

Exploiting restricted transitions in Quasi-Birth-and-Death processes

Juan F. Pérez and Benny Van Houdt

Performance Analysis of Telecommunication Systems Research Group

Department of Mathematics and Computer Science

University of Antwerp - IBBT

Middelheimlaan 1, B-2020 Antwerp, Belgium

Email: {juanfernando.perez, benny.vanhoudt}@ua.ac.be

Abstract—In this paper we consider Quasi-Birth-and-Death (QBD) processes where the upward (resp. downward) transitions are restricted to occur only from (resp. to) a subset of the phase space. This property is exploited to reduce the computation time to find the matrix R or G of the process. The reduction is done through the definition of a censored process which can be of the M/G/1- or GI/M/1-type. The approach is illustrated through examples that show the applicability and benefits of making use of the additional structure. The examples also show how these special structures arise naturally in the analysis of queuing systems. Even more substantial gains can be realized when we further restrict the class of QBD processes under consideration.

I. INTRODUCTION

Quasi-Birth-and-Death (QBD) processes are a generalization of simple birth-and-death processes where the addition of a second dimension, called the *phase*, allows the representation of more general systems. The second dimension typically describes a random environment or the state of the arrival and service processes in the case of queueing systems. QBD Markov chains (MCs) were introduced in [1] and studied in [2] as part of the more general class of GI/M/1-type MCs. The main feature of a QBD process is that its stationary probability vector, if it exists, has a matrix-geometric form, such that it can be expressed as function of a boundary probability vector and a *rate* matrix R . Many algorithms have been proposed for finding this matrix [2], [3], [4], which is the minimal nonnegative solution of a quadratic matrix equation.

In many cases the state space of the second dimension may be large enough to require long computation times to find the matrix R , even when relying on quadratically convergent algorithms such as Logarithmic Reduction [3] and Cyclic Reduction [4]. Provided that the matrices characterizing the QBD MC have a certain structure, Grassmann and Tavakoli [5] show how these computation times can be reduced. In particular, they considered the case where upward transitions can only occur in a certain subset of the phase space, using the resulting structure to accelerate the time per iteration of the linearly-convergent U-based method [6]. Their approach can also take into account the case where the downward transition can lead only to a subset of the phase space. In this paper we consider the same case as in [5], restricting the upward

(resp. downward) transitions to occur only from (resp. to) a subset of the phase space, called \mathcal{S}^+ . However, our approach differs from [5] in that we define a new process by observing the QBD MC only when the phase variable is in the set \mathcal{S}^+ . The resulting processes are of the M/G/1- and GI/M/1-type, depending on whether the restriction is on the downward or upward transitions, respectively. Since the block size of the new process is of the same size as \mathcal{S}^+ , the reduction in computation time is more significant when the ratio between the total size of the phase space and the size of \mathcal{S}^+ is larger. Additionally, once the parameters of the new process are computed we can use a quadratically convergent algorithm like Cyclic Reduction [4] to find the matrix R_+ of the new process, which is shown to be a submatrix of the R matrix of the original process, while the remaining entries are retrieved by solving a Sylvester matrix equation. Finally, we also show how additional restrictions on the transition probabilities lead to further reductions on the computation times.

There are many examples of queueing systems where the restricted transitions considered here arise. In an overflow queue [7] the arrivals only occur when the first queue is full. Hence the upward transitions are restricted to take place in those states corresponding to a full first queue. This case is similar to that when the inter-arrival times follow an Erlang distribution, since the arrivals can only happen in the states related to the last phase of the Erlang distribution. On the other hand, priority queues [8], [9] illustrate the case of restricted downward transitions. Here the downward transitions are associated to the service completion of a low-priority customer and these transitions only occur in those states where there are no high-priority customers. Another case is the QBD MC used in [10] to compute the waiting time distribution of a type- k customer in an MMAP[K]/PH[K]/ c ($c = 1, 2$) queue, where the downward transitions can only lead to a small subset of the phase space. Some of these examples are used to illustrate the behavior of the proposed approach compared to the traditional methods and to the approach in [5]. The method proposed here has been implemented in MATLAB and will be available online as part of the SMCSolver tool [11].

The paper is organized as follows: Section II gives a brief review of QBD processes. In sections III and IV we describe how to exploit restricted downward and upward transitions in

these processes. Section V provides some numerical examples of queueing systems with this kind of transitions and includes a comparison with a number of existing computational methods.

II. QBDs WITH RESTRICTED TRANSITIONS

A discrete-time QBD MC can be defined as a two-dimensional process $\{(N_t, X_t), t \geq 0\}$, where N_t is called the level variable and takes values on \mathbb{N} . The phase variable X_t takes values on the set $\{1, 2, \dots, m_0\}$ or $\{1, 2, \dots, m\}$ depending on whether the level is equal to or greater than 0. The level variable can only increase or decrease its value by one at each time epoch and these transition probabilities are level-independent. Therefore the QBD MC has a transition matrix P of the form

$$P = \begin{bmatrix} B_1 & B_2 & 0 & 0 & \dots \\ B_0 & A_1 & A_2 & 0 & \dots \\ 0 & A_0 & A_1 & A_2 & \dots \\ \vdots & \ddots & \ddots & \ddots & \ddots \end{bmatrix},$$

where B_1 and A_1 are square matrices of size m_0 and m , respectively. The matrices B_1 and B_2 hold the transition probabilities from level 0 to levels 0 and 1, respectively, and the matrix B_0 contains the transition probabilities from level 1 to level 0. Similarly, the matrices A_0 , A_1 and A_2 carry the transition probabilities from level i to level $i-1$, i and $i+1$, respectively, for $i > 0$. The key when computing the steady state probability vector $\pi = [\pi_0, \pi_1, \pi_2, \dots]$ of P , if it exists, is to find the minimal nonnegative solution R of the matrix equation

$$R = A_2 + RA_1 + R^2A_0. \quad (1)$$

The vectors π_i can then be computed as $\pi_i = \pi_1 R^i$, for $i > 1$, where $[\pi_0, \pi_1]$ is the solution of the boundary equation

$$[\pi_0, \pi_1] \begin{bmatrix} B_1 & B_2 \\ B_0 & A_1 + RA_0 \end{bmatrix} = [\pi_0, \pi_1].$$

Another way to find the matrix R is from $R = A_1(I - A_0 - A_1G)^{-1}$, where G is the minimal nonnegative solution of the matrix equation

$$G = A_0 + A_1G + A_0G^2. \quad (2)$$

The method introduced in this paper aims at computing either the matrix R or G , from which the stationary probability vector can be obtained. Moreover, our approach is actually independent of the behavior of the QBD near the boundary at level 0. Therefore a more general boundary behavior can be assumed as long as the QBD shows a repeating structure (matrices A_0 , A_1 and A_2) from a given level onward.

