VU Research Portal

An approximation approach for the deviation matrix of continuous-time Markov processes with application to Markov decision theory

Heidergott, B.F.; Hordijk, A.; Leder, N.

published in

Operations Research
2010

DOI (link to publisher)

[10.1287/opre.1090.0786](https://doi.org/10.1287/opre.1090.0786)

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Heidergott, B. F., Hordijk, A., & Leder, N. (2010). An approximation approach for the deviation matrix of continuous-time Markov processes with application to Markov decision theory. *Operations Research*, 58, 918-932. <https://doi.org/10.1287/opre.1090.0786>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

An Approximation Approach for the Deviation Matrix of Continuous-Time Markov Processes with Application to Markov Decision Theory

Nicole Leder

Department of Mathematics, University of Hamburg, Hamburg 20146, Germany, nicole.leder@studium.uni-hamburg.de

Bernd Heidergott

Department of Econometrics and Operations Research, and Tinbergen Institute, Vrije Universiteit Amsterdam, Amsterdam 1081 HV, The Netherlands, bheidergott@feweb.vu.nl

Arie Hordijk

Mathematical Institute, Leiden University, Leiden 2300 RA, The Netherlands, hordijk@math.leidenuniv.nl

We present an update formula that allows the expression of the deviation matrix of a continuous-time Markov process with denumerable state space having generator matrix Q^* through a continuous-time Markov process with generator matrix Q . We show that under suitable stability conditions the algorithm converges at a geometric rate. By applying the concept to three different examples, namely, the M/M/1 queue with vacations, the M/G/1 queue, and a tandem network, we illustrate the broad applicability of our approach. For a problem in admission control, we apply our approximation algorithm to Markov decision theory for computing the optimal control policy. Numerical examples are presented to highlight the efficiency of the proposed algorithm.

Subject classifications: Markov decision processes; deviation matrix; algorithm; probability; Markov processes; queues; algorithms.

Area of review: Stochastic Models.

History: Received September 2008; revision received July 2009; accepted September 2009. Published online in *Articles in Advance* February 26, 2010.

1. Introduction

Continuous-time Markov processes are a common tool for analyzing complex systems such as telecommunication networks, computer systems, or call centers. Typical models are the M/Ph/1 queue, where Ph indicates a phase-type distribution used for approximating a general service time distribution, or a $\cdot$ /M/1 tandem network. Whereas the stationary distribution of these systems is known, the deviation matrix that yields the so-called value function, an important tool in Markov decision theory, is only known for some special queues that are in essence variations of the M/M/c/L queue; see Koole (1998). In fact, the deviation matrix can only be computed explicitly for Markovian queues with a one-dimensional state space. Unfortunately, many important systems have a multidimensional state space, and the deviation matrix defies computation.

In this paper we provide an update formula that allows us to approximately compute the deviation matrix of models with a two-dimensional state space, and we illustrate our approach by the M/M/1 queue with vacations, the M/Ph/1 queue, and a tandem network. Starting point of our three examples is the M/M/1 queue, for which the deviation matrix can be obtained in a closed form, see Koole (1998).

In order to approximate the deviation matrix of the, say, M/Ph/1 queue, we initially enlarge the state space of the M/M/1 queue by transient states, so that the M/M/1 and the M/Ph/1 queue are defined on the same two-dimensional state space. As we show in this paper, the deviation matrix of an M/M/1 queue with additional transient states can still be computed explicitly. By inserting this adjusted deviation matrix into our update formula, we approximately compute the unknown deviation matrix of the M/Ph/1 queue.

Series expansions for Markov processes, such as our update formula, go back to Schweitzer (1968), where discrete-time, finite-state Markov chains have been studied; see also Heidergott et al. (2007). Our update formula is derived for continuous-time Markov processes on denumerable state space, without assuming uniformizability. The key contribution of our paper is that we can deal with two-dimensional infinite state spaces, which is a breakthrough in the computation of the deviation matrix for processes like the M/Ph/1 queue, for which it had not been possible to give an explicit representation. We complement our concept by the approximation approach for general distributions via phase-type distributions in order to extend our algorithm to the M/G/1 queue. To illustrate the relevance of the update formula for problems in control, we

will embed our approach into Markov decision processes (MDP). More specifically, we compute the average optimal admission policy for a call center in which each admitted customer raises a certain reward while the number of customers in the system causes certain holding cost.

The paper is organized as follows. In §2 basic properties of the deviation matrix are presented and v -norm ergodicity is introduced. In §3 the update formula for the deviation matrix is established and afterwards applied to three different examples in §4. The application to admission control is provided in §5. Eventually, we identify topics of further research in §6. The more-technical results and derivations of the respective deviation matrices are provided in an online companion available at <http://or.journal.informs.org/>.

2. Preliminaries

2.1. Basic Properties of the Deviation Matrix

Throughout this paper we will denote the infinitesimal generator of a continuous-time Markov process $\mathcal{X}$ that exists on a denumerable state space S by Q and its stationary distribution by π . The associated transition probability matrix with elements $p(i, j; t)$ representing the probability to go from state i to j within the time t will be denoted by $P(t)$. The ergodic projector, i.e., a matrix with all rows equal to π , will be represented by Π . Furthermore, we elementwise define by

$$d(i, j) = \int_0^\infty (p(i, j; t) - \pi(j)) dt, \quad i, j \in S, \quad (1)$$

the deviation matrix D of $\mathcal{X}$ that is said to exist whenever all integrals in (1) are finite. According to Coolen-Schrijner and van Doorn (2002) and Heidergott et al. (2009), these matrices satisfy the properties summarized in the following lemma.

LEMMA 1. *If it exists, the ergodic matrix Π of a continuous-time Markov process $\mathcal{X}$ with infinitesimal generator Q satisfies*

(i) $A\Pi = 0$ for all conservative matrices $A \in \mathbb{R}^{S \times S}$, and even $\Pi Q = 0$,

(ii) $B\Pi = \Pi$ for all stochastic matrices $B \in \mathbb{R}^{S \times S}$.

If the corresponding deviation matrix exists, it holds that

(iii) $\Pi D = 0$,

(iv) $-QD = -DQ = I - \Pi$,

(v) $D\mathbf{1} = 0$,

where I denotes the $S \times S$ identity matrix and $\mathbf{1}$ the column-vector with all S entries equal to 1.

In this paper we focus on continuous-time processes. It is, however, worth noting that our analysis carries over to the so-called sampled, or subordinate, discrete-time chain by using uniformization theory, see Kijima (1997) for details. Moreover, any discrete-time Markov chain $\tilde{\mathcal{X}}$ with transition matrix $\tilde{P}$ can be translated into a continuous-time process on the same state space S with generator matrix

$Q = \tilde{P} - I$, with I the identity matrix of appropriate size. Furthermore it holds for the associated ergodic projectors $\tilde{\Pi} = \Pi$ and the deviation matrices $\tilde{D} = D$. Hence, all formulas presented in this paper in terms of Q , D , and Π can be translated to a discrete-time chain.

2.2. Geometric Ergodicity

The main tool for our analysis is the weighted supremum norm, also called v -norm, denoted by $\|\cdot\|_v$, where v is some vector with elements $v_i \geq 1$ for all $i \in S$, and for any $w \in \mathbb{R}^S$

$$\|w\|_v \stackrel{\text{def}}{=} \sup_{i \in S} \frac{|w(i)|}{v(i)}.$$

For a matrix $A \in \mathbb{R}^{S \times S}$ the v -norm is given by

$$\|A\|_v \stackrel{\text{def}}{=} \sup_{i, \|w\|_v \leq 1} \frac{\sum_{j=1}^S |A(i, j)w(j)|}{v(i)},$$

which implies

$$\max_{j \in S} |A(i, j)| \leq \|A\|_v v(i), \quad i \in S.$$

Note that v -norm convergence to 0 implies elementwise convergence to 0. With the help of the above concepts, v -geometric ergodicity (also called v -normed ergodicity) of the transition matrix $P(t)$ of a continuous-time Markov process $\mathcal{X}$ can be introduced as follows.

DEFINITION 1. The Markov process $\mathcal{X}$ is v -geometric ergodic if $c < \infty$ and $\beta < 1$ exist such that

$$\|P(t) - \Pi\|_v \leq c\beta^t,$$

for all $t \geq 0$.

Note that

$$\|D\|_v \leq \int_0^\infty \|P(t) - \Pi\|_v dt \leq c \int_0^\infty \beta^t dt = -\frac{c}{\ln(\beta)}$$

and it is straightforward to check that geometric v -norm ergodicity implies existence of $\|D\|_v$ by assuring the finiteness of its elements. Unfortunately, geometric v -norm ergodicity is almost impossible to check in a direct way. One of the reasons is that $P(t)$ is in general not known in explicit form. Therefore, we use a different representation for the deviation matrix, which follows directly from the properties of Π and D presented in Lemma 1

$$D = (\Pi - Q)^{-1} - \Pi. \quad (2)$$

3. Series Representation of Denumerable Markov Processes

In this section we will present the basic formula of our concept. In §3.1 we derive the series expansion of the deviation matrix, and in §3.2 we present the associated sufficient condition assuring its convergence. Furthermore, in §3.3 we present the algorithm that allows establishment of a precision ϵ_{dev} up to which the deviation matrix D^* shall be approximated.

