_i j \Rightarrow \forall x < i \quad \forall u \in U: \rho_\pi(x, u) \leq \rho_{\pi^{i,j}}(x, u).$$

4.1. Examples of MDPs with the coupling property. In this section, we present a “one-dimensional” MDP, and prove that it satisfies the coupling property in the expected average cost model. We begin with a controlled nonabsorbing random walk. We then continue with a one-dimensional MDP that corresponds to a discrete-time controlled $M/M/1$ queue.

4.1.1. A controlled nonabsorbing random walk. A controlled nonabsorbing random walk is a simple example of an MDP that satisfies the coupling property. We formally describe it below.

The MDP has n states $\{0, \dots, n - 1\}$. For $i < n - 1$, there is a transition from state i to state $i + 1$ with probability $P(i + 1 | i, j) \in (0, 1)$. For $i > 0$, there is a transition from state i to state $i - 1$ with probability $P(i - 1 | i, j) = 1 - P(i + 1 | i, j)$. For $i = 0$, there is a self-loop $P(0 | 0, j) = 1 - P(1 | 0, j)$, and similarly, for state $n - 1$, there is a self-loop $P(n - 1 | n - 1, j) = 1 - P(n - 2 | n - 1, 0)$.

We assume that all $P(i + 1 | i, j)$ transition probabilities are positive. Hence the MDP is irreducible.

The linear orders \leq_i are defined as follows for each state $i \geq 1$:

$$j' \leq_i j'' \Leftrightarrow P(i - 1 | i, j') \leq P(i - 1 | i, j'').$$

Namely, the transition from state i to its left neighbor $i - 1$ is not more likely under the action j' than under the action j'' . The linear order \leq_i is defined arbitrarily for $i = 0$.

The proof of the following lemma appears in Appendix B.

LEMMA 4.1. *The coupling property holds for a controlled nonabsorbing random walk.*

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

4.1.2. A controlled discrete-time $M/M/1$ queue. We now consider a discrete-time version of a controlled $M/M/1$ queue obtained from a continuous-time controlled $M/M/1$ queue by a technique called uniformization (see Appendix A). A discrete controlled $M/M/1$ queue is similar to the controlled nonabsorbing random walk with the addition of self-loops in each state. Formally, the set of states is $\{0, \dots, n-1\}$. For $i < n-1$, there is a transition from state i to state $i+1$ with probability $P(i+1 | i, j) \in (0, 1)$. For $i > 0$, there is a transition from state i to state $i-1$ with probability $P(i-1 | i, j) \in (0, 1)$. In addition, for every state i , there is a self-loop with probability $P(i | i, j)$. Assumption 2.1 holds by the reduction from the continuous $M/M/1$ queue.

We assume that the actions do not affect the arrival rates, hence the probabilities $P(i+1 | i, j)$ do not depend on the action j . Moreover, the reduction from an $M/M/1$ queue implies that, for all states i , the transitions from state i to state $i+1$ have the same probability. We therefore denote $P(i+1 | i, j)$ simply by q . This means that the control only affects the service rates, and hence only the probabilities $P(i-1 | i, j)$ and $P(i | i, j)$ depend on the action j .

For each state $i \geq 1$, the linear order \leq_i in the discrete controlled $M/M/1$ queue is defined as follows:

$$j' \leq_i j'' \Leftrightarrow P(i-1 | i, j') \leq P(i-1 | i, j'').$$

We prove the following lemma for the expected average cost model. The same lemma can be proved if the control affects the arrival rates and does not affect the service rate.

The proof of the following lemma appears in Appendix B.

LEMMA 4.2. *The coupling property holds for the controlled discrete-time $M/M/1$ queue.*

5. Structure of optimal policies.

5.1. Deterministic policies. Given a policy π , let I_π denote the set of pairs (i, j) for which $\pi(i, j) > 0$. These pairs define columns of the matrix A . Let B_π denote the submatrix of A consisting of the projection of A to the columns in I_π . Let ρ_π denote the occupation measure corresponding to the policy π . Let $\tilde{\rho}_\pi$ denote the vector obtained by projecting ρ_π to coordinates in I_π .

The next proposition proves that, under Assumption 2.1, the mapping $\pi \mapsto \rho_\pi$ between deterministic policies and the corresponding occupation measure is one to one.

PROPOSITION 5.1. *If π is a deterministic policy for $\text{CMDP}(\alpha)$, then (i) $\tilde{\rho}_\pi$ is the unique solution of the equations $B_\pi \cdot \tilde{\rho} = b$, and (ii) the rank of B_π is n .*

PROOF. Part (ii) follows from part (i). We now prove part (i). By the definition of I_π , if $(x, u) \notin I_\pi$, then $\rho_\pi(x, u) = 0$. Hence $A \cdot \rho_\pi = b$ if and only if $B_\pi \cdot \tilde{\rho}_\pi = b$. In the model of discounted cost, the matrix B_π is invertible by Gersgorin's theorem (Horn and Johnson [7]), hence uniqueness follows.

In the model of expected average cost, if the MDP satisfies the Assumption 2.1, then by the Peron-Frobenius theorem (Horn and Johnson [7]), the system $B_\pi \cdot \tilde{\rho}_\pi = b$ has a unique solution, and the proposition follows. \square

5.2. Properties of optimal policies. The following theorem, proved for the various cost models in de Ghellinck [4], d'Epenoux [5], Derman [6], and Manne [12], states that, if $\text{CMDP}(\alpha)$ is feasible, there always exists an optimal policy that is deterministic or strictly 1-randomized. The theorem is stated in terms of the occupation measure (i.e., the optimal solution of the $\text{LP}(\alpha)$). This theorem and its proof are an extension of the theorem that every MDP has an optimal policy that is deterministic.

THEOREM 5.1. *If $\text{LP}(\alpha)$ is feasible, then there exists an optimal solution ρ^* of $\text{LP}(\alpha)$ that is deterministic or strictly 1-randomized.*

PROOF. The rank of the constraints in $\text{LP}(\alpha)$ is at most $n+1$. This implies that in every BFS there are at most $n+1$ nonzero variables. Fix an optimal BFS ρ^* . By Assumption 2.1, $\sum_u \rho^*(x, u) > 0$, for each state x . Hence, for each state x , except perhaps for one, $\rho(x, u)$ is positive for exactly one action, and the theorem follows. \square

5.3. Policies along an edge. Let π^0 and π^1 denote two deterministic policies that disagree in a single state. Let $\pi^q \triangleq q \cdot \pi^1 + (1-q) \cdot \pi^0$, for $0 \leq q \leq 1$. Note that π^q is a strictly 1-randomized policy if $0 < q < 1$. We say that a policy π agrees with the zeros of policy π^* if $\pi(x, u) = 0$ whenever $\pi^*(x, u) = 0$.

Let $A^{x,u}$ denote the column of A corresponding to $x \in X$ and $u \in U$. Complementary slackness implies the following optimality condition.

