

The Linear Program Approach in Multi-Chain Markov Decision Processes Revisited

EITAN ALTMAN

INRIA, Centre Sophia Antipolis, 06565 Valbonne Cedex, France

FLOS SPIEKSMAN

Institute of Mathematics & Computer Science, University of Leiden, P.O. Box 9512,
2300RA Leiden, The Netherlands

Abstract: Linear Programming is known to be an important and useful tool for solving Markov Decision Processes (MDP). Its derivation relies on the Dynamic Programming approach, which also serves to solve MDP. However, for Markov Decision Processes with several constraints the only available methods are based on Linear Programs. The aim of this paper is to investigate some aspects of such Linear Programs, related to multi-chain MDPs. We first present a stochastic interpretation of the decision variables that appear in the Linear Programs available in the literature. We then show for the multi-constrained Markov Decision Process that the Linear Program suggested in [9] can be obtained from an equivalent unconstrained Lagrange formulation of the control problem. This shows the connection between the Linear Program approach and the Lagrange approach, that was previously used only for the case of a single constraint [3, 14, 15].

Key Words: Multi-chain Markov Decision Processes, average cost criterion, state-action frequencies, deviation measure, linear programming, lagrange formulation.

Introduction

Markov Decision Processes (MDP) have been used extensively in the past to model and solve problems in data communications, computer networks, production etc. . . We consider in this note MDPs where the optimization criteria are of expected time average type, with finite state and action spaces. A basic method for solving such problems has been Dynamic Programming. Derman [7] obtained a Linear Program formulation from the dynamic program and investigated its dual program DLP. Using DLP and another Linear Program he was able to obtain an optimal control policy. Later on, Hordijk and Kallenberg obtained an optimal policy directly from DLP [8]. One advantage of using this new approach was that it could be extended to Constrained Markov Decision Processes, as Hordijk and Kallenberg did in [9], unlike the Dynamic Programming approach.

The first issue of this paper is to present a stochastic meaning to the two different kinds of decision variables that appear in the Linear Programs of [9].

The first kind is related to the expected frequencies of pairs of states and actions (see e.g. [7]). These frequencies are also known as “occupation measure” (see e.g. [4]) or expected empirical probabilities of state-action pairs. The precise definition of these frequencies is given in the next Section. Derman [7] considered the set of all frequencies that are achieved by different policies, and showed that it is equal to the set of achievable decision variables of the first kind. The properties of these frequencies have been studied in [7], [8] and [9].

However, the stochastic meaning of the second type of decision variables was unknown till now, to the best of our knowledge, except for stationary policies, where the decision variables were related to the deviation matrix by Kallenberg [10].

We note that for MDPs with a sample-path constraint (instead of the weaker constraint on the expected cost), an ε -optimal policy can be obtained by using Linear Programs that do not involve the second type of decision variables, see [12].

The importance in understanding the stochastic meaning of the decision variables is that it may enable us to obtain a Linear Program for solving MDP's directly, without requiring to go through the dual problem obtained from the Dynamic Program. This may be crucial in case duality problems may arise, such as the duality gap in infinite Linear Programs. Such a direct approach to obtain the Linear Problem is illustrated in [1]. There Altman and Schwartz obtained a Linear Program for solving constrained MDPs with a countable state space. However, they make strong ergodic assumptions (such as the unichain assumption) under which only decision variables of the first kind are needed. In the case of an infinite state space, it seems that understanding the meaning of all decision variables may enable to obtain Linear Programs with a general multichain structure. This is especially important in the constrained case where no alternative methods are known for solving the problem.

The second issue in this paper is to present an alternative derivation of the Linear Program used for solving the constrained Markov Decision problem [9]. Our derivation is based on a Lagrange formulation which generalizes the one used in the case of a single constraint [3] (later extended in [14], [15] to the countable state space) and is related to the approach used by Borkar in his book [5]. Our new derivation seems to be more natural and straightforward than the previous method [9], where the paradigm was to introduce the LP and then to prove that the optimal policy and optimal value of the control problem are appropriately related to the optimal solution and value of the LP.

After reviewing in Section 2 the Linear Program for solving MDPs, we introduce in Section 3 the biased deviation measure for arbitrary policies and show that it corresponds to the second kind of decision variables in the following sense. The sets of pairs of frequencies and biased deviation measures obtained by any arbitrary policy form a feasible solution for the Linear Pro-

gram. In Section 4 we then present some properties of the deviation measures. In Section 5 we present the alternative derivation of the LP used for the constrained case.

1 Model and Assumptions

Consider a Markov decision process with a finite state space $X = \{0, 1, \dots, N\}$ and a finite action space A . Without loss of generality, we assume that in any state x all actions in A are available. The probability to go from state x to state y given that action a is used, is given by the transition probability P_{xay} . A policy u in the policy space U is described as $u = \{u_1, u_2, \dots\}$, where u_t , applied at time epoch t , is a probability measure over A conditioned on the whole history of actions and state prior to t , as well as the state at time t . Given an initial distribution β on X , each policy u induces a probability measure denoted by P_β^u on the space of sample paths of states and actions (which serves as the canonical sample space Ω). The corresponding expectation operator is denoted by E_β^u . On this probability space are defined the state and action processes, $X_t, A_t, t = 1, 2, \dots$.

A Markov policy $u \in U(M)$ is characterized by the dependence of u_{t+1} on the current state and the time only. A stationary policy $g \in U(S)$ is characterized by a single conditional probability measure $p_{\cdot|x}^g$ over A , so that $p_{\lambda|x}^g = 1$; under g , X_t becomes a Markov chain with stationary transition probabilities, given by $P_{xy}^g = \sum_{a \in A} p_{a|x}^g P_{xay}$. The class of stationary deterministic policies $U(SD)$ is a subclass of $U(S)$, and every $g \in U(SD)$ is characterized by a mapping $g: X \rightarrow A$, so that $p_{\cdot|x}^g = \delta_{g(x)}(\cdot)$ is concentrated at the point $g(x)$ for each x .

For any given policy $g \in U(S)$ we denote by $v(g)$ the number of closed classes in the Markov chain generated by policy g . We write $B(g)$ for a subset of X that contains precisely one state from each closed class. Clearly $|B(g)| = v(g)$. $\mathcal{T}(g)$ denotes the set of transient state and $\mathcal{R}_l(g), l = 1, \dots, v(g)$ the closed classes.

$${}_M P_{\delta_x}^g \{X_t = y\} := \begin{cases} P_{\delta_x}^g \{X_t = y, X_s \notin M, 2 \leq s < t\}, & t \geq 2 \\ \delta_x(y), & t = 1 \end{cases}$$

stands for the taboo transition probabilities with taboo set M . The probabilities $F_M(\delta_x; g)$ of reaching set M , when the chain initially starts in state x , are equal to

$$F_M(\delta_x; g) = \sum_{m \in M} \sum_{t=1}^{\infty} \sum_{y \in X/M} {}_M P_{\delta_x}^g \{X_t = y\} P_{ym}^g$$

for $x \notin M$, and 1 for $x \in M$.