3.1. The Update Formula

Let Q and Q^* be the infinitesimal generators of two different continuous-time Markov processes. Denote by Π , Π^* , D , and D^* the respective ergodic matrices and deviation matrices. We know from Lemma 1(iv) that it holds

$$-QD = I - \Pi.$$

Adding Q^*D on both sides of the above equation gives us

$$(Q^* - Q)D = Q^*D + I - \Pi, \quad (3)$$

which we multiply by D^* to get

$$D^*(Q^* - Q)D = \underbrace{D^*Q^*D}_{=\Pi^* - I} + \underbrace{D^*}_{=0} - \underbrace{D^*\Pi}_{=0}.$$

By solving this equation for D^* , we obtain

$$D^* = (I - \Pi^*)D + D^*(Q^* - Q)D. \quad (4)$$

Inserting (4) repeatedly into its right-hand side yields

$$D^* = (I - \Pi^*)D \sum_{k=0}^n ((Q^* - Q)D)^k + D^*((Q^* - Q)D)^{n+1}, \quad (5)$$

which we separate into the series approximation of degree n

$$H(n) = (I - \Pi^*)D \sum_{k=0}^n ((Q^* - Q)D)^k \quad (6)$$

and the corresponding remainder term

$$R(n) = D^*((Q^* - Q)D)^{n+1}.$$

Equation (5) provides a useful tool to approximately compute the deviation matrix of a process $\mathcal{X}^*$ whenever the remainder term tends to zero for $n \rightarrow \infty$. Sufficient conditions for the convergence of the remainder term, and therefore the convergence of the series approximation, will be provided in the following section. We will show later on that by finding a finite N and an associated constant $\delta_N < 1$ such that it holds

$$\|((Q^* - Q)D)^N\|_v < \delta_N,$$

we assure that $\|D^* - H(n)\|_v$ tends to zero at an exponential rate. Provided that $\|f\|_v < \infty$, it holds that

$$|(D^*f)_i - (H(n)f)_i| \leq \|D^* - H(n)\|_v \|f\|_v \inf_{i \in S} v(i); \quad (7)$$

see Heidergott et al. (2009) for details. In applications, $\inf_{i \in S} v(i)$ is typically equal to 1, and therefore (7) reduces to

$$|(D^*f)_i - (H(n)f)_i| \leq \|D^* - H(n)\|_v \|f\|_v.$$

From (7) we obtain that the convergence of the update formula (6) directly implies that $H(n)f$ tends to D^*f exponentially fast for any finite-normed f . Now suppose that f is a cost function bounded for some appropriate v . Then, finding an approximation for D^* yields an approximation of the so-called value function D^*f , which is a basic tool in MDPs.

REMARK 1. Common synonyms of the value function are the *bias* (see, e.g., Feinberg and Schwartz 2002, Puterman 1994), the *performance potential* (see e.g., Cao 2007, Zhang et al. 2008), or the *relative cost* (see e.g., Bertsekas 2005). According to Feinberg and Schwartz (2002) and Koole (1998), the i th entry of the value function can be interpreted as the total difference in costs between starting in state i and the stationary version. This interpretation corresponds to the definition of the deviation matrix given by (1), where it can easily be seen that by multiplying (1) with a cost vector f , one compares the costs arising over the entire time frame caused by starting in a certain state i given by $\sum_{j \in S} p(i, j; t)f(j)$ with the stationary cost $\sum_{j \in S} \pi(j)f(j)$.

3.2. Convergence of the Series Expansion

To ensure the efficiency of our approximation for the deviation matrix, we demand geometric-fast convergence of the remainder $R(n)$ to zero as n tends to ∞ . As we will show in a subsequent lemma, geometric-fast convergence of $R(n)$ to zero is implied by the following condition.

[C] *There exists a finite number N such that we can find $\delta_N \in (0, 1)$, which satisfies*

$$\|((Q^* - Q)D)^N\|_v < \delta_N,$$

and we set

$$c_{\delta_N}^v \stackrel{\text{def}}{=} \frac{1}{1 - \delta_N} \left\| \sum_{k=0}^{N-1} ((Q^* - Q)D)^k \right\|_v.$$

The factor $c_{\delta_N}^v$ in condition [C] allows establishment of an upper bound for the remainder term that is independent of D^* . Denote by $T(k) = (I - \Pi^*)D((Q^* - Q)D)^k$ the k th element of the series in (6). Then we can summarize the main conditions in the following lemma.

LEMMA 2. *Under [C] it holds that for any $v \geq 1$:*

- (i) $\|R(k-1)\|_v \leq c_{\delta_N}^v \|T(k)\|_v$ for any k .
- (ii) $\lim_{k \rightarrow \infty} H(k) = (I - \Pi^*)D \sum_{n=0}^{\infty} ((Q^* - Q)D)^n = D^*$.
- (iii) $\rho \in \mathbb{R}$ and $\delta < 1$ exist such that $\|((Q^* - Q)D)^k\|_v < \rho \delta^k$ for all k .
- (iv) For all k it holds that $\|T(k)\|_v < \rho \delta^k \|(I - \Pi^*)D\|_v$, with ρ and δ as in (iii).

These results have been introduced in Heidergott et al. (2009) for finite-state discrete-time Markov chains and can be extended to continuous-time Markov processes on denumerable state space using the same line of arguments.

3.3. The Algorithm

With Lemma 2 we arrive at the following numerical approach. First we search for N such that $1 > \delta_N \stackrel{\text{def}}{=} \|((Q^* - Q)D)^N\|_v$. In words, we establish the minimal power of $((Q^* - Q)D)$ that yields geometrical convergence of $H(n)$. Then, we choose a precision ϵ_{dev} up to which we

want to approximate D^* . The algorithm computes the elements $T(k)$ of $H(k)$ until our upper bound for $R(k)$, given by $c_{\delta_N}^v \|((Q^* - Q)D)^{k+1}\|_v$, drops below ϵ_{dev} . We can now describe an algorithm that yields an approximation for D^* with ϵ_{dev} precision.

Algorithm 1

Choose precision $\epsilon_{dev} > 0$. Set $k = 1$ so that $T(1) = (I - \Pi^*)D(Q^* - Q)D$ and $H(0) = (I - \Pi^*)D$.

Step 1: Find N such that $\|((Q^* - Q)D)^N\|_v < 1$. Set $\delta_N = \|((Q^* - Q)D)^N\|_v$ and compute

$$c_{\delta_N}^v = \frac{1}{1 - \delta_N} \left\| \sum_{k=0}^{N-1} ((Q^* - Q)D)^k \right\|_v.$$

Step 2: If $c_{\delta_N}^v \|T(k)\|_v < \epsilon_{dev}$, the algorithm terminates and $H(k-1)$ yields the desired approximation. Otherwise, go to Step 3.

Step 3: Set $H(k) = H(k-1) + T(k)$. Set $k := k + 1$ and $T(k) = T(k-1)(Q^* - Q)D$. Go to Step 2.

If condition [C] holds, then the above algorithm terminates geometrically fast; see Lemma 2. In case S is finite, condition [C] and (ii) in Lemma 2 are equivalent, which allows us to empirically check whether condition [C] is satisfied. Moreover, for finite S all norms are equivalent with respect to norm ergodicity and, without loss of generality, we take $v \equiv \mathbf{1}$, with $\mathbf{1}$ the vector with all elements equal to one, for the algorithm.

If we again look at the value function D^*f instead of simply approximating D^* , we have to slightly adjust Algorithm 1 because the error term is no longer just $\|D^* - H(n)\|$, but according to (7) the remainder becomes $\|D^* - H(n)\|_v \|f\|_v \inf_{i \in S} v(i)$. To achieve precision ϵ_{value} for the approximation of the value function D^*f by $H(n)f$, we have to choose ϵ_{dev} in Algorithm 1 equal to $\epsilon_{value} / (\|f\|_v \inf_{i \in S} v(i))$, because this implies

$$\begin{aligned} |(D^*f)_i - (H(n)f)_i| &\leq \|D^* - H(n)\|_v \|f\|_v \inf_{i \in S} v(i) \\ &< \frac{\epsilon_{value}}{\|f\|_v \inf_{i \in S} v(i)} \|f\|_v \inf_{i \in S} v(i) = \epsilon_{value}. \end{aligned}$$

For given precision ϵ_{value} and by setting $v \equiv \mathbf{1}$ as we will do in the finite state-space examples presented in this paper, we obtain

$$\epsilon_{dev} = \frac{\epsilon_{value}}{\sup_{i \in S} f(i)}, \quad (8)$$

where $\|f\|_1 = \sup_{i \in S} |f(i)|$ follows from the definition of the v -norm.

