An aggregation-based method for the exact analysis of a class of GI/G/1-type processes*

Alma Riska
Seagate Research
Pittsburgh, PA
Alma.Riska@seagate.com

Evgenia Smirni Gianfranco Ciardo
College of William and Mary
Williamsburg, VA
{esmirni, ciardo}@cs.wm.edu

ABSTRACT

We present an aggregation-based algorithm for the exact analysis of Markov chains with GI/G/1-type pattern in their repetitive structure, i.e., chains that 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 faults/repairs that accept jobs from multiple exogenous sources. Our method provides exact computation of the steady state probabilities, and allows computation of performance measures of interest including the system queue length or any of its higher moments, the exact probability of system failures and repairs, and consequently a host of performability measures. Our algorithm also applies to systems that are purely of the M/G/1-type or the GI/M/1-type, or their intersection, i.e., quasi-birth-death processes.

Keywords: Markov chains; GI/G/1-type processess;M/G/1-type processes; GI/M/1-type processes; matrix-analytic techniques; stochastic complementation; reliability analysis.

Introduction

During the past two decades, significant effort has been put into the development of modeling tools that can capture the behavior of modern computer and communication systems. In many cases, the behavior of such systems can be captured by M/G/1-type or GI/M/1-type Markov chains, and their generalizations (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.

Here, we study a class of CTMCs that show a GI/G/1-type pattern in their repetitive structure and exhibit *both* M/G/1-type and GI/M/1-type patterns, and cannot be solved with existing techniques. Such chains occur when modeling open systems that accept customers from multiple exogenous sources (thus the existence of bulk arrivals) and are also subject to failures and repairs (since the system may empty-out in a single step when a catastrophic failure occurs or only parts of it may be operational when a non-catastrophic failure occurs). The state space[1] $\mathcal{S}$ of a GI/G/1 CTMC can be partitioned into a finite "boundary" set $\mathcal{S}^{(0)} = \{s_1^{(0)}, \ldots, s_m^{(0)}\}$ and a countably infinite sequence of finite "level" sets $\mathcal{S}^{(j)} = \{s_1^{(j)}, \ldots, s_n^{(j)}\}$, $j \geq 1$.

The generator matrix can accordingly be block-partitioned as

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

(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 "ˆ" 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. Since the same local and forward blocks appear from the third row block on, $RowSum(\widehat{\mathbf{B}}^{(j)}) + RowSum(\mathbf{B}^{(j-1)} + \cdots + \mathbf{B}^{(1)})$ 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 $RowSum(\mathbf{B}^{(1)} + \cdots + \mathbf{B}^{(j)})$, for any finite j, is bounded by $RowSum(-\mathbf{L} - \sum_{j=1}^{\infty} \mathbf{F}^{(j)})$. However, $\mathbf{L}^{(1)}$ can differ from $\mathbf{L}$ in the diagonal, hence $RowSum(\widehat{\mathbf{B}}^{(1)})$ might be different from $RowSum(\widehat{\mathbf{B}}^{(j)}) + RowSum(\mathbf{B}^{(j-1)} + \cdots + \mathbf{B}^{(1)})$, 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 $\boldsymbol{\pi}^{(0)} \in \mathbb{R}^m$ and $\boldsymbol{\pi}^{(j)} \in \mathbb{R}^n$, for $j \geq 1$. 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 [2], 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, forward jumps are allowed only to the next level) and the M/G/1-type CTMCs (the converse: $\widehat{\mathbf{B}}^{(j)}$ and $\mathbf{B}^{(j)}$ are zero for $j \geq 2$, backward jumps are allowed only to the previous level). Neuts proposed the elegant *matrix geometric* solution [7] for the former, while the latter can be solved using one of the *matrix analytic*-based methods [3, 8, 9]. 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 solution is preferred because it is simpler, more widely-known, and at least as efficient). In practice, such CTMCs often arise when ob-

*This work was partially supported by the National Science Foundation under grants EIA-9974992, CCR-0098278, ACI-0090221, and ACI-0203971.