Let $c: \mathbf{X} \times \mathbf{A} \rightarrow \mathcal{R}$, be a (real valued) cost function and define the average costs associated with a policy u and with an initial distribution β on \mathbf{X} :

$$C_\beta(u) = \overline{\lim}_{t \rightarrow \infty} \frac{1}{t} E_\beta^u \left[\sum_{s=1}^t c(X_s, A_s) \right] \tag{1.2}$$

Denote by OP the problem of finding a policy u that minimizes $C_\beta(u)$ for a given initial distribution β . Let C_β be the optimal value of OP. A policy that achieves $C_\beta(u) = C_\beta$ is said to be optimal for OP. Denote U_β the set of all such policies.

Derman [7] has shown that the cost in (1.2) can be represented in terms of the following *state action frequencies*: let $\bar{f}^T(\beta, u)$ is the matrix whose elements are given by $\bar{f}^T(y, a; \beta, u) :=$ the average expected frequency of state y and action a till time T when using policy u and initial probability distribution on the state space is $\beta(\cdot)$, i.e.

$$\bar{f}^T(y, a; \beta, u) = \frac{1}{T} \sum_{s=1}^T P_\beta^u(X_s = y, A_s = a)$$

Let $\bar{F}(\beta, u)$ be the set of accumulation points of $\bar{f}^T(\beta, u)$. A generic element of $\bar{F}(\beta, u)$ will be denoted by $\bar{f}(\beta, u)$.

Given a class of policies U' , define $L_\beta(U') := \bigcup_{u \in U'} \bar{F}(\beta, u)$. The set of frequencies obtained by all policies is $L_\beta := \bigcup_{u \in U} \bar{F}(\beta, u)$.

The following notation is used below: $1\{A\}$ is the indicator function of the set A and $\delta_a(x)$ is the Kronecker delta function. We denote by \bar{B} the closure of a set B , and $|B|$ is the cardinality of the set (we shall use this notation for finite sets only, in which case $|B|$ is the number of elements in B). For two matrices ζ, Q of appropriate dimensions, $\zeta \cdot Q$ stands for summation over common indices (scalar product).

2 Linear Program for Solving MDPs

One method of solving OP is based on the solution of a LP which we present below. The importance of this method lies in the fact that it also enables to deal with optimization problems with additional constraints, where other methods (based on dynamic programming) are not applicable.

Given the initial distribution β over \mathbf{X} , define Π_β to be the set of $\{(z, \zeta)\}$, $z, \zeta \in \mathcal{R}^{|\mathbf{X} \times \mathbf{A}|}$, that satisfy

$$\sum_{y \in \mathbf{X}} \sum_{a \in \mathbf{A}} (\delta_y(v) - P_{yav})z(y, a) = 0, \quad \forall v \in \mathbf{X} \tag{2.1.a}$$

$$\sum_{a \in \mathbf{A}} z(v, a) + \sum_{y \in \mathbf{X}} \sum_{a \in \mathbf{A}} (\delta_y(v) - P_{yav})\zeta(y, a) = \beta(v), \quad \forall v \in \mathbf{X} \tag{2.1.b}$$

$$z \geq 0 \tag{2.1.c}$$

$$\zeta \geq 0 \tag{2.1.d}$$

Remark: Every $z(\cdot, \cdot) \in \Pi_\beta$ satisfies $\sum_{y,a} z(y, a) = 1$. This can be seen by summing equation (2.1.b) over all $v \in \mathbf{X}$.

Consider the following Linear Program:

LP: Find $z, \zeta \in \mathcal{R}^{|\mathbf{X} \times \mathbf{A}|}$, that minimize $c \cdot z$ subject to $(z, \zeta) \in \Pi_\beta$.

LP is related to OP in the following way (see [8, 10]). Given any $(z, \zeta) \in \Pi_\beta$ define the stationary policy $g(z, \zeta)$ by

$$p_{a|y}^{g(z, \zeta)} = \begin{cases} \frac{z(y, a)}{\sum_a z(y, a)}, & \text{if } \sum_a z(y, a) > 0 \\ \frac{\zeta(y, a)}{\sum_a \zeta(y, a)}, & \text{if } \sum_a z(y, a) = 0 \text{ and } \sum_a \zeta(y, a) > 0 \\ \text{arbitrary,} & \text{otherwise.} \end{cases}$$

Lemma 1:

- (i) *The optimal values of OP and of LP are equal.*
- (ii) *Suppose that (z^*, ζ^*) is any optimal solution of LP. Then $g(z^*, \zeta^*)$ is optimal for OP.*

The proof of the Lemma is given in [8, 10], except that the average cost in these references is defined differently, namely, the limsup and the summation are interchanged. That these results apply to our case follows from the following observation. Let

$$\bar{C}_\beta(u) := \overline{\lim}_{t \rightarrow \infty} t^{-1} E_\beta^u \left[\sum_{s=1}^t c(X_s, A_s) \right]$$

$$\bar{C}(\beta, u) := \sum_{y \in \mathbf{X}} \beta(y) \overline{\lim}_{t \rightarrow \infty} t^{-1} E_\beta^u \left[\sum_{s=1}^t c(X_s, A_s) | X_1 = y \right]$$

$$\underline{C}_\beta(u) := \lim_{t \rightarrow \infty} t^{-1} E_\beta^u \left[\sum_{s=1}^t c(X_s, A_s) \right]$$

$$\underline{C}(\beta, u) := \sum_{y \in \mathbf{X}} \beta(y) \lim_{t \rightarrow \infty} t^{-1} E_\beta^u \left[\sum_{s=1}^t c(X_s, A_s) | X_1 = y \right].$$

Then,

$$\bar{C}(\beta, u) \geq \bar{C}_\beta(u) \geq \underline{C}_\beta(u) \geq \underline{C}(\beta, u)$$

for every policy u , and

$$\bar{C}(\beta, u) = \bar{C}_\beta(u) = \underline{C}_\beta(u) = \underline{C}(\beta, u)$$

for every stationary policy u . Since OP is known to have a stationary deterministic optimal policy, $\inf_u \bar{C}(\beta, u)$ may be considered and the results of [8, 10] can be used.