PROPOSITION 5.2. *Let ρ and w denote feasible solutions of $\text{LP}(\alpha)$ and the dual-linear program (DLP), respectively. The following two conditions are equivalent: (1) ρ and w are optimal and (2) For every $x \in X$ and $u \in U$, either $\rho(x, u) = 0$ or the dual constraint is tight (i.e., $w^T \cdot A^{x,u} = c(x, u)$).*

PROPOSITION 5.3 (ZADOROJNIY AND SHWARTZ [25]). *Let π^* denote an optimal policy for $\text{CMDP}(\alpha^*)$. Let π denote a policy that agrees with the zeros of π^* . Then, π is an optimal policy for $\text{CMDP}(D(\pi))$.*

PROOF. Let $\rho^* = \rho_{\pi^*}$ and $\rho = \rho_{\pi}$. Note that $\rho^*(x, u) = 0$ implies that $\rho(x, u) = 0$. Let w^* denote a dual-optimal solution of $\text{LP}(\alpha^*)$. By Proposition 5.2, it follows that, for every (x, u) , either $\rho^*(x, u) = 0$ or the dual constraint is tight (i.e., $(w^*)^T \cdot A^{x,u} = c(x, u)$). Note that w^* is also a feasible solution of the DLP corresponding to $\text{CMDP}(D(\pi))$. It follows that ρ and w^* also satisfy the optimality condition, and hence, by Proposition 5.2, ρ is optimal, as required. \square

PROPOSITION 5.4. *For every two policies π' and π'' such that $D(\pi') < D(\pi'')$, there exists a policy π such that $D(\pi') < D(\pi) < D(\pi'')$.*

PROOF. Denote $\pi^q \triangleq q \cdot \pi'' + (1 - q) \cdot \pi'$. Define $D(q) \triangleq D(\pi^q)$. Since $D(\pi)$ is continuous in π (Zadorojniy and Schwartz [25]), it follows that $D(q)$ is continuous in q . It follows that the image of $D(q)$ over the interval $[0, 1]$ contains the interval $[D(\pi'), D(\pi'')]$. \square

Proposition 5.3 and the proof of Proposition 5.4 imply the following.

COROLLARY 5.1 (ZADOROJNIY AND SHWARTZ [25]). *If π^{q^*} is an optimal policy for $\text{CMDP}(\alpha^*)$ and $q^* \in (0, 1)$, then, for each $\alpha \in [\inf_{0 \leq q \leq 1} D(\pi^q), \sup_{0 \leq q \leq 1} D(\pi^q)]$, there exists $q_\alpha \in [0, 1]$ such that π^{q_α} is an optimal policy for $\text{CMDP}(\alpha)$.*

Consider the strictly 1-randomized policy $\pi^{1/2} = (\pi^0 + \pi^1)/2$. Then $I_{\pi^{1/2}}$ is the set of pairs (i, j) for which $\pi^{1/2} > 0$.

Let $B(d)$ denote the $(n + 1) \times (n + 1)$ square matrix obtained by first augmenting the matrix A with the row d^T , and then projecting the augmented matrix on the columns in $I_{\pi^{1/2}}$.

LEMMA 5.1. *The following three conditions are equivalent:*

- (i) $D(\pi^0) = D(\pi^1)$.
- (ii) $B(d)$ is not of full rank.
- (iii) $D(\pi^q) = D(\pi^0)$, for all $q \in [0, 1]$.

PROOF. (i) \Rightarrow (ii). Fix $\alpha = D(\pi^0)$. The occupation measures ρ_{π^0} and ρ_{π^1} (induced by the deterministic policies π^0 and π^1 , respectively) are distinct feasible solutions of $\text{LP}(\alpha)$. Hence, both $\tilde{\rho}_{\pi^0}$ and $\tilde{\rho}_{\pi^1}$ are distinct solutions of the system of equations $B(d) \cdot \tilde{\rho} = \begin{pmatrix} b \\ \alpha \end{pmatrix}$. This implies that $B(d)$ is not of full rank.

(ii) \Rightarrow (iii). Both policies π^0 and π^1 induce occupation measures that are feasible solutions of LP. By convexity, for every $q \in [0, 1]$, the occupation measure ρ_{π^q} is also a feasible solution of LP. Since B has rank n , if $B(d)$ is not of full rank, the last row (corresponding to the constraint $d^T \cdot \rho = \alpha$) depends on the other rows. Hence, every occupation measure ρ that is a feasible solution of LP and whose support is contained in $I_{\pi^{1/2}}$ has the same cost $d^T \cdot \rho$. This implies that $D(\pi^q) = D(\pi^0)$, for all $q \in [0, 1]$, as required. Finally, the implication (iii) \Rightarrow (i) is trivial, and the lemma follows. \square

PROPOSITION 5.5. *If $D(\pi^0) \neq D(\pi^1)$, then $C(\pi^q)$ is linear in $D(\pi^q)$ over the range $q \in [0, 1]$.*

PROOF. We consider the following two cases:

(i) Suppose $B(d)$ is of full rank. In the model of discounted cost, $B(d)$ is an $(n + 1) \times (n + 1)$ square matrix, and thus invertible. Hence $\tilde{\rho}_{\pi^q} = B(d)^{-1} \cdot (b, D(\pi^q))^T$. Therefore, $C(\pi^q) = \tilde{c} \cdot \tilde{\rho}_{\pi^q} = \tilde{c} \cdot B(d)^{-1} \cdot (b, D(\pi^q))^T$, and $C(\pi^q)$ is linear in $D(\pi^q)$, as required. In the model of expected average cost, one needs to remove first a dependent row from $B(d)$ to make it square and thus invertible.

(ii) If $B(d)$ is not of full rank, then by Lemma 5.1, $D(\pi^0) = D(\pi^1)$, a contradiction. \square

PROPOSITION 5.6. Fix a value of α . Consider the set of policies $E_{(0,1)} \triangleq \{\pi^q: 0 < q < 1\}$. Exactly one of the following cases holds:

- (i) Every policy in $E_{(0,1)}$ is an optimal policy of $\text{CMDP}(\alpha)$.
- (ii) No policy in $E_{(0,1)}$ is an optimal policy of $\text{CMDP}(\alpha)$.
- (iii) Exactly one policy in $E_{(0,1)}$ is an optimal policy of $\text{CMDP}(\alpha)$.

PROOF. If $B(d)$ is of full rank, then by Proposition 5.5, either exactly one policy in $E_{(0,1)}$ is an optimal policy of $\text{CMDP}(\alpha)$ or no policy in $E_{(0,1)}$ is an optimal policy of $\text{CMDP}(\alpha)$.

If $B(d)$ is not of full rank, then by Lemma 5.1, $D(\pi^q) = D(\pi^0)$, for $q \in [0, 1]$, and thus either every policy in $E_{(0,1)}$ is a feasible policy of $\text{CMDP}(\alpha)$ or no policy in $E_{(0,1)}$ is a feasible policy of $\text{CMDP}(\alpha)$. By Proposition 5.3, if one policy in $E_{(0,1)}$ is an optimal policy of $\text{CMDP}(\alpha)$, then every policy in $E_{(0,1)}$ is optimal as well. \square

In the following lemmas, we abbreviate, and refer to a policy π as optimal if it is an optimal policy of $\text{CMDP}(D(\pi))$.

