

Exact analysis of a class of GI/G/1-type performability models*

Alma Riska Evgenia Smirni Gianfranco Ciardo

Department of Computer Science

College of William and Mary

Williamsburg, VA 23187-8795, USA

{riska,esmirni,ciardo}@cs.wm.edu

Abstract

In this paper we present an exact decomposition algorithm for the analysis of Markov chains with GI/G/1-type pattern in their repetitive structure. Such processes exhibit *both* M/G/1-type and GI/M/1-type patterns and cannot be solved with existing techniques. Markov chains with a GI/G/1 pattern result when modeling open systems with failures and repairs that accept jobs from multiple exogenous sources: a single failure can empty the system of jobs, while a single batch arrival can add many jobs to the system. Our method provides exact computation of the stationary probabilities, consequently of performance measures such as the system queue length or any of its higher moments, as well as of the probability of the system being in various failure

*This work has been supported by the National Science Foundation under grant no. EIA-9974992, by the National Aeronautics and Space Administration under NASA Grant NAG-1-2168, and by a William and Mary Summer Research Grant.

states, thus of performability measures. We formulate the conditions under which our approach is feasible and illustrate its applicability via a performability analysis of a parallel computer system.

Keywords: Modeling (performability analysis); Markov processes (GI/G/1-type processes; M/G/1-type processes; I/M/1-type processes; matrix analytic techniques; stochastic complementation).

1 Introduction

During the past two decades, significant effort has been put into the development of modeling techniques that can capture the behavior of modern computer and communication systems [10]. In many cases, M/G/1-type or GI/M/1-type Markov chains have been shown to be appropriate tools (we assume continuous time Markov chains, or CTMCs, but our discussion applies just as well to discrete time Markov chains, or DTMCs). CTMCs that model such processes have an infinite state space with a finite one-dimensional repetitive pattern.

In this paper, we study a class of CTMCs that show a GI/G/1-type pattern in their repetitive structure, that is, they exhibit *both* M/G/1-type and GI/M/1-type patterns, and cannot be solved by existing techniques [4]. Such chains occur when modeling open systems that accept customers from multiple exogenous sources (thus the existence of bulk arrivals) and are subject to failures and repairs (thus the system may become empty when a catastrophic failure occurs, or only parts of it may become non-operational if the failure is non-catastrophic).

The state space¹ \mathcal{S} of a GI/G/1-type CTMC can be partitioned into a finite “boundary”

¹We use calligraphic letters to indicate sets (e.g., \mathcal{A}), lower case boldface Roman or Greek letters to indicate row vectors (e.g., \mathbf{a} , $\boldsymbol{\alpha}$), and upper case boldface Roman letters to indicate matrices (e.g., \mathbf{A}). We use superscripts in parentheses or subscripts to indicate family of related entities (e.g., $\mathcal{A}^{(1)}$, $\mathbf{A}_{\mathcal{A}_1}$). Vector

set $\mathcal{S}^{(0)} = \{s_1^{(0)}, \dots, s_m^{(0)}\}$ and a countably infinite sequence of finite “level” sets $\mathcal{S}^{(j)} = \{s_1^{(j)}, \dots, s_n^{(j)}\}$, for $j \geq 1$. The infinitesimal generator matrix can accordingly be block-partitioned as

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

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

Since $RowSum(\mathbf{Q}) = \mathbf{0}$ and only the diagonal of \mathbf{Q} can contain negative entries, the infinite sets of matrices $\{\widehat{\mathbf{F}}^{(j)} : j \geq 1\}$ and $\{\mathbf{F}^{(j)} : j \geq 1\}$ must be summable. For presentation simplicity, we assume that these matrices have finite representations obeying the geometric expressions $\widehat{\mathbf{F}}^{(j)} = \widehat{\mathbf{A}}^{j-1}\widehat{\mathbf{F}}$ and $\mathbf{F}^{(j)} = \mathbf{A}^{j-1}\mathbf{F}$, for $j \geq 1$, where the spectral radii of $\widehat{\mathbf{A}}$ and \mathbf{A} are strictly less than one, to ensure that the infinite sums $\sum_{j=0}^{\infty} \widehat{\mathbf{A}}^j = (\mathbf{I} - \widehat{\mathbf{A}})^{-1}$ and $\sum_{j=0}^{\infty} \mathbf{A}^j = (\mathbf{I} - \mathbf{A})^{-1}$ exist. Since the same local and forward blocks appear from the third row block on, $RowSum(\widehat{\mathbf{B}}^{(j)}) + RowSum(\mathbf{B}^{(j-1)} + \dots + \mathbf{B}^{(1)})$ must have the same value for all $j \geq 2$. This implies that $\widehat{\mathbf{B}}^{(j)}$ is increasingly smaller, unless of course $\mathbf{B}^{(j)}$ is zero from some j on, and that the infinite set $\{\mathbf{B}^{(j)} : j \geq 1\}$ is also summable, since

and matrix elements are indicated using square brackets (e.g., $\mathbf{a}[1]$, $\mathbf{A}[1, 2]$), and we extend the notation to subvectors or submatrices by allowing sets of indices to be used instead of single indices (e.g., $\mathbf{a}[\mathcal{A}]$, $\mathbf{A}[\mathcal{A}, \mathcal{B}]$). $RowSum(\cdot)$ indicates the diagonal matrix whose entry in position (r, r) is the sum of the entries on the r^{th} row of the argument (which can be a rectangular matrix). $Norm(\cdot)$ indicates the matrix equal to its argument except that all nonzero rows are normalized to sum to one. $\mathbf{0}$ and $\mathbf{1}$ indicate a row vector or a matrix of 0's, or a row vector of 1's, of the appropriate dimensions, respectively.

$RowSum(\mathbf{B}^{(1)} + \dots + \mathbf{B}^{(j)})$, for any finite j , is bounded by $RowSum(-\mathbf{L} - (\mathbf{I} - \mathbf{A})^{-1}\mathbf{F})$. However, $\mathbf{L}^{(1)}$ can differ from \mathbf{L} in the diagonal, hence $RowSum(\widehat{\mathbf{B}}^{(1)})$ might be different from $RowSum(\widehat{\mathbf{B}}^{(j)}) + RowSum(\mathbf{B}^{(j-1)} + \dots + \mathbf{B}^{(1)})$, for $j \geq 2$.

We are interested in the computation of the stationary probability vector $\boldsymbol{\pi}$ solution of $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$, where $\boldsymbol{\pi}$ can be partitioned into $\boldsymbol{\pi}^{(0)} \in \mathbf{R}^m$ and $\boldsymbol{\pi}^{(j)} \in \mathbf{R}^n$, for $j \geq 1$. Of course, since $\boldsymbol{\pi}$ is infinite, in practice we compute $\boldsymbol{\pi}^{(j)}$ only up to a sufficiently large j , or an aggregate measure of the form $\sum_{j=0}^{\infty} \boldsymbol{\pi}^{(j)} \boldsymbol{\rho}^{(j)T}$, where $\boldsymbol{\rho}^{(j)}$ is a vector expressing the *reward rates* for the states in $\mathcal{S}^{(j)}$.

While GI/G/1-type CTMCs do not have a known solution algorithm in general [4], two special cases do: the GI/M/1-type CTMCs (where $\widehat{\mathbf{F}}^{(j)}$ and $\mathbf{F}^{(j)}$ are zero for $j \geq 2$, that is, only forward jumps to the next level are allowed), for which Neuts proposed the matrix geometric method [12], and the M/G/1-type CTMCs (the converse: $\widehat{\mathbf{B}}^{(j)}$ and $\mathbf{B}^{(j)}$ are zero for $j \geq 2$, only backward jumps to the previous level are allowed), for which Neuts proposed the matrix analytic method [11]. The intersection of these two cases, where \mathbf{Q} is block-tridiagonal, is the class of quasi-birth-death (QBD) processes, which can in principle be solved by either method (the matrix geometric method is preferred because it is simpler, more widely-known, and at least as efficient). In practice, such CTMCs often arise when embedding an open system with a single infinite-waiting-room queue at either the arrival or the service completion times.

The solution of M/G/1 and GI/M/1-type CTMCs requires an auxiliary matrix, \mathbf{G} or \mathbf{R} , respectively, typically computed through iterative methods [11, 12]. Several algorithms have been proposed for the efficient computation of \mathbf{R} [5] or \mathbf{G} [8]. See also [6] for a discussion of the fundamental aspects of matrix methods and their use for the solution of QBDs.

In this work, we propose a decomposition approach that extends both the applicability and the efficiency of matrix geometric and matrix analytic methods through an intelligent *partitioning* of the repetitive portion of the state space into subsets, according to its con-

nectivity. This partitioning allows us to define smaller CTMCs with a solvable structure, obtained from the original CTMC through the use of stochastic complementation, an exact decomposition technique that can be used to study the conditional stationary behavior in individual portions of a larger CTMC. To obtain the solution of the original process, the results of the analysis of these portions are coupled back together [9].

The main contribution of our work is the definition of an exact decomposition algorithm that provides *exact* analytic results for CTMCs that exhibit *both* M/G/1-type and GI/M/1-type patterns, i.e., CTMCs with infinitesimal generator having the block structure of Eq. 1, although with certain restrictions on the nonzero patterns of the blocks that will be made clear in the following sections. We do not consider instead the identification of the appropriate partitioning of the state space required to apply our method, this is a graph partitioning problem and is the subject of future work. A second contribution is the application of our algorithm to the performability analysis of computer systems.

In the following, we review the necessary background required to introduce our algorithm: matrix analytic methods and stochastic complementation (Sect. 2). Then, we formally present the proposed algorithm (Sect. 3) and we give an example of a system subject to failures and repairs that can be analyzed with our algorithm (Sect. 4). Finally, we summarize our findings and discuss future work (Sect. 5).

2 Background

We briefly review the basic terminology used to describe the processes we consider, as well as previous results on the solution of M/G/1 and GI/M/1-type processes. We also present some results on stochastic complementation needed to combine the results from the study of individual portions of the overall CTMCs.

