



ETAQA-MG1: an efficient technique for the analysis of a class of M/G/1-type processes by aggregation[☆]

Gianfranco Ciardo^{*}, Weizhen Mao, Alma Riska¹, Evgenia Smirni

Department of Computer Science, College of William and Mary, P.O. Box 8795, Williamsburg, VA 23187-8795, USA

Received 13 October 2000; received in revised form 10 November 2003

Available online 24 January 2004

Abstract

We extend the ETAQA approach, initially proposed for the efficient numerical solution of a class of quasi-birth–death processes, to a more complex class of M/G/1-type Markov processes where arbitrary forward transitions are allowed but backward transitions must be to a single state to the previous level. The new technique reduces the exact solution of this class of M/G/1-type models to that of a finite linear system. We demonstrate the utility of our method by describing the exact computation of an extensive class of Markov reward functions that include the expected queue length or its higher moments. We also provide an algorithm that finds an appropriate state reordering satisfying our applicability conditions, if one such order exists. We illustrate the method, discuss its complexity and numerical stability, and present comparisons with other traditional techniques.

© 2003 Elsevier B.V. All rights reserved.

Keywords: M/G/1-type Markov chain; Matrix-analytic solution; Bulk arrival

1. Introduction

We consider Markov chains on an infinite state space having an M/G/1-type structure. Such processes are often the modeling tool of choice for modern computer and communication systems [13]. As a consequence, considerable effort has been placed into the development of exact analysis techniques for them. In the continuous-time case, the infinitesimal generator of such processes is upper block Hessenberg

[☆] A preliminary version of this paper was presented at the Third Meeting on the Numerical Solution of Markov Chains, Zaragoza, Spain, 1999 [2]. This work was supported by National Science Foundation under grants no. EIA-9974992, CCR-0098278 and ACR-0090221, by the National Aeronautics and Space Administration under NASA Grant NAG-1-2168, and by a William and Mary Summer Research Grant.

^{*} Corresponding author. Fax: +1-757-221-1717.

E-mail addresses: ciardo@cs.wm.edu (G. Ciardo), wm@cs.wm.edu (W. Mao), alma.riska@seagate.com (A. Riska), esmirni@cs.wm.edu (E. Smirni).

¹ Present address: Seagate Research, 1251 Waterfront Place, Pittsburgh, PA 15222, USA.

with a repetitive structure, and matrix-analytic methods have been proposed for their solution [15]. The key in the matrix-analytic solution is the computation of an auxiliary matrix, \mathbf{G} . Similarly, for Markov chains of the GI/M/1-type, which have a lower block Hessenberg form, matrix-geometric solutions have been proposed [14]. Again, the key is the computation of an auxiliary matrix, \mathbf{R} . The traditional solution algorithms compute the stationary probability vector with a recursive function based on \mathbf{G} (for the case of M/G/1-type processes) or \mathbf{R} (for the case of GI/M/1-type processes). Various iterative algorithms for the efficient computation of \mathbf{G} (e.g., the work of Meini [12]) or \mathbf{R} (e.g., the work of Latouche [8]) have been proposed. See also [9] for a discussion of the fundamental aspects of matrix-analytic methods.

Our research differs from the above works in that we restrict our attention to a family of M/G/1-type processes with a specific form, for which “returns” from a higher “level” of states to the immediately lower level are always directed toward a single state. For such a subclass, the computation of the matrix \mathbf{G} is trivial, but the remaining solution steps using traditional methods are still expensive. We instead recast the problem into that of solving a finite linear system in $m + 2n$ unknowns, where m is the number of states in the boundary portion of the process and n is the number of states in each of the repetitive levels of the state space, and are able to obtain exact results.

Since the method that we propose applies only to a subclass of M/G/1-type processes, we further investigate an algorithm to repartition the state space such that the condition for backward transitions to a single state hold. This algorithm is guaranteed to find such a repartitioning if one exists. Thus, our technique does not apply to all M/G/1-type processes, but we argue that it applies to an important subclass. For example, any queueing system with a Coxian service process and bulk Poisson arrival process results in an M/G/1-type chain with backward transitions directed to a single state, and it is a known fact that the class of Coxian distributions is dense [4] and can approximate any distribution, given enough phases. Also, [3] presents a survey of applications from the literature whose Markov chains exhibit the required structure.

Our approach is an extension of the ETAQA method we introduced for the efficient solution of quasi-birth–death (QBD) processes with matrix-geometric form [3]. The proposed methodology uses basic results for Markov chains. We assume that the state space \mathcal{S} is partitioned into $\mathcal{S}^{(0)} = \{s_1^{(0)}, \dots, s_m^{(0)}\}$, containing m “boundary” states and, $\mathcal{S}^{(j)}$ for $j \geq 1$, each containing n “repetitive” states, $\mathcal{S}^{(j)} = \{s_1^{(j)}, \dots, s_n^{(j)}\}$. Then, we exploit the structure of the repetitive portion of the chain and instead of evaluating the probability distribution of *all* states in the chain, we calculate the *aggregate* probability of being in each of the n equivalence classes $\mathcal{T}_i = \{s_i^{(j)} : j \geq 2\}$ for $1 \leq i \leq n$ (corresponding to a specific partition of the repetitive portion of the chain, see Fig. 1), as well as the stationary probability of each state in $\mathcal{S}^{(0)} \cup \mathcal{S}^{(1)}$.

The new ETAQA-MG1 approach is both exact and efficient for the computation of a class of Markov reward functions that include moments of the queue length. Indeed, our method does not explicitly compute the *entire* stationary probability vector but only the part of it that is related to the boundary

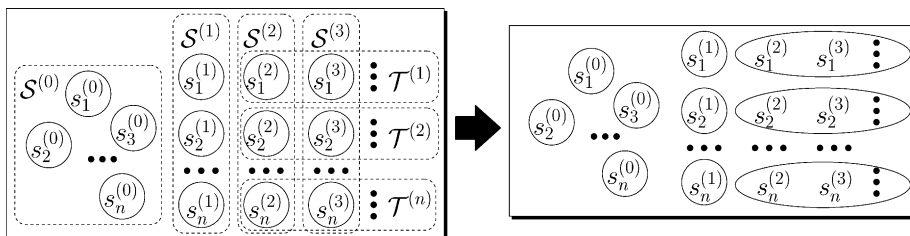


Fig. 1. Aggregation of an infinite \mathcal{S} into a finite number of states.

states, and the aggregate probability vector for the repetitive states. The traditional matrix-analytic method uses a recursive formula for the computation of the stationary probability vector. If this formula entails subtractions, then there is the possibility of numerical instability [15,17]. The preferred method for the recursive computation of the stationary probability vector uses Ramaswami’s recursive formula [17], which entails only additions. Thus, we compare experimentally the computation and storage complexity and the stability of our methodology with that of the matrix-analytic method based on Ramaswami’s formula.

Our paper is organized as follows. Section 2 contains terminology and related work. In Section 3, we present the basic theorem that extends ETAQA to M/G/1-type processes. We demonstrate how the methodology can be used for the computation of Markov reward functions in Section 4. We continue by showing the applicability of ETAQA-MG1 to bounded bulk arrivals in Section 5. In Section 6, we compare the computation and storage complexity of ETAQA-MG1 with the matrix-analytic method. We briefly refer in Section 7 to an algorithm that can repartition an infinitesimal generator into a form such that ETAQA-MG1 can be applied and elaborate on this algorithm in Appendix A. Section 8 presents a computer system modeling application that can be analyzed using the ETAQA-MG1 approach. Section 9 discusses the numerical stability of our approach. Finally, Section 10 summarizes our contributions.

2. Background

We briefly review the terminology used to describe the class of processes we consider. We restrict ourselves to the case of continuous-time Markov chains (hence we refer to the infinitesimal generator matrix \mathbf{Q}), but the theory can just as well be applied to the discrete case.

Neuts [14] defines various classes of infinite-state Markov chains with a repetitive structure. In all cases, the state space is partitioned into $\mathcal{S}^{(0)} = \{s_1^{(0)}, \dots, s_m^{(0)}\}$ and $\mathcal{S}^{(j)} = \{s_1^{(j)}, \dots, s_n^{(j)}\}$ for $j \geq 1$. For the class of M/G/1-type Markov chains, the infinitesimal generator \mathbf{Q} is block-partitioned as

$$\mathbf{Q} = \begin{bmatrix} \hat{\mathbf{L}}^{(0)} & \hat{\mathbf{F}}^{(1)} & \hat{\mathbf{F}}^{(2)} & \hat{\mathbf{F}}^{(3)} & \hat{\mathbf{F}}^{(4)} & \hat{\mathbf{F}}^{(5)} & \dots \\ \hat{\mathbf{B}} & \mathbf{L} & \mathbf{F}^{(1)} & \mathbf{F}^{(2)} & \mathbf{F}^{(3)} & \mathbf{F}^{(4)} & \dots \\ \mathbf{0} & \mathbf{B} & \mathbf{L} & \mathbf{F}^{(1)} & \mathbf{F}^{(2)} & \mathbf{F}^{(3)} & \dots \\ \mathbf{0} & \mathbf{0} & \mathbf{B} & \mathbf{L} & \mathbf{F}^{(1)} & \mathbf{F}^{(2)} & \dots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{B} & \mathbf{L} & \mathbf{F}^{(1)} & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \tag{1}$$

we use the letters “L”, “F”, and “B” according to whether they describe “local”, “forward”, and “backward” transition rates, respectively, and “^” for matrices related to $\mathcal{S}^{(0)}$.

For the solution of M/G/1-type processes, several recursive algorithms exist [5,12,15]. Here, we outline Ramaswami’s recursive formula [17], which provides a stable calculation for the values of $\boldsymbol{\pi}^{(j)}$, the stationary probability vector for states in $\mathcal{S}^{(j)}$:

$$\forall j \geq 1, \quad \boldsymbol{\pi}^{(j)} = - \left(\boldsymbol{\pi}^{(0)} \hat{\mathbf{S}}^{(j)} + \sum_{l=1}^{j-1} \boldsymbol{\pi}^{(l)} \mathbf{S}^{(j-l)} \right) \mathbf{S}^{(0)-1}, \tag{2}$$

where $\hat{\mathbf{S}}^{(j)}$ and $\mathbf{S}^{(j)}$ are defined as follows:

$$\hat{\mathbf{S}}^{(j)} = \sum_{l=j}^{\infty} \hat{\mathbf{F}}^{(l)} \mathbf{G}^{l-j}, \quad j \geq 1, \quad \mathbf{S}^{(j)} = \sum_{l=j}^{\infty} \mathbf{F}^{(l)} \mathbf{G}^{l-j}, \quad j \geq 0 \text{ (letting } \mathbf{F}^{(0)} = \mathbf{L})$$

and where, in turn, \mathbf{G} is the solution of the matrix equation

$$\mathbf{B} + \mathbf{L}\mathbf{G} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)} \mathbf{G}^{j+1} = \mathbf{0}. \tag{3}$$

Several iterative algorithms for the computation of \mathbf{G} exist [5,15], the most efficient being cyclic-reduction [12]. See [18] for cases where \mathbf{G} can be explicitly defined and does not require any calculation.

Given the above expression for $\pi^{(j)}$ and the normalization condition, a unique vector $\pi^{(0)}$ can be obtained by solving the following system of m equations

$$\pi^{(0)} \left[\left(\hat{\mathbf{L}}^{(0)} - \hat{\mathbf{S}}^{(1)} \mathbf{S}^{(0)-1} \hat{\mathbf{B}} \right) \diamond \mathbf{1}^T - \left(\sum_{j=1}^{\infty} \hat{\mathbf{S}}^{(j)} \right) \left(\sum_{j=0}^{\infty} \mathbf{S}^{(j)} \right)^{-1} \mathbf{1}^T \right] = [\mathbf{0} | \mathbf{1}], \tag{4}$$

where the symbol “ \diamond ” indicates that we discard one (any) column of the corresponding matrix, since we added a column representing the normalization condition. Once $\pi^{(0)}$ is known, we can then iteratively compute $\pi^{(j)}$ for $j = 1, 2, \dots$, stopping when the accumulated probability mass is close to one.