LEMMA 5.2. Let $q^* \in (0, 1)$. If π^{q^*} is an optimal strictly 1-randomized policy, then the function $D(q) \triangleq D(\pi^q)$ is strictly monotone in the interval $q \in [0, 1]$.

PROOF. The function $D(q)$ is continuous because the policy π^q is continuous in q , and $D(\pi)$ is continuous in π . If $D(q)$ is not strictly monotone, then there exist $q' < q''$ such that $D(q') = D(q'')$. By Proposition 5.3, each of the policies $\pi^{q'}$ and $\pi^{q''}$ is optimal for $\text{CMDP}(\alpha)$, where $\alpha = D(q')$. By the uniqueness assumption (Assumption 3.1), neither $\pi^{q'}$ or $\pi^{q''}$ is deterministic. Hence $0 < q' < q'' < 1$.

Let ρ' (resp. ρ'') denote the occupation measure that corresponds to the policy $\pi^{q'}$ (resp. $\pi^{q''}$). We first prove that $\rho' \neq \rho''$. Assume that π^0 and π^1 disagree in state s , and without loss of generality, assume that $\pi^0(s) = 0$ and $\pi^1(s) = 1$. By Assumption 2.1, both occupation measures ρ' and ρ'' assign positive probability to state s . However, the ratios $\rho'(s, 0)/\rho'(s, 1) \neq \rho''(s, 0)/\rho''(s, 1)$.

On the other hand, since the support of ρ' and ρ'' are equal, it follows that the bases corresponding to ρ' and ρ'' are the same. Hence ρ' and ρ'' are different solutions of the system $\tilde{B} \cdot \rho = \begin{pmatrix} b \\ \alpha \end{pmatrix}$, where \tilde{B} is the basis matrix. We consider two cases. If \tilde{B} is invertible, then immediately we have a contradiction. If \tilde{B} is not invertible, then by Lemma 5.1, $D(\pi^0) = D(\pi^1) = \alpha$. Therefore, both π^0 and π^1 are feasible policies of $\text{CMDP}(\alpha)$. On the other hand, both π^0 and π^1 are optimal, hence $C(\pi^0) = C(\pi^1)$, a contradiction to the uniqueness assumption (Assumption 3.1). \square

LEMMA 5.3. Let $\pi' \neq \pi''$ denote two distinct optimal policies of $\text{CMDP}(\alpha')$ and $\text{CMDP}(\alpha'')$, respectively. If π' and π'' are deterministic or strictly 1-randomized, then $\alpha' \neq \alpha''$.

PROOF. Assume for the sake of contradiction that $D(\pi') = D(\pi'')$. Recall that by definition $\alpha' = D(\pi')$ and $\alpha'' = D(\pi'')$. Since both π' and π'' are optimal, it follows that $C(\pi') = C(\pi'')$. If either π' or π'' is deterministic, then the lemma follows from the uniqueness assumption. If both policies are strictly 1-randomized, then let π' (resp. π'') be a convex combination of two deterministic policies π^0 and π^1 (resp. τ^0 and τ^1). By Lemma 5.2, D increases along the edge between π^0 and π^1 (resp. τ^0 and τ^1). Without loss of generality, $D(\tau^0) \leq D(\pi^0) \leq D(\pi')$. It follows that Assumption 3.1 is violated for $\alpha = D(\pi^0)$. \square

5.4. Graph representation.

DEFINITION 5.1 (POLICY GRAPH). The policy graph is a graph $G = (V, E)$, where V is the set of deterministic policies, and E is the set of pairs of neighboring deterministic policies (i.e., policies that disagree in exactly one state).

In the case of two actions $k = 2$, the policy graph is isomorphic to the n dimensional hypercube. In the general case, the policy graph is isomorphic to the Cartesian product of n copies of the complete graph over k vertices.

We consider the edge (π^0, π^1) between neighboring deterministic policies as a representation of all convex combinations $\pi^q = (1 - q) \cdot \pi^0 + q \cdot \pi^1$ of π^0 and π^1 . In such a case, we say that π^q belongs to the edge (π^0, π^1) .

Let Γ denote the set of deterministic or strictly 1-randomized feasible policies for $\text{CMDP}(\alpha)$, for all values of α . Let $\Gamma^* \subseteq \Gamma$ denote the subset of optimal policies in Γ . In light of Proposition 5.3, Γ^* consists of vertices (i.e., deterministic policies) and edges (i.e., deterministic and strictly 1-randomized policies).

LEMMA 5.4. The set Γ^* is a path in the policy graph G .

PROOF. Let G^* denote the subgraph of G that consists of the vertices and edges in Γ^* . The proof consists of the following two stages: (1) prove that G^* is connected and (2) prove that the degree of every vertex in G^* is at most two.

Denote the connected components of G^* by U_1, U_2, \dots, U_s . By continuity, the image of the function $D(\cdot)$ over each connected component is an interval. Denote the image of U_i by I_i . By Lemma 5.3, the intervals I_1, \dots, I_s are pairwise disjoint. By Proposition 5.4, the union of the intervals $I_1 \cup \dots \cup I_s$ is an interval. To avoid a contradiction, we conclude that G^* contains only a single connected component. Hence G^* is connected, as required.

If the degree of a vertex π is at least 3, consider three edges in Γ^* that are incident to π . By Lemma 5.2, $D(\pi)$ is strictly monotone as one travels along each of these edges incident to v . Moreover, for at least two edges, the slope of $D(\pi)$ as one approaches v has the same sign; namely, monotone increasing (or decreasing). Two such edges in Γ^* contain two optimal policies $\pi' \neq \pi'' \in \Gamma^*$ such that $D(\pi') = D(\pi'')$. This contradicts Lemma 5.3. \square

The next corollary follows from Lemmas 5.3 and 5.4.

COROLLARY 5.2. $D(\pi)$ is strictly monotone along the path Γ^* .

6. A general algorithm. In this section, we present a general algorithm for computing optimal policies of irreducible MDPs that satisfy the uniqueness property. Although we cannot prove that the running time of this algorithm is polynomial in general, in the next section, we prove strong polynomiality of a variant when the coupling property holds.

6.1. Geometric interpretation of the algorithm. The algorithm is based on Lemma 5.4 that states that the set Γ^* of optimal deterministic and strictly 1-randomized policies form a path in the policy graph. Consider the polytope P generated by the deterministic occupation measures. We introduce a cost vector d . Let P_α denote the intersection of P with the hyperplane $d^T \cdot \rho = \alpha$. Let α_{\min} (resp. α_{\max}) denote the minimum (resp. maximum) value of α for which P_α is not empty. For each $\alpha \in [\alpha_{\min}, \alpha_{\max}]$, the polytope P_α contains a single occupation measure $\rho(\alpha)$ that corresponds to a policy $\pi(\alpha) \in \Gamma^*$.