Many iterative algorithms have been developed to solve equations (1) and (2), including quadratically convergent algorithms such as Logarithmic Reduction [3] and Cyclic Reduction (CR) [4]. However, a large block size m may turn the solution of these equations into a lengthy task, as each iteration requires $O(m^3)$ time. In this paper we consider two special cases where the structure of the matrices A_0 and A_2 can be exploited to speed-up the computation of the matrix

G or R . In both cases we consider a partition of the set $\{1, \dots, m\}$ into two sets: \mathcal{S}^+ containing the first r phases, and \mathcal{S}^- containing the remaining $m-r$ phases. Using this partition the matrices A_i , for $i = \{0, 1, 2\}$, can be written as

$$A_i = \begin{bmatrix} A_i^{++} & A_i^{+-} \\ A_i^{-+} & A_i^{--} \end{bmatrix}, \quad (3)$$

where A_i^{++} and A_i^{--} are square matrices of size r and $m-r$, respectively. In Section III we consider the case where downward transitions can only occur to a state with phase in \mathcal{S}^+ , hence the matrix A_0 has only $r \ll m$ nonzero columns such that it can be written as

$$A_0 = \begin{bmatrix} A_0^{++} & 0 \\ A_0^{-+} & 0 \end{bmatrix}. \quad (4)$$

When the set \mathcal{S}^+ contains only one phase the matrix G can be computed explicitly without the need of resorting to iterative algorithms [12]. Furthermore, this particular case has also been exploited to compute performance measures in an efficient manner without computing all the terms of the vector π [13]. In this paper we consider the more general case where the cardinality of \mathcal{S}^+ is greater than one, meaning that the matrix G is not known explicitly from the parameters of the QBD.

The analogous case where upward transitions only occur in a state with phase in \mathcal{S}^+ is treated in Section IV. In this case the matrix A_2 has only $r \ll m$ nonzero rows, i.e.

$$A_2 = \begin{bmatrix} A_2^{++} & A_2^{+-} \\ 0 & 0 \end{bmatrix}. \quad (5)$$

This structure was analyzed by Grassmann and Tavakoli in [5], where it was exploited to reduce the computation time per iteration in the so-called U-algorithm [6], which computes a matrix U such that $R = A_2(I - U)^{-1}$. The algorithm starts with $U_0 = A_1$ and iteratively computes $U_{k+1} = A_1 + A_2(I - U_k)^{-1}A_0$, such that the iterates converge to the actual value of the matrix U . Even though the approach proposed in [5] provides an important computational gain per iteration, the number of iterations required may be large since this is a linearly convergent algorithm [14]. In Section V we consider an example with the structure described by (5) and compare the performance of our approach with the one proposed in [5]. The Grassmann and Tavakoli method can also be adapted to the case where the matrix A_0 has the form in (4).

A. Markov chains of the M/G/1- and GI/M/1-type

An M/G/1-type MC [15] can be seen as a generalization of a QBD MC, where the level is allowed to increase its value by more than one in a single transition. Therefore, the transition matrix \bar{P} of an M/G/1-type MC is of the form

$$\bar{P} = \begin{bmatrix} \bar{B}_0 & \bar{B}_1 & \bar{B}_2 & \bar{B}_3 & \dots \\ \bar{A}_0 & \bar{A}_1 & \bar{A}_2 & \bar{A}_3 & \dots \\ & \bar{A}_0 & \bar{A}_1 & \bar{A}_2 & \dots \\ & & \bar{A}_0 & \bar{A}_1 & \dots \\ 0 & & & \ddots & \ddots \end{bmatrix},$$

where \bar{A}_i , for $i \geq 0$, and \bar{B}_i , for $i \geq 0$, are nonnegative matrices in $\mathbb{R}^{b \times b}$ such that $\sum_{i=0}^{+\infty} \bar{A}_i$ and $\sum_{i=0}^{+\infty} \bar{B}_i$ are stochastic. A numerically stable method to find the stationary probability vector of this MC is Ramaswami's formula [16], which depends on the matrix \bar{G} , that is the minimal non-negative solution of

$$\bar{G} = \sum_{i=0}^{\infty} \bar{A}_i \bar{G}^i. \quad (6)$$

The quadratically-convergent Cyclic Reduction algorithm can also be applied to solve this equation.

On the other hand, a GI/M/1-type MC [2] can be seen as a QBD where the chain is allowed to decrease several levels in a single transition. The transition matrix for this MC is therefore given by

$$\hat{P} = \begin{bmatrix} \hat{B}_0 & \hat{A}_0 & & & 0 \\ \hat{B}_1 & \hat{A}_1 & \hat{A}_0 & & \\ \hat{B}_2 & \hat{A}_2 & \hat{A}_1 & \hat{A}_0 & \\ \hat{B}_3 & \hat{A}_3 & \hat{A}_2 & \hat{A}_1 & \dots \\ \vdots & \vdots & \ddots & \ddots & \ddots \end{bmatrix},$$

where \hat{A}_i , $i \geq 0$, and \hat{B}_i , $i \geq 0$ are nonnegative matrices in $\mathbb{R}^{b \times b}$ such that $\sum_{i=0}^n \hat{A}_i + \hat{B}_n$ is stochastic for all $n \geq 0$. In this case the stationary probability vector can be computed as $\pi_i = \pi_1 \hat{R}^i$, where \hat{R} is the minimal non-negative solution to

$$\hat{R} = \sum_{i=0}^{\infty} \hat{R}^i \hat{A}_i. \quad (7)$$

To solve this equation we first compute the dual process, which is of the M/G/1-type, which allows us to use the quadratically convergent CR algorithm. There are two different duals that can be used for this purpose. A brief description of both is included in Appendix A. Here we have assumed a particular boundary behavior for both the M/G/1- and GI/M/1-type MCs. A more general boundary can be assumed since our results are related to the behavior of the MCs away from the boundary, which is described by the \bar{A}_i or the \hat{A}_i matrices.

III. QBDS WITH RESTRICTED DOWNWARD TRANSITIONS

In this section we describe how the special structure of the matrix A_0 can be exploited to compute the matrix G . Consider the case where the matrix A_0 has only $r \ll m$ nonzero columns as shown in Equation (4). The (i, j) -th entry of the matrix G holds the probability that the first visit to level $k-1$ occurs by visiting state $(k-1, j)$, starting from state (k, i) , for $k > 1$ [12]. Since the downward transitions can only occur to the first r states of any level, the G matrix has the structure

$$G = \begin{bmatrix} G_+ & 0 \\ G_0 & 0 \end{bmatrix},$$

where G_+ (resp. G_0) is an $r \times r$ (resp. $(m-r) \times r$) matrix. The computation of G_+ and G_0 will be split in two steps such that for $r \ll m$ the total computation time can be significantly reduced.

A. Computing G_+

To compute G_+ we define a new process by observing the QBD MC only when the phase variable is in the set S^+ . In the original process any transition to a lower level triggers the phase to a state in S^+ , therefore the new process can only move one level down at each transition. On the other hand, the original process can move several levels upward while the phase is in S^- , i.e., between two visits to S^+ . Therefore the new process can move several levels up in one transition but only one level down. Hence, the new process is of the M/G/1-type and its behavior away from the boundary is characterized by a set of $r \times r$ matrices $(\bar{A}_i)_{i \geq 0}$. The minimal nonnegative solution \bar{G} of Equation (6) is actually equal to the matrix G_+ . This follows from the definition of the matrix \bar{G} as the first passage probability to the state $(k-1, j)$ starting from state (k, i) in the new process, and the fact that in the original process the downward transitions can only lead to S^+ . Hence, to compute the matrix G_+ we first need to determine the $r \times r$ blocks $(\bar{A}_i)_{i \geq 0}$ and then solve Equation (6).