The Lemma motivates our interest in studying the properties of LP, as these may relate to properties of Markov Decision Processes. The Lemma also has extensions to constrained Markov Decision Problems (see [9]). The aim of this paper is to study more about the properties of the decision variables of LP and their relation to quantities that characterize policies.

Denote $\Pi_\beta^z := \{z: \exists \zeta \text{ s.t. } (z, \zeta) \in \Pi_\beta\}$.

It is known [8] that $\Pi_\beta^z = L_\beta$. This means that for any policy u , $\bar{f}(\beta, u) \in \Pi_\beta^z$. On the other hand, for each $z \in \Pi_\beta^z$ there exists a policy u such that $\bar{f}(\beta, u) = z$.

However, little is known about the “physical” meaning of the variable ζ . For $u \in U(S)$ Kallenberg [10, p. 109] gives an explicit expression $\zeta = \zeta(\beta, u)$ such that $(\bar{f}(\beta, u), \zeta(\beta, u)) \in \Pi_\beta$. The reason that the meaning of the relation between the decision variables of LP and the control problem was not clear (for all policies) before is related to the fact that LP was not obtained directly from the control problem, but rather as the dual problem of another LP, related to dynamical programming. In the following sections we present an interpretation of the variables ζ that is related to the control problem.

3 The Deviation Measure

With relation to some $\bar{f}(\beta, u) \in \bar{F}(\beta, u)$ we define the biased total occupation matrix $\bar{r}^T(\beta, u)$ whose elements are given by $\bar{r}^T(y, a; \beta, u) := \sum_{s=1}^T [P_\beta^u \{X_s = y,$

$A_s = a\} - \bar{f}(y, a; \beta, u)$. Define the average biased occupation matrix $\bar{r}^T(\beta, u)$ whose elements are given by $\bar{r}^T(y, a; \beta, u) := \frac{1}{T} \sum_{t=1}^T \bar{r}^t(y, a; \beta, u)$. Let t_n be a subsequence of t along which $\bar{f}^{t_n}(\beta, u) \rightarrow \bar{f}(\beta, u)$. Pick a further subsequence s_n of t_n along which some (possibly infinite) limit $\tilde{r}(\beta, u) := \lim_{n \rightarrow \infty} \bar{r}^{s_n}(\beta, u)$ exists. Let $\tilde{R}(\bar{f}, \beta, u)$ denote the set of all such limits corresponding to $\bar{f}(\beta, u)$. We call $\tilde{r}(\beta, u)$ the *deviation measure*.

The following Lemma relates the quantities $\bar{f}(\beta, u)$ and $\tilde{r}(\beta, u)$ to the decision variables in LP, or more precisely to the variables that determine the set Π_β . By appropriately adding a bias factor, we then obtain in Theorem 3 a stronger characterization.

Lemma 2: Under any policy $u \in U$ and initial distribution β , if for some $\bar{f}(\beta, u) \in \bar{F}(\beta, u)$, the elements of some $\tilde{r}(\beta, u) \in \tilde{R}(\bar{f}, \beta, u)$ are finite then the tuple $(\bar{f}(\beta, u), \tilde{r}(\beta, u))$ satisfies equations (2.1.a), (2.1.b) and (2.1.c).

Proof: We shall use the following:

$$P\{X_t = y\} = E[P\{X_t = y | H_{t-1}\}] = \sum_{a,z} P\{X_{t-1} = z, A_{t-1} = a\} P_{zay} \tag{3.1}$$

By averaging we obtain:

$$\frac{1}{t} \sum_{s=2}^t P_\beta^u\{X_s = y\} = \frac{1}{t} \sum_{s=2}^t \sum_{a,z} P_\beta^u\{X_{s-1} = z, A_{s-1} = a\} P_{zay} . \tag{3.2}$$

Interchanging the order of summation in the right side of (3.2), we get after some algebra

$$\sum_{y,a} \bar{f}(z, a; \beta, u) [\delta_y(z) - P_{zay}] = 0 \tag{3.3}$$

for any accumulation point $\bar{f}(\beta, u)$, which proves that (2.1.a) is satisfied. Moreover, we clearly have

$$\bar{f}(y, a; \beta, u) \geq 0 . \tag{3.4}$$

This proves that (2.1.c) is satisfied.

Using (3.1) and (3.3) we get

$$\begin{aligned}
& \sum_{s=2}^t \sum_a [P_{\beta}^u \{X_s = y, A_s = a\} - \bar{f}(y, a; \beta, u)] \\
&= \sum_{s=2}^t \sum_{z,a} [(P_{\beta}^u \{X_{s-1} = z, A_{s-1} = a\} - \bar{f}(z, a; \beta, u)) P_{zay}] \tag{3.5}
\end{aligned}$$

We then obtain by definition of \bar{r}^t :

$$\sum_a \bar{r}^t(y, a; \beta, u) + \sum_a \bar{f}(y, a; \beta, u) - \beta(y) = \sum_{z,a} \bar{r}^{t-1}(z, a; \beta, u) P_{zay} . \tag{3.6}$$

Taking the time average of (3.6) we obtain:

$$\begin{aligned}
& \sum_a \frac{1}{s-1} \sum_{t=2}^s \bar{r}^t(y, a; \beta, u) + \sum_a \bar{f}(y, a; \beta, u) - \beta(y) \\
&= \sum_{z,a} P_{zay} \frac{1}{s-1} \sum_{t=2}^s \bar{r}^{t-1}(z, a; \beta, u) . \tag{3.7}
\end{aligned}$$

Equation (3.7) can then be rewritten as:

$$\begin{aligned}
& \sum_a \left(\frac{s}{s-1} \bar{r}^s(y, a; \beta, u) \right) - \frac{1}{s-1} \sum_a \bar{r}^1(y, a; \beta, u) + \sum_a \bar{f}(y, a; \beta, u) - \beta(y) \\
&= \sum_{z,a} \bar{r}^{s-1}(z, a; \beta, u) P_{zay} . \tag{3.8}
\end{aligned}$$

Note that the term $\frac{1}{s-1} \bar{r}^1(y, a; \beta, u)$ converges to zero as $s \rightarrow \infty$. With the sequence s_n defined above Lemma 2, note that we have

$$\lim_{n \rightarrow \infty} \bar{r}^{s_n-1}(\beta, u) = \bar{r}(\beta, u) \tag{3.9}$$

which follows from the fact that

$$\begin{aligned}
& \frac{t}{t-1} \bar{r}^t(y, a; \beta, u) - \bar{r}^{t-1}(y, a; \beta, u) = \frac{1}{t-1} \bar{r}^t(y, a; \beta, u) \\
&= \frac{t}{t-1} [\bar{f}^t(y, a; \beta, u) - \bar{f}(y, a; \beta, u)] \rightarrow 0 \tag{3.10}
\end{aligned}$$

as $t \rightarrow \infty$, along the subsequence s_n . We thus obtain from (3.8)

$$\sum_a \bar{f}(y, a; \beta, u) + \sum_{z,a} \bar{r}(z, a; \beta, u) [\delta_y(z) - P_{zay}] = \beta(y) \tag{3.11}$$

and hence also (2.1.b) is satisfied. Note that (3.11) is obtained for any subsequence s_n of t_n . ■

Remark: If the limit $\bar{r}(\beta, u) := \lim_{t \rightarrow \infty} r^t(\beta, u)$ exists then (3.11) can be obtained directly from (3.6) with $\bar{r}(\beta, u)$ replacing $\bar{r}(\beta, u)$.