3. Extending ETAQA to M/G/1-type processes

We consider the following structure in the infinitesimal generator matrix \mathbf{Q} :

$$\mathbf{Q} = \begin{bmatrix} \hat{\mathbf{L}}^{(0)} & \hat{\mathbf{F}}^{(1)} & \hat{\mathbf{F}}^{(2)} & \hat{\mathbf{F}}^{(3)} & \hat{\mathbf{F}}^{(4)} & \dots & \dots \\ \hat{\mathbf{B}} & \mathbf{L}^{(1)} & \mathbf{F}^{(1)} & \mathbf{F}^{(2)} & \mathbf{F}^{(3)} & \dots & \dots \\ \mathbf{0} & \mathbf{B} & \mathbf{L} & \mathbf{F}^{(1)} & \mathbf{F}^{(2)} & \dots & \dots \\ \mathbf{0} & \mathbf{0} & \mathbf{B} & \mathbf{L} & \mathbf{F}^{(1)} & \dots & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \tag{5}$$

which differs from that of Eq. (1) only in block $\mathbf{L}^{(1)}$, at the intersection of the rows and columns for $\mathcal{S}^{(1)}$, which is not restricted to be equal to \mathbf{L} . This allows us to have local transitions from $\mathcal{S}^{(1)}$ to $\mathcal{S}^{(1)}$ with a pattern other than those from $\mathcal{S}^{(j)}$ to $\mathcal{S}^{(j)}$, $j \geq 2$, and a total rate from each state $s_i^{(1)}$ to the states in $\mathcal{S}^{(0)}$ other than the total rate from state $s_i^{(j)}$ to the states in $\mathcal{S}^{(j-1)}$, i.e., the row sums of $\hat{\mathbf{B}}$ and \mathbf{B} can be different. For \mathbf{Q} to be an infinitesimal generator, the infinite sets of matrices $\{\hat{\mathbf{F}}^{(j)}: j \geq 1\}$ and $\{\mathbf{F}^j: j \geq 1\}$ must be summable. If we also partition the stationary probability vector satisfying $\pi\mathbf{Q} = \mathbf{0}$

as $\boldsymbol{\pi} = [\boldsymbol{\pi}^{(0)}, \boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}, \dots]$, with $\boldsymbol{\pi}^{(0)} \in \mathbb{R}^m$ and $\boldsymbol{\pi}^{(j)} \in \mathbb{R}^n$ for $j \geq 1$, we can then write $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$ as

$$\begin{cases} \boldsymbol{\pi}^{(0)}\hat{\mathbf{L}}^{(0)} + \boldsymbol{\pi}^{(1)}\hat{\mathbf{B}} = \mathbf{0}, \\ \boldsymbol{\pi}^{(0)}\hat{\mathbf{F}}^{(1)} + \boldsymbol{\pi}^{(1)}\mathbf{L}^{(1)} + \boldsymbol{\pi}^{(2)}\mathbf{B} = \mathbf{0}, \\ \boldsymbol{\pi}^{(0)}\hat{\mathbf{F}}^{(2)} + \boldsymbol{\pi}^{(1)}\mathbf{F}^{(1)} + \boldsymbol{\pi}^{(2)}\mathbf{L} + \boldsymbol{\pi}^{(3)}\mathbf{B} = \mathbf{0}, \\ \boldsymbol{\pi}^{(0)}\hat{\mathbf{F}}^{(3)} + \boldsymbol{\pi}^{(1)}\mathbf{F}^{(2)} + \boldsymbol{\pi}^{(2)}\mathbf{F}^{(1)} + \boldsymbol{\pi}^{(3)}\mathbf{L} + \boldsymbol{\pi}^{(4)}\mathbf{B} = \mathbf{0}, \\ \vdots \end{cases} \tag{6}$$

3.1. Conditions for stability

We briefly review the conditions that enable us to assert that the CTMC described by the infinitesimal generator \mathbf{Q} in (5) is stable, that is, admits a probability vector satisfying $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$ and $\boldsymbol{\pi}\mathbf{1}^T = 1$.

First, observe that the matrix $\tilde{\mathbf{Q}} = \mathbf{B} + \mathbf{L} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)}$ is an infinitesimal generator, since it has zero row sums and non-negative off-diagonal entries. If $\tilde{\mathbf{Q}}$ is irreducible, any state $s_i^{(j)}$ in the original process can reach any state $s_{i'}^{(j')}$ for $2 \leq j \leq j'$ and $1 \leq i, i' \leq n$, without having to go through states in $\mathcal{S}^{(l)}$, $l < j$. In this case, for “large values of j ”, the conditional probability of being in $s_i^{(j)}$ given that we are in $\mathcal{S}^{(j)}$ tends to $\tilde{\boldsymbol{\pi}}_i$, where $\tilde{\boldsymbol{\pi}}$ is the unique solution of $\tilde{\boldsymbol{\pi}}\tilde{\mathbf{Q}} = \mathbf{0}$, subject to $\tilde{\boldsymbol{\pi}}\mathbf{1}^T = 1$, that is, the stationary solution of the ergodic CTMC having $\tilde{\mathbf{Q}}$ as the infinitesimal generator.

For a proof of this statement in the case when only the first p matrices $\mathbf{F}^{(1)}, \dots, \mathbf{F}^{(p)}$ can be nonzero, simply observe that, if the process is in one of the states of $\mathcal{S}^{(j)}$, the last $j/p - 1$ state transitions (at least) *must* have been taken according to the rates specified by \mathbf{B}, \mathbf{L} , and $\mathbf{F}^{(1)}, \dots, \mathbf{F}^{(p)}$. Thus, for large j , the conditional probability of being in $s_i^{(j)}$ tends to the probability of being in state i in the CTMC with generator $\tilde{\mathbf{Q}}$. For the case where $\mathbf{F}^{(j)}$ can be nonzero for arbitrarily large j , recall that its rates must tend to zero faster than $1/j$ (since these matrices are summable), thus a “truncation” of this “exact” process to a sufficiently large maximum value p for the forward jumps will behave in essentially the same way. In other words, in the case of unbounded forward jumps, the last $j/p - 1$ transitions will have been taken according to the rates in $\tilde{\mathbf{Q}}$ with a probability, dependent on p , which can be made arbitrarily close to one.

Then, the M/G/1-type process is stable as long as, for large values of j , the forward drift from $\mathcal{S}^{(j)}$ is less than the backward drift from it:

$$\tilde{\boldsymbol{\pi}} \left(\sum_{j=1}^{\infty} j\mathbf{F}^{(j)} \right) \mathbf{1}^T < \tilde{\boldsymbol{\pi}}\mathbf{B}\mathbf{1}^T.$$

This stability condition can be verified numerically, and it is easy to see that it is equivalent to the one given by Neuts in [15], $\tilde{\boldsymbol{\pi}}\boldsymbol{\beta} < 1$, where, in our terminology, the column vector $\boldsymbol{\beta}$ is given by $\boldsymbol{\beta} = (\mathbf{L} + \sum_{j=1}^{\infty} j\mathbf{F}^{(j)})\mathbf{1}^T$. As in the scalar case, the condition where $\tilde{\boldsymbol{\pi}}\boldsymbol{\beta}$ is exactly equal to 1 results in a null-recurrent CTMC. If $\tilde{\mathbf{Q}}$ is instead reducible, this stability condition must be applied to each subset of $\{1, \dots, n\}$ corresponding to a recurrent class in the CTMC described by $\tilde{\mathbf{Q}}$.

3.2. Main theorem

As done in [3], we require that \mathbf{B} contains nonzero entries only in one column, which we assume to be the last one, column n , without loss of generality. The main idea behind our approach is to transform the infinitely countable set of linear equations $\boldsymbol{\pi} \mathbf{Q} = \mathbf{0}$, subject to $\boldsymbol{\pi} \mathbf{1}^T = 1$, into $m + 2n$ equations in $\boldsymbol{\pi}^{(0)}$, $\boldsymbol{\pi}^{(1)}$, and the new aggregate vector of n unknowns $\boldsymbol{\pi}^{(*)} = \sum_{j=2}^{\infty} \boldsymbol{\pi}^{(j)}$. This finite system can then be solved using ordinary numerical techniques, i.e., either direct methods such as Gaussian elimination, if $m + 2n$ is not too large, or indirect iterative methods such as Gauss–Seidel, which can better exploit the sparsity of the system, thus are applicable even when $m + 2n$ is of the order 10^4 or 10^5 , provided the sparsity of the original matrices describing \mathbf{Q} is sufficiently high, as is normally the case. Once $\boldsymbol{\pi}^{(0)}$, $\boldsymbol{\pi}^{(1)}$, and $\boldsymbol{\pi}^{(*)}$ have been obtained, very general measures of interest can be computed by solving further linear systems, as discussed in Section 4.

We now show the derivation of our $m + 2n$ equations.

- (i) The normalization condition offers one equation:

$$\boldsymbol{\pi}^{(0)} \mathbf{1}^T + \boldsymbol{\pi}^{(1)} \mathbf{1}^T + \sum_{j=2}^{\infty} \boldsymbol{\pi}^{(j)} \mathbf{1}^T = 1. \tag{7}$$

- (ii) m equations can be obtained from the first line in (6):

$$\boldsymbol{\pi}^{(0)} \hat{\mathbf{L}}^{(0)} + \boldsymbol{\pi}^{(1)} \hat{\mathbf{B}} = \mathbf{0}. \tag{8}$$

- (iii) $n - 1$ equations can be obtained from the second line in (6), which defines n equations, fortunately only the last one actually containing a contribution from $\boldsymbol{\pi}^{(2)}$, due to the structure we require in \mathbf{B} . This is of fundamental importance, since $\boldsymbol{\pi}^{(2)}$ is not one of our variables. Hence, we consider only the first $n - 1$ equations and write as

$$\boldsymbol{\pi}^{(0)} \hat{\mathbf{F}}_{1:m,1:n-1}^{(1)} + \boldsymbol{\pi}^{(1)} \mathbf{L}_{1:n,1:n-1}^{(1)} = \mathbf{0}. \tag{9}$$

- (iv) Another $n - 1$ equations are obtained as follows. First, we sum the remaining lines in (6):

$$\boldsymbol{\pi}^{(0)} \sum_{j=2}^{\infty} \hat{\mathbf{F}}^{(j)} + \boldsymbol{\pi}^{(1)} \sum_{j=1}^{\infty} \mathbf{F}^{(j)} + \left(\sum_{j=2}^{\infty} \boldsymbol{\pi}^{(j)} \right) \left(\mathbf{L} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)} \right) + \sum_{j=3}^{\infty} \boldsymbol{\pi}^{(j)} \mathbf{B} = \mathbf{0}.$$

Then, since $\sum_{j=3}^{\infty} \boldsymbol{\pi}^{(j)} = \boldsymbol{\pi}^{(*)} - \boldsymbol{\pi}^{(2)}$, we again remove the explicit mention of $\boldsymbol{\pi}^{(2)}$ by considering only the first $n - 1$ equations:

$$\boldsymbol{\pi}^{(0)} \sum_{j=2}^{\infty} \hat{\mathbf{F}}_{1:m,1:n-1}^{(j)} + \boldsymbol{\pi}^{(1)} \sum_{j=1}^{\infty} \mathbf{F}_{1:n,1:n-1}^{(j)} + \boldsymbol{\pi}^{(*)} \left(\mathbf{L}_{1:n,1:n-1} + \sum_{j=1}^{\infty} \mathbf{F}_{1:n,1:n-1}^{(j)} \right) = \mathbf{0}. \tag{10}$$

(v) For the last equation, consider the flow balance equations between the states in $\bigcup_{l=0}^j \mathcal{S}^{(l)}$ and those in $\bigcup_{l=j+1}^{\infty} \mathcal{S}^{(l)}$ for $j \geq 1$,

$$\left\{ \begin{array}{l} \pi^{(0)} \sum_{l=2}^{\infty} \hat{\mathbf{F}}^{(l)} \mathbf{1}^T + \pi^{(1)} \sum_{l=1}^{\infty} \mathbf{F}^{(l)} \mathbf{1}^T = \pi^{(2)} \mathbf{B} \mathbf{1}^T, \\ \pi^{(0)} \sum_{l=3}^{\infty} \hat{\mathbf{F}}^{(l)} \mathbf{1}^T + \pi^{(1)} \sum_{l=2}^{\infty} \mathbf{F}^{(l)} \mathbf{1}^T + \pi^{(2)} \sum_{l=1}^{\infty} \mathbf{F}^{(l)} \mathbf{1}^T = \pi^{(3)} \mathbf{B} \mathbf{1}^T, \\ \vdots \\ \pi^{(0)} \sum_{l=j+1}^{\infty} \hat{\mathbf{F}}^{(l)} \mathbf{1}^T + \pi^{(1)} \sum_{l=j}^{\infty} \mathbf{F}^{(l)} \mathbf{1}^T + \pi^{(2)} \sum_{l=j-1}^{\infty} \mathbf{F}^{(l)} \mathbf{1}^T + \dots + \pi^{(j)} \sum_{l=1}^{\infty} \mathbf{F}^{(l)} \mathbf{1}^T = \pi^{(j+1)} \mathbf{B} \mathbf{1}^T, \\ \vdots \end{array} \right. \quad (11)$$

and sum these equations by parts:

$$\pi^{(0)} \sum_{j=2}^{\infty} \sum_{l=j}^{\infty} \hat{\mathbf{F}}^{(l)} \mathbf{1}^T + \pi^{(1)} \sum_{j=1}^{\infty} \sum_{l=j}^{\infty} \mathbf{F}^{(l)} \mathbf{1}^T + \sum_{j=2}^{\infty} \pi^{(j)} \sum_{h=1}^{\infty} \sum_{l=h}^{\infty} \mathbf{F}^{(l)} \mathbf{1}^T = \sum_{j=2}^{\infty} \pi^{(j)} \mathbf{B} \mathbf{1}^T,$$

which can be rewritten as

$$\pi^{(0)} \sum_{j=2}^{\infty} (j-1) \hat{\mathbf{F}}^{(j)} \mathbf{1}^T + \pi^{(1)} \sum_{j=1}^{\infty} j \mathbf{F}^{(j)} \mathbf{1}^T + \pi^{(*)} \left(\sum_{j=1}^{\infty} j \mathbf{F}^{(j)} - \mathbf{B} \right) \mathbf{1}^T = \mathbf{0}. \quad (12)$$

The following theorem states that these equations are sufficient to compute $\pi^{(0)}$, $\pi^{(1)}$, and $\pi^{(*)}$.