To specify the blocks $(\bar{A}_i)_{i \geq 0}$ let the (i, j) -th entry of the $(m-r) \times r$ matrix K_i hold the probability that, given that the original process starts in state (k, i) , with $i \in S^-$, its first transition to a state with phase in S^+ occurs to the state $(k+l, j)$, for $j \in S^+$, $k > 1$ and $l \in \{-1, 0, 1, \dots\}$. Hence, the matrices $(K_i)_{i \geq -1}$ are given by

$$\begin{aligned} K_{-1} &= (I - A_1^{-})^{-1} A_0^{-+}, \\ K_0 &= (I - A_1^{-})^{-1} (A_1^{-+} + A_2^{-} K_{-1}), \\ K_1 &= (I - A_1^{-})^{-1} (A_2^{-+} + A_2^{-} K_0), \\ K_i &= (I - A_1^{-})^{-1} A_2^{-} K_{i-1}, \quad i \geq 2. \end{aligned} \quad (8)$$

To define K_{-1} we observe that the chain starts in level k and spends some time in the states of this level with phase in S^- . Afterward the chain has to move to a state $(k-1, j)$, with $j \in S^+$. The only other possible state that the chain could visit after its sojourn in level k , avoiding states with phase in S^+ , is to move to a state in level $k+1$ and phase in S^- . However, for the chain to visit level $k-1$ it first has to go back from level $k+1$ to level k , and this can only be done through a state with phase in S^+ . Therefore, this path is not possible if the first state with phase in S^+ to be visited must be in level $k-1$. The definition of the other matrices can be understood in a similar manner. Now we can define the blocks $(\bar{A}_i)_{i \geq 0}$ in terms of the matrices $(K_i)_{i \geq -1}$ as

$$\begin{aligned} \bar{A}_0 &= A_0^{++} + A_1^{+-} K_{-1}, \\ \bar{A}_1 &= A_1^{++} + A_1^{+-} K_0 + A_2^{+-} K_{-1}, \\ \bar{A}_2 &= A_2^{++} + A_1^{+-} K_1 + A_2^{+-} K_0, \\ \bar{A}_i &= A_1^{+-} K_{i-1} + A_2^{+-} K_{i-2}, \quad i \geq 3. \end{aligned} \quad (9)$$

To define \bar{A}_0 we see that the transition from a state (k, i) to a state $(k-1, j)$, with $i, j \in S^+$, can only occur in two ways: either the chain goes directly to $(k-1, j)$ with transition matrix A_0^{++} ; or it moves first to a state in level k with phase in S^- and, after a sojourn in these states, it moves downward avoiding other states in S^+ (with transition matrix $A_1^{+-} K_{-1}$). A transition to level $k+1$ is not allowed since the chain cannot

return to $k - 1$ without passing through a state in level k with phase in \mathcal{S}^+ . The other matrices can be defined similarly. Notice, to compute the matrices \bar{A}_i it suffices to store two K_i matrices at a time. The $r \times r$ matrices \bar{A}_i are sequentially computed from $i = 0$ to c , where c is the smallest positive integer such that $\sum_{i=0}^c \bar{A}_i e > (1 - \epsilon)e$, with e a column vector of ones and $\epsilon = 10^{-14}$. These blocks can then be used to compute the matrix G_+ using the CR algorithm [4].

B. Computing G_0

Given the structure of the matrices A_0 and G we can rewrite Equation (2) as

$$\begin{bmatrix} G_+ & 0 \\ G_0 & 0 \end{bmatrix} = \begin{bmatrix} A_0^{++} & 0 \\ A_0^{-+} & 0 \end{bmatrix} + \begin{bmatrix} A_1^{++} & A_1^{+-} \\ A_1^{-+} & A_1^{--} \end{bmatrix} \begin{bmatrix} G_+ & 0 \\ G_0 & 0 \end{bmatrix} + \begin{bmatrix} A_2^{++} & A_2^{+-} \\ A_2^{-+} & A_2^{--} \end{bmatrix} \begin{bmatrix} G_+^2 & 0 \\ G_0 G_+ & 0 \end{bmatrix}. \quad (10)$$

Extracting the lower-left block we find

$$G_0 - (I - A_1^{--})^{-1} A_2^{--} G_0 G_+ = (I - A_1^{--})^{-1} (A_0^{-+} + A_1^{-+} G_+ + A_2^{-+} G_+^2), \quad (11)$$

which is a Sylvester matrix equation [17], [18] of the type $AXB + X = E$, which can be solved in $O((m - r)^3)$ time with the Hessenberg-Schur method proposed in [17]. A brief description of this method is included in Appendix B together with a discussion on some additional considerations that influence the computation time of G_0 .

C. Restricted downward transitions and $A_2^{--} = 0$

Let the matrix A_0 have the structure shown in Equation (4). Additionally, assume that upward transitions from states with phase in \mathcal{S}^- take the process to a state with phase in \mathcal{S}^+ , i.e., the matrix A_2 has the form

$$A_2 = \begin{bmatrix} A_2^{++} & A_2^{+-} \\ A_2^{-+} & 0 \end{bmatrix}.$$

With this structure, the maximum number of upward transitions between two visits to \mathcal{S}^+ is two, since an upward transition from \mathcal{S}^- must end in \mathcal{S}^+ . Therefore the reduced process of the M/G/1-type, constructed by observing the original process when the phase is in \mathcal{S}^+ , has only four nonzero blocks defined as

$$\begin{aligned} \bar{A}_0 &= A_0^{++} + A_1^{+-} (I - A_1^{--})^{-1} A_0^{-+}, \\ \bar{A}_1 &= A_1^{++} + A_1^{+-} (I - A_1^{--})^{-1} A_1^{-+} \\ &\quad + A_2^{+-} (I - A_1^{--})^{-1} A_0^{-+}, \\ \bar{A}_2 &= A_2^{++} + A_1^{+-} (I - A_1^{--})^{-1} A_2^{-+} \\ &\quad + A_2^{+-} (I - A_1^{--})^{-1} A_1^{-+}, \\ \bar{A}_3 &= A_2^{+-} (I - A_1^{--})^{-1} A_2^{-+}. \end{aligned}$$

The definition of these blocks can be obtained directly from the equations (8) and (9) as follows: $A_2^{--} = 0$ implies that $K_i = 0$ for $i \geq 2$, which therefore means that $A_i = 0$ for $i \geq 3$. Additionally, the fact that $A_2^{--} = 0$ also simplifies the expressions for K_0 and K_1 , which are used in the definition of

the matrices \bar{A}_1 , \bar{A}_2 and \bar{A}_3 . This additional structure reduces both the time to compute the blocks and the time to find G_+ using CR. Additionally, to find G_0 we consider again Equation (10) and by extracting its lower-left block we find

$$G_0 = (I - A_1^{--})^{-1} (A_0^{-+} + A_1^{-+} G_+ + A_2^{-+} G_+^2).$$

Therefore, there is no need for solving a Sylvester matrix equation, as was done before, as G_0 can be determined directly from G_+ and other already computed matrices. With this additional constraint the problem of finding the $m \times m$ matrix G is replaced by the determination of just four $r \times r$ matrices and the solution of Equation (6) using these smaller matrices.