[1] Throughout this exposition we use calligraphic letters to indicate sets ($\mathcal{A}$), lower case boldface roman or greek letters to indicate vectors ($\mathbf{a}$, $\boldsymbol{\alpha}$), and upper case boldface roman letters to indicate matrices ($\mathbf{A}$). We use superscripts in parentheses or subscripts to indicate family of related entities ($\mathcal{A}^{(1)}$, $\mathbf{A}_{\mathcal{A}_1}$). Vector and matrix elements are indicated using square brackets ($\mathbf{a}[1]$, $\mathbf{A}[1,2]$), and we extend the notation to subvectors or submatrices by allowing sets of indices to be used instead of single indices ($\mathbf{a}[\mathcal{A}]$, $\mathbf{A}[\mathcal{A}, \mathcal{B}]$).

serving an open system with a single infinite-waiting-room queue at either the arrival or the service completion times.

Our contribution

We propose a decomposition approach that extends both the applicability and the efficiency of matrix analytic methods through an intelligent *partitioning* of the repetitive portion of the state space into subsets, according to its connectivity. 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 [5].

More specifically, our main contribution is a decomposition algorithm that can be used for the performability analysis of complex systems modeled as GI/G/1-type processes. Since such processes can capture both the normal behavior (arrivals and service) and the reliability aspects (failures and repairs) of system operation, they provide the necessary means to model a system's performability. In particular, GI/G/1-type processes are the tool of choice if one needs to analyze the user-perceived performance of a computer system. In such cases, the normal operation of the system corresponds to the M/G/1-type pattern of the GI/G/1-type model, while software or hardware failures correspond to the GI/M/1-type pattern of the GI/G/1-type model. Our method exploits this structure by reflecting a key performance component of the state (the number of customers in the system) in the level j of the sets $\mathcal{S}^{(j)}$, while the remaining component (the status of the service and repair processes, and the inter-arrival and service distributions, if not memoryless) are captured *within* the states of a level set. Only the latter components truly affect the solution complexity of our approach.

With regard to degradable behavior, we observe that the *reward structure* assigned to the CTMC states is very flexible, since it allows to define the reward, or "usefulness", of a state $s_i^{(j)}$ in terms of both its index i, describing the failure status of each component, and of the level j, describing the present workload in the system.

Since general GI/G/1-type models do not have a known exact solution [2], they are usually studied through approximations based on QBD processes or complex eigenvalue methods [1]. Informally, a QBD approximation truncates the "arbitrary forward and backward jump" behavior, but the resulting process is orders of magnitude larger than with the approach we propose, thus its solution is much more computationally expensive. Our methodology based on aggregation/decomposition techniques is applicable to an important subset of GI/G/1-type processes and provides instead *exact* solutions. Since we approach the problem in a divide-and-conquer way, we ensure that the subproblems are smaller in size and computationally more efficient to solve than the approximated original problem.

Overall approach

In this section, we present the high-level idea of our algorithm and outline the the structure that the CTMCs must have to be solved with the proposed technique. We observe that the pattern of interaction among states of a CTMC with infinitesimal generator $\mathbf{Q}$ given by Eq.(1) is the union of the patterns for GI/M/1-type and M/G/1-type processes, thus it is more general than either. Based on this observation, we propose the following solution steps (Fig. 1):

1. Partition the union of the level sets $\mathcal{S} = \bigcup_{j=1}^{\infty} \mathcal{S}^{(j)}$ into two disjoint sets $\mathcal{U}$ and $\mathcal{L}$ such that $\mathcal{U}$ captures the GI/M/1-type behavior of $\mathbf{Q}$ and $\mathcal{L}$ captures the M/G/1-type behavior of $\mathbf{Q}$.