REMARK 2. In case $\|f\|_1$ is large, the precision ϵ_{dev} used in Algorithm 1 might become prohibitively small, i.e., the value of the associated optimal degree k is such that computing $H(k)$ becomes numerical inefficient. In this case it is advisable to adjust v such that $\|f\|_v$, as well as $\|((Q^* - Q)D)\|_v$, becomes sufficiently small.

4. Applications of the Update Formula

In this section we will present various examples illustrating the broad applicability of our update formula. As the basic system that serves as the approximation analogue, we use the M/M/1 queue for which, beside Q and Π , the deviation matrix D also is well known (see Koole 1998). We will add transient states to the basic system so that we can still compute D in a closed form, but also get a sufficiently fast convergence of the series expansion. Starting with an M/M/1 queue with vacations in §4.1, we proceed with the M/G/1 queue in §4.2 and provide the approximation of the deviation matrix of a tandem network in §4.3. We support our theoretical computations by several numerical examples in §§4.1.2, 4.2.3, and 4.3.2.

4.1. Approximating the Deviation Matrix of an M/M/1 Queue with Vacations

4.1.1. The M/M/1 Queue with Vacations. Our first example is a slight variation of the M/M/1 queue for which the deviation matrix already cannot be given in closed form so far. Similar to the basic system, arrivals enter the queue according to a Poisson- λ process and are served by one server with exponential- μ distributed service time. However, whenever the queue is empty—that is, no customer is waiting for service—the server will go on an exponential- α distributed vacation. Another possible interpretation is that of the server to shut down for a maintenance period. Such a process $\mathcal{X}^*$ exists on the state space $S\{(i, j) \in \mathbb{N}_0 \times \{0; 1\}\}$, where i denotes the number of customers in the system and j indicates if the server is on vacation ($=1$) or not ($=0$). $\mathcal{X}^*$ has a generator matrix Q^* with entries

$$q^*(i, j, k, l) = \begin{cases} \lambda & i = k - 1, j = l \\ \mu & i = k + 1, k > 0, j = l = 0 \\ \mu & i = 1, k = 0, j = 0, l = 1 \\ \alpha & i = k, j = 1, l = 0 \\ -(\lambda + \mu) & i = k, j = l = 0 \\ -(\lambda + \alpha) & i = k, j = l = 1 \\ 0 & \text{otherwise,} \end{cases}$$

and a stationary distribution

$$\begin{aligned} \pi^*(k, l) &= \begin{cases} \left(\frac{(\lambda/\mu)^k - (\lambda/(\lambda + \alpha))^k}{1 - \mu/(\lambda + \alpha)} \right) \pi^*(0, 1) & l = 0 \\ \left(\frac{\lambda}{\lambda + \alpha} \right)^k \pi^*(0, 1) & k > 0, l = 1 \\ \frac{\alpha(\mu - \lambda)}{\mu(\alpha + \lambda)} & k = 0, l = 1, \end{cases} \end{aligned}$$

(see Chao and Zhao 1997) where the stability of the queue and therefore the existence of π^* , is assured by $\lambda/\mu < 1$ and $\alpha > 0$. By its generator matrix Q with entries

$$q(i, j, k, l) = \begin{cases} \lambda & i = k - 1, j = l \\ \mu & i = k + 1, j = l = 0 \\ \alpha & i = k, j = 1, l = 0 \\ -(\lambda + \mu) & i = k, j = l = 0 \\ -(\lambda + \alpha) & i = k, j = l = 1 \\ 0 & \text{otherwise,} \end{cases}$$

we define an approximation system $\mathcal{X}$ in which all states $(j, 1)$ are transient while the ergodic class of states $(j, 0)$ acts like an M/M/1 queue. The stationary distribution of $\mathcal{X}$ is

$$\pi(k, l) = \begin{cases} \left(1 - \frac{\lambda}{\mu}\right) \left(\frac{\lambda}{\mu}\right)^k & l = 0 \\ 0 & l = 1, \end{cases}$$

and for the deviation matrix it holds

$$d(i, j, k, l) = \begin{cases} \frac{(\lambda/\mu)^{\max\{k-i, 0\}} - (i+k+1)(1-\lambda/\mu)(\lambda/\mu)^k}{\mu(1-\lambda/\mu)} & j = l = 0 \\ \frac{\lambda^{k-i}}{(\lambda + \alpha)^{k-i+1}} & k \geq i, j = l = 1 \\ \frac{1}{\mu(1-\lambda/\mu)} - \frac{(\mu + (k+i)\alpha)(\lambda/\mu)^k}{\mu\alpha} & \\ - \left(\frac{1}{\mu(1-\lambda/\mu)} + \frac{\mu(1-\lambda/\mu)(\lambda/(\lambda+\alpha))^{k-i+1} - \alpha(\lambda/\mu)^{k-i+1}}{\mu(1-\lambda/\mu)(\lambda-\mu+\alpha)} \right) & j = 1, l = 0 \\ \cdot \mathbf{1}_{\{i-1 < k\}} & \\ 0 & \text{otherwise.} \end{cases}$$

Because its entries are not obvious and, to the best of our knowledge, have not been presented in the literature so far, we provide the corresponding derivations in the online companion associated with this paper. By inserting the above expressions for Q , Q^* , Π^* , and D into (6), we can apply our algorithm to approximate D^* .

4.1.2. Numerical Approximation of the Deviation Matrix of an M/M/1 Queue with Vacations. We will illustrate our algorithm now by applying it to an M/M/1 queue with arrival rate $\lambda = 1$, service rate $\mu = 2$, and a vacation rate $\alpha = 4$, and restrict the maximal system capacity to

100 customers. First we compute the power N for which it holds $\delta_N = \|((Q^* - Q)D)^N\|_v < 1$ and get $N = 2$. The corresponding constant c_{δ_N} is $c_{\delta_N} = 2.1164$. Now we compute the optimal k for which the upper bound of the remainder term becomes smaller than the desired precision ϵ_{dev} , which we choose to be 0.1. According to Figure 1, the respective value is $k_{opt} = 4$.

The v -norm bound on the remainder term is shown on the left-hand side of Figure 1 and the v -norm of the true error is shown on the right-hand side of the figure. As can be seen, the bound on the remainder term becomes rather sharp after $k = 4$. Note that the bound on the remainder term suggests $k_{opt} = 4$, whereas the true error is already less than 0.1 for $k = 2$.

4.2. Approximating the Deviation Matrix of an M/G/1 Queue

A means of keeping the Markovian structure of the M/M/1 queue and also generalizing, say, the service time distribution, is to approximate the general service time distribution by a phase-type distribution, such as the Cox or hyperexponential distribution. It is well-known that phase-type distributions allow an arbitrarily close approximation of any general service time distribution G , see Asmussen (1987).

Our approach to the approximate computation of the deviation matrix of an M/G/1 queue is as follows. First, an appropriate phase-type distribution, say Ph, is determined that sufficiently approximates G . To this M/Ph/1 queue, we then determine an M/M/1 system with additional transient states and compute the stationary distribution and deviation matrix of this related but simpler queue. Applying our update formula, it is now possible to approximate the deviation matrix of the M/Ph/1 queue via the deviation matrix of the simple queue up to a precision ϵ_{dev} , which in turn yields an approximation of the deviation matrix of the M/G/1 queue.

Because it is common knowledge and shall not be the focus of this paper we discuss details of the phase-type distribution approximation approach and motivate our choice of the M/Cox/1 and M/Hyp/1 model as basic approximation of the M/G/1 queue in the online companion. In §4.2.1 we derive the update formula for the M/Cox/1 and in §4.2.2 for the M/Hyp/1 model.

A queueing system with one server serving with a phase-type distribution known as the M/Ph/1 system can be described on a state space $S = \{(i, j) \in \mathbb{N}_0 \times \mathbb{N}_0\}$ where the first entry describes the number of customers in the system and the second entry denotes the phase in which the server is currently serving. Customers enter the system with rate λ and will be served according to a first-come-first-served discipline. The various phase-type service time distributions differ in the number of phases and the order according to which a customer passes the phases. Each single phase is exponentially distributed. Typical phase-type distributions are the Cox, Erlang, and Hyperexponential distribution; see

Figure 1. Convergence of the upper bound on remainder and the deviation matrix.

Figure 2. For reasons of clarity, we will restrict our computations in this paper to phase-type distributions with three phases. It is understood that by following the same line of arguments as used in this paper, it is easy to extend all formulas to more general phase-type distributions. Note that for three phases the Erlang distribution equals the Cox distribution with $p_1 = p_2 = 1$. For sake of completeness, we summarize the density functions of the Cox and Hyperexponential distribution in Table 1 in the online companion.