D. Restricted downward and upward transitions

Now we assume that the matrices A_0 and A_2 of the QBD have the structure described in equations (4) and (5), respectively. In this case the process obtained by observing the QBD when the phase is in the set \mathcal{S}^+ is again a QBD with parameters

$$\begin{aligned} \bar{A}_0 &= A_0^{++} + A_1^{+-} (I - A_1^{--})^{-1} A_0^{-+}, \\ \bar{A}_1 &= A_1^{++} + A_1^{+-} (I - A_1^{--})^{-1} A_1^{-+} \\ &\quad + A_2^{+-} (I - A_1^{--})^{-1} A_0^{-+}, \\ \bar{A}_2 &= A_2^{++} + A_2^{+-} (I - A_1^{--})^{-1} A_1^{-+}. \end{aligned}$$

To obtain these expressions, in addition to the simplifications due to $A_2^{--} = 0$ explained above, we notice that K_1 becomes zero since both A_2^{-+} and A_2^{--} are equal to zero. Hence \bar{A}_3 also becomes zero and the resulting process is again a QBD (of a smaller block size). Moreover, the matrix G_0 is given by

$$G_0 = (I - A_1^{--})^{-1} (A_0^{-+} + A_1^{-+} G_+).$$

The reduction in computation time is evident since now it is enough to find the solution to Equation (2) with matrices of size r instead of m . The number of matrix multiplications required to compute the blocks of the QBD process and the matrix G_0 is fixed and small compared to the solution of Equation (2).

IV. QBDS WITH RESTRICTED UPWARD TRANSITIONS

We now turn to the case where the matrix A_2 has only $r \ll m$ nonzero rows as in Equation (5), restricting the upward transitions to occur only when the phase variable is in \mathcal{S}^+ . In a QBD the (i, j) -th entry of the rate matrix R from Equation (1) can be interpreted as the expected number of visits to the state $(k + 1, j)$ starting from state (k, i) before visiting any other state at level k [2]. To visit a state in level $k + 1$ starting from level k , while avoiding level k , the first transition must take the chain from level k to level $k + 1$. However, due to the structure of A_2 , no upward transition can be made if the phase variable is in \mathcal{S}^- . Hence the last $m - r$ rows of the matrix R are equal to zero, and R can be written as

$$R = \begin{bmatrix} R_+ & R_0 \\ 0 & 0 \end{bmatrix},$$

where R_+ and R_0 are matrices of size $r \times r$ and $r \times (m - r)$, respectively. In a similar way as in the previous case, we define

a new process by observing the original QBD MC when the phase variable is in \mathcal{S}^+ . In this case the level cannot increase in the phases outside \mathcal{S}^+ , but it can decrease several levels between two visits to \mathcal{S}^+ . Therefore, the new process is a Markov chain of the GI/M/1-type. Using this process we can find the matrices R_+ and R_0 separately, as shown next.

A. Computing R_+

The behavior of the censored process, obtained by observing the original QBD MC when the phase is in \mathcal{S}^+ , is characterized away from the boundary by a set of $r \times r$ matrices $(\hat{A}_i)_{i \geq 0}$. Let \hat{R} be the minimal nonnegative solution of the Equation (7). Then the (i, j) -th entry of the matrix \hat{R} can be interpreted as the expected number of visits to state $(k+1, j)$, starting from state (k, i) , before the first return to level k [2], for $(i, j) \in \mathcal{S}^+$ and $k > 1$. This is the same interpretation as the (i, j) -th entry of R_+ ; therefore $R_+ = \hat{R}$. To find \hat{R} we first need to specify the blocks $(\hat{A}_i)_{i \geq 0}$, which is done in terms of the matrices $(W_{-i})_{i \geq 0}$.

Let the entry (i, j) of the $(m-r) \times r$ matrix W_{-l} be the probability that, given that the original process starts in state (k, i) with $i \in \mathcal{S}^-$, its first transition to a state with phase in the set \mathcal{S}^+ occurs in the state $(k-l, j)$, for $j \in \mathcal{S}^+$, $k > l \geq 0$. Hence, the matrices $(W_{-i})_{i \geq 0}$ are given by

$$\begin{aligned} W_0 &= (I - A_1^{--})^{-1} A_1^{+-}, \\ W_{-1} &= (I - A_1^{--})^{-1} (A_0^{++} + A_0^{--} W_0), \\ W_{-i} &= (I - A_1^{--})^{-1} A_0^{--} W_{-(i-1)}, \quad i \geq 2. \end{aligned}$$

The blocks $(\hat{A}_i)_{i \geq 0}$ can be defined in terms of the matrices $(W_{-i})_{i \geq 0}$ as

$$\begin{aligned} \hat{A}_0 &= A_2^{++} + A_2^{+-} W_0, \\ \hat{A}_1 &= A_1^{++} + A_1^{+-} W_0 + A_2^{+-} W_{-1}, \\ \hat{A}_2 &= A_0^{++} + A_0^{+-} W_0 + A_1^{+-} W_{-1} + A_2^{+-} W_{-2}, \\ \hat{A}_i &= A_0^{+-} W_{-i+2} + A_1^{+-} W_{-i+1} + A_2^{+-} W_{-i}, \quad i \geq 3. \end{aligned}$$

The blocks \hat{A}_i are computed from $i = 0$ to c , where c is the smallest positive integer such that $\sum_{i=0}^c \hat{A}_i e > (1 - \epsilon)e$. In this case it suffices to keep track of the three matrices $\{W_{-i+2}, W_{-i+1}, W_{-i}\}$ when computing the matrix A_i . As stated before, we need to compute the dual process of the GI/M/1-type MC characterized by $(\hat{A}_i)_{i \geq 0}$ in order to apply the CR algorithm. We use the dual relationship to compute the M/G/1-type blocks and, after solving a matrix equation of the type (6), retrieve R_+ from the G matrix of the dual. Since there are two different duals that can be used (see Appendix A), we consider both alternatives and compare their performance in Section V.

B. Computing R_0

By writing Equation (1) in block form and extracting the upper-right corner, we find

$$\begin{aligned} R_0 - R_+ R_0 A_0^{--} (I - A_1^{--})^{-1} = \\ (A_2^{+-} + R_+ A_1^{+-} + R_+^2 A_0^{+-}) (I - A_1^{--})^{-1}. \end{aligned} \quad (12)$$

This is also a Sylvester matrix equation of the type $AXB + X = E$, which can be solved in $O((m-r)^3)$ time using the Hessenberg-Schur method proposed in [17] (see Appendix B).