2.1 M/G/1-type processes

The infinitesimal generator matrix \mathbf{Q} of a CTMC having an M/G/1-type structure, as defined by Neuts [11], is a special case of Eq. 1:

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

Various recursive algorithms for the stationary solution of such chains exist [4, 8, 12]. Here, we outline one that provides a stable calculation for the values of $\boldsymbol{\pi}^{(j)}$. Using Ramaswami's recursive formula [13] we have:

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

where $\widehat{\mathbf{S}}^{(j)} = \sum_{l=j}^{\infty} \widehat{\mathbf{F}}^{(l)} \mathbf{G}^{l-j}$, for $j \geq 1$, $\mathbf{S}^{(j)} = \sum_{l=j}^{\infty} \mathbf{F}^{(l)} \mathbf{G}^{l-l}$, for $j \geq 0$ (letting $\mathbf{F}^{(0)} \equiv \mathbf{L}$), and \mathbf{G} can be computed as the solution of the matrix equation² $\mathbf{0} = \mathbf{B} + \mathbf{L}\mathbf{G} + \sum_{j=1}^{\infty} \mathbf{F}^{(j)} \mathbf{G}^{j+1}$. Several iterative algorithms for the computation of \mathbf{G} exist [4, 12], a particularly efficient one is based on Toeplitz matrices [8]. See [14] for cases where \mathbf{G} can be explicitly defined and does not require any calculation.

Given the above definition of $\boldsymbol{\pi}^{(j)}$ and the normalization condition, $\boldsymbol{\pi}^{(0)}$ can be calculated as the unique solution of the linear system in m variables

$$\boldsymbol{\pi}^{(0)} \left[\widehat{\mathbf{L}}^{(0)\diamond} - \widehat{\mathbf{S}}^{(1)} \mathbf{S}^{(0)-1} \widehat{\mathbf{B}}^{\diamond} \mid \mathbf{1}^T - \sum_{j=1}^{\infty} \widehat{\mathbf{S}}^{(j)} \left(\sum_{j=0}^{\infty} \mathbf{S}^{(j)} \right)^{-1} \mathbf{1}^T \right] = [\mathbf{0} \mid 1],$$

² $\mathbf{G}[k, l]$ represents the probability of first entering $\mathcal{S}^{(j-1)}$ through state $s_l^{(j-1)}$, starting from $\mathcal{S}^{(j)}$ in state $s_k^{(j)}$.

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

2.2 GI/M/1-type and QBD processes

The infinitesimal generator of a CTMC of the GI/M/1-type, also defined by Neuts in [11], is another special case of the structure of Eq. 1:

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

A further special case of the above structure is that of a QBD CTMC, whose infinitesimal generator can be block-partitioned as:

$$\mathbf{Q} = \begin{bmatrix} \widehat{\mathbf{L}}^{(0)} & \widehat{\mathbf{F}} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots \\ \widehat{\mathbf{B}} & \mathbf{L} & \mathbf{F} & \mathbf{0} & \mathbf{0} & \dots \\ \mathbf{0} & \mathbf{B} & \mathbf{L} & \mathbf{F} & \mathbf{0} & \dots \\ \mathbf{0} & \mathbf{0} & \mathbf{B} & \mathbf{L} & \mathbf{F} & \dots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{B} & \mathbf{L} & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}.$$

While the QBD case falls under both the M/G/1 and the GI/M/1-type case, it is most commonly associated with GI/M/1-type matrices because both can be solved using the well-

known matrix geometric approach [11]:

$$\forall j \geq 1, \boldsymbol{\pi}^{(j)} = \boldsymbol{\pi}^{(1)} \mathbf{R}^{j-1}, \quad (2)$$

where, in the GI/M/1-type case, \mathbf{R} is the solution of the matrix equation

$$\mathbf{F} + \mathbf{R}\mathbf{L} + \sum_{j=1}^{\infty} \mathbf{R}^{j+1} \mathbf{B}^{(j)} = \mathbf{0} \quad (3)$$

and can be computed using iterative numerical algorithms³. Then, together with the normalization condition $\boldsymbol{\pi}^{(0)} \mathbf{1}^T + \boldsymbol{\pi}^{(1)} \sum_{j=1}^{\infty} \mathbf{R}^{j-1} \mathbf{1}^T = 1$, which we can rewrite as $\boldsymbol{\pi}^{(0)} \mathbf{1}^T + \boldsymbol{\pi}^{(1)} (\mathbf{I} - \mathbf{R})^{-1} \mathbf{1}^T = 1$, we obtain the linear system in $m + n$ variables

$$[\boldsymbol{\pi}^{(0)} | \boldsymbol{\pi}^{(1)}] \left[\begin{array}{c|c|c} \mathbf{1}^T & \widehat{\mathbf{L}}^{(0)\diamond} & \widehat{\mathbf{F}}^{(1)} \\ \hline (\mathbf{I} - \mathbf{R})^{-1} \mathbf{1}^T & \left(\sum_{j=1}^{\infty} \mathbf{R}^{j-1} \widehat{\mathbf{B}}^{(j)\diamond} \right) & \mathbf{L} + \sum_{j=1}^{\infty} \mathbf{R}^j \mathbf{B}^{(j)} \end{array} \right] = [1 | \mathbf{0}] \quad (4)$$

which yields a unique solution for $\boldsymbol{\pi}^{(0)}$ and $\boldsymbol{\pi}^{(1)}$. For $j \geq 2$, $\boldsymbol{\pi}^{(j)}$ can be obtained numerically from Eq. 2, but many useful performance metrics such as expected system utilization, throughput, or queue length can be computed exactly in explicit form using $\boldsymbol{\pi}^{(0)}$, $\boldsymbol{\pi}^{(1)}$, and \mathbf{R} alone.

For QBD processes, Eq. 3 reduces to the quadratic matrix equation $\mathbf{F} + \mathbf{R}\mathbf{L} + \mathbf{R}^2 \mathbf{B} = \mathbf{0}$.

2.3 Stochastic complementation

In this section, we briefly outline the concept of stochastic complementation [9] and focus on results needed to derive our method. While [9] introduces the concept of stochastic complementation for DTMCs with finite state spaces we define it instead for the infinite case, a straightforward extension, and state the results in terms of CTMCs.

Partition the state space \mathcal{S} of an ergodic CTMC with infinitesimal generator matrix \mathbf{Q} and stationary probability vector $\boldsymbol{\pi}$, satisfying $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$, into two disjoint subsets, \mathcal{A} and $\overline{\mathcal{A}}$.

³ $\mathbf{R}[k, l]$ represents the expected time spent in state $s_l^{(j+1)}$ after a transition out of state $s_k^{(j)}$ and before entering a state in $\mathcal{S}^{(j)}$, measured using $-1/\mathbf{L}[k, k]$, the holding time in state $s_k^{(j)}$, as the time unit.

Definition 1 [9] (**Stochastic complement**) The stochastic complement of \mathcal{A} is

$$\mathbf{A} = \mathbf{Q}[\mathcal{A}, \mathcal{A}] + \mathbf{Q}[\mathcal{A}, \bar{\mathcal{A}}](-\mathbf{Q}[\bar{\mathcal{A}}, \bar{\mathcal{A}}])^{-1}\mathbf{Q}[\bar{\mathcal{A}}, \mathcal{A}], \quad (5)$$

where $(-\mathbf{Q}[\bar{\mathcal{A}}, \bar{\mathcal{A}}])^{-1}[r, c]$ represents the mean time spent in state $c \in \bar{\mathcal{A}}$, starting from state $r \in \bar{\mathcal{A}}$, before reaching any state in \mathcal{A} , and $((-\mathbf{Q}[\bar{\mathcal{A}}, \bar{\mathcal{A}}])^{-1}\mathbf{Q}[\bar{\mathcal{A}}, \mathcal{A}])[r, c']$ represents the probability that, starting from $r \in \bar{\mathcal{A}}$, we enter \mathcal{A} through state c' . \square

The stochastic complement \mathbf{A} is the infinitesimal generator of a new CTMC which mimics the original CTMC but “skips over” states in $\bar{\mathcal{A}}$. The following theorem formalizes this concept.

Theorem 1 [9] The stochastic complement \mathbf{A} of \mathcal{A} is an infinitesimal generator, and is irreducible if \mathbf{Q} is. If $\boldsymbol{\alpha}$ is its stationary probability vector satisfying $\boldsymbol{\alpha}\mathbf{A} = \mathbf{0}$, then $\boldsymbol{\alpha} = \text{Norm}(\boldsymbol{\pi}[\mathcal{A}])$. \square

In other words, $\boldsymbol{\alpha}$ gives the individual probability of being in each state of \mathcal{A} for the original CTMC, conditioned on being in \mathcal{A} . This implies that the stationary probability distribution $\boldsymbol{\alpha}$ of the stochastic complement differs from the corresponding portion of the stationary distribution of the original CTMC $\boldsymbol{\pi}[\mathcal{A}]$ only by a constant $a = \boldsymbol{\pi}[\mathcal{A}]\mathbf{1}^T$, which represents the probability of being in \mathcal{A} in the original CTMC. The value a is known as the *coupling factor* of the stochastic complement.

Stochastic complementation can be used as a divide-and-conquer strategy: to this end, computing the stationary distribution of each stochastic complement must be easier than computing the stationary distribution of the original CTMC directly. Since the inverse $(-\mathbf{Q}[\bar{\mathcal{A}}, \bar{\mathcal{A}}])^{-1}$ is often a full matrix, its computation can be costly. However, there are cases where we can take advantage of the special structure of the CTMC and avoid this computation. To consider these cases, rewrite the definition of stochastic complement in

Eq. 5 as

$$\mathbf{A} = \mathbf{Q}[\mathcal{A}, \mathcal{A}] + \text{RowSum}(\mathbf{Q}[\mathcal{A}, \bar{\mathcal{A}}]) \mathbf{Z} \quad \text{where} \quad \mathbf{Z} = \text{Norm}(\mathbf{Q}[\mathcal{A}, \bar{\mathcal{A}}]) (-\mathbf{Q}[\bar{\mathcal{A}}, \bar{\mathcal{A}}])^{-1} \mathbf{Q}[\bar{\mathcal{A}}, \mathcal{A}]. \quad (6)$$

The r^{th} diagonal element of $\text{RowSum}(\mathbf{Q}[\mathcal{A}, \bar{\mathcal{A}}])$ represents the rate at which the set \mathcal{A} is left from its r^{th} state to reach any of the states in $\bar{\mathcal{A}}$, while the r^{th} row of \mathbf{Z} , which sums to one, specifies how this rate should be redistributed over the states in \mathcal{A} when the process eventually reenters it.