4.2.1. Approximation of the Deviation Matrix of the M/Cox(3)/1 Queue. In this section we will apply the update formula to derive the deviation matrix of the queue-length process of the M/Cox(3)/1 queue, where we denote the queue-length process by $\mathcal{X}^*$, see Figure 3 for a depiction of the queueing model. The state space of this queue is $S = \{(i, j) \in \mathbb{N}_0 \times \{1, 2, 3\}\}$. Customers arrive to the system with rate λ . The service is separated into three phases. After a first μ_1 -exponentially distributed time the customer leaves the system immediately with probability $1 - p_1$, or will be served in a second phase with probability p_1 . The second service phase lasts a μ_2 -exponentially distributed

time, and the customer leaves the system with probability $1 - p_2$ or will be served in a final exponentially- μ_3 distributed phase with probability p_2 . The generator matrix Q^* of $\mathcal{X}^*$ is given by

$$q^*(i, j, k, l) = \begin{cases} \lambda & i = k - 1, j = l \\ (1 - p_1) \cdot \mu_1 & i = k + 1, j = l = 1 \\ p_1 \cdot \mu_1 & i = k, j = 1, l = 2 \\ (1 - p_2) \cdot \mu_2 & i = k + 1, j = 2, l = 1 \\ p_2 \cdot \mu_2 & i = k, j = 2, l = 3 \\ \mu_3 & i = k + 1, j = 3, l = 1 \\ - \sum_{(r,s) \neq (i,j)} q^*(i, j, r, s) & i = k \wedge j = l \\ 0 & \text{otherwise,} \end{cases}$$

Figure 2. (a) Cox(r)-, (b) Erlang(r)-, and (c) Hyperexponential(r)-distributions.

Figure 3. M/Cox(3)/1 queue.

and in accordance with Neuts (1994), Riska and Smirni (2002) the stationary distribution of $\mathcal{X}^*$ is given by $\pi^*(k) = (\pi^*(k, 1), \pi^*(k, 2), \pi^*(k, 3))$ with

$$\pi^*(k) = \begin{cases} \frac{\mu_1 \mu_2 \mu_3 - \lambda(\mu_2 \mu_3 + p_1 \mu_1 \mu_3 + p_1 p_2 \mu_1 \mu_2)}{\mu_1 \mu_2 \mu_3} & k = 0 \\ \pi^*(0) \cdot (1, 0, 0) \cdot R^k & k \geq 1, \end{cases}$$

where

$$R = \frac{\lambda}{(1-p_1)(\lambda^2 + \lambda \mu_3) + (1-p_1 p_2)\lambda \mu_2 + \mu_2 \mu_3} \cdot \begin{pmatrix} \frac{(\lambda + \mu_2)(\lambda + \mu_3)}{\mu_1} & p_1(\lambda + \mu_3) & p_1 p_2 p_3 \\ \frac{\lambda(\lambda + p_2 \mu_2 + \mu_3)}{\mu_1} & \lambda + \mu_3 & p_2 \mu_2 \\ \frac{\lambda(\lambda + \mu_2)}{\mu_1} & p_1 \lambda & (1-p_1)\lambda + \mu_2 \end{pmatrix}.$$

Existence of the stationary distribution follows from

$$\frac{\lambda(\mu_2 \mu_3 + p_1 \mu_1 \mu_3 + p_1 p_2 \mu_1 \mu_2)}{\mu_1 \mu_2 \mu_3} < 1.$$

Now we apply our update formula by first defining a second process $\mathcal{X}$ on the same state space S . Using the same parameters as before and letting p_1 and p_2 be equal to 0, we obtain a queueing system for which all states $(i, 2)$ and $(i, 3)$, with $1 \leq i$, (i.e., service is in the second or third phase), are transient states, whereas the remaining states form an ergodic class. As soon as this process enters the ergodic class, it acts like an M/M/1 queue with arrival rate λ and service rate μ_1 . The process $\mathcal{X}$ has a generator

matrix

$$q(i, j, k, l) = \begin{cases} \lambda & i = k-1, j = l \\ \mu_1 & i = k+1, j = l = 1 \\ \mu_2 & i = k+1, j = 2, l = 1 \\ \mu_3 & i = k+1, j = 3, l = 1 \\ -\sum_{(r,s) \neq (i,j)} q(i, j, r, s) & i = k, j = l \\ 0 & \text{otherwise.} \end{cases}$$

The stationary distribution of $\mathcal{X}$ resembles the one of an M/M/1 queue apart from the fact that its entries for transient states are 0, more specifically,

$$\pi(k, l) = \begin{cases} \left(1 - \frac{\lambda}{\mu_1}\right) \left(\frac{\lambda}{\mu_1}\right)^k & l = 1 \\ 0 & l = 2, 3, \end{cases}$$

where the condition $\lambda/\mu_1 < 1$ assures the existence of the stationary distribution. As the last component of our update formula, we present the deviation matrix of $\mathcal{X}$ with the respective derivations provided in the online companion.

$d(i, j, k, l)$

$$d(i, j, k, l) = \begin{cases} \frac{(\lambda/\mu_1)^{\max\{k-i, 0\}} - (i+k+1)(1-\lambda/\mu_1)(\lambda/\mu_1)^k}{\mu_1(1-\lambda/\mu_1)} & j = l = 1 \\ \frac{\lambda^{k-i}}{(\lambda + \mu_2)^{k-i+1}} & k \geq i, j = l = 2 \\ \frac{\lambda^{k-i}}{(\lambda + \mu_3)^{k-i+1}} & k \geq i, j = l = 3 \\ \frac{1}{\mu_1(1-\lambda/\mu_1)} - \frac{(\mu_1 + (k+i)\mu_2)(\lambda/\mu_1)^k}{\mu_1 \mu_2} & \\ - \left(\frac{1}{\mu_1(1-\lambda/\mu_1)} + \frac{\mu_1(1-\lambda/\mu_1)(\lambda/(\lambda + \mu_2))^{k-i+1} - \mu_2(\lambda/\mu_1)^{k-i+1}}{\mu_1(1-\lambda/\mu_1)(\lambda - \mu_1 + \mu_2)} \right) & j = 2, l = 1 \\ \frac{1}{\mu_1(1-\lambda/\mu_1)} - \frac{(\mu_1 + (k+i)\mu_3)(\lambda/\mu_1)^k}{\mu_1 \mu_3} & \\ - \left(\frac{1}{\mu_1(1-\lambda/\mu_1)} + \frac{\mu_1(1-\lambda/\mu_1)(\lambda/(\lambda + \mu_3))^{k-i+1} - \mu_3(\lambda/\mu_1)^{k-i+1}}{\mu_1(1-\lambda/\mu_1)(\lambda - \mu_1 + \mu_3)} \right) & j = 3, l = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Inserting the above expressions for Q , Q^* , Π^* , and D into (6), we obtain an approximation for the deviation matrix of an M/Cox(3)/1 queue.

4.2.2. Approximation of the Deviation Matrix of the M/Hyp(3)/1 Queue. Another well-known phase-type distribution is the hyperexponential one that, when applied to the service time of a queueing system, is reflected by the denotation M/Hyp(3)/1. In such a system the server provides service with an exponentially distributed service time, either with probability p_1 at rate μ_1 , with probability p_2 at rate μ_2 , or with probability $1 - p_1 - p_2$ at rate μ_3 , see Figure 4. This system has generator matrix

$$q^*(i, j, k, l) = \begin{cases} \lambda & i = k - 1, j = l \\ p_1 \cdot \mu_1 & i = k + 1, j = l = 1 \\ p_2 \cdot \mu_1 & i = k + 1, j = 1, l = 2 \\ (1 - p_1 - p_2) \cdot \mu_1 & i = k + 1, j = 1, l = 3 \\ p_1 \cdot \mu_2 & i = k + 1, j = 2, l = 1 \\ p_2 \cdot \mu_2 & i = k + 1, j = l = 2 \\ (1 - p_1 - p_2) \cdot \mu_2 & i = k + 1, j = 2, l = 3 \\ p_1 \cdot \mu_3 & i = k + 1, j = 3, l = 1 \\ p_2 \cdot \mu_3 & i = k + 1, j = 3, l = 2 \\ (1 - p_1 - p_2) \cdot \mu_3 & i = k + 1, j = l = 3 \\ - \sum_{(r,s) \neq (i,j)} q^*(i, j, r, s) & i = k, j = l \\ 0 & \text{otherwise,} \end{cases}$$

and stationary distribution

$$\pi^*(k) = \begin{cases} \frac{\mu_1 \mu_2 \mu_3 - \lambda(p_1 \mu_2 \mu_3 + p_2 \mu_1 \mu_3 + (1 - p_1 - p_2) \mu_1 \mu_2)}{\mu_1 \mu_2 \mu_3} & k = 0 \\ \pi^*(0) \cdot (p_1, p_2, 1 - p_1 - p_2) \cdot R^k & k \geq 1, \end{cases}$$

Figure 4. M/Hyp(3)/1 queue.