C. Restricted upward transitions and $A_0^{--} = 0$

When the matrix A_2 of the QBD MC has only r nonzero rows as in Equation (5) and additionally the block A_0^{--} is equal to zero, we can further improve the new algorithm in a manner analogue to Section III-C. We omit the details due to the similarity between both cases.

V. EXAMPLES

In this section we consider two continuous-time queueing systems in which the structures analyzed in the previous sections arise naturally (where a standard uniformization argument is applied to transform the problem to discrete time). First we present a priority queue with two customer classes that can be modeled as a QBD process with restricted downward transitions. The case of a QBD process with restricted upward transitions is illustrated with an overflow queue. In both cases we compare the times required to compute the matrix R or G using the full-size QBD and the approach proposed in this paper. For the overflow queue we also compare with the approach proposed in [5].

In the remainder of this section we provide a brief description of the continuous-time Markovian Arrival Processes (MAPs) and Phase-Type (PH) distribution [12] as both are used in the case studies presented below. A MAP is characterized by the parameters (n, D_0, D_1) , where n is a positive integer, and D_0 and D_1 are $n \times n$ matrices. This process is driven by an underlying MC with generator matrix $D = D_0 + D_1$, where D_1 and D_0 contain the intensities associated to transitions with and without arrivals, respectively. The off-diagonal entries of D_0 and all the entries of D_1 must be non-negative, while the diagonal entries of D_0 must be negative and such that $(D_0 + D_1)e = 0$. Let γ be the stationary distribution of the underlying MC, i.e., a $1 \times n$ vector such that $\gamma D = 0$ and $\gamma e = 1$. The arrival rate of the MAP is given by $\gamma D_1 e$. This process can be generalized by introducing markings to discriminate among different types of customers. In the forthcoming examples it is enough to consider a marked MAP (MMAP) with two types of customers. In addition to the parameters n and D_0 , the MMAP is characterized by the matrices D_1 and D_2 , which hold the transition intensities associated with an arrival of type 1 and 2, respectively. In this case the underlying MC has generator matrix $D = D_0 + D_1 + D_2$ and the arrival rate of customers of type i is equal to $\gamma D_i e$, for $i = 1, 2$.

A PH distribution is characterized by the triple (n, α, T) , where n is a positive integer, α is a $1 \times n$ vector and T a square matrix of size n . A PH distribution describes the absorption time in an MC where the states $\{1, \dots, n\}$ are transient and an additional state, say $n+1$, is absorbing. The initial probability distribution of the transient states is given by α , while T is the sub-generator matrix of these states. Therefore, the j -th entry of the vector $t = -T e$ holds the absorption rate in state

j , for $1 \leq j \leq n$. The cumulative distribution function of a PH variable is given by $F(x) = 1 - \alpha \exp(Tx)e$, for $x \geq 0$. For further reference recall that the Kronecker product of the matrices A and B , denoted $A \otimes B$, is the block matrix with block (i, j) equal to $A_{ij}B$. The Kronecker sum $A \oplus B$ is defined as $A \otimes I + I \otimes B$, where I is an identity matrix of appropriate dimension.

A. Priority Queue

Our first example is a continuous-time priority queue with two classes of customers. Class-1 customers have preemptive priority over class-2 customers. Therefore, customers of class 2 can only be served if there are no class-1 customers in the queue and the service of a class-2 customer is interrupted if a customer of class 1 arrives. The high-priority arrivals are described by a MAP characterized by (m_a^1, C_0^1, C_1^1) while the MAP of the low-priority arrivals has parameters (m_a^2, C_0^2, C_1^2) . These two processes can be combined in a single marked MAP with parameters $D_0 = C_0^1 \oplus C_0^2$, $D_1 = C_1^1 \otimes I$ and $D_2 = I \otimes C_1^2$. The service times of class-1 (resp. class-2) customers follow a PH distribution with parameters (m_s^1, α, T) (resp. (m_s^2, β, S)). To model this queue as a QBD with restricted downward transitions we take the level as the number of low-priority customers in the queue, and assume a finite buffer of size C for the class-1 customers. This assumption places no restriction in the analysis since this buffer can be dimensioned such that the blocking probability of a high-priority customer is below a certain threshold, allowing us to truncate its infinite size. Given the preemptive nature of the priority queue, this can be done using a QBD MC that ignores the low-priority customers. The second dimension of the QBD therefore holds the number of class-1 customers, the phase of the arrival process and the phase of the customer in service. If there is a class-1 customer in service, the service phase includes both the current phase of that customer and the starting phase of the next class-2 customer, and vice versa. This means that the block size is $m = (C + 1)\bar{m}$, where $\bar{m} = m_a^1 m_a^2 m_s^1 m_s^2$. With some minor abuse of notation, the QBD blocks are given by

$$A_0 = \begin{bmatrix} I \otimes s\beta \otimes I & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix},$$

$$A_1 = \begin{bmatrix} (D_0 \oplus S) \otimes I & D_1 \otimes I & 0 & \dots & 0 \\ I \otimes t\alpha & D_0 \oplus T & D_1 \otimes I & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots & D_0 \oplus T & D_1 \otimes I \\ 0 & 0 & \dots & I \otimes t\alpha & D \oplus T \end{bmatrix},$$

$$A_2 = \begin{bmatrix} D_2 \otimes I & 0 & \dots & 0 \\ 0 & D_2 \otimes I & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots & D_2 \otimes I \end{bmatrix},$$

where $D = D_0 + D_1$, $t = -Te$ and $s = -Se$. Since low-priority service completions can only occur when there are no high-priority customers in queue, downward transitions are limited to occur when the process is in one of the first $r = \bar{m}$ phases and such transitions trigger the process to the same set of phases. Therefore the structure of A_0 can be exploited as shown in Section III.

For the numerical results shown next we consider a high-priority buffer of size $C = 49$ and, for both customer classes, hyper-exponential service times with mean one and squared coefficient of variation (SCV) equal to two. The parameters of the service distribution are computed using the moment-matching method in [19], that results in a PH representation of order 2. The arrival processes are built using the method in [20], [21] that allows the matching of the first two moments of the inter-arrival distribution and the decay rate of the autocorrelation function γ with a MAP of size 2. In this case both MAPs have the same mean, fixed by the load ρ , and SCV equal to five. For this queue the load is given by $\rho = \lambda_1/\mu_1 + \lambda_2/\mu_2$, where λ_i and μ_i are the arrival and service rates of type- i customers, respectively, for $i = 1, 2$. Since the service rates are equal to one and the arrival rates are equal, then $\lambda_1 = \lambda_2 = \rho/2$. We consider two scenarios, in the first the inter-arrival times are independent ($\gamma = 0$), while in the second γ is equal to 0.9. With this set of parameters the block size is 800 while the number of nonzero columns in A_0 is 16.