Lemma 1 (Single entry) If \mathcal{A} can be entered from $\bar{\mathcal{A}}$ only through a single state $c \in \mathcal{A}$, the matrix \mathbf{Z} defined in Eq. 6 is trivially computable: it is a matrix of zeros except for its c^{th} column, which contains all ones. \square

Definition 2 (Pseudo-stochastic complement) The pseudo-stochastic complement of \mathcal{A} is

$$\mathbf{A} = \mathbf{Q}[\mathcal{A}, \mathcal{A}] + \mathbf{Q}[\mathcal{A}, \bar{\mathcal{A}}] \mathbf{1}^T \text{Norm}(\boldsymbol{\pi}[\bar{\mathcal{A}}] \mathbf{Q}[\bar{\mathcal{A}}, \mathcal{A}]),$$

where $\boldsymbol{\pi}$ is the stationary distribution for the states of the original CTMC.

Theorem 2 The pseudo-stochastic complement \mathbf{A} of \mathcal{A} is an infinitesimal generator and is irreducible if \mathbf{Q} is. If $\boldsymbol{\alpha}$ is its stationary probability vector satisfying $\boldsymbol{\alpha} \mathbf{A} = \mathbf{0}$, then $\boldsymbol{\alpha} = \text{Norm}(\boldsymbol{\pi}[\mathcal{A}])$. **Proof:** See Appendix A. \square

Comparing the definitions of stochastic and pseudo-stochastic complements, we see that the former is naturally based on considering all paths starting and ending in \mathcal{A} (given the interpretation of the inverse as an infinite sum over all possible paths), while the latter makes use of the conditional stationary probability vector for $\bar{\mathcal{A}}$. What is interesting, though, is that, even if the two complements are described by different matrices, Theorems 1 and 2 imply that they *have the same stationary probability vector* $\boldsymbol{\alpha}$. The intuition behind this

property is given by the stochastic meaning of the matrix $-\mathbf{Q}[\overline{\mathcal{A}}, \overline{\mathcal{A}}]^{-1}$ used in the formulation of the stochastic complement. Its entry in row r and column c represents the expected amount of time spent in state $c \in \overline{\mathcal{A}}$ starting from state $r \in \overline{\mathcal{A}}$, before entering \mathcal{A} , a quantity that is of course related to the stationary probability vector $\pi[\overline{\mathcal{A}}]$ used in the formulation of the pseudo-stochastic complement.

3 Our approach

The interactions between states of a CTMC with infinitesimal generator \mathbf{Q} given by Eq.(1) exhibit the patterns for both GI/M/1-type and M/G/1-type processes, thus they are more general than either. Our approach is based on decomposition: informally, we relegate the GI/M/1-type and the M/G/1-type behaviors into different CTMCs which not only must be individually solvable by known methods, but also, just as important, must provide results that are meaningful to the original process. We stress that we can do this even if, in the original CTMC, there are states with arbitrarily long forward and backward transitions to other states, as long as \mathbf{Q} satisfies the conditions outlined in Section 3.1. At a high level, these are our solution steps:

1. We determine a “gate” state g in $\mathcal{S}^{(0)}$, a covering $\mathcal{U}^{(0)}$ (for “upper”) and $\mathcal{L}^{(0)}$ (for “lower”) of $\mathcal{S}^{(0)}$, with $\mathcal{U}^{(0)} \cup \mathcal{L}^{(0)} = \mathcal{S}^{(0)}$ and $\mathcal{U}^{(0)} \cap \mathcal{L}^{(0)} = \{g\}$, and a partition of each level set $\mathcal{S}^{(j)}$ into two disjoint sets $\mathcal{U}^{(j)}$ and $\mathcal{L}^{(j)}$. These allow us to capture the GI/M/1-type behavior of \mathbf{Q} with an “upper CTMC” having state space $\mathcal{U} = \bigcup_{j=0}^{\infty} \mathcal{U}^{(j)}$, and the M/G/1-type behavior of \mathbf{Q} with a “lower CTMC” having state space $\mathcal{L} = \bigcup_{j=0}^{\infty} \mathcal{L}^{(j)}$.
2. We define the rates for the upper CTMC using stochastic complementation in the special setting of Lemma 1. Then, we solve it using the matrix geometric method [12], obtaining both the matrix \mathbf{R} and the conditional stationary probabilities of the states




Figure 1: The overall idea of our approach.

in \mathcal{U} for the original CTMC.

3. Using the matrix \mathbf{R} just computed, we apply pseudo-stochastic complementation to define the rates for the lower CTMC. Then, we solve it using the matrix analytic method [8, 12].
4. Finally, we “couple” the two solutions and obtain an expression for the stationary probability vector of the original CTMC. Unlike standard applications of stochastic complementation, where the state spaces of the stochastic complements being solved are normally disjoint, our upper and lower CTMCs share the state g , and this greatly simplifies the coupling.

The overall idea is illustrated in Fig. 1. We now describe in detail the key steps of our approach.

3.1 Determining the upper and lower sets

The partition of each $\mathcal{S}^{(j)}$, $j \geq 1$, must be consistent with a partition of the set of repetitive state indices $\{1, \dots, n\}$ into \mathcal{N}_u and \mathcal{N}_l , that is: $\mathcal{U}^{(j)} = \{s_i^{(j)} : i \in \mathcal{N}_u\}$ and $\mathcal{L}^{(j)} = \{s_i^{(j)} : i \in \mathcal{N}_l\}$. Recall that g belongs to both $\mathcal{U}^{(0)}$ and $\mathcal{L}^{(0)}$, hence to both \mathcal{U} and \mathcal{L} . We then use the notation $\mathcal{U}^{(0)}$, $\mathcal{L}^{(0)}$, \mathcal{U} , and \mathcal{L} to indicate the analogous sets without g , that is, $\mathcal{U}^{(0)} = \mathcal{U}^{(0)} \setminus \{g\}$, $\mathcal{L}^{(0)} = \mathcal{L}^{(0)} \setminus \{g\}$, $\mathcal{U} = \mathcal{U} \setminus \{g\}$, and $\mathcal{L} = \mathcal{L} \setminus \{g\}$. The covering of $\mathcal{S}^{(0)}$,

or equivalently its partition into $\mathcal{U}^{(0)}$, $\mathcal{L}^{(0)}$, and $\{g\}$, is completely determined by the choice of g : $\mathcal{L}^{(0)}$ is the set of states in $\mathcal{S}^{(0)}$ that can be reached from $\mathcal{L}^{(1)}$ without visiting g , while $\mathcal{U}^{(0)} = \mathcal{S}^{(0)} \setminus \mathcal{L}^{(0)}$ (there are no constraints on transitions from \mathcal{U} to $\mathcal{L}^{(0)}$). The definition of these two partitions, hence of the gate g , must satisfy the following conditions (see Fig. 2):

- Forward transitions are allowed from $\mathcal{U}^{(j)}$ to $\mathcal{U}^{(j+1)}$, $j \geq 0$. Backward transitions are allowed from $\mathcal{U}^{(j)}$, $j \geq 1$, to any lower level $\mathcal{U}^{(k)}$, $k < j$, and to $\mathcal{S}^{(0)}$. Note that, in Fig. 2, the set $\mathcal{S}^{(0)}$, which is exactly equal to $\mathcal{U}^{(0)} \cup \mathcal{L}^{(0)}$, is represented separately from them, in grey, simply to stress the fact that arcs reaching $\mathcal{S}^{(0)}$ or leaving from it can reach or leave any state in $\mathcal{U}^{(0)}$ or $\mathcal{L}^{(0)}$.
- Forward transitions are allowed from $\mathcal{S}^{(0)}$ to any $\mathcal{L}^{(j)}$ and from any $\mathcal{L}^{(j)}$ toward any $\mathcal{L}^{(k)}$, $k > j$. Backward transitions are allowed from $\mathcal{L}^{(j)}$ to $\mathcal{L}^{(j-1)}$, $j \geq 1$.
- Local transitions (not shown) are allowed within each $\mathcal{U}^{(j)}$ or $\mathcal{L}^{(j)}$, $j \geq 0$.
- Transitions from $\mathcal{U}^{(j)}$ to any $\mathcal{L}^{(k)}$, $k \geq 1$ are allowed.
- Transitions from $\mathcal{U}^{(0)}$ to $\mathcal{L}^{(0)}$ are allowed.
- Transitions from g are allowed to any state except those in $\bigcup_{j=2}^{\infty} \mathcal{U}^{(j)}$.
- No other transition is allowed. In particular, there is no direct transitions from \mathcal{L} to \mathcal{U} .

These requirements imply that any path from \mathcal{L} to \mathcal{U} must visit g . Furthermore, they imply not only that $\mathbf{Q}[\mathcal{U}, \mathcal{U}]$ has a GI/M/1-type structure (possibly QBD) and $\mathbf{Q}[\mathcal{L}, \mathcal{L}]$ has a M/G/1-type structure (possibly QBD), but also that their structure is preserved even after applying stochastic complementation to them, as discussed in the following sections.