with

$$R = \lambda \cdot (\lambda^2(p_1 \mu_1 + p_2 \mu_2 - p_2 \mu_3) + (1 - p_1)(\lambda \mu_2 \mu_3 + \lambda^2 \mu_3) + (1 - p_2) \lambda \mu_1 \mu_3 + (p_1 + p_2) \lambda \mu_1 \mu_2 + \mu_1 \mu_2 \mu_3)^{-1} \times \begin{pmatrix} p_1 \lambda (\lambda + \mu_2) + (1 - p_2) \lambda \mu_3 + p_2 \lambda \mu_2 + \mu_2 \mu_3 \\ p_1 \lambda (\lambda + \mu_3) \\ p_1 \lambda (\lambda + \mu_2) \\ p_2 \lambda (\lambda + \mu_3) \\ p_2 \lambda (\lambda + \mu_1) \\ (1 - p_1 - p_2)(\lambda^2 + \lambda \mu_2) \\ (1 - p_1 - p_2) \mu_1 + (1 - p_1) \lambda^2 - p_2 \lambda \\ (1 - p_1 - p_2) \lambda^2 + (1 - p_1) \lambda \mu_2 \\ + (1 - p_2) \lambda \mu_1 + \mu_1 \mu_2 \end{pmatrix}.$$

Existence of this stationary distribution follows from

$$\frac{\lambda(p_1 \mu_2 \mu_3 + p_2 \mu_1 \mu_3 + (1 - p_1 - p_2) \mu_1 \mu_2)}{\mu_1 \mu_2 \mu_3} < 1.$$

Analogously to §4.2.1, we can approximate the general system by the system with $p_1 = 1$ and $p_2 = 0$, i.e., every new arrival will be served at rate μ_1 . It can be easily seen that such a Markov process $\mathcal{X}$ has the same generator matrix, stationary distribution, and deviation matrix as the process used to approximate the deviation matrix of the M/Cox(3)/1 queue. Hence, we just have to exchange Q^* and Π^* in (6) and we obtain the update formula for the deviation matrix of the M/Hyp(3)/1 queue.

4.2.3. Numerical Approximation of the Deviation Matrix of an M/LogN/1 Queue. One of the most common examples for real-world queueing systems are call centers in which statistical analyses provide evidence for lognormal distributed service times, see, for example, Brown et al. (2005). The density of the lognormal distribution is given by

$$f^{\text{LogN}}(x) = \frac{1}{x \sqrt{2\pi\sigma^2}} \exp\left[-\frac{(\ln[x] - \mu)^2}{2\sigma^2}\right],$$

with μ and σ the mean and standard deviation of the variable's logarithm. Now we can determine the variables of the phase distributions by applying the EM algorithm. In the following we will apply the EMpht programme (for a manual, see Olsson 1998) for fitting distributions.

Suppose that $\mathcal{X}^{\text{LogN}}$ is the queue-length process of an M/LogN/1 queue with customers arriving according to a

Figure 5. Density functions of the lognormal, Cox(3), and Hyp(3) distributions.

Poisson arrival stream with rate $\lambda = 2$ and being served by a server having lognormal distributed service times with parameters $\mu = 0.48$ and $\sigma^2 = 1.40$. By using the EM algorithm we find an M/Cox(3)/1 queue with $p_1 = 0.50$, $p_2 = 0.30$, $\mu_1 = 0.50$, $\mu_2 = 9.00$, and $\mu_3 = 10.00$ as approximation, or, alternatively, an M/Hyp(3)/1 queue having parameters $p_1 = 0.07$, $p_2 = 0.43$, $\mu_1 = 0.45$, $\mu_2 = 0.45$, and $\mu_3 = 0.45$. By comparing the three density functions in Figure 5, we can already see that the density of the Cox distribution fits better to the original lognormal distribution. Additionally, the fitted hyperexponential distribution with all phases having equal service behaves like an exponential distribution. Nevertheless, the resulting M/Hyp(3)/1 differs from the simple M/M/1 queue due to the enlarged state space.

In our further computations we will restrict the waiting room of our queue to 49 places, which gives us an M/LogN/1/50 queue. For such a queueing system, we apply our algorithm now and first compute the respective powers

denoted by N^{Cox} and N^{Hyp} , respectively, for which it holds $\delta_N = \|((Q^* - Q)D)^N\|_v < 1$. It turns out that $N^{Cox} = 1$, $N^{Hyp} = 2$.

The corresponding constants, denoted by $c_{\delta_N}^{Hyp}$ and $c_{\delta_N}^{Cox}$, respectively, are

$$c_{\delta_N}^{Cox} = 3.5714, \quad c_{\delta_N}^{Hyp} = 9.2667.$$

Now it remains to identify the optimal k for which the upper bound of the remainder term becomes smaller than our desired precision of 0.1. According to Figure 6, the respective values are

$$k_{opt}^{Cox} = 5, \quad k_{opt}^{Hyp} = 2.$$

These findings are supported by Figure 7, in which we plot the norm convergence of $H(k-1)$ given by the update formula in (6) to the true deviation matrix D^* , which, in case of a finite state space, can be computed directly by (2).

4.3. Approximation of the Deviation Matrix of a Tandem Network

4.3.1. The Tandem Network. Another common example in queueing theory for which the deviation matrix cannot be computed directly is the tandem queue displayed in Figure 8(a). In such a system, customers enter a first queue according to a Poisson process with rate λ , and after an exponentially- μ_1 distributed service time are routed to a second queue where they are served by a second exponential server with rate μ_2 . This system has generator matrix

$$q^*(i, j, k, l) = \begin{cases} \lambda & i = k-1, j = l \\ \mu_1 & i = k+1, j = l-1 \\ \mu_2 & i = k, j = l+1 \\ -(\lambda + \mu_1 + \mu_2) & i = k, j = l \\ 0 & \text{otherwise,} \end{cases}$$

Figure 6. Convergence of upper bounds on remainders for Cox(3) and Hyp(3) distribution.

Figure 7. Norm convergence of the deviation matrices of Cox(3) and Hyp(3) distribution.

and stationary distribution

$$\pi^*(k, l) = \left(1 - \frac{\lambda}{\mu_1}\right) \left(\frac{\lambda}{\mu_1}\right)^k \left(1 - \frac{\mu_1}{\mu_2}\right) \left(\frac{\mu_1}{\mu_2}\right)^l,$$

with its existence assured by the condition

$$\frac{\lambda}{\mu_1} < 1 \quad \text{and} \quad \frac{\mu_1}{\mu_2} < 1.$$

The deviation matrix of such a tandem network cannot be computed directly using common tools so that we will apply our update formula to approximate it. First we define a second process $\mathcal{X}$ on the same state space $S = \{(i, j) \in \mathbb{N}_0 \times \mathbb{N}_0\}$, where i and j denote the number of customers queueing in the first and second queue, respectively. We model the first queue as an M/M/1 queue with exponential- λ distributed interarrival and exponential- μ_1 distributed service times. We furthermore do not allow for any arrivals in queue 1 as long as there are customers in the second queue, which is defined as an M/M/1 system without arrivals but with exponential- μ_2 service times. That is, all states (i, j) with $j > 0$ are transient while the states $(i, 0)$ form the system's ergodic class (see Figure 8(b)). Such a system is determined by its generator matrix

$$q(i, j, k, l) = \begin{cases} \lambda & i = k - 1, j = l = 0 \\ \mu_1 & i = k + 1, j = l \\ \mu_2 & i = k, j = l + 1 \\ -(\lambda + \mu_1 + \mu_2) & i = k, j = l = 0 \\ -(\mu_1 + \mu_2) & i = k, j = l > 0 \\ 0 & \text{otherwise,} \end{cases}$$

and has the stationary distribution

$$\pi(k, l) = \begin{cases} \left(1 - \frac{\lambda}{\mu_1}\right) \left(\frac{\lambda}{\mu_1}\right)^k & l = 0 \\ 0 & l > 0, \end{cases}$$

where the condition $\lambda/\mu_1 < 1$ assures its existence. The deviation matrix of $\mathcal{X}$ is provided by

$$d(i, j, k, l) = \begin{cases} \frac{(\lambda/\mu_1)^{\max\{k-i, 0\}} - (i+k+1)(1-\lambda/\mu_1)(\lambda/\mu_1)^k}{\mu_1(1-\lambda/\mu_1)} & j = l = 1 \\ \frac{(i-k+j-l)!}{(i-k)!(j-l)!} \frac{\mu_1^{i-k} \mu_2^{j-l}}{(\mu_1 + \mu_2)^{-(i-k+j-l+1)}} & j \geq l > 0, i \geq k \\ \frac{\mu_1 \mu_2 - ((\mu_1 - \lambda) \mu_2 (2k+1) + j(\mu_1 - \lambda)^2)(\lambda/\mu_1)^k}{\mu_1 \mu_2 (\mu_1 - \lambda)} & \\ - \frac{\mu_1^i \mu_2^j (\lambda/\mu_1)^k}{(\mu_1 + \mu_2)^{i+j}} \sum_{r=0}^i \binom{i-r+j-1}{j-1} \left(\frac{\mu_1 + \mu_2}{\mu_1}\right)^r & \\ \cdot \left(\frac{r-k}{\mu_1} + \left(\frac{\mu_1/\lambda}{\lambda - \mu_1}\right)^r \cdot \mathbf{1}_{\{r < k\}}\right) & j > 0, l = 0 \\ 0 & \text{otherwise,} \end{cases}$$

Figure 8. (a) Tandem queue, (b) independent manipulated M/M/1 queues.