In Table I we show the time required to compute the matrix G using the full-size QBD with the CR algorithm (QBD-CR), the time to compute the M/G/1-type blocks (BI), the number of those blocks, the time to compute the matrix G_+ with CR (MG1-CR) and the time to solve the Sylvester matrix equation to get G_0 (Sylv). The total computation time using the reduced process is shown in column MG1, and the last column has the ratio between the columns QBD and MG1. Clearly, the M/G/1-type based method outperforms the full-size approach, which can take 4 to 8 times longer to compute G . Also, when the load ρ increases both methods require more computation time, particularly the CR algorithm for the QBD and the computation of the M/G/1-type blocks. A large load has two major effects: first, it increases the rate of upward transitions per time unit; second, since the set \mathcal{S}^+ includes only the phases in which there are no high-priority customers in queue, a larger load increases the likelihood of having long sojourn times in \mathcal{S}^- . These two effects together imply that the number of blocks to compute increases and the CR algorithm requires more time to find G_+ . In contrast, the Hessenberg-Schur method to solve Equation (11) shows a more stable behavior, not being directly affected by the load of the queue.

Table II contains the same information as the previous one, but in this scenario the arrival processes are highly autocorrelated, with decay rate of the autocorrelation function $\gamma = 0.9$. As can be observed, the correlation, together with the load, has a large effect on the number of M/G/1-type blocks that describe the reduced process and therefore on the time required to compute those blocks and to find G_+ . On the other

Table I
COMPUTATION TIMES (SEC) FOR THE PRIORITY QUEUE WITH $\gamma = 0$

ρ	QBD-CR	BI	# BI	MG1-CR	Sylv	MG1	Ratio
0.1	34.34	1.38	19	0.06	5.14	6.58	5.22
0.2	34.36	1.55	33	0.13	6.81	8.49	4.05
0.3	39.34	1.78	50	0.14	5.16	7.08	5.56
0.4	39.36	2.08	73	0.25	3.53	5.86	6.72
0.5	44.36	2.5	103	0.28	6.84	9.63	4.61
0.6	44.42	3.08	141	0.75	6.84	10.67	4.16
0.7	49.45	3.75	192	0.72	6.86	11.33	4.37
0.8	49.44	4.72	260	1.19	5.19	11.1	4.46
0.9	59.48	5.97	347	2.91	3.55	12.42	4.79
1	139.78	7.72	466	3.34	5.19	16.25	8.6

Table II
COMPUTATION TIMES (SEC) FOR THE PRIORITY QUEUE WITH $\gamma = 0.9$

ρ	QBD-CR	BI	# BI	MG1-CR	Sylv	MG1	Ratio
0.1	34.44	1.39	20	0.06	5.2	6.66	5.17
0.2	39.44	1.61	38	0.13	6.89	8.63	4.57
0.3	39.42	1.94	64	0.25	3.55	5.73	6.87
0.4	44.45	2.44	105	0.27	3.55	6.25	7.11
0.5	49.47	3.33	169	0.78	5.22	9.33	5.3
0.6	54.44	4.53	270	1.3	5.24	11.06	4.92
0.7	59.5	6.8	431	3.02	5.2	15.02	3.96
0.8	64.5	10.48	680	11.19	5.23	26.91	2.4
0.9	69.47	16.42	1059	13.13	5.25	34.8	2
1	109.61	26.33	1602	45.39	3.58	75.3	1.46

hand, the correlation has little effect on the computation times of the full-size QBD, as well as on the time to find G_0 with the Hessenberg-Schur method. Therefore, the reduced process still offers a reduction in computation times but this gain is affected by the system parameters. A similar behavior is observed in the next example.

In addition to the computation times, it is relevant to consider the behavior of the approach introduced in this paper in terms of the residual error. Let the infinity norm of an $n \times m$ matrix K be given by $\|K\|_\infty = \max_{i=1}^n \sum_{j=1}^m K_{ij}$. Let \tilde{G} be the matrix that solves Equation (2) obtained with the approach of Section III. Then the residual error is defined as

$$\|\tilde{G} - A_0 + A_1\tilde{G} + A_0\tilde{G}^2\|_\infty,$$

which gives a measure of the goodness of \tilde{G} as a solution for Equation (2). In all the instance considered here the residual error was always below 10^{-14} , revealing the good behavior of the approach proposed here. This behavior is to be expected since the algorithms on which our method relies (Cyclic Reduction and the Hessenberg-Schur method for the Sylvester equation) are numerically stable. A similar result in terms of the residual error holds for the example in the next section.

B. Overflow Queue

The second example is an overflow queueing system consisting of two queues. The arrival process to the first queue is a MAP characterized by (m_a, D_0, D_1) . Customers arriving at the first queue are attended in FCFS order by a single server

with service times following a PH distribution characterized by the parameters (m_s^1, α, T) . This queue has a finite buffer of size C and a customer that arrives at a full buffer is sent to the second queue. The second queue receives *only* overflow arrivals from the first queue and attends them in FCFS order with a single server. The service times in this queue follow a PH distribution with parameters (m_s^2, β, S) . Hence, the arrival process at the second queue can be described by a MAP with parameters (m_o, C_0, C_1) given by $m_o = (C + 1)m_a m_s^1$,

$$C_0 = \begin{bmatrix} D_0 \otimes I & D_1 \otimes I & 0 & \dots & 0 & 0 \\ I \otimes t\alpha & D_0 \oplus T & D_1 \otimes I & \dots & 0 & 0 \\ 0 & I \otimes t\alpha & D_0 \oplus T & \dots & 0 & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & D_0 \oplus T & D_1 \otimes I \\ 0 & 0 & 0 & \dots & I \otimes t\alpha & D_0 \oplus T \end{bmatrix},$$

$$C_1 = \begin{bmatrix} 0 & \dots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & 0 & D_1 \otimes I \end{bmatrix},$$

where $t = -Te$. Assuming an infinite buffer at the second queue, we can model the queueing system as a QBD where the level describes the number of customers in the second queue. The second dimension holds the phase of the current customer in service and the phase of the arrival process in the second queue. The parameters of the QBD are $A_0 = I \otimes s\beta$, $A_1 = C_0 \oplus S$, $A_2 = C_1 \otimes I$, with $s = -Se$. In this case, the restricted upward transitions are a result of the overflow process, as can be seen in the structure of C_1 , where the arrivals to the second queue can only occur in the last $m_a m_s^1$ phases. As a result, the inclusion of a separate arrival stream directed to the second queue would suppress this structure. The block size in this case is $m = m_o m_s^2$ and the number of nonzero rows in A_2 is $r = m_a m_s^1 m_s^2$.

As with the previous example, we make use of the moment-matching methods in [19], [20], [21] to obtain PH and MAP representations of the service and arrival processes, respectively. The arrival process at the first queue has arrival rate and SCV equal to five while the service time has mean one and SCV equal to two. Therefore the first queue is heavily loaded and many customers are overflowed to the second queue. The arrival rate at the second queue (λ_2) is the arrival rate of the MAP with parameters (C_0, C_1) . Therefore for a given load at the second queue (ρ_2) the service rate at this queue is fixed by the relation $\rho_2 = \lambda_2 / \mu_2$. In this queue the service times have SCV equal to two, as in the first queue. The results are presented for different values of ρ_2 and a buffer size of 100 in the first queue. With these parameters the block size is $m = 808$ while the number of nonzero rows in A_2 is $r = 8$.