In practice, a suitable gate might exist in $\mathcal{S}^{(1)}$ but not in $\mathcal{S}^{(0)}$; in this case, we simply redefine a new $\mathcal{S}^{(0)}$ as the union of the original sets $\mathcal{S}^{(0)}$ and $\mathcal{S}^{(1)}$ and re-index all levels

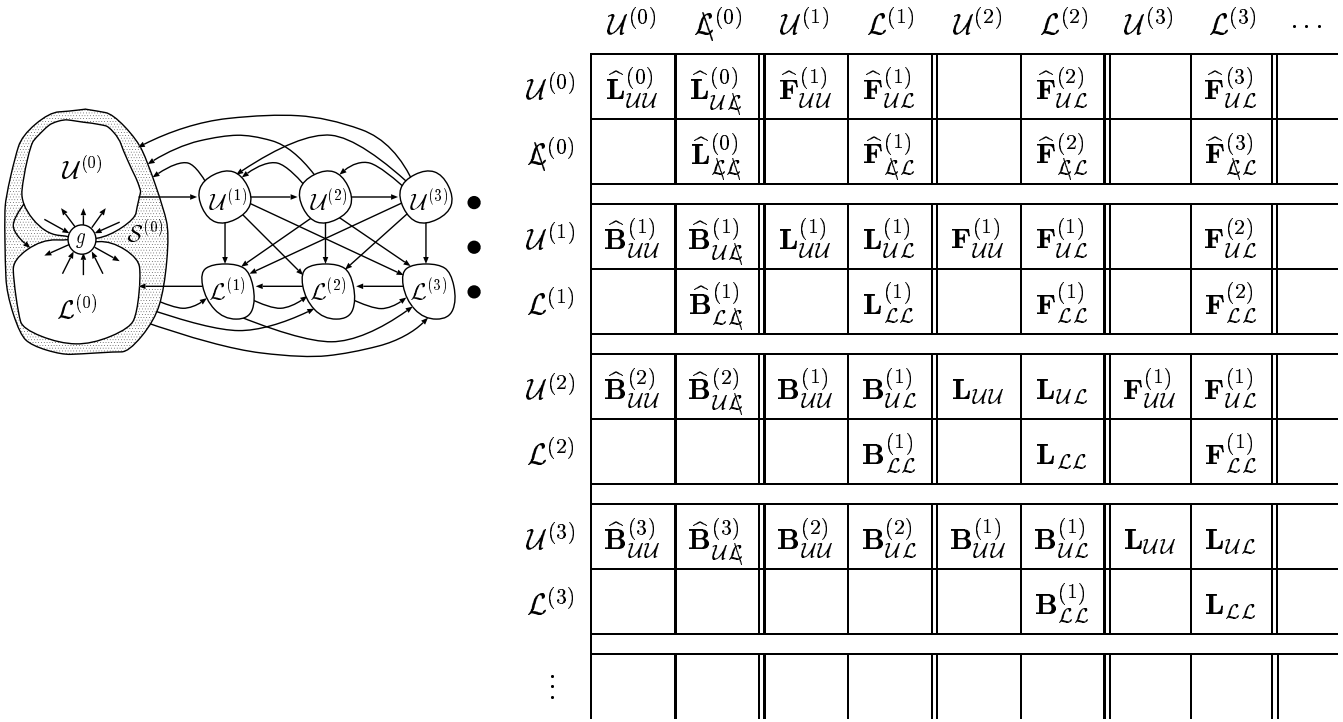


Figure 2: Two-level gated interaction and the corresponding nonzero pattern in matrix \mathbf{Q} .

accordingly. Also, note that this partition is not uniquely defined in general. For example, any infinite subchain having a strictly QBD pattern can be considered as a M/G/1-type or a GI/M/1-type portion, hence its states could be considered either as belonging to either \mathcal{U} or \mathcal{L} . In such cases, we will consider them to be in \mathcal{L} . If we cannot find a partition satisfying the required conditions, our approach is not applicable, and we must stop. Otherwise, we can proceed with the following steps.

3.2 Defining and solving the upper CTMC

The upper CTMC is obtained as the stochastic complement of \mathcal{U} . To compute its rates efficiently through stochastic complementation, the single-entry condition must be satisfied (Lemma 1): this is one of the reasons for requiring the existence of the gate g which all paths from \mathcal{K} to \mathcal{U} must visit (the other being that the application of pseudo-stochastic complementation for the lower CTMC is also simplified). The blocks composing \mathbf{Q} in Eq. 1

can be block-partitioned according to the upper and lower sets, as shown in Fig. 2 on the right (where non-zero blocks are labeled with a shorthand notation).

When we apply stochastic complementation to \mathcal{U} , any transition from \mathcal{U} to \mathcal{K} is simply rerouted to the gate state g . The infinitesimal generator of this new CTMC is then

$$\mathbf{Q}_{\mathcal{U}} = \begin{bmatrix} \tilde{\mathbf{L}}_{\mathcal{U}\mathcal{U}}^{(0)} & \hat{\mathbf{F}}_{\mathcal{U}\mathcal{U}}^{(1)} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots \\ \tilde{\mathbf{B}}_{\mathcal{U}\mathcal{U}}^{(1)} & \mathbf{L}_{\mathcal{U}\mathcal{U}}^{(1)} & \mathbf{F}_{\mathcal{U}\mathcal{U}}^{(1)} & \mathbf{0} & \mathbf{0} & \cdots \\ \tilde{\mathbf{B}}_{\mathcal{U}\mathcal{U}}^{(2)} & \mathbf{B}_{\mathcal{U}\mathcal{U}}^{(1)} & \mathbf{L}_{\mathcal{U}\mathcal{U}} & \mathbf{F}_{\mathcal{U}\mathcal{U}}^{(1)} & \mathbf{0} & \cdots \\ \tilde{\mathbf{B}}_{\mathcal{U}\mathcal{U}}^{(3)} & \mathbf{B}_{\mathcal{U}\mathcal{U}}^{(2)} & \mathbf{B}_{\mathcal{U}\mathcal{U}}^{(1)} & \mathbf{L}_{\mathcal{U}\mathcal{U}} & \mathbf{F}_{\mathcal{U}\mathcal{U}}^{(1)} & \cdots \\ \tilde{\mathbf{B}}_{\mathcal{U}\mathcal{U}}^{(4)} & \mathbf{B}_{\mathcal{U}\mathcal{U}}^{(3)} & \mathbf{B}_{\mathcal{U}\mathcal{U}}^{(2)} & \mathbf{B}_{\mathcal{U}\mathcal{U}}^{(1)} & \mathbf{L}_{\mathcal{U}\mathcal{U}} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix},$$

where the block $\tilde{\mathbf{L}}_{\mathcal{U}\mathcal{U}}^{(0)}$ is the same as $\hat{\mathbf{L}}_{\mathcal{U}\mathcal{U}}^{(0)}$ except in column g , which has also the contribution obtained by applying Eq. 6 and Lemma 1,

$$\tilde{\mathbf{L}}_{\mathcal{U}\mathcal{U}}^{(0)} = \hat{\mathbf{L}}_{\mathcal{U}\mathcal{U}}^{(0)} + \text{RowSum}(\hat{\mathbf{L}}_{\mathcal{U}\mathcal{K}}^{(0)})\mathbf{1}^T\mathbf{e}_g + \text{RowSum}\left(\sum_{l=1}^{\infty}\hat{\mathbf{F}}_{\mathcal{U}\mathcal{L}}^{(l)}\right)\mathbf{1}^T\mathbf{e}_g$$

(\mathbf{e}_g is a row vector of zeros except for a 1 in position g),

$$\tilde{\mathbf{B}}_{\mathcal{U}\mathcal{U}}^{(1)} = \hat{\mathbf{B}}_{\mathcal{U}\mathcal{U}}^{(1)} + \text{RowSum}(\hat{\mathbf{B}}_{\mathcal{U}\mathcal{K}}^{(1)})\mathbf{1}^T\mathbf{e}_g + \text{RowSum}\left(\mathbf{L}_{\mathcal{U}\mathcal{L}}^{(1)} + \sum_{l=1}^{\infty}\mathbf{F}_{\mathcal{U}\mathcal{L}}^{(l)}\right)\mathbf{1}^T\mathbf{e}_g,$$

and, for $j \geq 2$,

$$\tilde{\mathbf{B}}_{\mathcal{U}\mathcal{U}}^{(j)} = \hat{\mathbf{B}}_{\mathcal{U}\mathcal{U}}^{(j)} + \text{RowSum}(\hat{\mathbf{B}}_{\mathcal{U}\mathcal{K}}^{(j)})\mathbf{1}^T\mathbf{e}_g + \text{RowSum}\left(\sum_{l=1}^{j-1}\mathbf{B}_{\mathcal{U}\mathcal{L}}^{(l)} + \mathbf{L}_{\mathcal{U}\mathcal{L}} + \sum_{l=1}^{\infty}\mathbf{F}_{\mathcal{U}\mathcal{L}}^{(l)}\right)\mathbf{1}^T\mathbf{e}_g.$$

The result is a GI/M/1-type process that we solve using the matrix geometric method outlined in Sect. 2.2. We can choose to discard the g^{th} column of $\tilde{\mathbf{L}}_{\mathcal{U}\mathcal{U}}^{(0)}$ and $\tilde{\mathbf{B}}_{\mathcal{U}\mathcal{U}}^{(j)}$ when applying Eq. 4, so the contribution of the stochastic complementation is actually captured by the diagonal of matrices $\mathbf{L}_{\mathcal{U}\mathcal{U}}^{(1)}$ and $\mathbf{L}_{\mathcal{U}\mathcal{U}}$, and no computation is required in practice to set up the linear system of Eq. 4.

Letting $\bar{\boldsymbol{\alpha}} = [\bar{\boldsymbol{\alpha}}^{(0)}, \bar{\boldsymbol{\alpha}}^{(1)}, \dots]$ be the stationary probability vector of the upper CTMC, we obtain \mathbf{R} , $\bar{\boldsymbol{\alpha}}^{(0)}$, and $\bar{\boldsymbol{\alpha}}^{(1)}$ by applying the matrix geometric method. Then, we can obtain

any $\bar{\alpha}^{(j)}$, for $j \geq 2$, using the relation $\bar{\alpha}^{(j)} = \bar{\alpha}^{(1)}\mathbf{R}^{(j-1)}$. Of course, $\bar{\alpha}$ also represents the probabilities of the states in \mathcal{U} for the overall CTMC conditioned on being in \mathcal{U} . These are needed to formulate the pseudo-stochastic complement of \mathcal{L} , in the next section.