Figure 9. Convergence of upper bounds on remainders and deviation matrices.

with the derivations presented in the online companion of this paper. Inserting Q , Q^* , Π^* , and D into (6), we obtain an approximation for the deviation matrix of the tandem network.

4.3.2. Numerical Approximation of the Deviation Matrix of a Tandem Network. Let $\mathcal{X}^*$ be a tandem network with Poisson-0.5 arrival stream and exponentially-2 and exponentially-10 distributed service times in server 1 and 2, respectively. Each queue is supposed to have a maximal capacity of 20 customers. For this system we first calculate N for which it holds $\delta_N = \|((Q^* - Q)D)^N\|_v < 1$, the associated constant c_{δ_N} , and the optimal k for which the remainder becomes smaller than our given precision $\epsilon = 0.1$, and get

$$N = 2, \quad c_{\delta_N} = 9.4665, \quad k_{opt} = 7$$

(see Figure 9). By lowering the service rate of the second server to $\mu_2 = 5$ and repeating our calculations, we get a slightly slower convergence

$$N = 4, \quad c_{\delta_N} = 26.9075, \quad k_{opt} = 13,$$

whereas an increase of the second service rate leads to an improvement of the respective parameters

$$N = 2, \quad c_{\delta_N} = 3.2592, \quad k_{opt} = 5$$

(see Figure 9). Even in the case of a relatively slow service, we can still approximate the deviation matrix by our update formula, but will need more summands to reach our desired precision. That is due to the fact that the formula is sensitive towards the traffic rates $\rho_1 = \lambda/\mu_1$ and $\rho_2 = \mu_1/\mu_2$. Nevertheless, even traffic rates close to 1 will allow for an application of the update formula. To reduce the number of approximation steps necessary to reach the desired precision, it remains, however, to look for related processes $\mathcal{X}$ that are closer to the tandem queue than our exemplary M/M/1 queues, but still satisfy the condition of a given deviation matrix D . This will be a topic of further research.

5. Application to Markov Decision Theory

We revisit the M/G/1 queue and extend the basic Markov model by a certain set of actions $\mathcal{A}$. A stationary (deterministic) policy, say f , prescribes in each state i the action $f(i)$ taken by the controller of the system. The infinitesimal generator depends on the chosen policy f , which is reflected by writing Q_f for the generator. Moreover, there is a cost vector c_f , where $c_f(i)$ is the cost rate in state i under policy f . For an introduction to Markov decision processes (MDP), we refer to Dynkin and Yushkevich (1979), Puterman (1994), Ross (1970), Tijms (1994). MDPs are used to derive the structure of the optimal policy and, moreover, to compute it. There are several cost criteria, e.g., discounted and average costs. In our application we will consider average expected costs over an infinite time horizon.

A basic reference on MDPs is the monograph Puterman (1994); applications to the control of queueing systems can be found in Ross (1970), Stidham (1985), Tijms (1994). There are numerous papers on MDPs. Papers in which the condition of geometric ergodicity is introduced and used are Dekker and Hordijk (1988, 1991, 1992), Dekker et al. (1994), and Hordijk and Yushkevich (1999, 2002). The papers Bather (1976), Guo (2007), Guo and Cao (2005), Guo and Hernandez-Lerma (2003), Guo and Liu (2001), Hordijk and van der Duyn Schouten (1983), Kakumanu (1975), and Yushkevich (1977) deal with continuous-time decision processes.

5.1. Admission Control

A basic model in the control of communication networks is the admission control of a single server queue. For an overview of the application of Markov decision models in the control of communication networks we refer to Altman (2002). We assume here that there is a reward rate r for

each customer admitted to the queue, and that there are holding costs $h(i)$ per unit time when there are i customers in the queue. For this model, the optimal policy for discounted as well as average costs is of threshold type, that is, there is a critical level, say L , such that customers are allowed when there are at most $L - 1$ customers in the queue and rejected when there are L or more. The literature on admission control separates into two groups, depending on the amount of information on which the controls are based. In case of full state information, the problem is approached as an MDP, a Markov decision drift process, or a Markov game; see Altman and Hordijk (1995), Hordijk and van der Duyn Schouten (1983), Stidham (1985). It is shown in Stidham (1985) that the average optimal policy is a critical-level policy. For the continuous-time process, this follows from results in Hordijk and van der Duyn Schouten (1983). In Hordijk and Spieksma (1989), the constrained admission control is considered. Therein, the throughput is maximized while the expected delay of jobs has an upper bound. In this case the optimal policy is of threshold type too, but randomizes the admission at the critical level. A similar result is proven for worst-case optimal control in Altman and Hordijk (1995) and Altman et al. (1997). The customer admission model in a rather general setting with no state information is studied in Chapter 4 of Altman et al. (2003). A balanced policy is optimal in this case. The computation of the optimal policy was hindered by the fact that the deviation matrix could not be computed. The analysis put forward in this paper overcomes this problem. For the admission control model studied in this paper, there is an equivalence between the discrete-time and continuous-time models (see Serfozo 1979). For further study, we also refer to the survey papers Arapostathis et al. (1993) for discrete-time case and the recent one Guo et al. (2006) and the references therein for the continuous-time case.

Using the same notation as before, but adding subscript L to reflect the current admission level, we can express the average long-term costs by

$$\phi \stackrel{\text{def}}{=} \pi_L c_L,$$

where c_L is the respective cost vector defined as

$$c_L(i) = h(i) - r\lambda \cdot \mathbf{1}_{\{i < L\}}, \quad i \in S.$$

We will then consider a certain admission policy L_1 superior to L if long-term costs are lower so that it holds

$$\Pi_{L_1} c_{L_1} - \Pi_L c_L < 0.$$

From (3) multiplied by Π_{L_1} , we get

$$\begin{aligned} \Pi_{L_1}(c_{L_1} - c_L) + \Pi_{L_1} c_L - \Pi_L c_L \\ = \Pi_{L_1}(c_{L_1} - c_L) + \Pi_{L_1}(Q_{L_1} - Q_L)D_L c_L < 0, \end{aligned}$$

and because Π_{L_1} has no negative entries, this inequality is implied by

$$(c_{L_1} - c_L) + Q_{L_1}D_L c_L - Q_L D_L c_L < 0.$$

From Lemma 1 we have $Q_L D_L = \Pi_L - I$ so that we obtain

$$(c_{L_1} - c_L) + Q_{L_1}D_L c_L + (I - \Pi_L)c_L < 0,$$

which simplifies to the basic Poisson inequality used in MDPs

$$c_{L_1} - \Pi_L c_L + Q_{L_1}D_L c_L < 0.$$

Then the policy improvement step on critical level L policy in state $i \geq L$ is admitting the customer if

$$-r\lambda + h(i) - \pi_L c_L + Q_{i+1}(i)D_L c_L < 0, \quad (9)$$

where $Q_{i+1}(i)$ is the i th row of the generator matrix for the increased threshold $i + 1$. Similarly, one better rejects a customer in state $i < L$ if

$$h(i) - \pi_L c_L + Q_i(i)D_L c_L < 0. \quad (10)$$

The following algorithm allows for the iterative computation of the optimal threshold L' .

Algorithm 2

Start with an initial policy, say level L_0 .

Step 1: Compute the largest $i > L_0$ for which (9) holds, say L_1 .

Step 2: Compute π_{L_1} and D_{L_1} and insert them into (10).

Step 3: Compute the smallest $i < L_1$ for which (10) holds, say L_2 .

Step 4: Compute π_{L_2-1} and D_{L_2-1} and insert them into (9).

Step 5: Repeat Steps 1 to 4 until there is no improvement.

Step 6: The algorithm ends when no further improvement can be found.