The algorithm assigns a zero-one cost vector d , so that $\alpha_{\min} = 0$ and $\alpha_{\max} = 1$. Moreover, it is trivial to find the optimal deterministic policy π such that $D(\pi) = 0$. Given a prefix of Γ^* ending in a deterministic policy π , the algorithm finds the next deterministic policy τ along Γ^* as follows. First, note that τ must be a neighbor of π . Namely, there exists a state i and an action j such that $\tau = \pi^{i,j}$. This limits the number of candidates for τ to nk . Second, by Corollary 5.2, $D(\tau) > D(\pi)$. Thus, if we depict the neighboring policies on a (C, D) -plane (see Figure 1), then τ is simply the policy with the smallest slope.

The algorithm ends when all neighbors τ of π satisfy $D(\tau) \leq D(\pi)$. Thus the algorithm has reached the last policy along Γ^* .




FIGURE 1. A prefix of Γ^* ends in a deterministic policy π .

Notes. The algorithm has to compute the next policy along Γ^* among the neighboring policies τ_1, τ_2, τ_3 . The costs $C(\cdot)$ and $D(\cdot)$ of each policy are depicted in the graph. The algorithm chooses the policy τ_3 because the segment between $(D(\pi), C(\pi))$ and $(D(\tau_3), C(\tau_3))$ has the smallest slope.

6.2. Notation. Given a deterministic policy π , we define the gradient $\nabla_{i,j}$ as follows:

$$\nabla_{i,j} \triangleq \frac{C(\pi^{i,j}) - C(\pi)}{D(\pi^{i,j}) - D(\pi)}. \quad (4)$$

The parameters in the definition of $\nabla_{i,j}$ can be easily computed as follows. Recall that B_π denotes the projection of the columns of the matrix A on the pairs in I_π (i.e., the basis matrix corresponding to the BFS ρ_π). For a vector ρ_π , the projection to the coordinates in I_π is denoted by $\tilde{\rho}_\pi$. Since π is a deterministic policy, by Proposition 5.1, the corresponding occupation measure ρ_π when projected to I_π is the unique solution for $B_\pi \cdot \tilde{\rho}_\pi = b$. Hence $C(\pi) = \tilde{c}_\pi \cdot \tilde{\rho}_\pi$, $D(\pi) = \tilde{d}_\pi \cdot \tilde{\rho}_\pi$, and the analogous computations hold for $C(\pi^{i,j})$ and $D(\pi^{i,j})$.

6.3. Algorithm description. The algorithm adds a new artificial cost function $D(\pi)$ specified by a cost vector $d \in \{0, 1\}^{nk}$. The MDP with the constraint $D(\pi) = \alpha$ is denoted by CMDP(α). In the linear programming formulation, LP(α) is the LP obtained by adding the constraint $d^T \cdot \rho = \alpha$ to LP. The algorithm computes the set Γ^* of optimal (deterministic or strictly 1-randomized) policies for CMDP(α), for every value of α . This set is found by computing Γ^* . Finally, an optimal policy for the MDP is chosen as a deterministic policy in Γ^* with minimum cost $C(\cdot)$.

A listing of the algorithm appears as Algorithm 1. In line 1, the algorithm assigns zero-one costs $d(i, j)$. For each state, one (arbitrary) action is assigned zero cost, and the other actions are assigned unit cost. In line 2, the initial policy is set. This policy simply chooses the zero cost action for each state. This initial policy π achieves the minimum value for $D(\pi)$. The path p begins with the initial policy as its starting point.

The algorithm builds the path p by adding a new edge in each iteration of the while loop. The last policy (vertex) added to p is denoted by π . In each iteration of the while loop, the path p is augmented by a new edge $(\pi, \pi^{i,j})$. In line 4, this new edge $(\pi, \pi^{i,j})$ is chosen such that $(i, j) = \arg \min\{\nabla_{i,j} \mid \forall (i, j) \text{ such that } D(\pi^{i,j}) > D(\pi)\}$. In line 5, the new edge is added to the path p . In line 6, the new end point of p is updated. In lines 7–8, the minimum cost policy along p is updated if necessary. In line 11, a minimum cost policy is returned.

Algorithm 1

A heuristic for finding an optimal policy for the MDP $\min\{c \cdot \rho \mid A \cdot \rho = b\}$. We assume that the MDP is irreducible and satisfies the uniqueness property.

(1) Define

$$d(i, j) \triangleq \begin{cases} 0 & \text{if } j = 0, \\ 1 & \text{otherwise.} \end{cases}$$

(2) Initialize

$$\begin{aligned} \pi &\leftarrow (0, \dots, 0) && \{\pi \text{ chooses the “zero” action in each state}\} \\ \text{opt} &\leftarrow \pi && \{\text{best policy so far}\} \\ p &\leftarrow \{\pi\} && \{\text{path } p \text{ starts with } \pi\} \end{aligned}$$

- (3) **while** exists (i, j) such that $D(\pi^{i,j}) > D(\pi)$, **do**
- (4) $(i, j) \leftarrow \arg \min\{\nabla_{i,j} \mid \forall (i, j) \text{ such that } D(\pi^{i,j}) > D(\pi)\}$
- (5) add the edge $(\pi, \pi^{i,j})$ to p
- (6) $\pi \leftarrow \pi^{i,j}$ $\{\pi^{i,j}$ becomes the current end point of $p\}$
- (7) **if** $C(\pi^{i,j}) < C(\text{opt})$, **then**
- (8) $\text{opt} \leftarrow \pi^{i,j}$ $\{\text{opt is the best policy so far}\}$
- (9) **end if**
- (10) **end while**
- (11) **return** opt

6.4. Correctness. We now prove that Algorithm 1 finds an optimal policy. To prove this, we prove that the algorithm computes Γ^* , the path of optimal solutions of LP(α) (for all values of α) in the policy graph.

THEOREM 6.1. *The path p computed by the Algorithm 1 equals Γ^* .*

PROOF. We prove by induction on the number of iterations of the while loop that p is a prefix of Γ^* in each iteration. Since the costs $d(x, u)$ are in $\{0, 1\}$, it follows that for every policy τ , $D(\tau) \geq 0$. Hence $\text{LP}(\alpha)$ is feasible only if $\alpha \geq 0$. Clearly, the initial policy $\pi_0 = (0, \dots, 0)$ satisfies $D(\pi_0) = 0$. We claim that the initial policy is the only policy with $D(\pi) = 0$. Consider an optimal policy $\pi \neq \pi_0$. Consider a state x and action u for which $\pi(x, u) > 0$ while $\pi_0(x, u) = 0$. By the Assumption 2.1, $\rho_\pi(x, u) > 0$. Since $d(x, u) = 1$, it follows that $D(\pi) > 0$. We conclude that the initial policy is optimal for $\alpha = 0$. Moreover, the initial policy is the end point of the path Γ^* with smallest cost $D(\cdot)$, and the induction basis holds.