3.3 Defining and solving the lower CTMC

To derive the infinitesimal generator $\mathbf{Q}_{\mathcal{L}}$ of the lower CTMC, with state space \mathcal{L} , we apply Definition 2. Recall that, now the gate state g belongs to \mathcal{L} , hence the complement of \mathcal{L} in \mathcal{S} is \mathcal{U} . The pseudo-stochastic complement of \mathcal{L} is then $\mathbf{Q}_{\mathcal{L}} = \mathbf{Q}[\mathcal{L}, \mathcal{L}] + \mathbf{Q}[\mathcal{L}, \mathcal{U}] \mathbf{1}^T \text{Norm}(\pi[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}]) = \mathbf{Q}[\mathcal{L}, \mathcal{L}] + \mathbf{Q}[\mathcal{L}, \mathcal{U}] \mathbf{1}^T \text{Norm}(\bar{\alpha}[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}])$, where the last equality holds because $\bar{\alpha}[\mathcal{U}]$ and $\pi[\mathcal{U}]$ differ only by a constant, which is irrelevant due to the normalization. Furthermore, only row g of $\mathbf{Q}[\mathcal{L}, \mathcal{U}]$ can have nonzero entries, thus $\mathbf{Q}_{\mathcal{L}}$ differs from $\mathbf{Q}[\mathcal{L}, \mathcal{L}]$ only in that row. In other words, only transition from g to \mathcal{U} must be rerouted to \mathcal{L} .

Since the block-partition defined in Fig. 2 places g with the upper set of states, more precisely in $\mathcal{U}^{(0)}$, we need to repartition the first two rows and columns of blocks in Fig. 2, to reflect the fact that now g belongs to $\mathcal{L}^{(0)}$ instead. In other words, we need to define the blocks

$$\begin{aligned} \widehat{\mathbf{B}}_{\mathcal{L}\mathcal{L}}^{(1)} &= \mathbf{Q}[\mathcal{L}^{(1)}, \mathcal{L}^{(0)}] & \widehat{\mathbf{L}}_{\mathcal{L}\mathcal{L}}^{(0)} &= \mathbf{Q}[\mathcal{L}^{(0)}, \mathcal{L}^{(0)}] & \widehat{\mathbf{F}}_{\mathcal{L}\mathcal{L}}^{(j)} &= \mathbf{Q}[\mathcal{L}^{(0)}, \mathcal{L}^{(j)}] \\ \widehat{\mathbf{B}}_{\mathcal{U}\mathcal{L}}^{(j)} &= \mathbf{Q}[\mathcal{U}^{(j)}, \mathcal{L}^{(0)}] & \widehat{\mathbf{L}}_{\mathcal{U}\mathcal{L}}^{(0)} &= \mathbf{Q}[\mathcal{U}^{(0)}, \mathcal{L}^{(0)}] & \widehat{\mathbf{F}}_{\mathcal{U}\mathcal{L}}^{(j)} &= \mathbf{Q}[\mathcal{U}^{(0)}, \mathcal{L}^{(j)}] \end{aligned}$$

Then, the infinitesimal generator of the lower CTMC with state space \mathcal{L} is

$$\mathbf{Q}_{\mathcal{L}} = \begin{bmatrix} \widetilde{\mathbf{L}}_{\mathcal{L}\mathcal{L}}^{(0)} & \widetilde{\mathbf{F}}_{\mathcal{L}\mathcal{L}}^{(1)} & \widetilde{\mathbf{F}}_{\mathcal{L}\mathcal{L}}^{(2)} & \widetilde{\mathbf{F}}_{\mathcal{L}\mathcal{L}}^{(3)} & \widetilde{\mathbf{F}}_{\mathcal{L}\mathcal{L}}^{(4)} & \cdots \\ \widehat{\mathbf{B}}_{\mathcal{L}\mathcal{L}}^{(1)} & \mathbf{L}_{\mathcal{L}\mathcal{L}}^{(1)} & \mathbf{F}_{\mathcal{L}\mathcal{L}}^{(1)} & \mathbf{F}_{\mathcal{L}\mathcal{L}}^{(2)} & \mathbf{F}_{\mathcal{L}\mathcal{L}}^{(3)} & \cdots \\ \mathbf{0} & \mathbf{B}_{\mathcal{L}\mathcal{L}}^{(1)} & \mathbf{L}_{\mathcal{L}\mathcal{L}} & \mathbf{F}_{\mathcal{L}\mathcal{L}}^{(1)} & \mathbf{F}_{\mathcal{L}\mathcal{L}}^{(2)} & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{B}_{\mathcal{L}\mathcal{L}}^{(1)} & \mathbf{L}_{\mathcal{L}\mathcal{L}} & \mathbf{F}_{\mathcal{L}\mathcal{L}}^{(1)} & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{B}_{\mathcal{L}\mathcal{L}}^{(1)} & \mathbf{L}_{\mathcal{L}\mathcal{L}} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix},$$

where $\tilde{\mathbf{L}}_{\mathcal{L}\mathcal{L}}^{(0)} = \hat{\mathbf{L}}_{\mathcal{L}\mathcal{L}}^{(0)} + \mathbf{e}_g^T \boldsymbol{\gamma}^{(0)}$ and $\tilde{\mathbf{F}}_{\mathcal{L}\mathcal{L}}^{(j)} = \hat{\mathbf{F}}_{\mathcal{L}\mathcal{L}}^{(j)} + \mathbf{e}_g^T \boldsymbol{\gamma}^{(j)}$, for $j \geq 1$, where the ‘‘correcting’’ row vectors $\boldsymbol{\gamma}^{(j)}$, whose computation we show next, represent the rerouted rates of going from g to $\mathcal{L}^{(j)}$ through \mathcal{U} , for $j \geq 0$.

The total rate at which the original CTMC leaves state g toward \mathcal{U} once it is in g is

$$u = \mathbf{Q}[g, \mathcal{U}] \mathbf{1}^T = \sum_{i \in \mathcal{M}^{(0)} \cup \mathcal{U}^{(1)}} \mathbf{Q}[g, i].$$

To compute $Norm(\bar{\boldsymbol{\alpha}}[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}])$ we obtain first

$$\begin{aligned} \bar{\boldsymbol{\alpha}}[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}^{(0)}] &= \bar{\boldsymbol{\alpha}}^{(0)}[\mathcal{U}^{(0)}]\hat{\mathbf{L}}_{\mathcal{M}\mathcal{L}}^{(0)} + \sum_{l=1}^{\infty} \bar{\boldsymbol{\alpha}}^{(l)}\hat{\mathbf{B}}_{\mathcal{U}\mathcal{L}}^{(l)}, \\ \bar{\boldsymbol{\alpha}}[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}^{(1)}] &= \bar{\boldsymbol{\alpha}}^{(0)}[\mathcal{U}^{(0)}]\hat{\mathbf{F}}_{\mathcal{M}\mathcal{L}}^{(1)} + \bar{\boldsymbol{\alpha}}^{(1)}\mathbf{L}_{\mathcal{U}\mathcal{L}}^{(1)} + \sum_{l=1}^{\infty} \bar{\boldsymbol{\alpha}}^{(l+1)}\mathbf{B}_{\mathcal{U}\mathcal{L}}^{(l)}, \\ \bar{\boldsymbol{\alpha}}[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}^{(j)}] &= \bar{\boldsymbol{\alpha}}^{(0)}[\mathcal{U}^{(0)}]\hat{\mathbf{F}}_{\mathcal{M}\mathcal{L}}^{(j)} + \sum_{l=1}^{j-1} \bar{\boldsymbol{\alpha}}^{(l)}\mathbf{F}_{\mathcal{U}\mathcal{L}}^{(j-l)} + \bar{\boldsymbol{\alpha}}^{(j)}\mathbf{L}_{\mathcal{U}\mathcal{L}} + \sum_{l=j+1}^{\infty} \bar{\boldsymbol{\alpha}}^{(l)}\mathbf{B}_{\mathcal{U}\mathcal{L}}^{(l-j)}, \quad \text{for } j > 1. \end{aligned}$$

Then, the normalization factor $f = \bar{\boldsymbol{\alpha}}[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}]\mathbf{1}^T$ is given by

$$\begin{aligned} &\bar{\boldsymbol{\alpha}}^{(0)}[\mathcal{U}^{(0)}] \left(\hat{\mathbf{L}}_{\mathcal{M}\mathcal{L}}^{(0)} \mathbf{1}^T + \sum_{j=1}^{\infty} \hat{\mathbf{F}}_{\mathcal{M}\mathcal{L}}^{(j)} \mathbf{1}^T \right) + \bar{\boldsymbol{\alpha}}^{(1)} \left(\hat{\mathbf{B}}_{\mathcal{U}\mathcal{L}}^{(1)} \mathbf{1}^T + \mathbf{L}_{\mathcal{U}\mathcal{L}}^{(1)} \mathbf{1}^T + \sum_{j=1}^{\infty} \mathbf{F}_{\mathcal{U}\mathcal{L}}^{(j)} \mathbf{1}^T \right) \\ &+ \sum_{j=2}^{\infty} \bar{\boldsymbol{\alpha}}^{(j)} \left(\hat{\mathbf{B}}_{\mathcal{U}\mathcal{L}}^{(j)} \mathbf{1}^T + \sum_{l=1}^{j-1} \mathbf{B}_{\mathcal{U}\mathcal{L}}^{(l)} \mathbf{1}^T + \mathbf{L}_{\mathcal{U}\mathcal{L}} \mathbf{1}^T + \sum_{l=1}^{\infty} \mathbf{F}_{\mathcal{U}\mathcal{L}}^{(l)} \mathbf{1}^T \right) \end{aligned}$$

and the correcting vectors $\boldsymbol{\gamma}^{(j)}$, for $j \geq 0$, can be computed as

$$\boldsymbol{\gamma}^{(j)} = u \, Norm(\bar{\boldsymbol{\alpha}}[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}])[\mathcal{L}^{(j)}] = u \bar{\boldsymbol{\alpha}}[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}^{(j)}] / f.$$

In practice, we must truncate any infinite summation involving the vectors $\bar{\boldsymbol{\alpha}}^{(l)}$, such as those involved in the computation of $\bar{\boldsymbol{\alpha}}[\mathcal{U}]\mathbf{Q}[\mathcal{U}, \mathcal{L}^{(j)}]$ or of the normalization factor f , to a value of l large enough. Finally, once $\mathbf{Q}_{\mathcal{L}}$ has been built, we can use the M/G/1 algorithm outlined in Sect. 2.1 to solve for its stationary probability vector $\boldsymbol{\alpha}$.