Theorem 1 (ETAQA-MG1). *Given an ergodic CTMC with infinitesimal generator \mathbf{Q} having the structure shown in (5), such that the first $n - 1$ columns of \mathbf{B} are null, $\mathbf{B}_{1:n,1:n-1} = \mathbf{0}$, and with stationary probability vector $\boldsymbol{\pi} = [\pi^{(0)}, \pi^{(1)}, \pi^{(2)}, \dots]$, the system of linear equations*

$$\mathbf{xM} = [1, \mathbf{0}], \quad (13)$$

where $\mathbf{M} \in \mathbb{R}^{(m+2n) \times (m+2n)}$ is defined as follows:

$$\mathbf{M} = \left[\begin{array}{c|c|c|c|c} \mathbf{1}^T & \hat{\mathbf{L}}^{(0)} & \hat{\mathbf{F}}_{1:m,1:n-1}^{(1)} & \left(\sum_{j=2}^{\infty} \hat{\mathbf{F}}^{(j)} \right)_{1:m,1:n-1} & \left(\sum_{j=2}^{\infty} (j-1) \hat{\mathbf{F}}^{(j)} \right) \mathbf{1}^T \\ \mathbf{1}^T & \hat{\mathbf{B}} & \mathbf{L}_{1:n,1:n-1}^{(1)} & \left(\sum_{j=1}^{\infty} \mathbf{F}^{(j)} \right)_{1:n,1:n-1} & \left(\sum_{j=1}^{\infty} j \mathbf{F}^{(j)} \right) \mathbf{1}^T \\ \mathbf{1}^T & \mathbf{0} & \mathbf{0} & \left(\sum_{j=1}^{\infty} \mathbf{F}^{(j)} + \mathbf{L} \right)_{1:n,1:n-1} & \left(\sum_{j=1}^{\infty} j \mathbf{F}^{(j)} - \mathbf{B} \right) \mathbf{1}^T \end{array} \right] \quad (14)$$

admits a unique solution $\mathbf{x} = [\pi^{(0)}, \pi^{(1)}, \pi^{(*)}]$, where $\pi^{(*)} = \sum_{j=2}^{\infty} \pi^{(j)}$.

	$\mathbf{v}^{[2]}$	$\mathbf{v}^{[m+2]}$	$\mathbf{v}^{[m+n+1]}$	$\mathbf{v}^{[m+2n+1]}$	$\mathbf{v}^{[m+3n+1]}$		$\mathbf{z}^{[1]}$
$\mathbf{v}^{[1]}$	through	through	through	through	through	...	through
	$\mathbf{v}^{[m+1]}$	$\mathbf{v}^{[m+n]}$	$\mathbf{v}^{[m+2n]}$	$\mathbf{v}^{[m+3n]}$	$\mathbf{v}^{[m+4n]}$		$\mathbf{z}^{[n-1]}$

$\mathbf{1}^T$	$\widehat{\mathbf{L}}^{(0)}$	$\widehat{\mathbf{F}}_{1:m,1:n-1}^{(1)}$	$\widehat{\mathbf{F}}^{(2)}$	$\widehat{\mathbf{F}}^{(3)}$	$\widehat{\mathbf{F}}^{(4)}$...	$\left(\sum_{j=2}^{\infty} \widehat{\mathbf{F}}^{(j)}\right)_{1:m,1:n-1}$
$\mathbf{1}^T$	$\widehat{\mathbf{B}}$	$\mathbf{L}_{1:n,1:n-1}^{(1)}$	$\mathbf{F}^{(1)}$	$\mathbf{F}^{(2)}$	$\mathbf{F}^{(3)}$...	$\left(\sum_{j=1}^{\infty} \mathbf{F}^{(j)}\right)_{1:n,1:n-1}$
$\mathbf{1}^T$	0	0	\mathbf{L}	$\mathbf{F}^{(1)}$	$\mathbf{F}^{(2)}$...	$\left(\sum_{j=1}^{\infty} \mathbf{F}^{(j)} + \mathbf{L}\right)_{1:n,1:n-1}$
$\mathbf{1}^T$	0	0	\mathbf{B}	\mathbf{L}	$\mathbf{F}^{(1)}$...	$\left(\sum_{j=1}^{\infty} \mathbf{F}^{(j)} + \mathbf{L}\right)_{1:n,1:n-1}$
$\mathbf{1}^T$	0	0	0	\mathbf{B}	\mathbf{L}	...	$\left(\sum_{j=1}^{\infty} \mathbf{F}^{(j)} + \mathbf{L}\right)_{1:n,1:n-1}$
$\mathbf{1}^T$	0	0	0	0	\mathbf{B}	...	$\left(\sum_{j=1}^{\infty} \mathbf{F}^{(j)} + \mathbf{L}\right)_{1:n,1:n-1}$
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots

$$\mathbf{z}^{[n]} = \sum_{j=1}^{\infty} \sum_{l=j}^{\infty} \sum_{i=1}^n \mathbf{v}^{[m+ln+i]} = \begin{bmatrix} \sum_{j=1}^{\infty} \sum_{l=j}^{\infty} \widehat{\mathbf{F}}^{(l+1)} \mathbf{1}^T \\ \sum_{j=1}^{\infty} \sum_{l=j}^{\infty} \mathbf{F}^{(l)} \mathbf{1}^T \\ \left(\mathbf{L} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)} + \sum_{j=1}^{\infty} \sum_{l=j}^{\infty} \mathbf{F}^{(l)}\right) \mathbf{1}^T \\ \left(\mathbf{B} + 2\left(\mathbf{L} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)}\right) + \sum_{j=1}^{\infty} \sum_{l=j}^{\infty} \mathbf{F}^{(l)}\right) \mathbf{1}^T \\ \left(2\mathbf{B} + 3\left(\mathbf{L} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)}\right) + \sum_{j=1}^{\infty} \sum_{l=j}^{\infty} \mathbf{F}^{(l)}\right) \mathbf{1}^T \\ \vdots \end{bmatrix} = \begin{bmatrix} \sum_{j=1}^{\infty} j \widehat{\mathbf{F}}^{(j+1)} \mathbf{1}^T \\ \sum_{j=1}^{\infty} j \mathbf{F}^{(j)} \mathbf{1}^T \\ \left(\sum_{j=1}^{\infty} j \mathbf{F}^{(j)} - \mathbf{B}\right) \mathbf{1}^T \\ \left(\sum_{j=1}^{\infty} j \mathbf{F}^{(j)} - \mathbf{B}\right) \mathbf{1}^T \\ \left(\sum_{j=1}^{\infty} j \mathbf{F}^{(j)} - \mathbf{B}\right) \mathbf{1}^T \\ \vdots \end{bmatrix}$$

Fig. 2. The column vectors used to prove linear independence.

Proof. The vector $[\boldsymbol{\pi}^{(0)}, \boldsymbol{\pi}^{(1)}, \sum_{j=2}^{\infty} \boldsymbol{\pi}^{(j)}]$ satisfies (7)–(10) and (12), hence it is a solution of (13). To show that this solution is unique, it is enough to prove that the rank of \mathbf{M} is $m + 2n$, i.e., its $m + 2n$ rows are linearly independent. Recall that, in the construction of \mathbf{M} , we use either columns from the infinitesimal generator \mathbf{Q} (i.e., the first block of columns directly as it is in \mathbf{Q} and the second block of columns in \mathbf{Q} after removing the column corresponding to the nonzero column of \mathbf{B}) or columns built using the remaining block columns of \mathbf{Q} (i.e., a block of columns obtained by summing these blocks after removing the column corresponding to the nonzero column of \mathbf{B} , and a column built as a linear combination of the above columns in \mathbf{Q} , including those corresponding to the nonzero column of \mathbf{B}), plus of course the normalization condition. The following formally shows that the resulting $m + 2n$ columns of \mathbf{M} are linearly independent.

Since \mathbf{Q} is ergodic, we know that the vector $\mathbf{1}^T$ and the set of vectors corresponding to all the columns of \mathbf{Q} except one (any one of them) are linearly independent. In particular, we choose to remove from \mathbf{Q} the n th column of the second block of columns in (5), corresponding to the transitions into state s_n^1 . The result is then the countably infinite set of linearly independent column vectors of \mathbb{R}^N , $\{\mathbf{v}^{[1]}, \mathbf{v}^{[2]}, \dots\}$, shown in Fig. 2. We then define n new vectors of \mathbb{R}^N , $\{\mathbf{z}^{[1]}, \dots, \mathbf{z}^{[n]}\}$, as follows (see Fig. 2):

- For $i = 1, \dots, n - 1$, let $\mathbf{z}^{[i]} = \sum_{j=1}^{\infty} \mathbf{v}^{[m+jn+i]}$, that is we sum the i th column of each block of columns, starting from the block corresponding to transitions into level $\mathcal{S}^{(2)}$. \mathbf{B} does not appear in the expression of these vectors because $\mathbf{B}_{1:n,1:n-1} = \mathbf{0}$.

- Let $\mathbf{z}^{[n]} = \sum_{j=1}^{\infty} \sum_{l=j}^{\infty} \sum_{i=1}^n \mathbf{v}^{[m+ln+i]}$. Fig. 2 shows the steps required to obtain a closed-form expression for $\mathbf{z}^{[n]}$, which are proved by recalling that, since \mathbf{Q} is an infinitesimal generator, $(\mathbf{B} + \mathbf{L} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)})\mathbf{1}^T = \mathbf{0}$, which also implies $(\mathbf{L} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)})\mathbf{1}^T = -\mathbf{B}\mathbf{1}^T$.

Then, we can show that the $m + 2n$ vectors $\{\mathbf{v}^{[1]}, \dots, \mathbf{v}^{[m+n]}, \mathbf{z}^{[1]}, \dots, \mathbf{z}^{[n]}\}$ are linearly independent, since the original set $\{\mathbf{v}^{[1]}, \mathbf{v}^{[2]}, \dots\}$ is linearly independent and the vectors $\{\mathbf{z}^{[1]}, \dots, \mathbf{z}^{[n]}\}$ are obtained as linear combinations of different subsets of vectors from $\{\mathbf{v}^{[m+n+1]}, \mathbf{v}^{[m+n+2]}, \dots\}$:

- Disjoint subsets of vectors are used to build $\{\mathbf{z}^{[1]}, \dots, \mathbf{z}^{[n-1]}\}$.
- $\mathbf{z}^{[n]}$ is built using vectors already used for $\{\mathbf{z}^{[1]}, \dots, \mathbf{z}^{[n-1]}\}$, but also vectors of the form $\mathbf{v}^{[m+jn]}$, which are not used to build any of the vectors in $\{\mathbf{z}^{[1]}, \dots, \mathbf{z}^{[n-1]}\}$.

Hence, the matrix having as column the vectors $\{\mathbf{v}^{[1]}, \dots, \mathbf{v}^{[m+n]}, \mathbf{z}^{[1]}, \dots, \mathbf{z}^{[n]}\}$ has rank $m + 2n$, which implies that we must be able to find $m + 2n$ linearly independent rows in it. Since row $m + jn + i$ is identical to row $m + n + i$ for $j > 1$ and $i = 1, \dots, n$, the first $m + 2n$ rows must be linearly independent. These are also the rows of our matrix \mathbf{M} , so the proof is complete. \square

4. Computing the measures of interest

We now consider the problem of obtaining stationary measures of interest once $\boldsymbol{\pi}^{(0)}$, $\boldsymbol{\pi}^{(1)}$, and $\boldsymbol{\pi}^{(*)}$ have been computed. We consider measures that can be expressed as the expected reward rate

$$r = \sum_{j=0}^{\infty} \sum_{i \in \mathcal{S}^{(j)}} \rho_i^{(j)} \pi_i^{(j)},$$

where $\rho_i^{(j)}$ is the reward rate of state $s_i^{(j)}$. For example, if we wanted to compute the expected queue length in steady state for a model where $\mathcal{S}^{(j)}$ contains the system states with j customers in the queue, we would let $\rho_i^{(j)} = j$, while, to compute the second moment of the queue length, we would let $\rho_i^{(j)} = j^2$.

Since our solution approach computes $\boldsymbol{\pi}^{(0)}$, $\boldsymbol{\pi}^{(1)}$, and $\boldsymbol{\pi}^{(*)} = \sum_{j=2}^{\infty} \boldsymbol{\pi}^{(j)}$, we rewrite r as

$$r = \boldsymbol{\pi}^{(0)} \boldsymbol{\rho}^{(0)T} + \boldsymbol{\pi}^{(1)} \boldsymbol{\rho}^{(1)T} + \sum_{j=2}^{\infty} \boldsymbol{\pi}^{(j)} \boldsymbol{\rho}^{(j)T},$$

where $\boldsymbol{\rho}^{(0)} = [\rho_1^{(0)}, \dots, \rho_m^{(0)}]$ and $\boldsymbol{\rho}^{(j)} = [\rho_1^{(j)}, \dots, \rho_n^{(j)}]$ for $j \geq 1$. Then, we must show how to compute the above summation without explicitly using the values of $\boldsymbol{\pi}^{(j)}$ for $j \geq 2$. We can do so if the reward rate of state $s_i^{(j)}$ for $j \geq 2$ and $i = 1, \dots, n$, is a polynomial of degree k in j with arbitrary coefficients $\mathbf{a}_i^{[0]}, \mathbf{a}_i^{[1]}, \dots, \mathbf{a}_i^{[k]}$:

$$\forall j \geq 2, \forall i \in \{1, 2, \dots, n\}, \quad \rho_i^{(j)} = \mathbf{a}_i^{[0]} + \mathbf{a}_i^{[1]}j + \dots + \mathbf{a}_i^{[k]}j^k. \tag{15}$$

The above equation allows the reward rate for each state $s_i^{(j)}$ in the set $\mathcal{S}^{(j)}$ to be a general polynomial function with arbitrary coefficient that may differ from those for other states in $\mathcal{S}^{(j)}$. We emphasize this because, in Section 7, the states within $\mathcal{S}^{(j)}$ may need to have different reward rates, and the method we describe below is general enough to handle these cases as well.