2. Use the well-known concept of stochastic complementation [5] to define two new processes (stochastic complements), one containing all states in $\mathcal{U}$ (plus a special "gate" state g) and one containing all states in $\mathcal{L}$ (plus state g).

3. Solve each new process using well-known techniques, and obtain the conditional stationary probabilities for all states in $\mathcal{U}\cup g$ (or $\mathcal{L}\cup g$) given that the original process $\mathbf{Q}$ is in $\mathcal{U}\cup g$ (or $\mathcal{L}\cup g$, respectively). In particular, the stochastic complement of the set $\mathcal{U} \cup g$ is a GI/M/1-type process that is solved with the matrix geometric method [7] and the stochastic complement of $\mathcal{L} \cup g$ is an M/G/1-type process that is solved using the matrix analytic method [3, 8].

4. Finally, "couple" the two solutions by scaling back the conditional state probabilities of all states in $\mathcal{U}\cup g$ and $\mathcal{L}\cup g$, and obtain the stationary probabilities for the original process.

The required block structure for $\mathbf{Q}$ is shown in Fig. 2. We observe a "two-level" repetitive structure, where each level set $\mathcal{S}^{(j)}$, $j \geq 1$, is partitioned into two disjoint classes denoted $\mathcal{U}^{(j)}$ (for "upper") and $\mathcal{L}^{(j)}$ (for "lower"). Let $\mathcal{U} = \bigcup_{j=1}^{\infty} \mathcal{U}^{(j)}$ and $\mathcal{L} = \bigcup_{j=1}^{\infty} \mathcal{L}^{(j)}$. The following list summarizes the interactions within each set and across the two sets.

- Within set $\mathcal{U}$, forward transitions are allowed from any $\mathcal{U}^{(j)}$ only toward the next level $\mathcal{U}^{(j+1)}$. Backward transitions are allowed from set $\mathcal{U}^{(j)}$ to any lower level sets $\mathcal{U}^{(k)}$, $k < j$.

- Within set $\mathcal{L}$, forward transitions are allowed from any $\mathcal{L}^{(j)}$ toward any higher level $\mathcal{L}^{(k)}$, $k > j$. Backward transitions are allowed from set $\mathcal{L}^{(j)}$ to $\mathcal{L}^{(j-1)}$ only.

- Local transitions (not shown) are allowed within each $\mathcal{U}^{(j)}$ and $\mathcal{L}^{(j)}$.

- Transitions from $\mathcal{U}^{(j)}$ to any $\mathcal{L}^{(k)}$, $k \geq 1$ are allowed.

- There is strictly no interaction from $\mathcal{L}$ toward $\mathcal{U}$ (except, of course, through the boundary portion $\mathcal{S}^{(0)}$).

- There is a special "gate" state g in $\mathcal{S}^{(0)}$ such that any path from $\mathcal{L}$ to $\mathcal{U}$ must visit state g. In practice, such a 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 "shift all levels to the left by one".

The gated structure for the interaction of $\mathcal{U}$ and $\mathcal{L}$ is critical for our algorithm, as it allows us to apply stochastic complementation in a special setting. Furthermore, in conjunction with the upper/lower interaction between sets $\mathcal{U}$ and $\mathcal{L}$, it ensures that (a) the stochastic complement of the upper set of states is a GI/M/1-type CTMC and (b) the stochastic complement of the lower set of states is a M/G/1-type CTMC.

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 is 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.

For more details on our methodology including the formulation of the conditions under which our approach is feasible as well as an example from the area of realiability analysis of parallel computer systems, we direct the interested reader to [10].