5.2. Numerical Example: Applying the Update Formula to Admission Control

Recall the example introduced in §4.2.3, but now with threshold L . The infinitesimal generators of the original system $\mathcal{X}_L^*$ and the approximation process $\mathcal{X}_L$ remain the same except for the fact that $q_L(i, j, i + 1, l) = 0$ for every $i \geq L$. Hence, we create transient states $(i, j) > L$. Then the stationary distribution $\pi_L^*(k, l)$ is equal to 0 whenever $k > L$. And for $k \leq L$ it resembles the one of a finite M/Ph/1/L. According to Neuts (1994), π_L^* of such a queue equals the denumerable one except for

$$\pi_L^*(k) = \begin{cases} \left((1, 0, 0) \left(\sum_{i=0}^{L-1} R^i + \lambda R^{L-1} S^{-1} \right) \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right)^{-1} & k = 0 \\ \pi^*(0)(1, 0, 0)\lambda R^{L-1} S^{-1} & k = L, \end{cases}$$

with

$$S^{-1} = \begin{pmatrix} \frac{1}{\mu_1} & \frac{p_1}{\mu_2} & \frac{p_1 p_2}{\mu_3} \\ 0 & \frac{1}{\mu_2} & \frac{p_2}{\mu_3} \\ 0 & 0 & \frac{1}{\mu_3} \end{pmatrix}$$

in case of the M/Cox(3)/1/L queue and

$$\pi_L^*(k) = \begin{cases} \left((p_1, p_2, 1 - p_1 - p_2) \left(\sum_{i=0}^{L-1} R^i + \lambda R^{L-1} S^{-1} \right) \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right)^{-1} & k = 0 \\ \pi^*(0)(p_1, p_2, 1 - p_1 - p_2) \lambda R^{L-1} S^{-1} & k = L, \end{cases}$$

with

$$S^{-1} = \begin{pmatrix} \frac{1}{\mu_1} & 0 & 0 \\ 0 & \frac{1}{\mu_2} & 0 \\ 0 & 0 & \frac{1}{\mu_3} \end{pmatrix}$$

for the M/Hyp(3)/1/L system, and for the stationary distribution of $\mathcal{X}_L$ we have

$$\pi(k, l) = \begin{cases} \frac{1 - \lambda/\mu_1}{1 - (\lambda/\mu_1)^{L+1}} \left(\frac{\lambda}{\mu_1} \right)^k & l = 1 \\ 0 & l = 2, 3. \end{cases}$$

Although in the case of a denumerable state space S a stability condition is needed to assure the existence of a stationary distribution, this is not needed for the finite state space. The derivation of the deviation matrix is for reasons of clarity provided in the online companion (see §6 of the online companion).

Suppose that the systems arrival rate is still $\lambda = 2.00$ and customers are served by one server having lognormal distributed service times with parameters $\mu = 0.48$ and $\sigma^2 = 1.40$. This leads to the approximation system of an M/Cox(3)/1 queue with $p_1 = 0.50$, $p_2 = 0.30$, $\mu_1 = 0.50$, $\mu_2 = 9.00$, and $\mu_3 = 10.00$ and an M/Hyp(3)/1 queue having parameters $p_1 = 0.07$, $p_2 = 0.43$, $\mu_1 = 0.45$, $\mu_2 = 0.45$, and $\mu_3 = 0.45$. Like in §4.2.3 the waiting room is restricted to a maximum capacity of 49 waiting customers. Let the holding costs of our example be defined as $h(i) = e^{0.01i} - 1$, and let each accepted customer create a reward of $r = 9$.

We know from the example in §4.2.3 that in case of the M/Cox(3)/1 queue we need $H(4)$ to approximate the deviation matrix sufficiently precisely, and for the M/Hyp(3)/1 system, $H(1)$ given by (6) has the right degree to compute

the deviation matrix up to a precision 0.1. However, because we no longer want to approximate D_L^* but the value function $D_L^* c_L$, we use (8) to assure the desired precision 0.1

$$\epsilon_{dev} = \frac{\epsilon_{value}}{\sup_{i \in S} |c_L(i)|} = \frac{0.1}{10.0} \approx 0.01,$$

and by applying Algorithm 1 get $H(5)$ for the Cox model and $H(1)$ in case of hyperexponential service times. Additionally, it is no longer sufficient to use (9) and (10), but we have to consider the two-dimensional state space. Hence, each state with i customers inside the system is reflected by three states—namely, $(i, 1)$, $(i, 2)$, and $(i, 3)$. Checking if (9) holds for a certain customer level i allows for several interpretations: (i) demanding that it holds for all service phases, which is quite conservative, (ii) demanding that the condition (9) is valid for only one service phase, which is too optimistic, (iii) weighting the phases by the stationary probability that the server serves in the respective phase. In the following, we chose approach (iii) and, using the PASTA property (Poisson arrivals see time averages), we obtain

$$\begin{aligned} & -r\lambda + h(i) - \pi_L^* c_L + \frac{1}{\pi_{i+1}^*(i, 1) + \pi_{i+1}^*(i, 2) + \pi_{i+1}^*(i, 3)} \\ & \cdot (\pi_{i+1}^*(i, 1) Q_{i+1}^*(i, 1) + \pi_{i+1}^*(i, 2) Q_{i+1}^*(i, 2) \\ & + \pi_{i+1}^*(i, 3) Q_{i+1}^*(i, 3)) H_L(k) c_L < 0 \end{aligned} \quad (11)$$

for (9) and replace (10) by

$$\begin{aligned} & h(i) - \pi_L^* c_L + \frac{1}{\pi_{i+1}^*(i, 1) + \pi_{i+1}^*(i, 2) + \pi_{i+1}^*(i, 3)} \\ & \cdot (\pi_{i+1}^*(i, 1) Q_L^*(i, 1) + \pi_{i+1}^*(i, 2) Q_L^*(i, 2) \\ & + \pi_{i+1}^*(i, 3) Q_L^*(i, 3)) H_L(k) c_L < 0. \end{aligned} \quad (12)$$

Now we start Algorithm 2 with the initial step assuming $L = 10$, i.e., no customers are accepted to enter the system if there are already 10 or more inside. The results of the following improvement steps are summarized in Table 1, where L_0 and ϕ_0 denote the initial values of L and ϕ , and L_i and ϕ_i denote the values of L and ϕ in the i th iteration of the algorithm. Although the M/Cox(3)/1 needs one improvement step fewer, both approximations suggest the same optimal policy level $L' = 4$. Recall that both models represent the same M/LogN/1 queue, so this result is not surprising. In principle, level L is dependent on the traffic rate ρ given by

$$\rho^{Cox} = \frac{\lambda(\mu_2 \mu_3 + p_1 \mu_1 \mu_3 + p_1 p_2 \mu_1 \mu_2)}{\mu_1 \mu_2 \mu_3} = 4.14$$

Table 1. Policy improvement steps.

	L_0	ϕ_0	L_1	ϕ_1	L_2	ϕ_2	L_3	ϕ_3
Cox(3)	10	-4.2449	3	-4.2802	4	-4.3004		
Hyp(3)	10	-3.9480	3	-3.9866	6	-3.9908	4	-4.0041

Figure 10. Comparison of thresholds depending on the traffic rate.

for M/Cox(3)/1 and

$$\rho^{Hyp} = \frac{\lambda(p_1\mu_2\mu_3 + p_2\mu_1\mu_3 + (1-p_1-p_2)\mu_1\mu_2)}{\mu_1\mu_2\mu_3} = 4.44$$

in the case of an M/Hyp(3)/1 queue. This rate compares the average number of arrivals within a certain time interval to the number of served customers. In our example we have a ρ that is somewhat larger than one so that in the case of an unrestricted entrance policy the system would overflow. Even in the case of a restriction this has to be small so that holding costs do not exceed the rewards. Figure 10 illustrates the dependence of the optimal threshold L on the traffic rate.

5.3. Extension to Denumerable State Spaces

In the previous section we documented the adaptability of our update formula to MDPs by presenting a numerical example on finite state space. Naturally, in such a case one can easily compute the real deviation matrix by (2). In the following we will discuss an application of our update formula to the denumerable state space. Consider the Coxian model and assume that a Cox distribution has been fitted to the general service time. As the next step, use the update formula (6) to approximate the deviation matrix of an M/Cox/1 system. Suppose that the series is only developed for the first n elements, that is, take

$$H^{Cox}(n) = (I - \Pi^{Cox})D \sum_{k=0}^n (I + (Q^{Cox} - Q)D)^k$$

as an approximation of D^{Cox} , where D , Q are obtained from the simple M/M/1 queue with transient states; §§4.2.1, 5.2,

and 6 of the online companion. Provided that n is small, for example, $n = 1$, inserting the given formulas into (11) and (12) one can easily determine the optimal threshold policy even for denumerable state space S . For values of n larger than, say 3, the analytical application of (6) is surely still possible but a lot less convenient because formulas become way too complex. However, because our current intention is no longer the exact approximation of the deviation matrix but the determination of the optimal policy, we do not necessarily need to compute an $H^{Cox}(n)$ for $n > 2$. For our numerical example, we have $k^{Cox} = 5$, and Table 2 displays the optimal policy L for our Cox model depending on the chosen approximation parameter n for $H^{Cox}(n)$. The table suggests that $H^{Cox}(2)$ already reveals the optimal threshold level L , and even $H^{Cox}(1)$ provides fairly cost-optimal policies, even though the deviation matrix of the M/Cox/1 queue is only roughly approximated.