Consequently, for any such policy u , $(\bar{f}(\beta, u), \bar{r}(\beta, u))$ satisfies (2.1.a), (2.1.b) and (2.1.c).

Note that the \bar{r} need not be nonnegative and hence equation (2.1.d) is not satisfied. In fact, for any $u \in U$ and initial distribution β , $\sum_{y,a} \bar{r}(y, a; \beta, u) = 0$. This follows from the fact that

$$\begin{aligned} \sum_{y,a} \bar{r}(y, a; \beta, u) &= \sum_{a,y} \lim_{n \rightarrow \infty} \frac{1}{s_n} \sum_{s=1}^{s_n} \sum_{t=1}^s \left[P_\beta^u(X_t = y, A_t = a) - \sum_a \bar{f}(y, a; \beta, u) \right] \\ &= \lim_{n \rightarrow \infty} \sum_{a,y} \frac{1}{s_n} \sum_{s=1}^{s_n} \sum_{t=1}^s \left[P_\beta^u(X_t = y, A_t = a) - \sum_a \bar{f}(y, a; \beta, u) \right] \\ &= \lim_{n \rightarrow \infty} \frac{1}{s_n} \sum_{s=1}^{s_n} \frac{1}{s_n} \sum_{t=1}^s [1 - 1] = 0 . \end{aligned}$$

Hence in order that all the \bar{r} be non-negative, so that we have in fact $(\bar{f}(\beta, u), \bar{r}(\beta, u)) \in \Pi_\beta$, another bias factor should be added. By doing so, (see Theorem 3 below), we obtain the main result of the paper, that states that for any policy u , the pair (frequencies, corresponding biased deviation) are feasible solutions for the LP in [8], i.e. belong to the set Π_β . This result was known to hold only for stationary policies, and the Theorem below generalizes it to any policy.

Let $T \subset \mathbf{X} \times \mathbf{A}$ be the set of pairs $\{y, a\}$ satisfying $\bar{f}(y, a; \beta, u) = 0$. Note that for any $(y, a) \in T$, $\bar{r}(y, a; \beta, u) \geq 0$. Let γ be any real number that satisfies

$$\gamma \geq \min_k \{ \bar{r}(y, a; \beta, u) + k \bar{f}(y, a; \beta, u) \geq 0 \text{ for all } (y, a) \} .$$

Then $\gamma > 0$. Define:

$$Y(y, a; \beta, u) := \bar{r}(y, a; \beta, u) + \gamma \bar{f}(y, a; \beta, u) .$$

Clearly $Y(y, a; \beta, u)$ are nonnegative for all state-action pairs. Hence $Y(\beta, u)$ satisfies equation (2.1.d). We thus obtain:

Theorem 3: Under any policy $u \in U$ and initial distribuion β , if the elements of $\bar{f}(\beta, u)$ are finite then $(\bar{f}(\beta, u), Y(\beta, u)) \in \Pi_\beta$.

Proof: The proof that equations (2.1.a) and (2.1.c) are satisfied remains unchanged. Equation (2.1.b) holds by (2.1.a) and (3.11). ■

An open problem is whether for any solution (z, ζ) of LP we can construct a policy u , such that $(z, \zeta) = (\bar{f}(\beta, u), Y(\beta, u))$. A step towards the answer is presented in Theorem 4, where we give a characterization of the polytope Π_β . To do so, we need some notation first.

For sake of convenience we will write $z(y), \zeta(y)$ and $\bar{f}(y; \beta, u)$ for $\sum_a z(y, a), \sum_a \zeta(y, a)$ and $\sum_a \bar{f}(y, a; \beta, u)$. From the context it will always be clear, whether these quantities have to be interpreted as vectors or as matrices.

Finally we have to introduce the deviation matrix, the entries of which are given formally by

$$D(y; \delta_x, g) = \lim_{\alpha \uparrow 1} \sum_{t=1}^{\infty} \alpha^{t-1} (P_{\delta_x}^g \{X_t = y\} - \bar{f}(y; \delta_x, g)) .$$

This can be rewritten as (cf. Spieksma [16] p. 82)

$$D(y; \delta_x, g) = \sum_{v \in X} (\delta_x(v) - \bar{f}(v; \delta_x, g)) \cdot \sum_{t=1}^{\infty} \sum_{w \in X} P_{\delta_v}^{g_2} \{X_t = w\} (\delta_w(y) - \bar{f}(y; \delta_w, g)) . \tag{3.12}$$

Theorem 4: Fix some initial distribution β . $(z, \zeta) \in \Pi_\beta \Leftrightarrow$ there are $g_1, g_2 \in U(\mathcal{S})$ and an initial distribution $\tilde{\beta}$, such that z, ζ as vectors on X satisfy (3.13) and $z(y, a) = z(y)P_{a|y}^{g_1}, \zeta(y, a) = \zeta(y)P_{a|y}^{g_2}, y \in X$.