The induction step is proved as follows. Let π denote the last policy added to p . Let $\pi^{i,j}$ denote the next policy added to p ; namely, $(i, j) \leftarrow \arg \min\{\nabla_{i,j} \mid \forall (i, j) \text{ such that } D(\pi^{i,j}) > D(\pi)\}$. Let $\pi^{\hat{i},\hat{j}}$ denote the next policy along Γ^* after π . We wish to prove that $(i, j) = (\hat{i}, \hat{j})$.

Assume for the sake of contradiction that $(i, j) \neq (\hat{i}, \hat{j})$. By Corollary 5.2, $D(\pi^{i,j}) > D(\pi)$. Let $D' = \min\{D(\pi^{\hat{i},\hat{j}}), D(\pi^{i,j})\}$. Since the cost $D(\tau)$ is a continuous function of the policy τ , the cost D' is obtained in two policies: (1) π_1 along the edge between π and (2) $\pi^{i,j}$ and π_2 along the edge between π and $\pi^{\hat{i},\hat{j}}$. For example, $\pi_1 = \pi^{i,j}$ and π_2 is a convex combination of π and $\pi^{\hat{i},\hat{j}}$. The policy π_2 is also an optimal policy. However, by Proposition 5.5 and the definition of (i, j) , it follows that $C(\pi_1) \leq C(\pi_2)$. This contradicts the optimality of π_2 (if $C(\pi_1) < C(\pi_2)$) or the uniqueness of the solution (if $C(\pi_1) = C(\pi_2)$). Hence $(i, j) = (\hat{i}, \hat{j})$, which completes the induction step.

We now prove that when the algorithm terminates, p cannot be augmented anymore, and hence p equals Γ^* , as required. Indeed, if p is a proper prefix of Γ^* , then the cost $D(\cdot)$ increases from π to the next deterministic policy in Γ^* . In this case, the algorithm would not have terminated yet because the set $\{(i, j): D(\pi^{i,j}) > D(\pi)\}$ is not empty. \square

COROLLARY 6.1. *Algorithm 1 computes an optimal policy of the MDP.*

PROOF. The MDP has an optimal policy π^* that is deterministic. This policy is also in Γ^* . By Theorem 6.1, π^* appears in the sequence of policies scanned by the algorithm. \square

Let $|\Gamma^*|$ denote the number of deterministic policies in Γ^* .

PROPOSITION 6.1. *The complexity of Algorithm 1 is $O(|\Gamma^*| \cdot n^3 \cdot k)$.*

PROOF. In each deterministic policy (vertex along Γ^*), the algorithm checks at most $n \cdot k$ options for the next policy. The running time of each check is dominated by matrix inversion. Matrix inversion is applied to a basis matrix that is obtained from an adjacent basis; namely, a change in a single column. By the Sherman-Morrison formula (Meyer [15]), the inverse matrix can be computed in time $O(n^2)$. Thus the complexity of the algorithm is $O(|\Gamma^*| \cdot n^3 \cdot k)$, as required. \square

7. A strongly polynomial algorithm: When coupling property holds. For every deterministic policy τ , let $\rho_{\min}(\tau) \triangleq \min\{\rho_\tau(i, \tau(i)): i \in X\}$. Similarly, $\rho_{\max}(\tau) \triangleq \max\{\rho_\tau(i, \tau(i)): i \in X\}$. Let $\rho_{\min} \triangleq \min_\tau \rho_{\min}(\tau)$ and $\rho_{\max} \triangleq \max_\tau \rho_{\max}(\tau)$, where the minimum and maximum are taken only over deterministic policies. Assumption 2.1 implies that $\rho_{\min} > 0$.

The algorithm uses a parameter R that satisfies

$$R \geq \frac{1 + \rho_{\max}}{\rho_{\min}}. \tag{5}$$

There is no need to precisely compute the right-hand side of Equation (5); instead, we use an upper bound based on Assumption 2.1 as follows. Obviously, $\rho_{\max} < 1$. In the expected average cost model, we lower bound ρ_{\min} by $\tilde{\rho}_{\min} = p_{\min}^n$, where p_{\min} is the minimum nonzero transition probability in the MDP. This lower bound holds simply by considering all paths of length n with nonzero transition probabilities to a given state. In the discounted cost model, we lower bound ρ_{\min} by $\tilde{\rho}_{\min} = (1 - \beta) \cdot \beta^{n-1} \cdot p_{\min}^{n-1}$. The algorithm uses the following value for R :

$$R \triangleq \max \left\{ \frac{2}{\tilde{\rho}_{\min}}, nk \right\}. \tag{6}$$

7.1. Algorithm description. A listing of the algorithm appears as Algorithm 2. The algorithm works under the additional assumption that the coupling property holds. The algorithm is a variation of Algorithm 1. The only difference is in the definition of the new artificial cost constraint $d^T \cdot \rho = \alpha$. This definition now relies on the linear orders \leq_i and the parameter R . The costs $d(i, j)$ are exponential functions of i and j .

In line 1, the algorithm sorts the actions in each state; namely, it computes the linear orders \leq_i over U for each state $i \in X$. In line 2, costs $d(i, j)$ are assigned to each pair $(i, j) \in X \times U$. In line 3, the initial policy is set. This policy simply chooses the first action (according to the order \leq_i) for each state i . This initial policy π achieves the minimum value for $D(\pi)$. The remaining lines are identical to corresponding lines in Algorithm 1.

Algorithm 2

A strongly polynomial algorithm for finding an optimal policy for the MDP $\min\{c \cdot \rho \mid A \cdot \rho = b\}$. We assume that the MDP is irreducible and satisfies both the uniqueness and coupling properties.

- (1) Sort the actions for each state $i \in X$ according to the order \leq_i . Let $j_0^i \leq_i j_1^i \leq_i \dots \leq_i j_{k-1}^i$ denote the actions sorted according to the order \leq_i .
- (2) Define $d(i, j_l^i) \triangleq R^{k \cdot (n-i) + l}$.
- (3) Initialize

$\pi \leftarrow (j_0^0, \dots, j_0^{n-1})$	{ π chooses the “first” action in each state}
$\text{opt} \leftarrow \pi$	{best policy so far}
- (4) **while** exists (i, j) such that $D(\pi^{i,j}) > D(\pi)$, **do**
- (5) $(i, j) \leftarrow \arg \min\{\nabla_{i,j} \mid \forall (i, j) \text{ such that } D(\pi^{i,j}) > D(\pi)\}$
- (6) $\pi \leftarrow \pi^{i,j}$
- (7) **if** $C(\pi^{i,j}) < C(\text{opt})$, **then**
- (8) $\text{opt} \leftarrow \pi^{i,j}$ {opt is the best policy so far}
- (9) **end if**
- (10) **end while**
- (11) **return** opt

7.2. Correctness. The following lemma is used both for proving the correctness and running time of Algorithm 2. Consider two neighboring deterministic policies π and τ . By Lemma 5.2, it follows that the cost $D(\cdot)$ is strictly monotone along the edge in the policy graph from π to τ . The following lemma determines whether $D(\cdot)$ increases or decreases along the edge (π, τ) . The lemma relies on both Assumption 2.1 and the coupling assumption.

For two actions j_1 and j_2 , we say that $j_1 <_i j_2$ if $j_1 \leq_i j_2$ and $j_1 \neq j_2$.