Table III shows the computation times in a similar fashion as the previous tables, with the main difference being that the column GM1-CR-R (resp. GM1-CR-B) includes the time to compute the blocks of the Ramaswami (resp. Bright) dual process and the time to solve the dual with the CR algorithm.

Table III
COMPUTATION TIMES (SEC) FOR THE OVERFLOW QUEUE WITH $C = 100$

ρ_2	QBD-CR	BI	# BI	GM1-CR-R	GM1-CR-B	Sylv	GM1-R	GM1-B	Ratio-R	Ratio-B
0.1	31.91	24.92	2791	17.11	8.14	2.25	44.28	35.31	0.72	0.90
0.2	37.19	12.38	1431	9.78	1.95	2.23	24.39	16.56	1.52	2.25
0.3	42.3	8.63	980	2.33	0.53	2.25	13.20	11.41	3.20	3.71
0.4	42.3	6.73	745	2.27	0.63	2.25	11.25	9.61	3.76	4.40
0.5	47.38	5.63	601	1.17	0.64	3.94	10.73	10.2	4.41	4.64
0.6	47.38	4.94	511	0.81	0.25	2.23	7.98	7.42	5.93	6.38
0.7	52.47	4.41	442	0.81	0.58	2.24	7.45	7.22	7.04	7.27
0.8	57.56	4.02	391	0.78	0.58	2.25	7.05	6.84	8.17	8.41
0.9	62.67	3.70	350	0.66	0.45	2.23	6.59	6.39	9.51	9.81
1	169.69	3.47	317	0.78	0.78	2.22	6.47	6.47	26.23	26.24

The columns GM1-R and GM1-B show the total computation times to find R using the two different duals, while the columns Ratio-R and Ratio-B hold the ratio between the QBD-CR and the GM1-R and GM1-B columns, respectively. Again, the load has an important effect on the computation times, but in this case the consequences are reversed. When the load is low the original process can make many downward transitions between two visits to the set S^+ , increasing the number of GI/M/1-type blocks. As before, a large number of blocks increases the computation time of the CR algorithm for the M/G/1-type MC (the dual process) but it has little effect on the solution of the Sylvester equation and the full-size QBD. For loads between 0.2 and 0.9 in this scenario, the solution of the full-size QBD may take between 2 and 10 times as long as the solution of the reduced process. When the load is one the process is null recurrent and the QBD-CR takes a much longer time than for lower loads. This effect can be reduced by using a shifting technique [14], resulting in times similar to those shown for loads up to 0.9. When comparing the two alternative duals, it is clear how the Bright dual outperforms the Ramaswami dual, being specially effective when the load is low, i.e., when the number of GI/M/1-type blocks is large. This effect is to be expected since for $\rho_2 < 1$ the GI/M/1-type MC is positive recurrent, and therefore the Ramaswami dual is transient while the Bright dual is positive recurrent (see [22]).

In Figure 1 we include, for the full-size QBD, the computation times of CR (QBD-CR), the original U-based algorithm (QBD-U) and the modified version of the U-based algorithm (QBD-GT) proposed by Grassman and Tavakoli [5] to exploit the special structure of A_2 . We also include the total time required to solve the reduced process using the Ramaswami dual (GM1-R) and the Bright dual (GM1-B). The scenario is the same as in the previous case with the only exception that the buffer size in the first queue is $C = 50$. This means that the block size is $m = 408$ while the number of nonzero rows in A_2 does not change. This reduction is done because of the long computation times experienced with the U-based method, as can be observed in the figure. From these results it is evident the substantial gain obtained by the QBD-GT method compared to the original QBD-U, which requires about 5 times as much computation time. In spite of this gain, the QBD-GT method performs better than the QBD-CR only for small




Figure 1. Computation times (sec) for the overflow queue with $C = 50$

values of ρ_2 . In contrast, the GI/M/1-type based approach performs better than CR on the full-size QBD, except for low values of ρ_2 . For the remaining part of the load range (except $\rho_2 = 1$) the QBD-CR takes up to 6 times as much time as the GI/M/1-type based approach. In this case the time to compute R is smaller using the Bright dual than the Ramaswami dual. The difference is significant for low loads, when the number of blocks is large (2800 for $\rho_2 = 0.1$) and it vanishes when the load is high. Therefore, the use of the Bright dual implies an important reduction in computation times in the range of the load that is more critical for the reduced process.

In both examples we have shown that the computation times can be substantially reduced by using the approach proposed in this paper. For each case the gain increases with the ratio m/r as expected, but it also depends on other factors related to the parameters of the system modeled. Even though for some cases the reduced-process approach may take longer than solving the full-size QBD, exploiting the structure of the matrices A_0 or A_2 may reduce the computation times substantially. To determine whether the reduced process can be useful for a particular system or not, attention must be paid to the expected sojourn times in S^- related to those in S^+ . If the sojourn times in S^- are too long compared with the sojourn times in S^+ , the reduced process will need many blocks to be described. This increases both the time required to compute the blocks and the time to find G_+ or R_+ . However, to analyze the performance

of a particular system it is usual to consider a broad range of conditions (load, variability, etc.), and it is likely that for a considerable part of this range the reduced process can provide important reductions in computation times. Additionally, the parameters of the system under analysis not only affect the reduced process but also the original process. This can be observed in our examples, where an increase in the load increased the computation times of the original QBD for both the overflow and the priority queue. An additional gain can be obtained for the GI/M/1-type case by using the Bright dual, which helps to reduce the computation time specially in those cases where the reduced process requires more time, i.e., when the number of blocks is large.

APPENDIX A DUAL PROCESSES

In this section we describe two dual relationships between discrete-time M/G/1- and GI/M/1-type processes. In both cases the dual process can be seen as the time-reverse of the original process with respect to an invariant measure [23], [24]. We consider the computation of an M/G/1-type MC as the dual of a GI/M/1-type MC, but the opposite relationship can be defined in a similar manner. The Ramaswami dual was introduced in [25] and its probabilistic interpretation given in [23]. Let the set of matrices $(A_i)_{i \geq 0}$ describe a GI/M/1-type MC, such that $A = \sum_{i=0}^{\infty} A_i$ is stochastic and irreducible. Then A is the transition matrix of a discrete-time MC with stationary probability vector α , i.e., $\alpha A = \alpha$ and $\alpha e = 1$. Then the Ramaswami dual is an M/G/1-type MC characterized by the set of matrices $(A_i^R)_{i \geq 0}$ given by $A_i^R = \Delta_R^{-1} A_i' \Delta_R$, where $\Delta_R = \text{diag}(\alpha)$. The G matrix of this process, denoted G_R , is related to the R matrix of the original process by $G_R = \Delta_R^{-1} R' \Delta_R$. Let $\rho(M)$ denote the spectral radius of a matrix M . Since the matrix G_R has the same eigenvalues as R , if the original GI/M/1-type MC is positive recurrent ($\rho(R) < 1$) the dual process is transient ($\rho(G_R) < 1$), and vice versa. The dual process will be null recurrent if and only if the original process is also null recurrent. In this case the dual process is the time-reverse process with respect to the invariant measure α .