$$\left\{ \begin{array}{l} z(y) = \bar{f}(y; \tilde{\beta}, g_1) , \quad \forall y \end{array} \right. \tag{3.13a}$$

$$\left\{ \begin{array}{l} \zeta(y) = \sum_x \left[\beta(x)D(y; \delta_x, g_2) \right. \end{array} \right. \tag{3.13b}$$

$$\left. \left. \begin{array}{l} + (\bar{f}(x; \beta, g_2) - z(x)) \sum_{t=1}^{\infty} P_{\delta_x}^{g_2} \{X_t = y\} + \gamma(x)\bar{f}(y; \delta_x, g_2) \right] , \quad \forall y \end{array} \right. \tag{3.13c}$$

$$\left\{ \begin{array}{l} \sum_x (\beta(x) - z(x))F_{\mathcal{A}_l(g_2)}(\delta_x, g_2) = 0 , \quad l = 1, \dots, v(g_2) \end{array} \right. \tag{3.13c}$$

$$\left\{ \begin{array}{l} \sum_x (\beta(x) - z(x)) \sum_{t=1}^{\infty} P_{\delta_x}^{g_2} \{X_t = y\} \geq 0 , \quad \forall y \in \mathcal{T}(g_2) , \end{array} \right. \tag{3.13d}$$

where γ is a vector on X with

$$\gamma(y) = \begin{cases} 0, & y \in \mathcal{F}(g_2) \\ c_l, & y \in \mathcal{R}_l(g_2), \end{cases}$$

with c_l a constant satisfying

$$c_l \geq \max_{y \in \mathcal{R}_l(g_2)} \frac{-\sum_x [\beta(x)D(y; \delta_x, g_2) + (\bar{f}(x; \beta, g_2) - z(x)) \sum_{i=1}^{\infty} P_{\delta_x}^{g_2} \{X_i = y\}]}{\sum_{x \in \mathcal{R}_l(g_2)} \bar{f}(y; \delta_x, g_2)},$$

$$l = 1, \dots, v(g_2).$$

Notice, that the vector γ is constant on each positive recurrent class in the Markov chain induced by g_2 . Equations (3.13c) and (3.13d) are necessary to ensure that we only get positive solutions ζ to system (3.13).

If $\bar{\beta} = \beta$ and $g_1 = g_2$, then $\zeta(y) = \sum_x [\beta(x)D(y; \delta_x, g_2) + \gamma(x)\bar{f}(y; \delta_x, g_2)]$. This is in fact the formula used by Kallenberg to show, that we can construct a feasible solution (x, ζ) of Π_β for any stationary policy.

Proof of Theorem 4: First choose $(z, \zeta) \in \Pi_\beta$. By the construction in [10] Π_β is not empty. Define two stationary policies g_1, g_2 in the following way:

$$p_{a|y}^{g_1} := \begin{cases} \frac{z(y, a)}{z(y)}, & \text{if } z(y) > 0 \\ \frac{\zeta(y, a)}{\zeta(y)}, & \text{if } z(y) = 0, \zeta(y) > 0 \\ \text{arbitrarily,} & \text{otherwise,} \end{cases}$$

$$p_{a|y}^{g_2} := \begin{cases} \frac{\zeta(y, a)}{\zeta(y)}, & \text{if } \zeta(y) > 0 \\ \frac{z(y, a)}{z(y)}, & \text{if } \zeta(y) = 0, z(y) > 0 \\ \text{arbitrarily,} & \text{otherwise.} \end{cases}$$

So, g_1 and g_2 only differ in states y with $z(y), \zeta(y) > 0$ and $z(y, a)/z(y) \neq \zeta(y, a)/\zeta(y)$, or possibly in states y with $z(y) = \zeta(y) = 0$.

By the construction of g_1 and g_2 and by (2.1a) (resp. (2.1b)) the vectors z and ζ satisfy the following linear equations.

$$z(y) - \sum_x z(x) P_{xy}^{g_1} = 0, \quad y \in \mathbf{X} \quad (3.14a)$$

$$z(y) + \zeta(y) - \sum_x \zeta(x) P_{xy}^{g_2} = \beta(y), \quad y \in \mathbf{X}. \quad (3.14b)$$

By the remark following (2.1), $\sum_y z(y) = 1$. Consequently z is an invariant probability measure for the Markov chain generated by policy g_1 and so $z = \tilde{f}(z, g_1)$. In fact, $z = \tilde{f}(\tilde{\beta}, g_1)$ for any initial distribution $\tilde{\beta}$ so that $\sum_x \tilde{\beta}(x) R_{i(g_1)} F(\delta_x, g_1) = \sum_{y \in R_i(g_1)} z(y)$.

Notice, that $\tilde{f}(\cdot; \delta_x, g_2)$ is an invariant measure under g_2 . Hence for any vector η on \mathbf{X} , $z, \tilde{\zeta}(\bullet) := \zeta(\bullet) + \sum_x \eta(x) \tilde{f}(\bullet; \delta_x, g_2)$ are a solution of (3.14) as well. Let $B(g_2)$ be a set of reference states in the Markov chain generated by g_2 . We write $b_l := \mathcal{R}_l(g_2) \cap B(g_2)$, $l = 1, \dots, v(g_2)$. As the vector η we choose $\eta(y) := 0$ for $y \in \mathcal{T}(g_2)$, and $\eta(y) := -\zeta(b_l) / \sum_{x \in \mathcal{R}_l(g_2)} \tilde{f}(b_l; \delta_x, g_2)$ for $y \in \mathcal{R}_l(g_2)$. Note that $\eta(x) \tilde{f}(y; \delta_x, g_2) = 0$ for $y \in \mathcal{R}_l(g_2)$, $x \notin \mathcal{R}_l(g_2)$. Thus, we have chosen η to be constant on each positive recurrent class and 0 otherwise, in such a way that $\tilde{\zeta}(b_l) = 0$. Plugging $\tilde{\zeta}$ into equation (3.14b) we get

$$\tilde{\zeta}(y) = \beta(y) - z(y) + \sum_{x \notin B(g_2)} \tilde{\zeta}(x) P_{xy}^{g_2}.$$

Iterating this we obtain

$$\begin{aligned} \tilde{\zeta}(y) &= \beta(y) - z(y) + \sum_{x \notin B(g_2)} \left[\beta(x) - z(x) + \sum_{v \notin B(g_2)} \tilde{\zeta}(v) P_{vx}^{g_2} \right] P_{xy}^{g_2} \\ &= \beta(y) - z(y) + \sum_{x \notin B(g_2)} (\beta(x) - z(x)) P_{xy}^{g_2} \\ &\quad + \sum_{v \notin B(g_2)} \tilde{\zeta}(v) P_{v_y}^{g_2} \{X_3 = y, X_2 \notin B(g_2)\} \\ &= \beta(y) - z(y) + \sum_{x \notin B(g_2)} (\beta(x) - z(x)) \sum_{i=2}^T P_{ix}^{g_2} \{X_i = y, X_s \notin B(g_2), 2 \leq s < i\} \\ &\quad + \sum_{v \notin B(g_2)} \tilde{\zeta}(v) P_{v_y}^{g_2} \{X_{T+1} = y, X_s \notin B(g_2), 2 \leq s \leq T\}. \end{aligned}$$