LEMMA 7.1. Let $\pi \neq \tau$ denote two neighboring deterministic policies, i.e., $\tau = \pi^{i,j}$. Then, $\pi(i) <_i \tau(i)$ if and only if $D(\pi) < D(\tau)$.

PROOF. We first prove that $\pi(i) <_i \tau(i)$ implies that $D(\pi) < D(\tau)$. The proof is based on a coupling argument and the exponential growth of $d(x, u)$, as x is closer to the initial state 0. By definition,

$$D(\tau) - D(\pi) = \sum_{x \in X, u \in U} d(x, u) \cdot (\rho_\tau(x, u) - \rho_\pi(x, u)). \tag{7}$$

We partition this difference into the following three parts:

$$\begin{aligned} \delta_1 &\triangleq \sum_{x < i} \sum_{u \in U} d(x, u) \cdot (\rho_\tau(x, u) - \rho_\pi(x, u)). \\ \delta_2 &\triangleq \sum_{x > i} \sum_{u \in U} d(x, u) \cdot (\rho_\tau(x, u) - \rho_\pi(x, u)). \\ \delta_3 &\triangleq \sum_{u \in U} d(i, u) \cdot (\rho_\tau(i, u) - \rho_\pi(i, u)). \end{aligned}$$

Since $\pi(i) <_i \tau(i)$, the coupling property states that $\rho_\pi(x, u) \leq \rho_\tau(x, u)$, for every $x < i$ and every $u \in U$. Therefore, $\delta_1 \geq 0$.

We now bound δ_2 as follows. Note that, for every $x > i$ and $u \in U$, it follows that $d(x, u) \leq R^{k \cdot (n-i)-1}$. Hence

$$\begin{aligned} \delta_2 &\geq \sum_{x>i} \sum_{u \in U} d(x, u) \cdot (0 - \rho_\pi(x, u)) \\ &\geq -\rho_{\max} \cdot (n-i) \cdot k \cdot R^{k \cdot (n-i)-1} \\ &> -R^{k \cdot (n-i)}, \end{aligned}$$

where the last line follows from $\rho_{\max} < 1$ and $R \geq kn > k \cdot (n-i)$.

We now bound δ_3 as follows. Denote the index of $\pi(i)$ and $\tau(i)$ in the order \leq_i , as $l(\pi)$ and $l(\tau)$, respectively. The assumption $\pi(i) <_i \tau(i)$ implies that $0 \leq l(\pi) < l(\tau) < k$. Since π and τ are deterministic, it follows that

$$\begin{aligned} \delta_3 &= d(i, j) \cdot \rho_\tau(i, j) - d(i, \pi(i)) \cdot \rho_\pi(i, \pi(i)) \\ &\geq R^{k \cdot (n-i)} \cdot (\rho_{\min} \cdot R^{l(\tau)} - \rho_{\max} \cdot R^{l(\pi)}) \\ &\geq R^{k \cdot (n-i)}, \end{aligned}$$

where the last line follows from $R \geq (1 + \rho_{\max})/\rho_{\min}$ and $l(\tau) \geq l(\pi) + 1$. Note that the second line requires that $\rho_\tau(i, j) > 0$. Indeed, Assumption 2.1 implies that $\rho_\tau(i, j) > 0$.

It follows that $\delta_1 + \delta_2 + \delta_3 > 0 - R^{k \cdot (n-i)} + R^{k \cdot (n-i)} = 0$, as required.

The converse direction is proved as follows. By contraposition, $D(\pi) < D(\tau)$ implies that $\pi(i) \leq_i \tau(i)$. We rule out equality (namely, $\pi(i) = \tau(i)$) since $\pi \neq \tau$ and $\tau = \pi^{i,j}$. \square

COROLLARY 7.1. *The initial policy $\pi_0 = (j_0^1, \dots, j_0^n)$ in Algorithm 2 is an optimal policy of CMDP(α_0) for $\alpha_0 \triangleq D(\pi_0)$. Moreover, LP(α) is not feasible for $\alpha < \alpha_0$.*

PROOF. Consider the policy τ of minimum cost $D(\cdot)$ in Γ^* . By Lemma 5.2, τ is a deterministic policy. Suppose, for the sake of contradiction, that $\tau \neq \pi_0$. Let i denote a state for which $\pi_0(i) \neq \tau(i)$. By the definition of π_0 , it follows that $\pi_0(i) <_i \tau(i)$. Let $\pi = \tau^{i, \pi_0(i)}$. Note that π and τ satisfy the premises of Lemma 7.1. Hence $D(\pi) < D(\tau)$, contradicting the minimality of $D(\tau)$. It follows that π_0 is the unique policy in Γ^* whose cost is $D(\pi_0)$, as required. \square

THEOREM 7.1. *Algorithm 2 returns an optimal deterministic policy of MDP if the coupling property holds.*

PROOF. The proof follows the proof of Theorem 6.1 and Corollary 6.1. The only modification, based on Corollary 7.1, is the justification that the initial policy is an end point of Γ^* . \square

7.3. Running time.

PROPOSITION 7.1. *If the coupling property holds, then $|\Gamma^*| \leq n \cdot k$.*

PROOF. By Corollary 5.2, the cost $D(\cdot)$ increases along Γ^* . Let π immediately precede τ along Γ^* . Suppose $\tau = \pi^{i,j}$. By Lemma 7.1, $\pi(i) <_i \tau(i)$. This implies that the length of the path is bounded by $n \cdot k$. \square

COROLLARY 7.2. *The complexity of Algorithm 2 is $O(k^2 \cdot n^4)$.*

PROOF. Follows directly from Propositions 6.1 and 7.1. \square

8. Discussion. We presented an algorithm for computing an optimal policy of an irreducible MDP. A variation of this algorithm runs in strongly polynomial time if the MDP satisfies the coupling property, e.g., a controlled discrete time $M/M/1$ queue. The algorithm is based on two assumptions: irreducibility and uniqueness.

The uniqueness property is shown in Lemma 3.1 to hold with high probability if the cost function is randomly perturbed. Therefore, if the MDP does not satisfy the uniqueness property, then the need for a random perturbation implies that the algorithm is a randomized ε -approximation algorithm.

The irreducibility assumption of the MDP (i.e., Assumption 2.1) is used several times in the proofs to show that the occupation measure is positive for every state. In the discounted cost model (i.e., $\beta < 1$), irreducibility can be replaced with the more relaxed assumption that, for every deterministic policy, every state is reachable from the initial state with positive probability.

Appendix A. Uniformization of a birth-death process. Uniformization (Beutler and Ross [2], Serfozo [19], Tijms [20]) is a technique that transforms a continuous-time Markov process into a discrete one, while

preserving the occupation measure. In the case of a controlled Markov process, uniformization also preserves the optimality of policies and their costs. We apply uniformization to transform a controlled continuous time birth-death process (Kleinrock [10]) (i.e., one-dimensional queue) to a discrete MDP.