Using the representation of Eq. (15) we write

$$\begin{aligned} \sum_{j=2}^{\infty} \boldsymbol{\pi}^{(j)} \boldsymbol{\rho}^{(j)\text{T}} &= \sum_{j=2}^{\infty} \boldsymbol{\pi}^{(j)} (\mathbf{a}^{[0]} + \mathbf{a}^{[1]} j + \dots + \mathbf{a}^{[k]} j^k)^{\text{T}} \\ &= \sum_{j=2}^{\infty} \boldsymbol{\pi}^{(j)} \mathbf{a}^{[0]\text{T}} + \sum_{j=2}^{\infty} j \boldsymbol{\pi}^{(j)} \mathbf{a}^{[1]\text{T}} + \dots + \sum_{j=2}^{\infty} j^k \boldsymbol{\pi}^{(j)} \mathbf{a}^{[k]\text{T}} \\ &= \mathbf{r}^{[0]} \mathbf{a}^{[0]\text{T}} + \mathbf{r}^{[1]} \mathbf{a}^{[1]\text{T}} + \dots + \mathbf{r}^{[k]} \mathbf{a}^{[k]\text{T}}, \end{aligned}$$

and the problem is reduced to the computation of $\mathbf{r}^{[l]} = \sum_{j=2}^{\infty} j^l \boldsymbol{\pi}^{(j)}$ for $l = 0, \dots, k$ (for an interpretation of $\mathbf{r}^{[l]}$, one can identify it, except for the missing term $\boldsymbol{\pi}^{(1)}$, with the l th vector moment of the random variable representing the level j at which the chain resides).

We show how $\mathbf{r}^{[k]}$ for $k > 0$, can be computed recursively, starting from $\mathbf{r}^{[0]}$, which is simply $\boldsymbol{\pi}^{(*)}$. We first define the following sums:

$$\hat{\mathbf{F}}^{[j]} = \sum_{l=j}^{\infty} \hat{\mathbf{F}}^{(l)}, \quad \mathbf{F}^{[j]} = \sum_{l=j}^{\infty} \mathbf{F}^{(l)}, \quad \hat{\mathbf{F}}^{[c,i]} = \sum_{j=1}^{\infty} (j+c)^i \hat{\mathbf{F}}^{(j)} \quad \text{and} \quad \mathbf{F}^{[c,i]} = \sum_{j=1}^{\infty} (j+c)^i \mathbf{F}^{(j)}.$$

Multiplying the equations in (6) from the second line on by the appropriate factor j^k results in

$$\begin{cases} 2^k \boldsymbol{\pi}^{(0)} \hat{\mathbf{F}}^{(1)} + 2^k \boldsymbol{\pi}^{(1)} \mathbf{L}^{(1)} + 2^k \boldsymbol{\pi}^{(2)} \mathbf{B} = \mathbf{0}, \\ 3^k \boldsymbol{\pi}^{(0)} \hat{\mathbf{F}}^{(2)} + 3^k \boldsymbol{\pi}^{(1)} \mathbf{F}^{(1)} + 3^k \boldsymbol{\pi}^{(2)} \mathbf{L} + 3^k \boldsymbol{\pi}^{(3)} \mathbf{B} = \mathbf{0}, \\ \vdots \end{cases} \tag{16}$$

Summing the equations in (16) by parts we obtain

$$\begin{aligned} &\underbrace{\boldsymbol{\pi}^{(0)} \sum_{j=1}^{\infty} (j+1)^k \hat{\mathbf{F}}^{(j)}}_{=\hat{\mathbf{F}}^{[1,k]}} + \underbrace{\boldsymbol{\pi}^{(1)} 2^k \mathbf{L}^{(1)} + \boldsymbol{\pi}^{(1)} \sum_{j=1}^{\infty} (j+2)^k \mathbf{F}^{(j)}}_{=\mathbf{F}^{[2,k]}} \\ &+ \sum_{h=2}^{\infty} \boldsymbol{\pi}^{(h)} \left[\sum_{j=1}^{\infty} (j+h+1)^k \mathbf{F}^{(j)} + (h+1)^k \mathbf{L} + h^k \mathbf{B} \right] = \mathbf{0}, \end{aligned}$$

which can be rewritten as

$$\begin{aligned} &\boldsymbol{\pi}^{(0)} \hat{\mathbf{F}}^{[1,k]} + \boldsymbol{\pi}^{(1)} (2^k \mathbf{L}^{(1)} + \mathbf{F}^{[2,k]}) + \sum_{h=2}^{\infty} \boldsymbol{\pi}^{(h)} \sum_{j=1}^{\infty} (j+h+1)^k \mathbf{F}^{(j)} + \sum_{l=0}^{k-1} \binom{k}{l} \mathbf{r}^{[l]} \mathbf{L} + \mathbf{r}^{[k]} (\mathbf{L} + \mathbf{B}) \\ &= \mathbf{0}. \end{aligned} \tag{17}$$

We can express the third term on the left-hand side of the above equation as

$$\begin{aligned} \sum_{h=2}^{\infty} \pi^{(h)} \sum_{j=1}^{\infty} (j+h+1)^k \mathbf{F}^{(j)} &= \sum_{h=2}^{\infty} \pi^{(h)} \sum_{l=0}^k \binom{k}{l} (h+1)^l \sum_{j=1}^{\infty} j^{k-l} \mathbf{F}^{(j)} \\ &= \sum_{h=2}^{\infty} \pi^{(h)} \sum_{l=0}^k \binom{k}{l} (h+1)^l \mathbf{F}^{[0,k-l]} = \sum_{l=0}^k \binom{k}{l} \sum_{h=2}^{\infty} \pi^{(h)} (h+1)^l \mathbf{F}^{[0,k-l]} \\ &= \sum_{l=0}^k \binom{k}{l} \sum_{t=0}^l \binom{l}{t} \sum_{h=2}^{\infty} h^t \pi^{(h)} \mathbf{F}^{[0,k-l]} = \sum_{l=0}^k \binom{k}{l} \sum_{t=0}^l \binom{l}{t} \mathbf{r}^{[t]} \mathbf{F}^{[0,k-l]} \\ &= \sum_{l=0}^{k-1} \binom{k}{l} \sum_{t=0}^l \mathbf{r}^{[t]} \mathbf{F}^{[0,k-l]} + \sum_{t=0}^{k-1} \binom{k}{t} \mathbf{r}^{[t]} \mathbf{F}^{[0,0]} + \mathbf{r}^{[k]} \mathbf{F}^{[0,0]} \end{aligned}$$

and substituting the third term in (17) with the above expression we obtain

$$\begin{aligned} \mathbf{r}^{[k]} (\mathbf{F}^{[0,0]} + \mathbf{L} + \mathbf{B}) &= - \left(\pi^{(0)} \hat{\mathbf{F}}^{[1,k]} + \pi^{(1)} (2^k \mathbf{L}^{(1)} + \mathbf{F}^{[2,k]}) + \sum_{l=0}^{k-1} \binom{k}{l} \mathbf{r}^{[l]} \mathbf{L} \right) \\ &\quad - \left(\sum_{l=0}^{k-1} \binom{k}{l} \sum_{t=0}^l \mathbf{r}^{[t]} \mathbf{F}^{[0,k-l]} + \sum_{t=0}^{k-1} \binom{k}{t} \mathbf{r}^{[t]} \mathbf{F}^{[0,0]} \right) \end{aligned}$$

which is a linear system of the form $\mathbf{r}^{[k]} (\mathbf{F}^{[0,0]} + \mathbf{L} + \mathbf{B}) = \mathbf{b}^{[k]}$ where $\mathbf{b}^{[k]}$ is an expression that can be computed from $\pi^{(0)}$, $\pi^{(1)}$, and the vectors $\mathbf{r}^{[0]}$ to $\mathbf{r}^{[k-1]}$. Since we know that the rank of $\mathbf{F}^{[0,0]} + \mathbf{L} + \mathbf{B}$ (alternatively, of $\sum_{j=1}^{\infty} \mathbf{F}^{(j)} + \mathbf{L} + \mathbf{B}$) is $n - 1$, we remove the equation corresponding to the last column, resulting in $n - 1$ equations

$$\mathbf{r}^{[k]} (\mathbf{F}^{[0,0]} + \mathbf{L})_{1:n,1:n-1} = \mathbf{b}^{[k]}_{1:n-1}. \tag{18}$$

One additional equation is then required. We obtain it from the equations in (12), again by multiplying them by the appropriate j^k and summing them:

$$\begin{aligned} \pi^{(0)} \sum_{j=2}^{\infty} j^k \underbrace{\sum_{h=j}^{\infty} \hat{\mathbf{F}}^{(h)} \mathbf{1}^T}_{=\hat{\mathbf{F}}^{[j]}} + \pi^{(1)} \sum_{j=1}^{\infty} (j+1)^k \underbrace{\sum_{h=j}^{\infty} \mathbf{F}^{(h)} \mathbf{1}^T}_{=\mathbf{F}^{[j]}} + \sum_{h=2}^{\infty} \pi^{(h)} \sum_{j=1}^{\infty} (j+h)^k \underbrace{\sum_{h=j}^{\infty} \mathbf{F}^{(h)} \mathbf{1}^T}_{=\mathbf{F}^{[j]}} \\ = \underbrace{\sum_{j=2}^{\infty} \pi^{(j)} j^k \mathbf{B} \mathbf{1}^T}_{=\mathbf{r}^{[k]} \mathbf{B} \mathbf{1}^T} \end{aligned}$$

which can be rewritten as

$$\pi^{(0)} \sum_{j=2}^{\infty} j^k \hat{\mathbf{F}}^{[j]} \mathbf{1}^T + \pi^{(1)} \sum_{j=1}^{\infty} (j+1)^k \mathbf{F}^{[j]} \mathbf{1}^T + \sum_{h=2}^{\infty} \pi^{(h)} \sum_{j=1}^{\infty} (j+h)^k \mathbf{F}^{[j]} \mathbf{1}^T = \mathbf{r}^{[k]} \mathbf{B} \mathbf{1}^T. \tag{19}$$

The third term of the above summation can then be expressed as

$$\begin{aligned} \sum_{h=2}^{\infty} \boldsymbol{\pi}^{(h)} \sum_{j=1}^{\infty} \sum_{l=0}^k \binom{k}{l} j^k h^{k-l} \mathbf{F}^{[j]} \mathbf{1}^T &= \sum_{h=2}^{\infty} \boldsymbol{\pi}^{(h)} \sum_{t=0}^k \binom{k}{t} h^{k-t} \sum_{j=1}^{\infty} j^k \mathbf{F}^{[j]} \mathbf{1}^T \\ &= \sum_{t=0}^k \binom{k}{t} \sum_{h=2}^{\infty} \boldsymbol{\pi}^{(h)} h^{k-t} \sum_{j=1}^{\infty} j^k \mathbf{F}^{[j]} \mathbf{1}^T = \sum_{t=0}^k \binom{k}{t} \mathbf{r}^{[k-t]} \sum_{j=1}^{\infty} j^k \mathbf{F}^{[j]} \mathbf{1}^T. \end{aligned}$$

Substituting the above in (19) we obtain

$$\begin{aligned} \mathbf{r}^{[k]} \left(\sum_{j=1}^{\infty} j^k \mathbf{F}^{[j]} - \mathbf{B} \right) \mathbf{1}^T \\ = - \left(\boldsymbol{\pi}^{(0)} \sum_{j=2}^{\infty} j^k \hat{\mathbf{F}}^{[j]} \mathbf{1}^T + \boldsymbol{\pi}^{(1)} \sum_{j=1}^{\infty} (j+1)^k \mathbf{F}^{[j]} \mathbf{1}^T + \sum_{t=1}^k \binom{k}{t} \mathbf{r}^{[k-t]} \sum_{j=1}^{\infty} j^k \mathbf{F}^{[j]} \right). \end{aligned} \tag{20}$$

The right-hand side of the above equation is an expression containing $\boldsymbol{\pi}^{(0)}$, $\boldsymbol{\pi}^{(1)}$, and the vectors $\mathbf{r}^{[0]}$ to $\mathbf{r}^{[k-1]}$. The $n \times n$ matrix

$$\left[(\mathbf{F}^{[0,0]} + \mathbf{L})_{1:n,1:n-1} \left| \left(\sum_{j=1}^{\infty} j^k \mathbf{F}^{[j]} - \mathbf{B} \right) \mathbf{1}^T \right. \right]$$

required to compute $\mathbf{r}^{[k]}$ has full rank (the proof is analogous to that of Theorem 1). Thus, the above result can be used to obtain any moment k , provided that the first $k - 1$ moments have been previously computed. As an example, we consider $\mathbf{r}^{[1]}$, which is used to compute measures such as the first moment of the queue length. In this case,

$$\begin{aligned} \mathbf{b}^{[1]} &= - \left(\boldsymbol{\pi}^{(0)} \sum_{j=1}^{\infty} (j+1) \hat{\mathbf{F}}^{(j)} + \boldsymbol{\pi}^{(1)} \left(2\mathbf{L}^{(1)} + \sum_{j=1}^{\infty} (j+2) \mathbf{F}^{(j)} \right) + \boldsymbol{\pi}^{(*)} \left(\mathbf{L} + \sum_{j=1}^{\infty} (j+1) \mathbf{F}^{(j)} \right) \right), \\ c^{[1]} &= - \left(\boldsymbol{\pi}^{(0)} \sum_{j=2}^{\infty} j \hat{\mathbf{F}}^{[j]} + \boldsymbol{\pi}^{(1)} \sum_{j=1}^{\infty} (j+1) \mathbf{F}^{[j]} + \boldsymbol{\pi}^{(*)} \sum_{j=1}^{\infty} j \mathbf{F}^{[j]} \right) \mathbf{1}^T. \end{aligned}$$

We conclude by observing that some measures might be infinite. For example, if the matrices $\mathbf{F}^{(j)}$ are summable but decrease only like $1/j^h$ for some $h > 1$, then the moments of order $h - 1$ or higher for the queue length do not exist (are infinite).