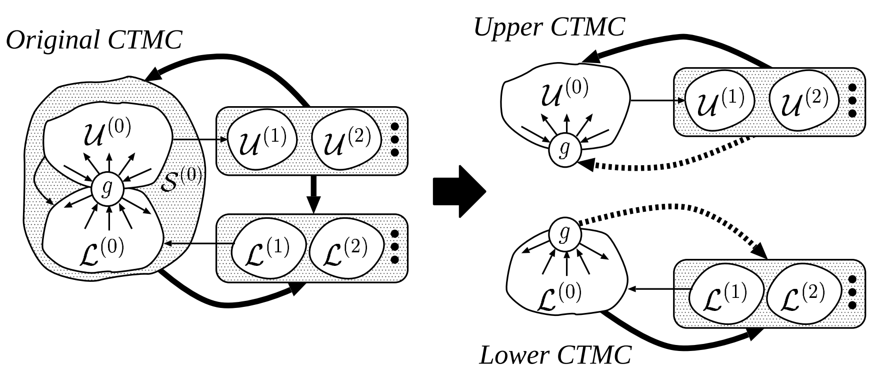

Figure 1: The overall idea of our approach.

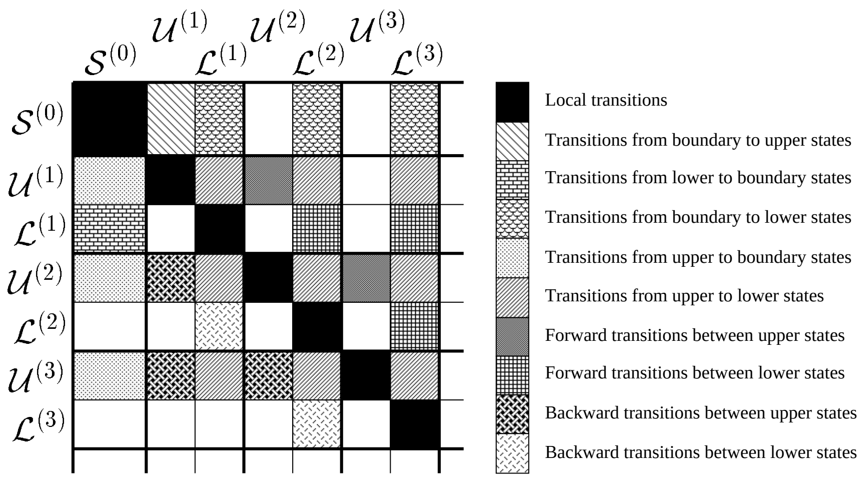

Figure 2: The nonzero pattern in our matrix Q.

1. REFERENCES

[1] H. R. Gail, S. L. Hantler, and B. A. Taylor. Use of characteristic roots for solving infinite state markov chains. In W. K. Grassman, editor, *Computational Probability*, pages 205–255, Boston, MA, 2000. Kluwer Academic Publishers.

[2] 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.

[3] 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.

[4] B. Meini. Solving M/G/1 type Markov chains: Recent advances and applications. *Stochastic Models*, 14(1 & 2):479–496, 1998.

[5] C. D. Meyer. Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems. *SIAM Review*, 31(2):240–271, June 1989.

[6] J. F. Meyer. On evaluating the performability of degradable computing systems. *IEEE Trans. Comp.*, C-29(8):720–731, Aug. 1980.

[7] M. F. Neuts. *Matrix-geometric Solutions in Stochastic Models*. Johns Hopkins University Press, Baltimore, MD, 1981.

[8] M. F. Neuts. *Structured Stochastic Matrices of M/G/1 Type and Their Applications*. Marcel Dekker, Inc., New York, NY, 1989.

[9] A. Riska and E. Smirni. Exact aggregate solutions for M/G/1-type Markov processes. In *Proceedings of ACM SIGMETRICS Conference*, pages 86–96, Marina del Rey, CA, June 2002.

[10] A. Riska, E. Smirni, and G. Ciardo. Exact analysis of a class of GI/G/1 performability models. *IEEE Trans. on Reliability* Dec. 2003, (to appear). Preprint available at http://www.cs.wm.edu/˜esmirni/docs/pubs.html.