Set $B(g_2)$ is reached with probability 1 from any initial state, when policy g_2 is used. So, taking the limit for $T \rightarrow \infty$ in the last equation yields

$$\begin{aligned} \tilde{\zeta}(y) &= \beta(y) - z(y) \\ &+ \sum_{x \notin B(g_2)} (\beta(x) - z(x)) \sum_{t=2}^{\infty} P_{\delta_x}^{g_2} \{X_t = y, X_s \notin B(g_2), 2 \leq s < t\} . \end{aligned} \quad (3.15)$$

We note that in (3.15), we interchanged limit with summation. This is justified since the second summation is finite. Indeed,

$$\begin{aligned} P_{\delta_x}^{g_2} \{X_t = y, X_s \notin B(g_2), 2 \leq s < t\} &\leq P_{\delta_x}^{g_2} \{X_s \notin B(g_2), 2 \leq s < t\} \\ &= P_{\delta_x}^{g_2} \left\{ \frac{1}{t-2} \sum_{s=2}^{t-1} \delta_{X(s)}(B(g_2)) = 0 \right\} . \end{aligned}$$

It follows from the Large deviation principle for Markov chains (see e.g. [6]) that the latter probability behaves asymptotically (for large t) as $\exp(-tI)$ where $I > 0$ is some constant (the rate function). This implies the finiteness of the second summation in (3.15).

For $y \notin B(g_2)$ the probabilities in the last expression in (3.15) equal the taboo transition probabilities for the taboo set $B(g_2)$. For $y \in B(g_2)$ they equal the first hitting probabilities of y . As a consequence we obtain the following expressions for $\tilde{\zeta}$.

$$\tilde{\zeta}(y) = \begin{cases} \sum_{x \notin B(g_2)} (\beta(x) - z(x)) \sum_{t=1}^{\infty} P_{\delta_x}^{g_2} \{X_t = y\} , & y \notin B(g_2) \\ \sum_x (\beta(x) - z(x)) F_y(\delta_x, g_2) , & y \in B(g_2) . \end{cases} \quad (3.16a)$$

$$(3.16b)$$

As $F_{b_1}(\delta_x, g_2) = F_{\mathcal{R}_1(g_2)}(\delta_x, g_2)$ for $x \in \mathbf{X}$, $F_{b_1}(\delta_x, g_2) = 1$ for $x \in \mathcal{R}_1(g_2)$ and $\tilde{\zeta}(b_1) = 0$ (by definition), (3.13c) follows directly from (3.16b).

Obviously (3.16a) equals 0 for $y \in B(g_2)$. Consequently, the first summation in (3.16a) can be extended to *all* states. Thus we get for $y \in \mathcal{R}_1(g_2)$:

$$\begin{aligned} \tilde{\zeta}(y) &= \sum_{x \in \mathbf{X}} (\beta(x) - z(x)) \sum_{t=1}^{\infty} P_{\delta_x}^{g_2} \{X_t = y\} \\ &\quad - (\beta(b_1) - z(b_1)) \sum_{t=2}^{\infty} P_{\delta_{b_1}}^{g_2} \{X_t = y\} \\ &= \sum_{x \in \mathbf{X}} (\beta(x) - z(x)) \sum_{t=1}^{\infty} P_{\delta_x}^{g_2} \{X_t = y\} - (\beta(b_1) - z(b_1)) \frac{\bar{f}(y; \delta_{b_1}, g_2)}{\bar{f}(b_1; \delta_{b_1}, g_2)} \end{aligned}$$

$$\begin{aligned}
 &= \sum_{v \in X} \beta(v) \sum_x (\delta_x(v) - \bar{f}(x; \delta_v, g_2)) \sum_t B(g_2) P_{\delta_x}^{g_2} \{X_t = y\} \\
 &\quad + \sum_{x \in X} (\bar{f}(x; \beta, g_2) - z(x)) \sum_t B(g_2) P_{\delta_x}^{g_2} \{X_t = y\} \\
 &\quad - \frac{\beta(b_1) - z(b_1)}{\bar{f}(b_1; \delta_{b_1}, g_2)} \bar{f}(y; \delta_{b_1}, g_2) \\
 &= \sum_{x \in X} \beta(x) D(y; \delta_x, g_2) + \sum_{x \in X} (\bar{f}(x; \beta, g_2) - z(x)) \sum_t B(g_2) P_{\delta_x}^{g_2} \{X_t = y\} \\
 &\quad + \left\{ \sum_{v \in X} \beta(v) \sum_{x \in X} (\delta_x(v) - \bar{f}(x; \delta_v, g_2)) \right. \\
 &\quad \cdot \sum_t \sum_{w \in X} B(g_2) P_{\delta_x}^{g_2} \{X_t = w\} F_{\mathcal{A}_t(g_2)}(\delta_w; g_2) \\
 &\quad \left. - \frac{\beta(b_1) - z(b_1)}{\bar{f}(b_1; \delta_{b_1}, g_2)} \right\} \bar{f}(y; \delta_{b_1}, g_2). \tag{3.17}
 \end{aligned}$$

For the second equality we use the renewal theorem, if $y \neq b_1$, and if $y = b_1$ we use $\sum_{t=2}^{\infty} B(g_2) P_{\delta_{b_1}}^{g_2} \{X_t = y\} = 1 = \bar{f}(b_1, \delta_{b_1}, g_2) / \bar{f}(b_1, \delta_{b_1}, g_2)$. For the 4th equality we use formula (3.12) for the deviation matrix. For $y \in \mathcal{T}(g_2)$, (3.17) holds since $\bar{f}(y; \delta_{b_1}, g_2) = 0$. The derivation of the other equalities is similar to the foregoing ones.

Recall that $\zeta(y) = \tilde{\zeta}(y) - \sum_{x \in X} \eta(x) \bar{f}(y; \delta_x, g_2)$. So, we choose the vector γ equal to

$$\gamma(y) := \begin{cases} -\eta(y) + \frac{1}{|\mathcal{A}_1(g_2)|} \left\{ \sum_{v \in X} \beta(v) \sum_{x \in X} (\delta_x(v) - \bar{f}(x; \delta_v, g_2)) \right. \\ \quad \times \sum_t \sum_{w \in X} B(g_2) P_{\delta_x}^{g_2} \{X_t = w\} F_{\mathcal{A}_t(g_2)}(\delta_w; g_2) - \frac{\beta(b_1) - z(b_1)}{\bar{f}(b_1; \delta_{b_1}, g_2)} \Big\}, & y \in \mathcal{A}_1(g_2) \\ -\eta(y) = 0, & y \in \mathcal{T}(g_2). \end{cases}$$

Then by virtue of (3.17) and our choice of η, z satisfies (3.13b) for this choice of γ . Moreover, $\gamma(y)$ is constant on each positive recurrent class. As $\zeta(y) \geq 0$ by assumption, we have obtained that γ satisfies the conditions of the Theorem. For transient states y , $\tilde{\zeta}(y) = \zeta(y) \geq 0$. Using expression (3.16a) for $\tilde{\zeta}(y)$ gives (3.13d).