5. Bounded bulk arrivals

We now consider processes with bulk arrivals of maximum size p . These can be solved using the original ETAQA [3] or matrix-geometric [14] methods, by merging every p levels into a single larger

level, but the results of this section show how to compute the solution without the corresponding increase in complexity (a factor of p for both execution time and storage in ETAQA, a factor of p^3 for execution time and p^2 for storage in the matrix-geometric method).

If we restrict the process so that it is allowed to jump forward by at most p levels, its infinitesimal generator \mathbf{Q} still has the structure of (5), except that $\hat{\mathbf{F}}^{(j)}$ and $\mathbf{F}^{(j)}$ are zero for $j > p$:

$$\mathbf{Q} = \begin{bmatrix} \hat{\mathbf{L}}^{(0)} & \hat{\mathbf{F}}^{(1)} & \hat{\mathbf{F}}^{(2)} & \dots & \hat{\mathbf{F}}^{(p)} & \mathbf{0} & \mathbf{0} & \dots \\ \hat{\mathbf{B}} & \mathbf{L}^{(1)} & \mathbf{F}^{(1)} & \dots & \mathbf{F}^{(p-1)} & \mathbf{F}^{(p)} & \mathbf{0} & \dots \\ \mathbf{0} & \mathbf{B} & \mathbf{L} & \dots & \mathbf{F}^{(p-2)} & \mathbf{F}^{(p-1)} & \mathbf{F}^{(p)} & \dots \\ \mathbf{0} & \mathbf{0} & \mathbf{B} & \dots & \mathbf{F}^{(p-3)} & \mathbf{F}^{(p-2)} & \mathbf{F}^{(p-1)} & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \tag{21}$$

We can then formulate the following lemma.

Lemma 1 (Bulk arrivals). *Given an ergodic CTMC with infinitesimal generator \mathbf{Q} having the structure shown in (21), such that the first $n - 1$ columns of \mathbf{B} are null, $\mathbf{B}_{1:n,1:n-1} = \mathbf{0}$, and with stationary probability vector $\boldsymbol{\pi} = [\boldsymbol{\pi}^{(0)}, \boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}, \dots]$, the system of linear equations*

$$\mathbf{x} \begin{bmatrix} \mathbf{1}^T & \hat{\mathbf{L}}^{(0)} & \hat{\mathbf{F}}_{1:m,1:n-1}^{(1)} & \left(\sum_{j=2}^p \hat{\mathbf{F}}^{(j)} \right)_{1:m,1:n-1} & \left(\sum_{j=2}^p (j-1)\hat{\mathbf{F}}^{(j)} \right) \mathbf{1}^T \\ \mathbf{1}^T & \hat{\mathbf{B}} & \mathbf{L}_{1:n,1:n-1}^{(1)} & \left(\sum_{j=1}^p \mathbf{F}^{(j)} \right)_{1:n,1:n-1} & \left(\sum_{j=1}^p j\mathbf{F}^{(j)} \right) \mathbf{1}^T \\ \mathbf{1}^T & \mathbf{0} & \mathbf{0} & \left(\sum_{j=1}^p \mathbf{F}^{(j)} + \mathbf{L} \right)_{1:n,1:n-1} & \left(\sum_{j=1}^p j\mathbf{F}^{(j)} - \mathbf{B} \right) \mathbf{1}^T \end{bmatrix} = [1, \mathbf{0}]$$

admits a unique solution $\mathbf{x} = [\boldsymbol{\pi}^{(0)}, \boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(*)}]$, where $\boldsymbol{\pi}^{(*)} = \sum_{i=2}^{\infty} \boldsymbol{\pi}^{(i)}$.

Proof. The steps of the proof are exactly the same as those of the theorem introduced in Section 3, hence they are omitted. □

The measure of interest r can also be derived as done in Section 4. Using the same definitions for r , ρ , and $\mathbf{r}^{[l]}$, we need to show how to express the vectors $\mathbf{r}^{[l]}$ for $l = 0, \dots, k$. Here, $\mathbf{r}^{[k]}$ can be computed recursively by solving the equation:

$$\mathbf{r}^{[k]} \left[\left(\sum_{j=1}^p \mathbf{F}^{(j)} + \mathbf{L} \right)_{1:n,1:n-1} \middle| \left(\sum_{j=1}^p j^k \mathbf{F}^{[j]} - \mathbf{B} \right) \mathbf{1}^T \right] = [\mathbf{b}^{[k]} | c^{[k]}],$$

where $\mathbf{b}^{[k]}$ and $c^{[k]}$ are computed using $\boldsymbol{\pi}^{(0)}$, $\boldsymbol{\pi}^{(1)}$, and $\mathbf{r}^{[0]}$ to $\mathbf{r}^{[k-1]}$, following the exact same steps as for the case of unbounded bulks described in Section 4, but adjusted so as to consider bounded bulks only.

6. Theoretical complexity

In this section, we present a detailed complexity comparison of ETAQA-MG1 with the matrix-analytic method outlined in Section 2 (in Section 8, we present actual run-times on an application). First, we note that, under our condition that \mathbf{B} has a single nonzero column, \mathbf{G} does not need to be computed [18], since it is a zero matrix except for the last column, which contains all ones. This observation is further confirmed by the probabilistic interpretation of \mathbf{G} : an entry (k, l) in \mathbf{G} expresses the conditional probability of the process first entering $\mathcal{S}^{(j-1)}$ through state l , given that it starts from state k of $\mathcal{S}^{(j)}$ [15]². Thus, for the special case of \mathbf{Q} that we consider in this paper, there is no cost associated for computing and storing \mathbf{G} .

To obtain $\boldsymbol{\pi}^{(0)}$, the matrix-analytic algorithm requires the computation of the inverse of $\mathbf{S}^{(0)}$. This inverse is in general a full matrix that must be stored for the duration of the entire computation, increasing the storage complexity for the matrix-analytic approach. During the construction of the system of linear equations of Eq. (4), we also need to multiply this full inverse by sparse $m \times n$ and $n \times m$ matrices, for an overall complexity of $O(n^3 + n\eta\{\hat{\mathbf{B}}, \sum_{j=1}^{\infty} \hat{\mathbf{F}}^{(j)}\})$, where $\eta\{\cdot\}$ denotes the total number of nonzero entries in the argument matrices. Finally, the resulting linear system is described by an $m \times m$ matrix, which is also full in general. Instead, for ETAQA-MG1, we only require the solution of a sparse system of $m + 2n$ linear equations, plus some sparse matrix summations and scalar multiplications.

Although there are closed form formulas for the first and the second moment of the queue length, based on the matrix-analytic approach [11], the entire stationary probability vector is in principle required to compute the measures of interest. Using Ramaswami's recursive formula, the vectors $\boldsymbol{\pi}^{(j)}$ can be computed only up to an index s such that $\sum_{j=0}^s \boldsymbol{\pi}^{(j)} \mathbf{1}^T$ is very close to one.

We also note that, in practice, use of either ETAQA-MG1 or the matrix-analytic method requires storing only a finite number of matrices $\hat{\mathbf{F}}^{(j)}$ and $\mathbf{F}^{(j)}$, enough to accurately describe the M/G/1-type process. In our study, we assume that matrices $\hat{\mathbf{F}}^{(j)}$ and $\mathbf{F}^{(j)}$ for $j > p$ are zero, as discussed in Section 5. However, ETAQA's superiority is more obvious in this case. The complexity of the matrix-analytic algorithm has now an additional cost due to the fact that the matrices $\hat{\mathbf{S}}^{(j)}$ and $\mathbf{S}^{(j)}$ do not have a closed-form expression, therefore they must be computed for all $j \geq 1$, up to the forward matrix truncation point p . Fortunately, computing these matrices involves only full vector operations, and their storage requires only a full vector each (plus the storage for the forward matrices themselves), due to the structure of \mathbf{G} and to the fact that, in our case, $\mathbf{G}\mathbf{G} = \mathbf{G}$. Since $\hat{\mathbf{S}}^{(j)}$ and $\mathbf{S}^{(j)}$ are zero for $j > p$, only p full vectors are additionally required to store these matrices, and, using Eq. (2), we only need to keep a sliding window of the $p + 1$ most recently computed vectors $\boldsymbol{\pi}^{(j)}$. ETAQA-MG1 requires instead only sparse matrix additions and the solution of a sparse system of $m + 2n$ linearly independent equations.

Table 1 summarizes the computational and storage complexity of ETAQA-MG1 and the matrix-analytic method. Since that ETAQA-MG1 is an aggregation-based technique that truly exploits the sparsity of the matrices describing the original infinitesimal generator, its storage and time complexity are superior to the matrix-analytic method. Note that the contribution of \mathbf{B} to the time complexity can be ignored, since it contains at most n nonzero entries.

² The probabilistic interpretation of \mathbf{G} is the same for both DTMCs and CTMCs. The interpretation in [15, p. 81] is consistent with the discussion in [9, p. 142], where CTMCs are taken into consideration.

Table 1
Computational and storage complexity of ETAQA-MG1 and the matrix-analytic method^a

	Time complexity	Additional storage
Computation of $\pi^{(0)}$ (matrix-analytic) or $\pi^{(0)}$, $\pi^{(1)}$, and $\pi^{(*)}$ (ETAQA)		
ETAQA-MG1	$O^L(\eta\{\hat{\mathbf{B}}, \hat{\mathbf{L}}^{(0)}, \mathbf{L}^{(1)}, \mathbf{L}, \sum_{j=1}^p \hat{\mathbf{F}}^{(j)}, \sum_{j=1}^p \mathbf{F}^{(j)}\})$	None
Matrix-analytic	$O^L(m^2) + O(n\eta\{\sum_{j=1}^p \hat{\mathbf{F}}^{(j)}, \hat{\mathbf{B}}\} + \eta\{\sum_{j=1}^p \mathbf{F}^{(j)}\} + n^3)$	$O(m^2 + n^2)$
First moment measures		
ETAQA-MG1	$O^L(\eta\{\mathbf{L}, \sum_{j=1}^p \mathbf{F}^{(j)}\})$	None
Matrix-analytic	$O(s(p\eta\{\sum_{j=1}^p \mathbf{F}^{(j)}\} + n^2) + p\eta\{\sum_{j=1}^p \hat{\mathbf{F}}^{(j)}\})$	$O(pn + n^2)$

^a $O^L(x)$ denotes the time complexity of solving a linear system described by x nonzero entries.

7. Repartitioning Q

The approach we introduced in Sections 3 and 5 is very efficient but its applicability is limited by the requirement that \mathbf{B} must have only one nonzero column. There are models where this condition on \mathbf{B} is not immediately satisfied, yet it can be achieved through an appropriate *repartitioning* of the states. For illustration simplicity, Fig. 3 shows an example of state repartitioning applied to a QBD process, but exactly the same idea would apply to that model structure if it had a more complex forward arc structure, i.e., if it were an M/G/1-type process.

Fig. 3 (top) shows a QBD process where the transitions from $\mathcal{S}^{(j)}$ to $\mathcal{S}^{(j-1)}$ have two destinations, $s_b^{(j-1)}$ and $s_c^{(j-1)}$. Clearly, the single column condition for \mathbf{B} is not immediately satisfied. Fig. 3 (bottom) shows

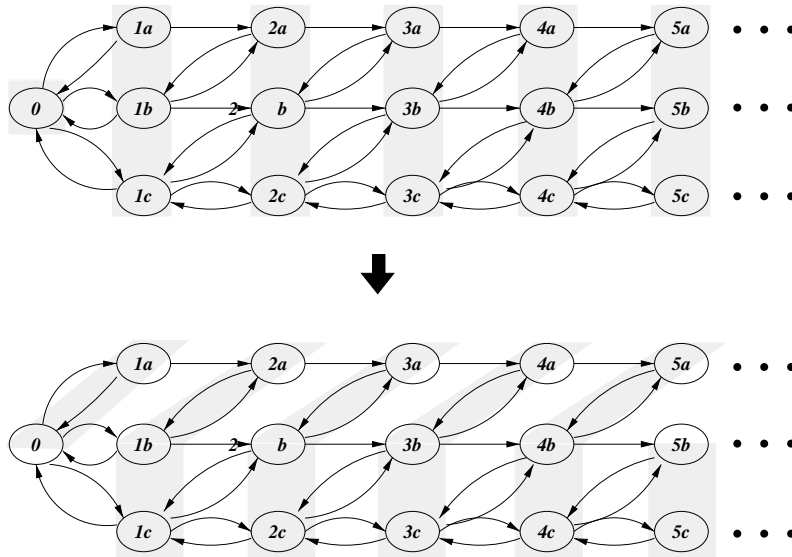


Fig. 3. State repartitioning.

instead that, by redefining the sets $\mathcal{S}^{(j)}$ in such a way that $\mathcal{S}^{(1)} = \{s_a^{(2)}, s_b^{(1)}, s_c^{(1)}\}$, $\mathcal{S}^{(2)} = \{s_a^{(3)}, s_b^{(2)}, s_c^{(2)}\}$, etc., the transitions from $\mathcal{S}^{(j)}$ to $\mathcal{S}^{(j-1)}$ now go to a single state, $s_c^{(j-1)}$. The condition on \mathbf{B} is then satisfied.