The queue dynamics are determined by service completions, which are independent, exponentially distributed random variables with rates $\mu(i, u)$ that depend on the state i and the action u , and an independent Poisson arrival process with rate λ (that does not depend on the state or the action). Equivalently, the interarrival times are exponentially distributed with rate λ .

We now use the following fact. Given two independent exponential random variables X, Y with rates a and b resp. and a number $c > a + b$, we can construct the two variables as follows. Choose a Poisson process with rate c . When the process jumps, throw an independent three-sided die. With probability a/c claim that X occurred. With probability b/c claim that Y occurred, and with probability $(c - a - b)/c$ neither occurred—in which case we repeat. This gives the same probabilistic behavior as the original variables.

Now, denote by $v_i(u)$ the rate by which the queue leaves state i if the selected action is u . To apply the previous argument, set $v_i(u) = \lambda + \mu(i, u)$ as the total rate in state i . Let $v = \max_{i,u} v_i(u)$. The transition rates of the uniformized process are now independent of the state and of the actions chosen in each state, and are equal to v . Once a transition occurred, its type is determined according to

$$P(j | i, u) = \begin{cases} \frac{\lambda}{v} & \text{for } j = i + 1, \\ \frac{\mu(i, u)}{v} & \text{for } j = i - 1, \\ 1 - \frac{\lambda + \mu(i, u)}{v} & \text{for } j = i, \end{cases} \quad (\text{A1})$$

with the appropriate modifications at $i = 0$ and $i = n - 1$. We have obtained a process that shares the probabilistic description of the original, except that self-loops were introduced. In the new process, times between events (including self-loops) are independent and identically distributed (exponential with rate v).

To obtain a discrete time process, we observe the process at jump times (including times of self-loops). Because interjump times are i.i.d., we can simply count the number of jumps, where the transition probabilities are determined by Equation (A1), and obtain a discrete-time process. The properties of the interjump times imply that, under any deterministic policy, the occupation measure of the original process agrees with that of the discrete-time process, and, in particular, all cost functionals agree as well.

Appendix B. Proof that the coupling property holds.

PROOF OF LEMMA 4.1. We prove the lemma for the expected average cost model. We use the following abbreviated notation. Fix a deterministic policy π . For each state x , let $\rho(x)$ denote the occupation measure for state x under the policy π ; namely, $\rho(x) \triangleq \sum_{u \in U} \rho_{\pi}(x, u)$. Let $\rho'(x) \triangleq \sum_{u \in U} \rho_{\pi^{i,j}}(x, u)$. Let $p(x) \triangleq P(x - 1 | x, \pi(x))$ and $q(x) \triangleq P(x + 1 | x, \pi(x))$. Similarly, let $p'(x) \triangleq P(x - 1 | x, \pi^{i,j}(x))$ and $q'(x) \triangleq P(x + 1 | x, \pi^{i,j}(x))$. For state 0, let $p(0) \triangleq P(0 | 0, \pi(0))$, and for state $n - 1$, let $q(n - 1) \triangleq P(n - 1 | n - 1, \pi(n - 1))$.

We claim that the following holds, for every state $x \geq 1$:

$$\rho(x) = \frac{q(x - 1) \cdot q(x - 2) \cdots q(0)}{p(x) \cdot p(x - 1) \cdots p(1)} \cdot \rho(0). \quad (\text{B1})$$

The proof of Equation (B1) is similar to the analytic solution of the equilibrium probabilities of continuous-time birth-death queuing systems (Kleinrock [10]). We prove Equation (B1) by induction on x . The basis for $x = 1$ is equivalent to $\rho(1) \cdot p(1) = q(0) \cdot \rho(0)$. Indeed, the first constraint of LP implies that

$$\rho(0) \cdot p(0) + \rho(1) \cdot p(1) = \rho(0) \cdot p(0) + \rho(0) \cdot q(0).$$

This is, in fact, the balance equation (Kleinrock [10], Tijms [21]) that compares the probability of transitions entering state 0 with the probabilities of the transitions emanating from state 0.

Assume that Equation (B1) holds for $x \leq k$. The induction step for $x = k + 1$ uses the LP constraint for state k (i.e., the balance equation for state k). Namely,

$$\rho(k - 1) \cdot q(k - 1) + \rho(k + 1) \cdot p(k + 1) = \rho(k)(p(k) + q(k)).$$

Rearranging, we obtain

$$\rho(k + 1) = \frac{1}{p(k + 1)} \cdot (\rho(k)(p(k) + q(k)) - \rho(k - 1) \cdot q(k - 1)).$$

Dividing $\rho(k)/\rho(k - 1)$, and substituting according to Equation (B1) gives $\rho(k) \cdot p(k) = \rho(k - 1) \cdot q(k - 1)$. Therefore,

$$\begin{aligned} \rho(k + 1) &= \frac{1}{p(k + 1)} \cdot (\rho(k) \cdot q(k)) \\ &= \frac{q(k) \cdot q(k - 1) \cdots q(0)}{p(k + 1) \cdot p(k) \cdots p(1)} \cdot \rho(0), \end{aligned}$$

which completes the proof of Equation (B1).

Our goal is to prove that if $\pi(i) \leq_i j$, then $\rho_{\pi}(x) \leq \rho_{\pi^{i,j}}(x)$, for every $x < i$. Let $\gamma(x) \triangleq (q(x - 1) \cdot q(x - 2) \cdots q(0)) / (p(x) \cdot p(x - 1) \cdots p(1))$. Similarly, let $\gamma'(x)$ denote the above ratio with respect to the policy $\pi^{i,j}$.

We claim that for every state x , $\gamma(x) \geq \gamma'(x)$. Indeed, for $x < i$, $\gamma(x) = \gamma'(x)$ because the ratio differs only when $x \geq i$. For $x = i$, it follows that $\gamma(x)/\gamma'(x) = p'(i)/p(i) \geq 1$ since $\pi(i) \leq_i j$. For $x > i$, it follows that $\gamma(x)/\gamma'(x) = (p'(i)/p(i)) \cdot (q(i)/q'(i)) \geq 1$.

Recall first that since ρ is an occupation measure, it follows that $\sum_{x \in X} \rho(x) = 1$. Hence, by Equation (B1),

$$\begin{aligned} 1 &= \sum_{x \in X} \rho(x) = \rho(0) \cdot \sum_{x \in X} \gamma(x) \\ 1 &= \sum_{x \in X} \rho(x) = \rho'(0) \cdot \sum_{x \in X} \gamma'(x). \end{aligned}$$

Since $\sum_{x \in X} \gamma(x) \geq \sum_{x \in X} \gamma'(x)$, it follows that $\rho(0) \leq \rho'(0)$. For every state $x < i$, we have $\gamma(x) = \gamma'(x)$, hence by Equation (B1), it follows that $\rho(x) \leq \rho'(x)$, as required. \square

PROOF OF LEMMA 4.2. We use the same notation as in the proof of Lemma 4.1. We claim that the following holds, for every state $x \geq 1$:

$$\rho(x) = \frac{q^x}{p(x) \cdot p(x - 1) \cdots p(1)} \cdot \rho(0). \tag{B2}$$

The proof of Equation (B2) is by induction on x . The basis for $x = 1$ is equivalent to $\rho(1) \cdot p(1) = q \cdot \rho(0)$. Indeed, it holds because of the balance equation (Kleinrock [10], Tijms [21]):

$$\rho(0) \cdot p(0) + \rho(1) \cdot p(1) = \rho(0) \cdot p(0) + \rho(0) \cdot q$$

that compares the probability of transitions entering state 0 with the probabilities of the transitions emanating from state 0. Note that this balance equation does not hold in the discounted cost model.