3.4 Combining the results from the upper and lower CTMCs

Once we computed the conditional probability vectors $\bar{\boldsymbol{\alpha}}$, of the states in \mathcal{U} , and $\boldsymbol{\alpha}$, of the states in \mathcal{L} , we can compute the overall probability distribution $\boldsymbol{\pi}$ of original CTMC.

Since g belongs to both \mathcal{U} and \mathcal{L} , we can scale $\bar{\alpha}$ and α so that the probability of g in both has the same value; then the entries for *any* state in \mathcal{U} or \mathcal{L} will be correctly scaled in relation to each other. For example, we can scale $\bar{\alpha}$ and α so that the entry for g assumes value 1 in both, i.e., divide $\bar{\alpha}$ and α by $\bar{\alpha}[g]$ and $\alpha[g]$, respectively. Then, the sum of the scaled entries corresponding to \mathcal{U} is $(1 - \bar{\alpha}[g])/\bar{\alpha}[g]$, while the sum of the scaled entries corresponding to \mathcal{L} is $(1 - \alpha[g])/\alpha[g]$. Thus, the overall sum of the scaled entries corresponding to \mathcal{U} , g , and \mathcal{L} is

$$s = \frac{1 - \bar{\alpha}[g]}{\bar{\alpha}[g]} + 1 + \frac{1 - \alpha[g]}{\alpha[g]} = \frac{\bar{\alpha}[g] + \alpha[g] - \bar{\alpha}[g]\alpha[g]}{\bar{\alpha}[g]\alpha[g]},$$

and we can express the entries of the stationary probability vector of the original CTMC as

$$\pi[i] = \begin{cases} \bar{\alpha}[i]/(\bar{\alpha}[g]s) & \text{if } i \in \mathcal{U} \\ \alpha[i]/(\alpha[g]s) & \text{if } i \in \mathcal{L} \\ 1/s & \text{if } i = g \end{cases}.$$

Of course, in practice, we compute only a finite portion of $\bar{\alpha}$ and α , hence of π .

3.5 Multiple upper or lower classes

In our exposition so far, we restricted ourselves to CTMCs where we can identify one upper and one lower set of states. However, it might be possible to further decompose the state space of certain models to fully exploit their beneficial structural characteristics.

For CTMCs with a strictly GI/M/1-type or M/G/1-type structure, this opportunity has been observed before. For example, if the repeating states $\mathcal{S}^{(j)}$ of a GI/M/1-type CTMC can be partitioned into two subsets $\mathcal{S}_a^{(j)}$ and $\mathcal{S}_b^{(j)}$ such that any path from $\mathcal{S}_b^{(j)}$ to $\mathcal{S}_a^{(j)}$ must visit $\mathcal{S}^{(0)}$, the probabilistic interpretation of \mathbf{R} implies that its submatrix corresponding to rows of “ b states” and columns of “ a states” is zero [10]. Analogously, the probabilistic interpretation of \mathbf{G} suggest that a similar property exists for the \mathbf{G} matrix in the case of M/G/1-type processes.




Figure 3: Multiple upper and lower sets.

Other types of simplifications are possible as well. For example, in [2, 3], we discuss the solution of special QBD and M/G/1-type CTMCs where the structure of the CTMC has another special property: all transitions from $\mathcal{S}^{(j)}$ to $\mathcal{S}^{(j-1)}$ are directed to a single return state $s_n^{(j-1)}$. This ensures that all the columns of matrix \mathbf{G} except for column n are zero, and allows for a very memory- and time-efficient computation of the expected measures of interest.

Here, we consider a special case of GI/G/1-type structure which can be seen as an extension of the one that can be managed by the algorithm we just introduced. Consider for example the structure shown in Fig. 3 on the left, where there are two sets of upper states, \mathcal{U}_a and \mathcal{U}_b , and two sets of lower states, \mathcal{L}_c and \mathcal{L}_d . This requires the identification of four gate states g_a , g_b , g_c , and g_d . The sets $\mathcal{U}_a^{(0)}$ and $\mathcal{U}_b^{(0)}$ contain any state that can be reached from g_a or g_b on a path to $\mathcal{U}_a^{(1)}$ or $\mathcal{U}_b^{(1)}$, respectively. Analogously, $\mathcal{L}_c^{(0)}$ and $\mathcal{L}_d^{(0)}$ contain any state that $\mathcal{L}_c^{(1)}$ or $\mathcal{L}_d^{(1)}$ can reach on a path to g_c or g_d , respectively. These four sets must be disjoint, except for the gate states themselves, since some or all of g_a , g_b , g_c , and g_d can coincide. $\mathcal{X}^{(0)}$ contains any remaining state: $\mathcal{X}^{(0)} = \mathcal{S}^{(0)} \setminus (\mathcal{U}_a^{(0)} \cup \mathcal{U}_b^{(0)} \cup \mathcal{L}_c^{(0)} \cup \mathcal{L}_d^{(0)})$. To extend our approach to this situation, we define:

1. Two upper GI/M/1-type CTMCs obtained using stochastic complementation on the

sets \mathcal{U}_a and \mathcal{U}_b , respectively.

2. One finite CTMC with state space $\mathcal{X}^{(0)} \cup \{A, B, g_c, g_d\}$ obtained using stochastic complementation on the set $\mathcal{U}_a \cup \mathcal{U}_b \cup \mathcal{X}^{(0)}$, then aggregating the set \mathcal{U}_a into the single macro-state A and the set \mathcal{U}_b into the single macro-state B .
3. Two lower M/G/1-type CTMCs obtained using pseudo-stochastic complementation on the sets \mathcal{L}_c and \mathcal{L}_d , respectively.

The order of the solution, however, is critical. First, we must solve the two upper CTMCs shown in portion (1) of the figure, to compute the conditional stationary probability vectors $\bar{\alpha}_a$ and $\bar{\alpha}_b$ of the states in \mathcal{U}_a and \mathcal{U}_b , respectively. These probabilities can then be used to compute the rates leaving the macro-states A and B in the CTMC with state space $\mathcal{X}^{(0)} \cup \{A, B, g_c, g_d\}$: for example the rate from A to g_c is $\bar{\alpha}_a \cdot \mathbf{Q}[\mathcal{U}_a, \mathcal{L}_c] \cdot \mathbf{1}^T$. After computing the stationary probability vector β for this finite CTMC, we can say that $[\beta[A]\bar{\alpha}_a | \beta[B]\bar{\alpha}_b]$ is proportional to the corresponding portion $\pi[\mathcal{U}_a \cup \mathcal{U}_b]$ of the stationary probability vector for the original CTMC; this is enough to compute the rates from g_c and g_d to \mathcal{L}_c and \mathcal{L}_d , respectively, using pseudo-stochastic complementation, for the CTMCs shown in portion (3) of the figure, since only states in \mathcal{U}_a and \mathcal{U}_b can reach states in \mathcal{L}_c and \mathcal{L}_d . Once these last two CTMCs are also solved, the solutions from the five CTMCs can be combined using the probability of the four gate states g_a, g_b, g_c , and g_d as reference points for the required scaling (while the finite CTMC does not explicitly have states g_a and g_b , their probabilities according to this CTMC can be obtained as $\beta[A]\bar{\alpha}_a[g_a]$ and $\beta[B]\bar{\alpha}_b[g_b]$).

Note that, unlike the special structures mentioned for the GI/M/1-type and the M/G/1-type CTMCs, which simply increase the efficiency of standard solution approaches, *the structure of Fig. 3 represents a true extension* which we can now solve.. This is because our approach with a single upper and a single lower set of states can be used to solve such structure only if g_a coincides with g_b and g_c coincides with g_d (in which case the role of the

single gate g could then be played indifferently by $g_a \equiv g_b$, provided $\mathcal{X}^{(0)}$ is considered part of the lower states, or by $g_c \equiv g_d$, provided $\mathcal{X}^{(0)}$ is considered part of the upper states). Of course, when our approach with a single upper and a single lower set can be applied, identifying further partitions of the upper set, the lower set, or both, is still nevertheless important, but simply because it allows a more efficient solution.

4 Application: parallel scheduling in the presence of failures

We now employ our method to solve a system that can be modeled as a GI/G/1-type CTMC. A popular way to allocate processors among competing applications in a parallel system environment is by space-sharing: processors are partitioned into disjoint sets and each application executes in isolation on one of these sets. Space-sharing can be done in a static, adaptive, or dynamic way. If a job requires a fixed number of processors for execution, this requires a static space-sharing policy [16]. Adaptive space-sharing policies [1] have been proposed for jobs that can configure themselves to the number of processors allocated by the scheduler at the beginning of their execution. Dynamic space-sharing policies [7] have been proposed for jobs that are flexible enough to allow the scheduler to reduce or increase the number of processors allocated to them as they run, in response to environment changes. Because of their flexibility, dynamic policies can offer optimal performance; however, they are also the most difficult to implement because they reallocate resources while applications are executing.

Modeling the behavior of scheduling policies in parallel systems often results in CTMCs with matrix geometric form [17]. In these models, neither failures nor arrivals from multiple exogenous sources are considered. To illustrate the applicability of our methodology, we

present instead a CTMC that models the behavior of a scheduling policy in a cluster environment subject to software and hardware failures. Our system is a cluster composed of two sub-clusters connected via a high speed medium (e.g., Gigabit Ethernet), while the nodes in each sub-cluster are connected via a lower speed switch (e.g., an ATM switch). For simplicity, we present a small system composed of two sub-clusters only, where a limited number of possible partitions is allowed, but our methodology readily applies to larger systems with multiple partitions. The arrival process is Poisson with parameter λ but each arrival may be a bulk of arbitrary size governed by a geometric distribution with parameter ρ . For clarity's sake, we only draw arcs corresponding to bulks of size one and two only, labeled with the rates $\lambda_1 = \lambda\rho$ and $\lambda_2 = \lambda(1 - \rho)\rho$, respectively.