In Appendix A, we detail an algorithm that accepts as input the (finite) block description of the (infinite) infinitesimal generator \mathbf{Q} , as given in Eq. (5), of an M/G/1-type process and then determines a repartitioning of the state space so that \mathbf{B} has a single nonzero column in the new partition, if such a partition exists. The computational complexity of this algorithm is $O(n^2 e \log e)$ and its space complexity is $O(n + e)$, where n is the size of the repetitive sets and e is at most the total number of nonzero entries in $\tilde{\mathbf{Q}} = \mathbf{B} + \mathbf{L} + \sum_{j=1}^{2n} \mathbf{F}^{(j)}$ in the original partition.

We observe that, as a result of this repartitioning, at most $n(n - 1)/2$ states might be shifted from the repetitive portion of the chain into $\mathcal{S}^{(0)}$. This is because, among the n sets of states $\mathcal{T}_i = \{s_i^{(j)} : j \geq 2\}$, $k_0 \geq 1$ of them will remain unchanged; $k_1 \leq n - k_0$ of them will be shifted one position to the left (i.e., state $s_i^{(j)}$ becomes state $s_i^{(j-1)}$), adding k_1 states to $\mathcal{S}^{(0)}$; $k_2 \leq n - k_0 - k_1$ of them will be shifted two positions to the left, adding $2k_2$ states to $\mathcal{S}^{(0)}$, and so on. In the worst case, this adds $1 + 2 + \dots + n - 1$ states to $\mathcal{S}^{(0)}$.

Furthermore, this change in the definition of the sets $\mathcal{S}^{(j)}$ does not affect our ability to compute the measures of interest. Assume that we have defined the reward rate $\rho_i^{(j)} = \mathbf{a}_i^{[0]} + \mathbf{a}_i^{[1]}j + \dots + \mathbf{a}_i^{[k]}j^k$ for state $s_i^{(j)}$ in the original partition of the state space. After repartitioning, the “old” state $s_i^{(j)}$ may now be state $\bar{s}_i^{(j-c)}$, for some fixed c (we use a “-” for quantities referring to the new partition). The same expected reward rate will still be computed, that is,

$$r = \sum_{j=0}^{\infty} \sum_{i \in \mathcal{S}^{(j)}} \rho_i^{(j)} \pi_i^{(j)} = \sum_{j=0}^{\infty} \sum_{i \in \bar{\mathcal{S}}^{(j)}} \bar{\rho}_i^{(j)} \bar{\pi}_i^{(j)},$$

provided the new reward rate for state $\bar{s}_i^{(j-c)}$ satisfies $\bar{\rho}_i^{(j-c)} = \rho_i^{(j)}$. In other words, we simply require that $\bar{\mathbf{a}}_i^{[0]} + \bar{\mathbf{a}}_i^{[1]}(j - c) + \dots + \bar{\mathbf{a}}_i^{[k]}(j - c)^k = \mathbf{a}_i^{[0]} + \mathbf{a}_i^{[1]}j + \dots + \mathbf{a}_i^{[k]}j^k$, and this can be obviously achieved through an appropriate definition of the coefficients $\bar{\mathbf{a}}_i^{[0]}, \dots, \bar{\mathbf{a}}_i^{[k]}$.

8. Application: multiprocessor scheduling

In this section, we describe an application that can be studied through the ETAQA-MG1 approach. We concentrate on presenting the CTMCs that models the application of interest, focusing on the form of the repeating matrix pattern, and we present numerical results to compare ETAQA-MG1 with the matrix-analytic approach.

Modeling the behavior of scheduling policies in parallel systems often results in CTMCs with matrix-geometric form [22]. Here, we present a CTMC that models the behavior of a class of dynamic scheduling policies where the reconfiguration cost is reduced by limiting how and when processor reallocations can occur [1]. Fig. 4 illustrates the CTMC modeling a dynamic scheduling policy that can only reduce the number of processors allocated to an application, and only immediately after a service completion. Fig. 4 shows the system behavior under bursty arrival conditions modeled as bulks of finite size (see [20] for an analysis of adaptive space-sharing policies under similar assumptions for the arrival process).

The system state is described by the number of applications that are waiting for service in the queue and the number of applications that are in service. $NwMs_r$ denotes that there are N applications in the queue,

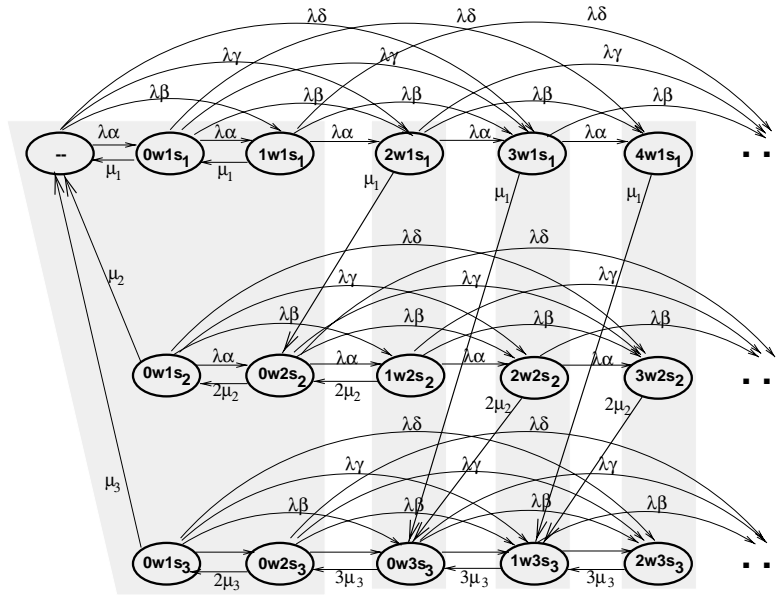


Fig. 4. The CTMC that models a dynamic multiprocessor scheduling policy.

waiting to be executed, and M applications in service, each having a fraction $1/r$ of the total number of processors allocated to it. Thus, the service rate of an application is μ_r , a function of r . The arrival rate for bulks is λ , and the probability that the size of the bulk is 1, 2, 3, or 4 is α , β , γ , and δ , respectively.

Our lemma for bounded bulk arrivals can be readily applied to study the performance of this policy. Since the behavior of this dynamic scheduling policy under different workloads is outside the scope of this paper, we do not present numerical results that illustrate the performance of the policy per se. Instead, we present the time for the solution needed by ETAQA-MG1 and the most efficient matrix-analytic method (as presented in Section 2) on a 1 GHz Pentium IV with 512 KB cache and 1 GB memory. For simplicity, Fig. 4 shows bulk arrivals of maximum size four and the multiprocessor partitioned in up to three classes, but our experiments consider maximum bulks of sizes 50 or 100, and several values for the maximum number of classes. Fig. 5(a) and (b) report the significant run-time improvements achieved when computing $\pi^{(0)}$,

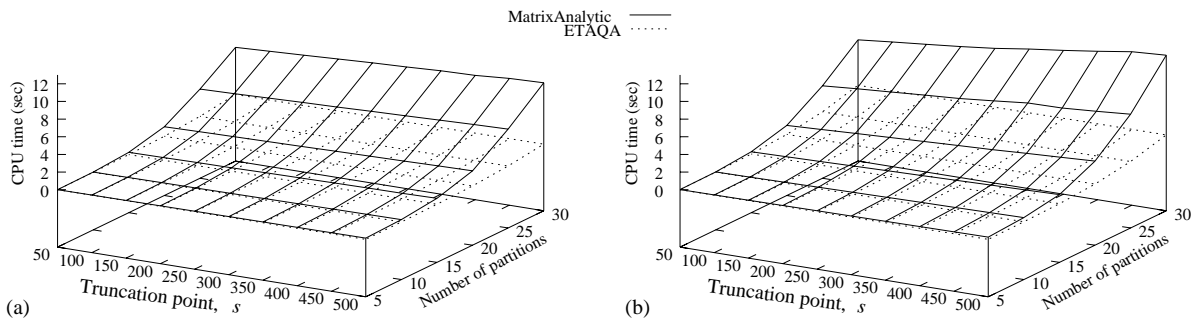


Fig. 5. Computation time (in s) for a process with bulk sizes of (a) 50 and (b) 100.

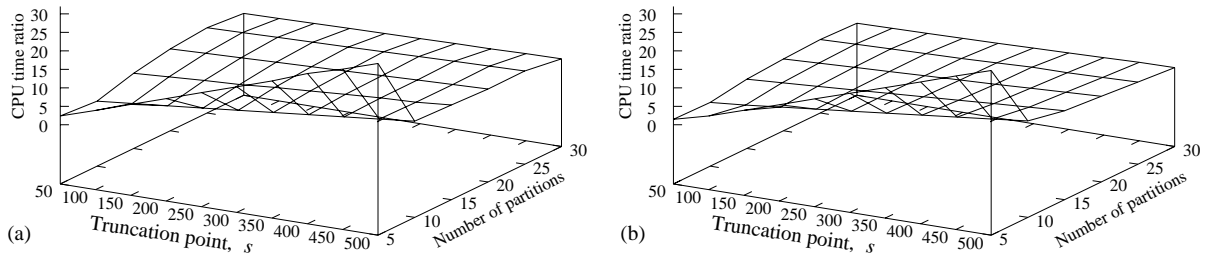


Fig. 6. Ratio of the matrix-analytic and ETAQA-MG1 computation times for a process with bulk sizes of 50 (a) or 100 (b).

$\pi^{(1)}$, and $\pi^{(*)}$ using ETAQA-MG1, as opposed to computing $\pi^{(0)}$, and $\pi^{(j)}$ for $1 \leq j \leq s$, using the matrix-analytic method, as a function of the number of multiprocessor partitions and the truncation point s in the computation of the probability vector, i.e., $\sum_{i=0}^s \pi^{(i)} \mathbf{1}^T$ is approximately equal to one.

To better observe the efficiency of ETAQA-MG1, Fig. 6(a) and (b) plot the relative computation time of the matrix-analytic method versus ETAQA-MG1 for the cases presented in Fig. 5(a) and (b), respectively. ETAQA-MG1 outperforms the matrix-analytic method in both small and large size problems. The efficiency of ETAQA-MG1 increases when the length of the computed probability vector for matrix-analytic increases and when the cardinality m of $\mathcal{S}^{(0)}$ is in the same range as the cardinality n of $\mathcal{S}^{(j)}$ for $j > 0$. In the case under study, $m = O(n^2)$, which explains higher computational gains for ETAQA-MG1 for partitions into small number of classes.

For the experiments of Figs. 5 and 6, we used an absolute test as a convergence criterion for the iterative solution of the systems of linear equations required by ETAQA-MG1 and the matrix-analytic method: we stop when the maximum entry in the vector defined as the difference of two successive iterates does not exceed 10^{-16} . As a confirmation of the appropriateness of this test, we observe that the entries corresponding to $\mathcal{S}^{(0)}$ and $\mathcal{S}^{(1)}$ for the probability vectors computed with the two methods match with an accuracy of 10^{-16} as well. We elaborate further on the numerical stability of our technique in the next section.

9. Numerical stability of ETAQA-MG1

ETAQA-MG1 is a numerical approach to solve M/G/1-type processes, and it entails matrix multiplications, additions, and subtractions, as well as the solution of systems of linear equations. This raises the issue of numerical stability both in the construction phase of the blocks forming matrix \mathbf{M} of Theorem 1 (or the analogous one of Lemma 1) and in the solution phase of the resulting linear systems.

We investigate the numerical stability of ETAQA-MG1 experimentally, adopting the definition of a “numerically stable algorithm” given in [23]: *an algorithm is stable if, for an almost exact input, it produces an almost exact output*. As a testbed, we choose a queueing system with bulk arrivals and Coxian service, and solve it with both ETAQA-MG1 and matrix-analytic/Ramaswami’s formula, which is known to be numerically stable.

For the case under study, the input consists of the parameters for the arrival process and for the Coxian service process, while the output consists of $\pi^{(0)}$, $\pi^{(1)}$, and $\pi^{(*)}$ (for the matrix-analytic method, we compute $\pi^{(*)}$ as $\sum_{j=1}^s \pi^{(j)}$). First, we solve the queueing system with both methods for a chosen set of input parameters, which we refer to as the unperturbed input. Then, we randomly perturb the values of the

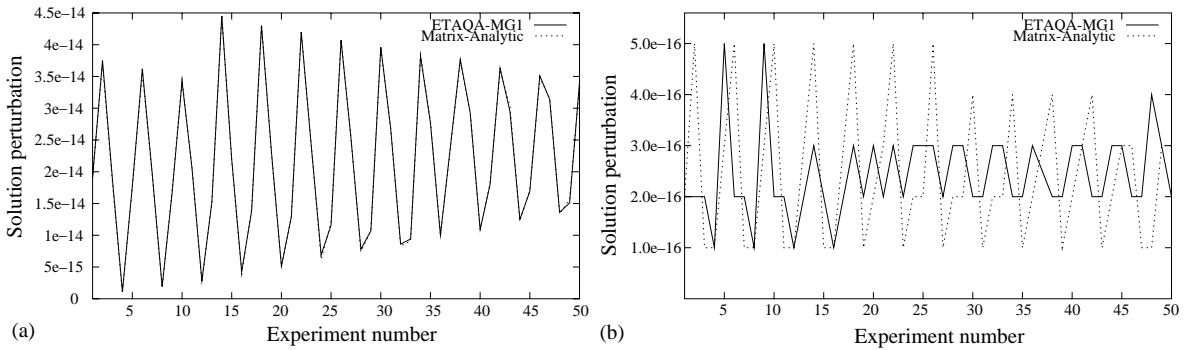


Fig. 7. Perturbation of the solution as result of the perturbation of the input for both ETAQA-MG1 and matrix-analytic method for 50 different experiments, with matrix size $n = 32$ and bulk size of $p = 64$, for perturbation magnitudes of (a) 10^{-12} and (b) 10^{-14} .

input parameters with uniformly distributed quantities in $(-\epsilon, \epsilon)$, where ϵ is a small constant, either 10^{-14} or 10^{-12} . We solve again each model with each method and compare the output with that obtained from solving the problem with unperturbed input using the same method. To increase the statistical confidence, we conduct 50 (differently) perturbed experiments.