Assume that Equation (B2) holds for $x \leq k$, the induction step for $x = k + 1$ uses the balance equation for state k . Namely,

$$\rho(k - 1) \cdot q + \rho(k) \cdot P(k | k, \pi(k)) + \rho(k + 1) \cdot p(k + 1) = \rho(k)(p(k) + P(k | k, \pi(k)) + q).$$

Rearranging, we obtain

$$\rho(k + 1) = \frac{1}{p(k + 1)} \cdot (\rho(k)(p(k) + q) - \rho(k - 1) \cdot q).$$

By dividing Equation (B2) for $\rho(k)$ and $\rho(k - 1)$, it follows that $\rho(k) \cdot p(k) = \rho(k - 1) \cdot q$. Therefore

$$\begin{aligned} \rho(k + 1) &= \frac{1}{p(k + 1)} \cdot (\rho(k) \cdot q) \\ &= \frac{q^{(k+1)}}{p(k + 1) \cdot p(k) \cdots p(1)} \cdot \rho(0), \end{aligned}$$

which completes the proof of Equation (B2).

Our goal is to prove that if $\pi(i) \leq_i j$, then $\rho_\pi(x) \leq \rho_{\pi^{i,j}}(x)$ for every $x < i$. Let $\gamma(x) \triangleq q^x / (p(x) \cdot p(x-1) \cdots p(1))$. Similarly, let $\gamma'(x)$ denote the above ratio with respect to the policy $\pi^{i,j}$.

We claim that for every state x , $\gamma(x) \geq \gamma'(x)$. Indeed, for $x < i$, $\gamma(x) = \gamma'(x)$ because the ratio differs only when $x \geq i$. For $x \geq i$, it follows that $\gamma(x)/\gamma'(x) = p'(i)/p(i) \geq 1$ since $\pi(i) \leq_i j$.

Recall first that since ρ is an occupation measure, it follows that $\sum_{x \in X} \rho(x) = 1$. Hence, by Equation (B2),

$$1 = \sum_{x \in X} \rho(x) = \rho(0) \cdot \sum_{x \in X} \gamma(x),$$
$$1 = \sum_{x \in X} \rho(x) = \rho'(0) \cdot \sum_{x \in X} \gamma'(x).$$

Since $\sum_{x \in X} \gamma(x) \geq \sum_{x \in X} \gamma'(x)$, it follows that $\rho(0) \leq \rho'(0)$. For every state $x < i$, we have $\gamma(x) = \gamma'(x)$, hence by Equation (B2), it follows that $\rho(x) \leq \rho'(x)$, as required. \square

Acknowledgments. The authors thank Boaz Patt-Shamir for many helpful conversations. Adam Schwartz holds The Julius M. and Bernice Naiman Chair in Engineering. His research was supported, in part, by the fund for promotion of research, and by the fund for promotion of sponsored research at Technion.

References

- [1] Altman, E. 1999. *Constrained Markov Decision Processes*. Chapman & Hall/CRC Press, Boca Raton, FL.
- [2] Beutler, F. J., K. W. Ross. 1987. Uniformization for semi-Markov decision processes under stationary policies. *J. Appl. Probab.* **24** 644–656.
- [3] Blondel, V. D., J. N. Tsitsiklis. 2000. A survey of computational complexity results in systems and control. *Automatica* **36** 1249–1274.
- [4] de Ghellinck, G. 1960. Les problemes de decisions sequentielles. *Cahiers Centre Etudes Rech. Operationnelle* **2** 161–179.
- [5] d'Epenoux, F. 1963. A probabilistic production and inventory problem. *Management Sci.* **10**(1) 98–108.
- [6] Derman, C. 1962. On sequential decisions and Markov chains. *Management Sci.* **9**(1) 16–24.
- [7] Horn, R. A., C. R. Johnson. 2005. *Matrix Analysis*. Cambridge University Press, Cambridge, UK.
- [8] Kallenberg, L. 2002. Finite state and action MDPs. E. Feinberg, A. Schwartz, eds. *Handbook of Markov Decision Processes: Methods and Applications*. Kluwer, Boston, 21–87.
- [9] Kitaev, M. Y., V. V. Rykov. 1995. *Controlled Queueing Systems*. CRC Press, New York.
- [10] Kleinrock, L. 1975. *Queueing Systems, Volume I: Theory*. John Wiley and Sons, New York.
- [11] Littman, M. L., T. L. Dean, L. P. Kaelbling. 1995. On the complexity of solving Markov decision problems. *Proc. 11th Conf. Uncertainty in Artificial Intelligence*, Morgan Kaufmann, San Francisco, 394–402.
- [12] Manne, A. S. 1960. Linear programming and sequential decisions. *Management Sci.* **6**(3) 259–267.
- [13] Mansour, Y., S. Singh. 1999. On the complexity of policy iteration. *Proc. 15th Conf. Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers, San Francisco.
- [14] Megiddo, N. 2006. Method for solving stochastic control problems of linear systems in high dimension. U.S. Patent 7,117,130, October 3, 2006.
- [15] Meyer, C. D. 2000. *Matrix Analysis and Applied Linear Algebra*. Society for Industrial Mathematics, Philadelphia.
- [16] Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, Inc., New York.
- [17] Ross, S. M. 2000. *Introduction to Probability Models*. Academic Press, Boston.
- [18] Schrijver, A. 2000. *Theory of Linear and Integer Programming*. John Wiley and Sons, New York.
- [19] Serfozo, R. F. 1979. An equivalence between continuous and discrete time Markov decision processes. *Oper. Res.* **27** 616–620.
- [20] Tijms, H. C. 1986. *Stochastic Modelling and Analysis: A Computational Approach*. John Wiley and Sons, Inc., New York.
- [21] Tijms, H. C. 1994. *Stochastic Models: An Algorithmic Approach*. John Wiley and Sons, Inc., New York.
- [22] Tseng, P. 1990. Solving H-horizon, stationary Markov decision problems in time proportional to $\log(H)$. *Oper. Res. Lett.* **9**(5) 287–297.
- [23] Yadin, M., P. Naor. 1967. On queueing systems with variable service capacities. *Naval Res. Logist. Q.* **14** 43–53.
- [24] Ye, Y. 2005. A new complexity result on solving the Markov decision problem. *Math. Oper. Res.* **30**(3) 733–749.
- [25] Zadorojnyi, A., A. Schwartz. 2006. Robustness of policies in constrained Markov decision processes. *IEEE Trans. Automatic Control* **51**(4) 635–638.