The system employs a space-sharing policy that allows up to two parallel jobs to run simultaneously. A parallel job may execute across the whole cluster (i.e., on both sub-clusters) or occupy only one sub-cluster: the service time is exponentially distributed with rate μ_2 in the former case, or μ_1 in the latter. The policy is as follows. Upon an arrival

- while there are no jobs in the system:
 - if the arrival is of a single job, that is, the bulk size equals one, the whole cluster is assigned to that job,
 - otherwise, if multiple jobs arrive simultaneously, that is, the bulk size is greater than one, two jobs are scheduled, one on each sub-cluster, while the remaining jobs in the bulk, if any, are queued;
- while there are already jobs in the system and one job running using the whole cluster:
 - if the bulk size equals one, the job is simply queued,
 - otherwise, if the bulk size is greater than one, the arriving jobs are enqueued, the current job is stopped and restarted on a single sub-cluster, and one of the queued

jobs is also started, on the other sub-cluster;

- while there are jobs in the system already and two jobs are running, each on one sub-cluster:
 - we simply enqueue arriving jobs, regardless of the bulk size.

Upon a completion (departure)

- of a job that was using the entire system:
 - if there is only one job waiting, it is assigned the entire system,
 - if there are multiple jobs waiting, two of them are assigned a sub-cluster each;
- of a job that was using only one sub-cluster:
 - one of the jobs in the queue, if any, is assigned the sub-cluster just released (if there are no jobs waiting, the other job running, if any, is *not* reassigned to the entire system).

The rationale behind these decisions is that, under bursty conditions, we would like to reach our goal of having smaller partitions quickly, even at the cost of killing a running job and rescheduling it on a smaller partition [15].

We consider the performance of our scheduling policy under the following failure scenarios. Each of the two sub-clusters can experience a local hardware failure independently of each other; in this case, only the affected sub-cluster must be brought down for repair, and the system can still accept arrivals. In addition, when a parallel job is assigned the entire cluster, it makes use of software whose execution can cause the entire system to crash. After such a failure, each sub-cluster must be brought down for reboot; consequently, the system's queue is flushed and no arrivals are accepted until repairs are completed. We assume that all event durations are exponentially distributed. The rate for a hardware failure is f_h , for




Figure 4: The CTMC for our performability study.

State	Description
0	Empty system, no failures
$0S$	Empty system, rebooting after a software failure
kW	$k \geq 0$ jobs in the queue, one job executing on the whole system
$0H$	0 jobs in the queue, one sub-cluster idle, the other running a job
kHH	$k \geq 0$ jobs in the queue, each sub-cluster running a job
$0F$	0 jobs in the queue, one sub-cluster idle, the other failed
kHF	$k \geq 0$ jobs in the queue, one sub-cluster running a job, the other failed
kFF	$k \geq 0$ jobs in the queue, both sub-cluster failed

Table 1: Meaning of the CTMC states.

a software failure is f_s , for a repair after a hardware failure is r_h , and for a repair (reboot) after a software failure is r_s .

Fig. 4 depicts a CTMC modeling the behavior of our system and Table 1 describes the meaning of the system states. The infinitesimal generator of this CTMC has a GI/G/1 struc-

ture consistent with the requirements of our solution algorithm, since we can immediately identify state 0 as the gate state, states kW plus $0S$ as the upper states, and states $0H$, kHH , $0F$, kHF and kFF as the lower states. Note that the model assumes that hardware failures do not occur during reboot; if this were not true, we could add a transition from $0S$ to $0F$ with rate $2f_h$, and the structure required by our approach would still be present, since there can be transitions from $\mathcal{U}^{(0)}$ (i.e., $\{0S\}$) to any state in \mathcal{L} .

The numerical values used for our model parameters are: $f_h = 10^{-6}..10^{-3}$, $f_s = 10 \cdot f_h$, $r_h = 10^{-3}$ or 10^{-1} , $r_s = 10 \cdot r_h$, $\mu_1 = 3.0$ or 6.0 , $\mu_2 = 1.8 \cdot \mu_1$, $\lambda = 10^{-5}..5$, and $\rho = 0.8$. Fig. 5(a) and 5(b) show the long-term availability of the system, that is, its ability of accepting arrivals, computed as the stationary probability of being in any of the states except $0S$ and kFF , for $k \geq 0$, for a given choice of r_h and two choices of μ_1 , and for a given choice of μ_1 and two choices of r_h , respectively, as a function of f_h and λ (note that the arrival rate λ does affect the availability, since it affects the probability of the system being in the kW states, hence of software failures). Fig. 5(c) and 5(d) are analogous, except that they focus on the probability on not being unavailable due to a software failure, that is, they plot the complement of the probability of being in state $0S$. Fig. 5(e) and 5(f) focus instead on the system power, defined as the ratio of system throughput over average job response time, also as a function of f_h and λ and for various choices of r_h and μ_1 . It is apparent that there is a significant correlation between the above parameters and the workload that the system can handle. Note that, in all plots, the missing data points correspond to parameter combinations for which the system becomes unstable.

5 Conclusion

In this paper we presented an exact decomposition algorithm for the analysis of a class of GI/G/1-type Markov chains with repetitive structure. Such chains appear when modeling