The results for perturbations of magnitude 10^{-12} and 10^{-14} , in Fig. 7(a) and (b), respectively, suggest that ETAQA-MG1 is a stable algorithm. In Fig. 8(a) and (b) we present the maximum difference between the solutions obtained from both methods for one of the perturbed experiments for different problem sizes, i.e., different sizes on the matrices of the process and different number of matrices involved.

The results presented in Figs. 7 and 8 correspond to cases where the entries within each of the input matrices are in the same range. To study the effect of perturbations on the more problematic “stiff” chains, we consider the same queueing system where entries differ up to six orders of magnitude. The results are in Figs. 9 and 10.

For stiff inputs, the resulting perturbation in the solution is higher than for non-stiff inputs. Again, however, ETAQA-MG1 performs similarly to the stable matrix-analytic approach. Thus, while we do not provide a theoretical proof of the numerical stability of ETAQA-MG1, the results of the numerical experiments of this section provide evidence that ETAQA-MG1 is a stable approach.

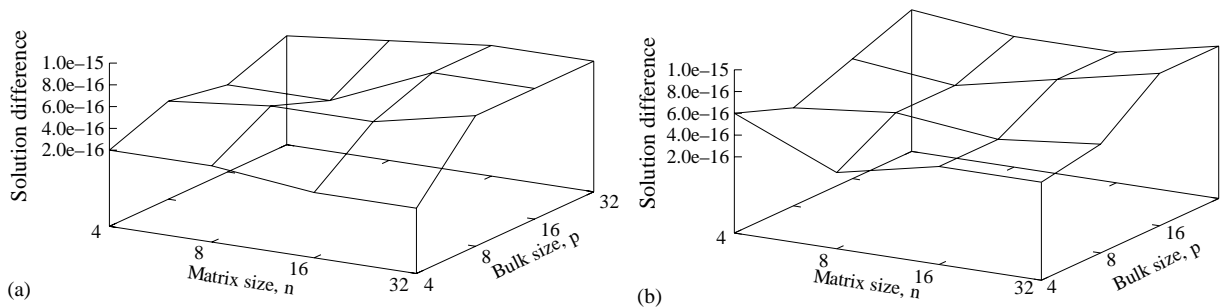


Fig. 8. Maximum difference in the solutions obtained from ETAQA-MG1 and the matrix-analytic method for one perturbation experiment as a function of different problem sizes, with perturbation magnitudes of (a) 10^{-12} and (b) 10^{-14} .

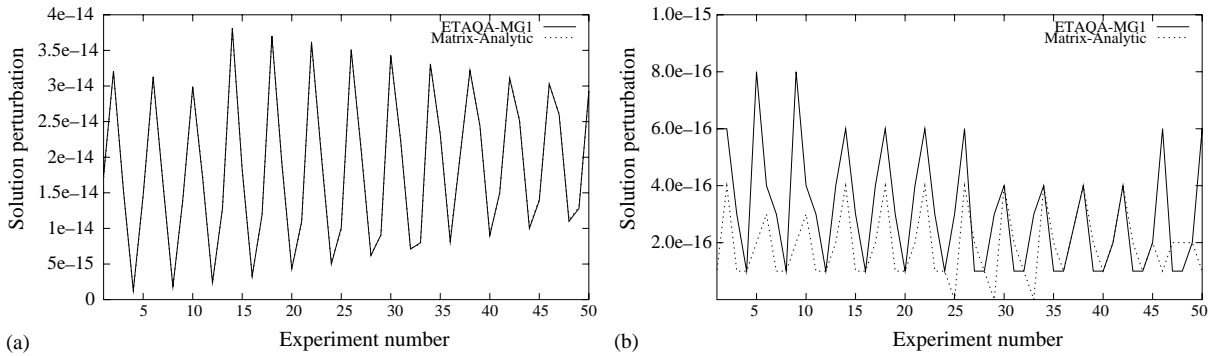


Fig. 9. Perturbation of the solution as result of the perturbation of the stiff input for both ETAQA-MG1 and the matrix-analytic method for 50 different experiments, with matrix size $n = 32$ and bulk size of $p = 64$, with perturbation magnitudes of (a) 10^{-12} and (b) 10^{-14} .

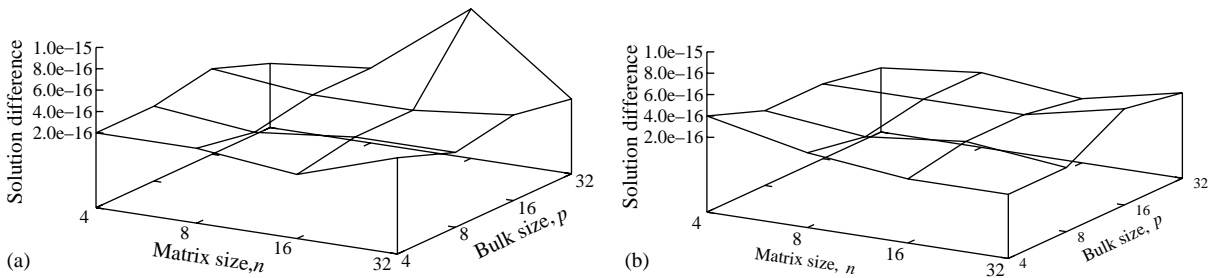


Fig. 10. Maximum difference in the solutions obtained from ETAQA-MG1 and the matrix-analytic method for one perturbation experiment with stiff input as a function of different problem sizes, with perturbation magnitudes of (a) 10^{-12} and (b) 10^{-14} .

10. Conclusions and future work

In this paper we presented an extension of the ETAQA approach to M/G/1-type processes. Our exposition focuses on the description of the extended ETAQA methodology and its application to efficiently compute the *exact* probabilities of the boundary states of the process and the aggregate probability distribution of the states in each of the equivalence classes corresponding to a specific partitioning of the remaining infinite portion of the state space. Although the method does not compute the probability distribution of all states, it still provides enough information for the “mathematically exact” computation of a wide variety of Markov reward functions.

We must emphasize that our treatment applies only to M/G/1-type processes where all transitions from one level to the previous one are directed to a single state (transitions within a level or toward higher levels are completely unrestricted). When this condition is satisfied, the solution is derived by solving a system of $m + 2n$ linear equations and provides significant savings over standard algorithms for the solution of general M/G/1-type processes. While the condition on \mathbf{B} is indisputably restrictive and might not hold on a given CTMC as stated, we also provide an algorithm that finds a repartition of the state space so that the condition becomes satisfied, if such a repartition exists.

We have recently release a software tool called MAMSolver (available at <http://www.cs.wm.edu/MAMSolver/>) that implements ETAQA-MG1 as well as the classic matrix analytic methodologies for the solution of MG1-type processes [19].

Acknowledgements

We thank the anonymous referees, whose comments greatly helped us to improve this manuscript. We also thank Prof. Andreas Stathopoulos for giving us advice on addressing the numerical stability of our method.

Appendix A. State repartitioning

A.1. Problem definition

We seek an algorithm that, given an infinitesimal generator \mathbf{Q} block-partitioned as in Eq. (5), corresponding to a state space initially partitioned as $\mathcal{S}^{(0)} \cup \mathcal{S}^{(1)} \cup \mathcal{S}^{(2)} \cup \dots$, with $|\mathcal{S}^{(0)}| = m$ and $|\mathcal{S}^{(j)}| = n$ for $j \geq 1$, discovers a new partition $\bar{\mathcal{S}}^{(0)} \cup \bar{\mathcal{S}}^{(1)} \cup \bar{\mathcal{S}}^{(2)} \cup \dots$, with $|\bar{\mathcal{S}}^{(0)}| = \bar{m}$ and $|\bar{\mathcal{S}}^{(j)}| = n$ for $j \geq 1$, i.e., possibly different-size boundary class but same-size repetitive classes, such that the new corresponding block-partition of Eq. (5) is such that \mathbf{B} contains a single nonzero column, if such a repartition exists, or returns a “no” otherwise.

If we just focus on the zero-nonzero pattern of the infinitesimal generator \mathbf{Q} , the result is a directed graph. The repeating pattern of \mathbf{Q} implies that there is an edge from $s_i^{(j)}$ to $s_{i'}^{(j)}$ iff there is an edge from $s_i^{(j')}$ to $s_{i'}^{(j')}$ for $i, i' = 1, \dots, n$ and $j, j' = 1, 2, \dots$ (since $\mathbf{L}^{(1)}$ can differ from \mathbf{L} , exceptions can occur when $j = 1$ or $j' = 1$, but this is inessential). We call $s_i^{(1)}, s_i^{(2)}, \dots$ for $i = 1, \dots, n$, *same-position* nodes. We call the index a corresponding to the nonzero column in \mathbf{B} the *anchor*. Furthermore, since \mathbf{Q} has an M/G/1-type structure, we know that there are no arcs from $s_i^{(j)}$ to $s_{i'}^{(j')}$ for $j' < j - 1$.

A.2. An algebraic approach

Since the subgraphs induced by $\mathcal{S}^{(1)}, \mathcal{S}^{(2)}, \dots$ in the original partition are all isomorphic and those defined by the new sets $\bar{\mathcal{S}}^{(1)}, \bar{\mathcal{S}}^{(2)}, \dots$ after the repartition are required to be isomorphic as well, and of the same size, any change we make to the definition of $\mathcal{S}^{(j)}$ should be made also to the other sets $\mathcal{S}^{(j')}$ for $j, j' \geq 1$. For example, if we shift one node from $\mathcal{S}^{(2)}$ to $\mathcal{S}^{(1)}$, the same-position node in $\mathcal{S}^{(j+1)}$ must also be shifted to $\mathcal{S}^{(j)}$ for $j \geq 2$. Applying the same moving action (shift left, shift right, or leave alone) to same-position nodes guarantees the isomorphism of the resulting subgraphs and that of the forward and backward edges between different subgraphs. As a result, we only need determine where to move the nodes of one of the repetitive subgraphs, say $\mathcal{S}^{(j)}$, to form the new partition. The other subgraphs follow the same pattern.

Consider this subgraph, which contains nodes $s_1^{(j)}, s_2^{(j)}, \dots, s_n^{(j)}$. After the repartition some nodes will be shifted left to subgraphs with indices smaller than j , some will go right to subgraphs with indices larger than j , and some will stay in their current home. Let $x_i^{(j)}$ be the index of the subgraph to which

node $s_i^{(j)}$ will be moved by the repartition. For the time being, suppose that we know the position a of the anchor. We then apply the following rules to establish a linear system of integer inequalities:

- For each edge with a starting point in $\mathcal{S}^{(j)}$, say $s_k^{(j)}$, to $s_l^{(j+r)}$, l other than the anchor a , we must ensure that this edge will not be a backward edge in the new partition, thus

$$x_k^{(j)} \leq x_l^{(j+r)} = x_l^{(j)} + r.$$

- For each edge with a starting point in $\mathcal{S}^{(j)}$, say $s_k^{(j)}$, to an anchor node $s_a^{(j+r)}$, we must ensure that it will not be a backward edge going back more than one level, thus

$$x_k^{(j)} \leq x_a^{(j+r)} + 1 = x_a^{(j)} + r + 1.$$

(in the above, $r = -1$, $r = 0$, or $r > 0$ if the edge is originally a backward, local, or forward edge, respectively).

We then obtain a system of inequalities in n unknowns, x_1, \dots, x_n , with each inequality of the form

$$x_k \leq x_l + c,$$

where c is an integer no smaller than -1 , and we have dropped the superscripts “ (j) ”, since no confusion can arise. The linear system can be simplified by removing redundant inequalities. That is, for any two variables x_k and x_l , if there are several inequalities in the form of $x_k \leq x_l + c$, only the one with the smallest c needs to be kept in the linear system. As a result of this simplification, the number of inequalities in the linear system, denoted as e , is at most the number of off-diagonal nonzero entries in matrix $\mathbf{B} + \mathbf{L} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)}$ (instead of the number of off-diagonal nonzero entries in \mathbf{B} , \mathbf{L} , and $\mathbf{F}^{(j)}$ for $j = 1, 2, \dots$). Later in Section A.4, we shall reduce this upper bound to e to an even smaller quantity. Thus, \mathbf{Q} has a new partition with a prespecified anchor a satisfying the condition on \mathbf{B} iff the linear system of e inequalities with n variables has an integer solution.

A.3. A necessary and sufficient condition

For the linear system, denoted by L_a , developed in the previous section under the assumption that a is the anchor, we can construct a graph, denoted by G_a , where there is a node for each variable and an edge from node x_l to node x_k with weight c iff $x_k \leq x_l + c$ is one of the inequalities in the system.

The sufficient and necessary condition for the linear system L_a to have an integer solution is that there are no cycles with negative total weight in graph G_a . This condition was observed by Pratt [16] and later further discussed by Shostak [21], hence the detailed proof is omitted here.

According to van Leeuwen [24], there are efficient algorithms to determine whether there is a negative-weight cycle in a directed graph. One of these algorithms [7,10] achieves $O(ne)$ worst-case time complexity, if G_a has n nodes and e edges. Thus we conclude that determining whether the linear system L_a has an integer solution takes $O(ne)$ time. Next, we need to show how to construct a solution if there is one.