We performed various numerical experiments, all of which suggest that the approximation $H^{Cox}(2)c_L$ already yields the optimal policy, despite the fact that the approximation $H^{Cox}(2)c_L$ differs in norm from the real value function given by D^*c_L by more than the desired precision ϵ_{value} . Hence, we expect that even if the convergence of $H^{Cox}(k)$ to D^* is too slow to work with $n = 2$, a series expansion of degree 2 is already sufficient to derive optimal policies. We conjecture that for denumerable state spaces the application of

$$H(2) = (I - \Pi^*)D(I + (Q^* - Q)D + ((Q^* - Q)D)^2)$$

in MDPs yields optimal policies. Further investigation into this matter is a topic of future research.

Table 2. Optimal policy L .

ρ	Exact	D^*c_L	Opt. appr.	$H^{Cox}(5)c_L$	$H^{Cox}(2)c_L$	$2H^{Cox}(1)c_L$
	L	ϕ	L	ϕ	L	ϕ
0.1	158	-0.4488	158	-0.4488	158	-0.4488
0.5	113	-2.2393	113	-2.2393	113	-2.2393
1.0	22	-4.0902	22	-4.0902	22	-4.0901
1.5	10	-4.2384	10	-4.2384	10	-4.2384

6. Conclusion

In this paper we presented an approach to efficiently approximate the deviation matrix of continuous-time Markov processes and applied it to three different types of examples. For a problem in admission control, we combined our approximation scheme with Markov decision theory and computed the optimal admission policy. Our three examples are all related to a certain kind of manipulation of the M/M/1 queue as the approximation system. In future research it will be of particular interest to extend these findings to more elaborate approximation models that increase the speed of convergence of our algorithm. Although our approach is the first step towards higher-dimensional state spaces, it will be of special interest to apply the concept to larger networks.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

Acknowledgments

The first author thanks the Faculty of Exact Sciences of the Vrije Universiteit Amsterdam for their hospitality and financial support.

References

- Altman, E. 2002. Applications of Markov decision processes in communication networks. E. Feinberg, A. Schwartz, eds. *Handbook of Markov Decision Processes*. Kluwer, Boston, 489–536.
- Altman, E., A. Hordijk. 1995. Zero-sum Markov games and worst-case optimal control of queueing systems. *Queueing Systems* **21** 415–447.
- Altman, E., B. Gaujal, A. Hordijk. 2003. *Discrete-Event Control of Stochastic Networks: Multimodularity and Regularity*. Springer Lecture Notes in Mathematics, Vol. 1829. Springer, New York.
- Altman, E., A. Hordijk, F. Spieksma. 1997. Contraction conditions for average and α -discount optimality in countable state Markov games with unbounded rewards. *Math. Oper. Res.* **22** 588–618.
- Arapostathis, A., V. Borkar, E. Fernández-Gaucherand, M. Ghosh, S. Marcus. 1993. Discrete-time controlled Markov processes with average cost criterion: A survey. *SIAM J. Control Optim.* **31** 282–344.
- Asmussen, S. 1987. *Applied Probability and Queues*. Springer, New York.
- Bather, J. 1976. Optimal stationary policies for denumerable Markov chains in continuous time. *Adv. Appl. Probab.* **8** 144–158.
- Bertsekas, D. 2005. *Dynamic Programming and Optimal Control*, 3rd ed. Athena Scientific, Nashua, NH.
- Brown, L., N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn, L. Zhao. 2005. Statistical analysis of a telephone call center: A queueing-science perspective. *J. Amer. Statist. Assoc.* **469** 36–50.
- Cao, X. 2007. *Stochastic Learning and Optimization: A Sensitivity-Based Approach*. Springer, New York.
- Chao, X., Y. Zhao. 1997. Analysis of multi-server queues with station and server vacations. *Eur. J. Oper. Res.* **110** 392–406.
- Coolen-Schrijner, P., E. van Doorn. 2002. The deviation matrix of a continuous-time Markov chain. *Probab. Engrg. Information Sci.* **16** 351–366.
- Dekker, R., A. Hordijk. 1988. Average, sensitive and Blackwell optimal policies in denumerable Markov decision chains with unbounded rewards. *Math. Oper. Res.* **13** 395–420.
- Dekker, R., A. Hordijk. 1991. Denumerable semi-Markov decision chains with small interest rates. *Ann. Oper. Res.* **28** 185–212.
- Dekker, R., A. Hordijk. 1992. Recurrence conditions for average and Blackwell optimality in denumerable state Markov decision chains. *Math. Oper. Res.* **17** 271–289.
- Dekker, R., A. Hordijk, F. Spieksma. 1994. On the relation between recurrence and ergodicity properties in denumerable Markov decision chains. *Math. Oper. Res.* **19** 539–559.
- Dynkin, E., A. Yushkevich. 1979. *Controlled Markov Processes*. Springer-Verlag, New York.
- Feinberg, E., A. Schwartz. 2002. *Handbook of Markov Decision Processes: Methods and Applications*. Kluwer, Boston.
- Guo, X. 2007. Continuous-time Markov decision processes with discounted rewards: The case of Polish spaces. *Math. Oper. Res.* **32** 73–87.
- Guo, X., X. Cao. 2005. Optimal control of ergodic continuous-time Markov chains with average sample-path rewards. *SIAM J. Control Optim.* **44** 29–48.
- Guo, X., O. Hernández-Lerma. 2003. Continuous-time controlled Markov chains with discounted rewards. *Acta Appl. Math.* **79** 195–216.
- Guo, X., K. Liu. 2001. A note on optimality conditions for continuous-time Markov decision processes with average cost criterion. *IEEE Trans. Automatic Control* **46** 1984–1989.
- Guo, X., O. Hernández-Lerma, T. Prieto-Rumeau. 2006. A survey of recent results on continuous-time Markov decision processes. *Sociedad de Estadística e Investigación Operativa TOP* **14** 177–261.
- Heidergott, B., A. Hordijk, N. Leder. 2009. Series expansions for continuous-time Markov processes. *Oper. Res.*, ePub ahead of print October 28, <http://or.journal.informs.org/cgi/content/abstract/opre.1090.0738v1>.
- Heidergott, B., A. Hordijk, M. van Uitert. 2007. Series expansions for finite-state Markov chains. *Probab. Engrg. Information Sci.* **21** 381–400.
- Hordijk, A., F. Spieksma. 1989. Constrained admission control to a queueing system. *Adv. Appl. Probab.* **21** 409–431.
- Hordijk, A., F. van der Duyn Schouten. 1983. Average optimal policies in Markov decision drift processes with applications to a queueing and a replacement model. *Adv. Appl. Probab.* **15** 274–303.
- Hordijk, A., A. Yushkevich. 1999. Blackwell optimality in the class of all policies in Markov decision chains with a Borel state and unbounded rewards. *Math. Methods Oper. Res.* **50** 421–448.
- Hordijk, A., A. Yushkevich. 2002. Blackwell optimality. E. Feinberg, A. Schwartz, eds. *Handbook of Markov Decision Processes*. Kluwer, Boston, 231–267.
- Kakumanu, P. 1975. Continuous time Markov decision processes with average return criterion. *J. Math. Anal. Appl.* **52** 173–188.
- Kijima, M. 1997. *Markov Processes for Stochastic Modeling*. Chapman & Hall, London.
- Koole, G. 1998. The deviation matrix of the M/M/1/ ∞ and M/M/1/N queue, with application to controlled queueing models. *Proc. 37th IEEE CDC*. IEEE Press, Tampa, FL, 56–59.
- Neuts, M. 1994. *Matrix-Geometric Solutions in Stochastic Models—An Algorithmic Approach*. Constable and Company, London.
- Olsson, M. 1998. The EMPht-programme. *Manual*. Gothenburg University, Gothenburg, Sweden.
- Puterman, M. 1994. *Markov Decision Processes*. John Wiley & Sons, New York.
- Riska, A., E. Smirni. 2002. M/G/1-type Markov processes: A tutorial. *Performance Evaluation of Complex Systems: Techniques and Tools* **2459** 315–325.
- Ross, S. 1970. *Applied Probability Models with Optimization Applications*. Holden-Day, San Francisco.
- Schweitzer, E. 1968. Perturbation theory and finite Markov chains. *J. Appl. Probab.* **5** 401–413.
- Serfozo, R. 1979. An equivalence between continuous and discrete time Markov decision processes. *Oper. Res.* **27** 616–620.
- Stidham, S., Jr. 1985. Optimal control to a queueing system. *IEEE Trans. Automatic Control* **30** 705–713.
- Tijms, H. 1994. *Stochastic Models, An Algorithmic Approach*. John Wiley & Sons, New York.
- Yushkevich, A. 1977. Controlled Markov models with countable state and continuous time. *Theory Probab. Its Appl.* **22** 215–235.
- Zhang, K., Y. Xu, X. Chen, X. Cao. 2008. Policy iteration based feedback control. *Automatica* **44** 1055–1061.