The proof of the reverse implication is straightforward. Here we have to use, that $\sum_v D(v; \delta_x, g_2)(\delta_y(v) - P_{vy}^{g_2}) = \delta_y(x) - \bar{f}(y; \delta_x, g_2)$ (cf. Kallenberg [10, p. 29], Spieksma [16, p. 68]). ■

4 Properties of the Deviation Measure

4.1 Bounds

Unlike the state-action frequencies, which sum to one under any arbitrary policy, the deviation measures are not bounded in the class of policies. This is demonstrated in the following examples.

Example 1: This example shows that $\tilde{r}(\beta, u)$, and hence $Y(\beta, u)$, may be unbounded in U . Assume that there exists a state v that is recurrent under all the stationary deterministic policies. Consider the following set of policies $\{u(t)\}_{t=1}^{\infty}$: let $u(t)$ be a policy that follows one stationary policy till time t and then switches to another one v . It follows that $\bar{F}(\beta, u(t)) = \{\bar{f}(\beta, v)\}$ and $\tilde{R}(\bar{f}, \beta, u(t)) = \{\tilde{r}(\beta, u(t))\}$ are singletons for all t ; moreover $|\tilde{r}(y, a; \beta, u(t))|$ grows like $O(t)$ and is thus unbounded. Note however that for any t , $\tilde{r}(y, a; \beta, u(t))$ is finite. This follows from the fact that for any stationary policy u and any initial distribution β , $\tilde{r}(y, a; \beta, u)$ is finite (and $\tilde{R}(\beta, u) = \{\tilde{r}(\beta, u)\}$ is a singleton).

Example 2: This example show that $\tilde{r}(\beta, u)$ may be infinite, and $Y(\beta, u)$ not defined. Let $X = \{x\}$, and $A = \{a, b\}$. Let u be the policy that chooses a at time t with probability $1/t$. Then $\bar{F}(x, a; x, u) = \{0\}$, and $\tilde{r}^t(x, a; x, u) = \sum_{s=1}^t t^{-1}$ tends to infinity as $t \rightarrow \infty$. Hence $\tilde{R}(x, a; x, u) = \{\infty\}$. Similarly, $\tilde{R}(x, b; x, u) = \{-\infty\}$, and $Y(x, u)$ can thus not be defined.

4.2 Uniqueness of the Deviation Measure

Lemma 5: Assume that under every deterministic stationary policy there exists one recurrent class (plus possibly a transient set). Let u and v be two policies such that

$$\bar{f}(\beta, v) = \bar{f}(\beta, u)$$

for some $\bar{f}(\beta, v) \in \bar{F}(\beta, v)$ and some $\bar{f}(\beta, u) \in \bar{F}(\beta, u)$. Assume moreover that

there are some finite $\tilde{r}(\beta, u) \in \tilde{R}(\bar{f}(\beta, u), \beta, u)$ and $\tilde{r}(\beta, v) \in \tilde{R}(\bar{f}(\beta, v), \beta, v)$ that satisfy

$$\frac{\tilde{r}(y, a; \beta, v)}{\sum_a \tilde{r}(y, a; \beta, v)} = \frac{\tilde{r}(y, a; \beta, u)}{\sum_a \tilde{r}(y, a; \beta, u)} =: \gamma_y^a$$

for all y and a . Then $\tilde{r}(y, a; \beta, u) = \tilde{r}(y, a; \beta, v)$.

Proof: Let P^* be the matrix whose components are given by: $P_{xy}^* = \sum_a \gamma_x^a P_{xay}$. It follows then from (3.11) that

$$0 = \bar{f}(y; \beta, u) + \sum_x \tilde{r}(z; \beta, u)[\delta_y(z) - P_{zy}^*] - \beta(y)$$

(where $\tilde{r}(z; \beta, u) := \sum_a \tilde{r}(z, a; \beta, u)$). Since the same holds for v too, we obtain by subtraction:

$$\sum_z [\tilde{r}(z; \beta, u) - \tilde{r}(z; \beta, v)][\delta_y(z) - P_{zy}^*] = 0$$

Hence $[\tilde{r}(\beta, u) - \tilde{r}(\beta, v)]$ is an eigenvector of the matrix $\{P_{zy}^*\}$ which corresponds to the eigenvalue 1. Since P^* is an irreducible stochastic matrix, it follows from the Perron–Frobenius Theorem (see e.g. [13, p. 22]) that there exists some constant c such that $\tilde{r}(z; \beta, u) - \tilde{r}(z; \beta, v) = c\pi^*(z)$ where π^* is the vector of stationary distribution that corresponds to P^* . But since $\tilde{r}(z; \beta, u)$ sum to 0, it follows that $c = 0$ from which the Lemma follows. ■

5 Linear Program for Constrained MDPs

In this Section we study the Linear Program method applied for the constrained Markov Decision Process [9]. An alternative method based on a Lagrange approach for solving such problems in the case of a single constraint was introduced in [3] (later extended to countable state space in [14], [15]).

The aim of this Section is to generalize the formulation of the Lagrange method to several constraints and then to use it for obtaining a new derivation of the Linear Program method of Hordijk and Kallenberg [9]. Our new derivation seems to be more natural and straightforward than the previous method [9], where the paradigm was to introduce the LP and then to prove that the optimal policy and optimal value of the control problem are appropriately related to the optimal solution and value of the LP.

To define the constrained problem, we define a set of K (real valued) cost function $d^k: X \times A \rightarrow \mathcal{R}, k = 1, \dots, K$, and define the average costs associated with a policy u and with an initial distribution β on X :

$$D_\beta^k(u) = \overline{\lim}_{t \rightarrow \infty} \frac{1}{t} - E_\beta^u \left[\sum_{s=1}^t d^k(X_s, A_s) \right]$$

For a given vector of real numbers $V_k, k = 1, \dots, K$ we define the constrained problem COP: find a policy $u \in U$ that minimizes $C_\beta(u)$ subject to $D_\beta^k(u) \leq V_k, k = 1, \dots, K$, for a given initial distribution β .

A policy that satisfies $D_\beta^k(u) \leq V_k, k = 1, \dots, K$ is said to be feasible. Let C_β be the optimal value of COP. A feasible policy u that achieves $C_\beta(u) = C_\beta$, is said to be optimal for COP.