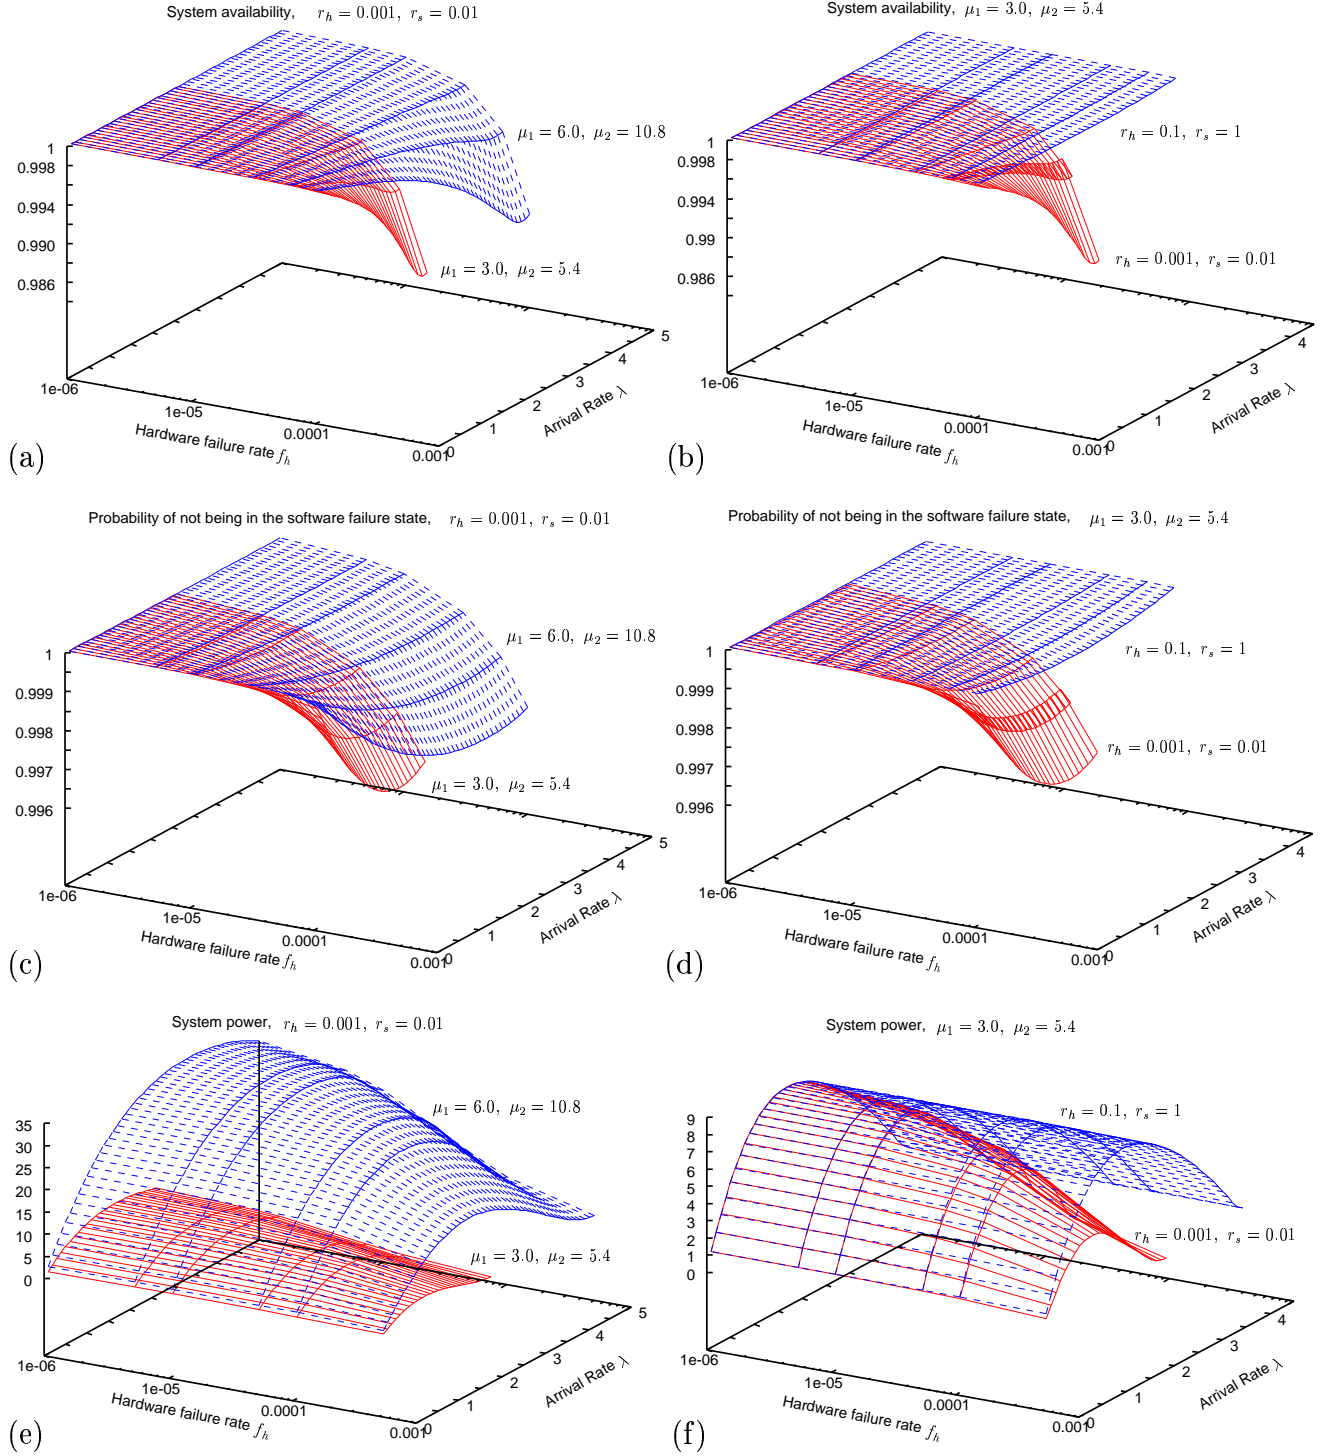


Figure 5: Performance results.

systems with failures and repairs that also accept arrivals from multiple exogenous sources.

In particular, we considered the performability of a parallel task scheduling policy in a cluster

system subject to failures and repairs.

The conditions required for the application of our algorithm include the presence of a special gate state, which imposes a certain structure on the graph describing the Markov chain. The identification of this gate state and the partition of the overall state space into the upper and lower sets is a graph partitioning problem, which will be the subject of future work.

In many cases, however, the nature of the system under examination may immediately guide us into identifying the possible state space partition. Then, our algorithm is easily applied and allows the modeler to compute stationary measures for Markov chains previously thought not solvable by analytical methods.

References

- [1] S.-H. Chiang, R. K. Mansharamani, and 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*, pages 33–44, May 1994.
- [2] G. Ciardo, A. Riska, and E. Smirni. An aggregation-based solution method for M/G/1-type processes. In B. Plateau, W. J. Stewart, and M. Silva, editors, *Numerical Solution of Markov Chains*, pages 21–40. Prensas Universitarias de Zaragoza, Zaragoza, Spain, Sept. 1999.
- [3] G. Ciardo and E. Smirni. ETAQA: an efficient technique for the analysis of QBD-processes by aggregation. *Performance Evaluation*, (36-37):71–93, 1999.

- [4] W. K. Grassmann and D. A. Stanford. Matrix analytic methods. In W. K. Grassmann, editor, *Computational Probability*, pages 153–204. Kluwer Academic Publishers, Boston, MA, 2000.
- [5] G. Latouche. Algorithms for infinite Markov chains with repeating columns. In C. Meyer and R. J. Plemmons, editors, *Linear Algebra, Markov Chains, and Queueing Models*, volume 48 of *IMA Volumes in Mathematics and its Applications*, pages 231–265. Springer-Verlag, 1993.
- [6] G. Latouche and V. Ramaswami. *Introduction to Matrix Geometric Methods in Stochastic Modeling*. ASA-SIAM Series on Statistics and Applied Probability. SIAM, Philadelphia PA, 1999.
- [7] C. McCann, R. Vaswani, and J. Zahorjan. A dynamic processor allocation policy for multiprogrammed shared memory multiprocessors. *ACM Transactions on Computer Systems*, 11(2):146–178, 1993.
- [8] B. Meini. Solving M/G/1 type Markov chains: Recent advances and applications. *Stochastic Models*, 14(1 & 2):479–496, 1998.
- [9] C. D. Meyer. Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems. *SIAM Review*, 31(2):240–271, June 1989.
- [10] R. Nelson. *Probability, Stochastic Processes, and Queueing Theory*. Springer-Verlag, 1995.
- [11] M. F. Neuts. *Matrix-geometric solutions in stochastic models*. Johns Hopkins University Press, Baltimore, MD, 1981.
- [12] M. F. Neuts. *Structured stochastic matrices of M/G/1 type and their applications*. Marcel Dekker, New York, NY, 1989.

- [13] V. Ramaswami. A stable recursion for the steady state vector in Markov chains of M/G/1 type. *Commun. Statist.– Stochastic Models*, 4:183–263, 1988.
- [14] V. Ramaswami and G. Latouche. A general class of Markov processes with explicit matrix-geometric solutions. *OR Spektrum*, 8:209–218, Aug. 1986.
- [15] E. Rosti, E. Smirni, G. Serazzi, L. W. Dowdy, and K. C. Sevcik. On processor saving scheduling policies for multiprocessor systems. *IEEE Trans. Comp.*, 47(2):178–189, Feb. 1998.
- [16] E. Smirni, E. Rosti, L. W. Dowdy, and G. Serazzi. A methodology for the evaluation of multiprocessor non-preemptive allocation policies. *Journal of Systems Architecture*, 44:703–721, 1998.
- [17] M. S. Squillante, F. Wang, and M. Papaefthymiou. Stochastic analysis of gang scheduling in parallel and distributed systems. *Perf. Eval.*, 27/28:273–296, 1996.

A Proof of Theorem 2

It is easy to see that \mathbf{A} is an infinitesimal generator: its off-diagonal entries are non-negative by construction, and its rows sum to zero because:

$$\begin{aligned}
 \mathbf{A}\mathbf{1}^T &= \left(\mathbf{Q}[\mathcal{A}, \mathcal{A}] + \mathbf{Q}[\mathcal{A}, \overline{\mathcal{A}}]\mathbf{1}^T \text{Norm}(\boldsymbol{\pi}[\overline{\mathcal{A}}]\mathbf{Q}[\overline{\mathcal{A}}, \mathcal{A}]) \right) \mathbf{1}^T \\
 &= \mathbf{Q}[\mathcal{A}, \mathcal{A}]\mathbf{1}^T + \mathbf{Q}[\mathcal{A}, \overline{\mathcal{A}}]\mathbf{1}^T \text{Norm}(\boldsymbol{\pi}[\overline{\mathcal{A}}]\mathbf{Q}[\overline{\mathcal{A}}, \mathcal{A}])\mathbf{1}^T \\
 &= \mathbf{Q}[\mathcal{A}, \mathcal{A}]\mathbf{1}^T + \mathbf{Q}[\mathcal{A}, \overline{\mathcal{A}}]\mathbf{1}^T \\
 &= \mathbf{Q}[\mathcal{A}, \mathcal{S}]\mathbf{1}^T = \mathbf{0}^T.
 \end{aligned}$$

The irreducibility of the pseudo stochastic complement of \mathcal{A} follows directly from its probabilistic interpretation and the fact that the original process is irreducible. We now

need to show that $\boldsymbol{\alpha} = \boldsymbol{\pi}[\mathcal{A}]/(\boldsymbol{\pi}[\mathcal{A}]\mathbf{1}^T)$ satisfies $\boldsymbol{\alpha}\mathbf{A} = \mathbf{0}$. Starting from the definition of pseudo stochastic complement, we have:

$$\begin{aligned}
\boldsymbol{\alpha}\mathbf{A} &= \frac{\boldsymbol{\pi}[\mathcal{A}]}{\boldsymbol{\pi}[\mathcal{A}]\mathbf{1}^T} \left(\mathbf{Q}[\mathcal{A}, \mathcal{A}] + \mathbf{Q}[\mathcal{A}, \overline{\mathcal{A}}]\mathbf{1}^T \text{Norm}(\boldsymbol{\pi}[\overline{\mathcal{A}}]\mathbf{Q}[\overline{\mathcal{A}}, \mathcal{A}]) \right) \\
&= \frac{1}{\boldsymbol{\pi}[\mathcal{A}]\mathbf{1}^T} \left(\boldsymbol{\pi}[\mathcal{A}]\mathbf{Q}[\mathcal{A}, \mathcal{A}] + \boldsymbol{\pi}[\mathcal{A}]\mathbf{Q}[\mathcal{A}, \overline{\mathcal{A}}]\mathbf{1}^T \frac{\boldsymbol{\pi}[\overline{\mathcal{A}}]\mathbf{Q}[\overline{\mathcal{A}}, \mathcal{A}]}{\boldsymbol{\pi}[\overline{\mathcal{A}}]\mathbf{Q}[\overline{\mathcal{A}}, \mathcal{A}]\mathbf{1}^T} \right) \\
&= \frac{1}{\boldsymbol{\pi}[\mathcal{A}]\mathbf{1}^T} \left(\boldsymbol{\pi}[\mathcal{A}]\mathbf{Q}[\mathcal{A}, \mathcal{A}] + \boldsymbol{\pi}[\overline{\mathcal{A}}]\mathbf{Q}[\overline{\mathcal{A}}, \mathcal{A}] \right) = \mathbf{0},
\end{aligned}$$

where the last two steps are obtained considering that $\boldsymbol{\pi}[\mathcal{A}]\mathbf{Q}[\mathcal{A}, \overline{\mathcal{A}}]\mathbf{1}^T = \boldsymbol{\pi}[\overline{\mathcal{A}}]\mathbf{Q}[\overline{\mathcal{A}}, \mathcal{A}]\mathbf{1}^T$ (they represent the flow from \mathcal{A} to $\overline{\mathcal{A}}$ and from $\overline{\mathcal{A}}$ to \mathcal{A} in steady state, respectively) and that $\boldsymbol{\pi}[\mathcal{A}]\mathbf{Q}[\mathcal{A}, \mathcal{A}] + \boldsymbol{\pi}[\overline{\mathcal{A}}]\mathbf{Q}[\overline{\mathcal{A}}, \mathcal{A}] = \mathbf{0}$ (since $\boldsymbol{\pi}$ is the stationary probability vector). \square