A.4. Solving the linear system

Hochbaum and Naor [6] presented an algorithm that finds a feasible integer solution to a system of monotone inequalities of the form $ax_k \leq bx_l + c$, where a and b are positive. Obviously, our inequalities,

$x_k \leq x_l + c$, are monotone. Thus, Hochbaum and Naor’s algorithm may be applied. However, the algorithm requires that each variable x_i has a lower bound x_i^{\min} and an upper bound x_i^{\max} , and its resulting complexity is $O(n^2 e \log e + e \sum_{i=1}^n (x_i^{\max} - x_i^{\min} + 1))$. To use this algorithm, we need to show that bounded variables can be assumed without changing the integer feasibility of our linear system. Specifically, we need to show that the original linear system with unbounded variables has an integer solution if and only if the linear system with bounded variables has an integer solution. The proof of the “if” part is trivial. We next focus on the proof of the “only if” part.

Consider any system of inequalities of the form $x_k \leq x_l + c$ with at least one integer solution. Without loss of generality, assume that, among the n variables, there is one variable taking the value of j in the solution, and all other variables have the values relative to j . For example, if $x_1 = j$, $x_2 = j + 1$, $x_3 = j$, and $x_4 = j - 1$ is a solution to a system with four variables, then for nodes in $\mathcal{S}^{(j)}$, $s_1^{(j)}$ and $s_3^{(j)}$ stay in $\mathcal{S}^{(j)}$, $s_2^{(j)}$ is moved to $\mathcal{S}^{(j+1)}$ by the repartition, and $s_4^{(j)}$ is moved to $\mathcal{S}^{(j-1)}$. We next argue that, if there is an integer solution (with reference parameter j) then there is an integer solution with $j - n < x_i < j + n$ for $i = 1, \dots, n$ (in other words, we never need to shift a state by n or more positions). To do so, we construct a new integer solution with the bounded variables based on the solution given.

Let \mathcal{V}_i be the set of all variables given the value of i in the initial solution. Clearly, $\mathcal{V}_j \neq \emptyset$ according to our assumption. Let a be the smallest index such that \mathcal{V}_a is non-empty and b be the largest index such that \mathcal{V}_b is non-empty. If $a \leq j - n$, there must be at least one empty set (“hole”) with index between a and j . Let r be the smallest such index. For variables in $\mathcal{V}_a \cup \dots \cup \mathcal{V}_{r-1}$, we increase their values by one. We argue that the resulting assignment is still a feasible solution. For an inequality $x_k \leq x_l + c$ with both variables in or not in $\mathcal{V}_a \cup \dots \cup \mathcal{V}_{r-1}$, the equality still holds under the new assignment, since the difference of the values of x_k and x_l is unchanged. For an inequality $x_k \leq x_l + c$ with one variable in $\mathcal{V}_a \cup \dots \cup \mathcal{V}_{r-1}$ and the other in \mathcal{V}_s for $s > r$, there are two cases. The easy case is $x_k \in \mathcal{V}_s$ and $x_l \in \mathcal{V}_a \cup \dots \cup \mathcal{V}_{r-1}$. Under the new assignment, the value of x_k is unchanged while the value of x_l is increased by one. The inequality still holds since the increased variable is on the right-hand side of \leq . The second case is $x_k \in \mathcal{V}_a \cup \dots \cup \mathcal{V}_{r-1}$ and $x_l \in \mathcal{V}_s$. Since there is at least one “hole” between \mathcal{V}_{r-1} and \mathcal{V}_s , we must have $r - 1 \leq s - 2$, thus $x_k \leq x_l - 2$ before the increase of x_k . Increasing the value of x_k by one, we get $x_k \leq x_l - 1 \leq x_l + c$. So the inequality remains true under the new assignment.

We continue increasing values of some variables until the smallest value in the solution is larger than $j - n$. To the variables holding the values greater than or equal to $j + n$ in the solution, we can decrease their values by one at a time using the similar method until all variables have values smaller than $j + n$ in the solution.

Using the procedure described above, we can bound the variables within the range between $j - n$ and $j + n$. This does not change the integer feasibility of our linear system. For any variable, the difference of the upper bound and lower bound is smaller than $2n$. Using Hochbaum and Naor’s algorithm, we achieve the time complexity of $O(n^2 e \log e + e \sum_{k=1}^n (2n + 1)) = O(n^2 e \log e + en(2n + 1)) = O(n^2 e \log e)$. Recall that e , the number of inequalities in the linear system, is at most the number of off-diagonal nonzero entries in matrix $\mathbf{B} + \mathbf{L} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)}$. Here we argue that e is actually at most the number of off-diagonal nonzero entries in matrix $\mathbf{B} + \mathbf{L} + \sum_{j=1}^{2n} \mathbf{F}^{(j)}$. This is because that when we impose the lower bound $j - n$ and upper bound $j + n$ to any variable x_i in the linear system, the difference between any two variables is made less than $2n$. Therefore, when we are simplifying the linear system after it is constructed using the method in Section A.2, those inequalities $x_k \leq x_l + c$ with $c \geq 2n$ can also be removed from the linear system.

A.5. Our algorithm and its complexity

We now outline an algorithm for the repartitioning problem. Given an initial partition $\mathcal{S}^{(0)}, \mathcal{S}^{(1)}, \dots$, for each possible anchor s_i for $i = 1, 2, \dots, n$, we construct a linear system using the method described in Section A.2 and its corresponding graph G_i . Then we determine whether G_i contains any negative-weight cycle. If not, Hochbaum and Naor's algorithm is called to find an integer solution, which is then converted to a new partition of \mathbf{Q} . Otherwise, we check the next node to see if it can be the anchor. If no node can be the anchor that can result in a repartition, the algorithm returns a negative answer.

Input: Graph \mathbf{Q} with partition $\mathcal{S}^{(0)}, \mathcal{S}^{(1)}, \dots$

Output: A new partition $\bar{\mathcal{S}}^{(0)}, \bar{\mathcal{S}}^{(1)}, \dots$, if there is one, or “No” otherwise

For i from 1 to n

 Construct the linear system with bounded variables, L_i , using s_i as the anchor

 Construct the corresponding graph, G_i

 If G_i does not have any negative-weight cycle

 Use Hochbaum and Naor's algorithm to construct an integer solution

 Return the corresponding partition

 End if

End for

Return “No”

In any case Hochbaum and Naor's algorithm is called by our algorithm at most once, while the negative-weight cycle detection algorithm may be called n times in the worst case. Therefore, the time complexity of our repartitioning algorithm is $O(n^2e + n^2e \log e) = O(n^2e \log e)$. Since the data structures involved are graphs with n nodes and e edges and a linear system with n unknowns and e inequalities, the space complexity of our repartitioning algorithm is $O(n + e)$.

References

- [1] S.-H. Chiang, R.K. Mansharamani, M.K. Vernon, Use of application characteristics and limited preemption for run-to-completion parallel processor scheduling policies, in: Proceedings of the 1994 Sigmetrics Conference on Measurement and Modeling of Computer Systems, May 1994, pp. 33–44.
- [2] G. Ciardo, A. Riska, E. Smirni, An aggregation-based solution method for M/G/1-type processes, in: B. Plateau, W.J. Stewart, M. Silva (Eds.), Numerical Solution of Markov Chains, Prensas Universitarias de Zaragoza, Zaragoza, Spain, September 1999, pp. 21–40.
- [3] G. Ciardo, E. Smirni, ETAQA: an efficient technique for the analysis of QBD-processes by aggregation, Perf. Eval. 36–37 (1999) 71–93.
- [4] D.R. Cox, A use of complex probabilities in the theory of stochastic processes, Proc. Camb. Phil. Soc. 51 (1955) 313–319.
- [5] W.K. Grassmann, D.A. Stanford, Matrix analytic methods, in: W.K. Grassmann (Ed.), Computational Probability, Kluwer Academic Publishers, Boston, MA, 2000, pp. 153–204.
- [6] D.S. Hochbaum, J. Naor, Simple and fast algorithms for linear and integer programs with two variables per inequality, SIAM J. Comp. 23 (6) (1994) 1179–1192.
- [7] D.B. Johnson, Algorithms for shortest paths, Ph.D. Thesis, Cornell University, 1973.
- [8] G. Latouche, Algorithms for infinite Markov chains with repeating columns, in: C. Meyer, R.J. Plemmons (Eds.), Linear Algebra, Markov Chains, and Queueing Models, vol. 48, IMA Volumes in Mathematics and its Applications, Springer, Berlin, 1993, pp. 231–265.

- [9] G. Latouche, V. Ramaswami, Introduction to Matrix Geometric Methods in Stochastic Modeling, ASA-SIAM Series on Statistics and Applied Probability, SIAM, Philadelphia PA, 1999.
- [10] E.L. Lawler, Combinatorial Optimization: Networks and Matroids, Holt, Rinehart and Winston, New York, 1976.
- [11] D.M. Lucantoni, An Algorithmic Analysis of a Communication Model with Retransmission of Flawed Messages, Pitman, Boston, 1983.
- [12] B. Meini, Solving M/G/1 type Markov chains: Recent advances and applications, Stoch. Mod. 14 (1–2) (1998) 479–496.
- [13] R. Nelson, Probability, Stochastic Processes, and Queueing Theory, Springer, Berlin, 1995.
- [14] M.F. Neuts, Matrix-geometric Solutions in Stochastic Models, Johns Hopkins University Press, Baltimore, MD, 1981.
- [15] M.F. Neuts, Structured Stochastic Matrices of M/G/1 Type and Their Applications, Marcel Dekker, New York, 1989.
- [16] V.R. Pratt, Two easy theories whose combination is hard, Technical Report, MIT, 1977.
- [17] V. Ramaswami, A stable recursion for the steady state vector in Markov chains of M/G/1 type, Comm. Stat. Stoch. Mod. 4 (1988) 183–263.
- [18] V. Ramaswami, G. Latouche, A general class of Markov processes with explicit matrix-geometric solutions, OR Spektrum 8 (1986) 209–218.
- [19] A. Riska, E. Smirni, MAMSolver: a matrix-analytic methods tool, in: T. Filed et al. (Eds.), Proceedings of the 12th International Conference on Modelling Techniques and Tools for Computer and Communication System Performance Evaluation (TOOLS 2002), London, UK, April 2002, Lecture Notes in Computer Science, Vol. 2324, Springer, Berlin, 2000, pp. 205–211.
- [20] E. Rosti, E. Smirni, G. Serazzi, L.W. Dowdy, K.C. Sevcik, On processor saving scheduling policies for multiprocessor systems, IEEE Trans. Comp. 47 (2) (1998) 178–189.
- [21] R. Shostak, Deciding linear inequalities by computing loop residues, J. ACM 28 (4) (1981) 769–779.
- [22] M.S. Squillante, F. Wang, M. Papaefthymiou, Stochastic analysis of gang scheduling in parallel and distributed systems, Perf. Eval. 27–28 (1996) 273–296.
- [23] L.N. Trefethen, D. Bau III, Numerical Linear Algebra, SIAM, Philadelphia, 1997.
- [24] J. van Leeuwen, Graph algorithms, in: Handbook of Theoretical Computer Science, MIT Press, Cambridge, MA, 1990.



Gianfranco Ciardo is a professor of computer science at the College of William and Mary, Williamsburg, VA. He has been a visiting professor at the University of Torino, Italy and at the Technical University of Berlin, Germany, and held research positions at HP Labs (Palo Alto, CA), Software Productivity Consortium (Herndon, VA), and CSELT (Torino, Italy). He received a Ph.D. degree in computer science from Duke University, Durham, North Carolina and a Laurea in Scienze dell' Informazione from the University of Torino, Italy. He is interested in theoretical research and tool building for logic and stochastic analysis of discrete-state models, symbolic model checking, performance and reliability evaluation of complex hardware/software systems, Petri nets, and Markov models. He is on the editorial board of IEEE Transactions on Software Engineering and was a keynote speaker at PNPM'01. He is a member of ACM and a senior member of IEEE.



Weizhen Mao received her Ph.D. degree in computer science from Princeton University in 1990 and her B.S. degree in computer science from Tsinghua University in China in 1986. She is currently an associate professor of computer science at the College of William and Mary in Williamsburg, VA. Her research areas include design and analysis of algorithms and combinatorial optimization. In particular, she is interested in approximation algorithms for NP-complete problems, competitive analysis of online algorithms, and theoretical problems arising from job scheduling, data broadcast scheduling, network routing, parallel computing, and computational molecular biology. She is a member of ACM.



Alma Riska received her Ph.D. in computer science from College of William and Mary, Williamsburg, VA, in 2002 and her diploma in computer science from the University of Tirana, Albania, in 1996. Currently, she is a research staff member of Seagate Research Center in Pittsburgh, Pennsylvania. Her research interests are on performance evaluation of computer systems, with emphasis on applying analytic techniques and detailed workload characterization in adaptive storage systems. She is a member of IEEE and ACM.



Evgenia Smirni is an associate professor at the College of William and Mary, Department of Computer Science, Williamsburg, VA. She received her diploma in computer engineering and informatics from the University of Patras, Greece, in 1987, and her M.S. and Ph.D. in computer science from Vanderbilt University in 1993 and 1995, respectively. From August 1995 to June 1997 she had a postdoctoral research associate position at the University of Illinois at Urbana-Champaign. Her research interests include analytic modeling, Markov chains, matrix analytic methods, scheduling and load balancing, resource allocation policies, parallel input/output, workload characterization, and modeling of distributed systems and applications. She is a member of ACM and the IEEE Computer Society.