Introduce the Lagrange minimax problem L_{minimax} :

$$\min_u \sup_{v \geq 0} \left[C_\beta(u) + \sum_k v_k (D_\beta^k(u) - V_k) \right]$$

Let C_β^{minimax} be the optimal value of L_{minimax} . It is easily seen that L_{minimax} is equivalent to COP (in the sense that the optimal policy and optimal value are the same).

Next we introduce the Linear Program used for calculating the optimal value and policy for COP [9]:

LPC: Find $z, \zeta \in \mathcal{R}^{X \times A}$, that minimize $c \cdot z$ subject to

$$(z, \zeta) \in \Pi_\beta,$$

$$d^k \cdot z \leq V_k$$

where Π_β is given in Section 2. It is shown in [9] that the optimal value of COP is equal to the optimal value of LPC, and an optimal Markovian policy is constructed from the optimal solution of LPC.

We show below that LPC can be derived from L_{minimax} . We first note that we may restrict the minimization in COP to the class of policies U^* under which $\bar{F}(\beta, u)$ is a singleton, see e.g. [9]. Note that this class contains the stationary policies.

Define L_{maximin} :

$$\sup_{v \geq 0} \min_{u \in U^*} \left[C_\beta(u) + \sum_k v_k (D_\beta^k(u) - V_k) \right].$$

Let $C_\beta^{maximin}$ be the optimal value of $L_{maximin}$. It can be shown that the min and sup can be interchanged, so that $C_\beta^{maximin} = C_\beta^{minimax}$. We leave the proof to the appendix. In order to obtain LPC we first consider the problem of $\min_{u \in U^*} [C_\beta(u) + \sum_k v_k (D_\beta^k(u) - V_k)]$ for a given v . It follows from [8] that an optimal stationary deterministic policy for that problem can be obtained and the optimal value is given by LP': find $\lambda_y, g_y, y \in X$, that

$$\max \left[\sum_y \beta(y) g_y - \sum_k v_k V_k \right] \quad (5.1)$$

s.t.

$$\lambda_y + g_y \leq c(y, a) + \sum_k v_k d^k(y, a) + \sum_v P_{yav} \lambda_v \quad \forall y, a \quad (5.2)$$

$$g_y \leq \sum_{v \in X} P_{yav} g_v \quad \forall y, a \quad (5.3)$$

We may conclude that the optimal value of COP is given by solving the LP: find $v_k, k = 1, \dots, K, \lambda_y$, and $g_y, y \in X$, that

$$\max_{g, \lambda, v \geq 0} \left[\sum_y \beta(y) g_y - \sum_k v_k V_k \right]$$

s.t.

$$\lambda_y + g_y \leq c(y, a) + \sum_k v_k d^k(y, a) + \sum_v P_{yav} \lambda_v \quad \forall y, a$$

$$g_y \leq \sum_{v \in X} P_{yav} g_v \quad \forall y, a$$

The dual of this LP is LPC.

Remark: The results of this Section extend easily to the discounted criteria, even with countable state space. They also extends easily to the countable case under a strong ergodic assumption (the unichain assumption). In both cases, there is no need of equation (5.3), and of (2.1.b) and (2.1.d). Thus we need not consider the decision variables ζ in the Linear Program (2.1). (See e.g. [1], [2], [16] for alternative derivations of such LPs).

Appendix

Interchanging the min and max in Section 5

We show in this appendix that $C_\beta^{\text{minimax}} = C_\beta^{\text{maximin}}$. Since we restrict to $u \in U^*$ it easily follows that

$$C_\beta(u) + \sum_k v_k(D_\beta^k(u) - V_k) = \sum_{y,a} \bar{f}(y, a; \beta, u) \left[c(y, a) + \sum_k v_k(d^k(y, a) - V_k) \right]$$

and hence

$$\begin{aligned} \min_u \sup_{v \geq 0} \left[C_\beta(u) + \sum_k v_k(D_\beta^k(u) - V_k) \right] \\ = \min_{z \in L_\beta(U^*)} \sup_{v \geq 0} \sum_{y,a} \bar{f}(y, a; \beta, u) \left[c(y, a) + \sum_k v_k(d^k(y, a) - V_k) \right] \end{aligned}$$

Since $L_\beta = L_\beta(U^*)$ is convex and compact (this is a straightforward generalization of [7] p. 93–94, [9] who prove it for β of the form of δ_x), and since the set $v \geq 0$ is convex and compact in the compactified space $[\mathcal{R} \cup \{\infty\}]^K$, it follows from the minimax Theorem (e.g. [11] p. 208) that the min and sup can be interchanged from which it finally follows that $C_\beta^{\text{minimax}} = C_\beta^{\text{maximin}}$.

Acknowledgement: We authors wish to thanks L.C.M. Kallenberg for fruitful discussions and helpful comments.

References

- [1] Altman E, Schwartz A (1991) Markov decision problems and state-action frequencies. SIAM J Control and Optimization 29/4:786–809
- [2] Altman E (1994) Denumerable constrained Markov decision problems and finite approximations. Math of OR 19:169–191
- [3] Beutler FJ, Ross KW (1985) Optimal policies for controlled Markov chains with a constraint. Math Anal Appl 112:236–252

- [4] Borkar VS (1988) A convex analytic approach to Markov decision processes. *Probab Th Rel Fields* 78:583–602
- [5] Borkar VS (1991) *Topics in controlled Markov chains*. Pitman
- [6] Dembo A, Zeitouni O (1993) *Large deviations techniques and applications*. Jones and Bartlett
- [7] Derman C (1970) *Finite state markovian decision processes*. Academic Press
- [8] Hordijk A, Kallenberg LCM (1979) Linear programing and Markov decision chains. *Management Science* 25/4:352–362
- [9] Hordijk A, Kallenberg LCM (1984) Constrained undiscounted stochastic dynamic programming. *Math of OR* 9:277–298
- [10] Kallenberg LCM (1983) *Linear programming and finite markovian control problems*. Math Centre Tracts 148 Amsterdam
- [11] Luenberger DG (1968) *Optimization by vector space methods*. John Wiley
- [12] Ross K, Varadarajan R (1991) Multichain Markov Decision Processes with a Sample Path Constraint: A Decomposition Approach. *MOR* 16/1:195–207
- [13] Seneta E (1981) *Non-negative martices and markov chains*. Springer-Verlag
- [14] Sennott LI (1991) Constrained discounted Markov decision chains. *Probability in the Engineering and Informational Sciences* 5:463–475
- [15] Sennott LI (1993) Constrained average cost markov decision chains. *Probability in the Engineering and Informational Sciences* 7:69–83
- [16] Spieksma F (1990) *Geometrically ergodic markov chains and the optimal control of queues*. Ph D thesis Leiden

Received: February 1994