We now turn to the Bright dual [24], which is defined as the time-reverse process with respect to a different invariant measure. If the GI/M/1-type MC is positive recurrent, the eigenvalue of maximum real part of R is $\eta = \rho(R) < 1$. It has been shown that the spectral radius of the matrix $\sum_{i=0}^{\infty} A_i \eta^i$ is equal to one [15]. Therefore there exists a positive vector w_η such that

$$w_\eta \left(\sum_{i=0}^{\infty} A_i \eta^i \right) = w_\eta.$$

The Bright dual is an M/G/1-type MC characterized by the matrices $(A_i^B)_{i \geq 0}$ defined as $A_i^B = \eta^{i-1} \Delta_B^{-1} A_i' \Delta_B$, where $\Delta_B = \text{diag}(w_\eta)$. The matrix R of the original GI/M/1-type MC and the matrix G_B of the dual process are related by $G_B = \eta^{-1} \Delta_B^{-1} R' \Delta_B$. In this case the eigenvalues of the matrix G_B are the eigenvalues of R divided by η . Hence

the spectral radius of G_B is equal to one and the dual process is positive recurrent [24]. The positive recurrent case is particularly relevant because it can make a difference in the time to compute the matrix R_+ . To find this matrix we first compute the dual of the reduced process, which is of the GI/M/1-type. When the process is positive recurrent, as in the examples shown in Section V for loads less than one, the Ramaswami dual will be transient while the Bright dual will be positive recurrent. As explained in detail in [22], the Bright dual can therefore reduce the computation times achieved by the Ramaswami dual considerably. This is confirmed numerically in Section V, especially when the load of the overflow queue is small, which results in a large number of blocks for the GI/M/1-type MC and a small value of η . The computation time for the reduced process increases with the number of blocks, but the gain that can be realized by using the Bright dual is larger when η is smaller [22]. Therefore, the Bright dual becomes especially useful in this case as it compensates the larger computation times caused by the number of blocks.

APPENDIX B SOLVING SYLVESTER MATRIX EQUATIONS

In this section we describe how to solve the matrix equations (11) and (12) using the Hessenberg-Schur decomposition proposed in [17]. As noted before, these are Sylvester matrix equations of the form $AXB + X = E$. Consider Equation (11) and let $n = m - r$, then X and E are $n \times r$ matrices, while A and B are square matrices of size n and r , respectively. The first step to solve this linear system is to find orthogonal matrices U and V such that $U'AU = P$ and $V'BV = R$, where P is an upper-Hessenberg matrix, R is a quasi-upper triangular matrix and $'$ denotes the transpose operator. A matrix P is upper-Hessenberg if its entries $P_{ij} = 0$ for $i > j + 1$. A quasi-upper triangular matrix, also called real Schur form, is block-triangular with 1×1 (resp. 2×2) diagonal blocks that correspond to the real (resp. complex) eigenvalues [26]. While the Hessenberg decomposition to obtain U can be done using Householder transformations, the real Schur decomposition to compute V makes use of the QR algorithm, see [26, Chapter 7]. Let $F = U'EV$ and $Y = U'XV$, then the linear system becomes $PYR + Y = F$. Therefore, to find Y_k , the k -th column of the matrix Y , we need to solve the system

$$P \sum_{j=1}^{\max(k+1,r)} R_{jk} Y_j + Y_k = F_k,$$

for $1 \leq k \leq r$. However, the quasi-upper triangular form of R greatly simplifies this system. For $k < r$ there are two possible cases, either $R_{k+1,k} = 0$ or not. If $R_{k+1,k} = 0$, then Y_k is the solution to the $n \times n$ Hessenberg system

$$(PR_{k,k} + I)Y_k = F_k - \sum_{j=1}^{k-1} R_{jk} P Y_j, \quad (13)$$

which can be solved in $O(n^2)$ time. On the other hand, $R_{k+1,k} \neq 0$ implies $R_{k+2,k+1} = 0$, and hence we need to solve

$$\begin{bmatrix} PR_{k,k} + I & PR_{k+1,k} \\ PR_{k,k+1} & PR_{k+1,k+1} + I \end{bmatrix} \begin{bmatrix} Y_k \\ Y_{k+1} \end{bmatrix} = \begin{bmatrix} \hat{F}_k^{k-1} \\ \hat{F}_{k+1}^{k-1} \end{bmatrix}, \quad (14)$$

where $\hat{F}_k^l = F_k - \sum_{j=1}^l R_{jk}PY_j$, for $1 \leq l \leq k-1$ and $1 \leq k \leq r$. This $2n \times 2n$ linear system is upper-triangular with two nonzero subdiagonals that can be solved in $O(n^2)$ time [17]. Notice, to determine Y_k it is necessary to know Y_1, \dots, Y_{k-1} . Therefore, the algorithm starts by computing the first (or first two) column(s), and then works forward until the last column of Y has been computed. After finding Y , the matrix X can be computed as $X = UYV'$.

It is possible to apply this procedure to either the original or the transpose $B'X'A' + X' = E'$ system. In the first case A is transformed into Hessenberg form and B into real Schur form, while the opposite happens in the second case. The choice directly affects the computation times since for a matrix of size b the Schur decomposition can be done in $10b^3$ operations, while it takes $\frac{5}{3}b^3$ operations to compute the Hessenberg decomposition using Householder transformations [17], [26]. Therefore, to solve Equation (11) it is better to use the original system since the Hessenberg decomposition is applied on the $n \times n$ matrix A , which is larger than the $r \times r$ matrix B under the assumption that $r \ll m$. On the other hand, to solve Equation (12) it is preferable to first transpose the system since in that case B is a $n \times n$ matrix given by $B = A_0^{-}(I - A_1^{-})^{-1}$.

An additional issue to take into account when solving equations (11) and (12) is the actual computation of the matrices A , B and E . Take for example Equation (11), where $A = (I - A_1^{-})^{-1}A_2^{-}$, $B = G_+$ and $E = (I - A_1^{-})^{-1}(A_0^{-} + A_1^{-}G_+ + A_2^{-}G_+)$. Although all the matrices involved are already computed it is still necessary to perform two matrix multiplications to determine A and E . In the examples shown in this paper the A_i^{-} blocks are sparse or even zero. In some cases however these blocks can be dense and therefore these matrix multiplications may require considerably more time. A way to avoid this is to solve the slightly different equation

$$(I - A_1^{-})G_0 - A_2^{-}G_0G_+ = A_0^{-} + A_1^{-}G_+ + A_2^{-}G_+^2,$$

which is a Sylvester matrix equation of the type $AXB + CX = E$. The procedure to solve this equation is very similar to the one shown above, but in this case the first step of the QZ algorithm [26] is applied to the pair (A, C) . As a result A is reduced to Hessenberg form while C is transformed into upper-triangular form. This, together with a reduction of B to quasi-upper triangular form, allows the solution of this equation in a similar way as done in (13) and (14). A detailed explanation can be found in [18]. Since the matrices A , B , C and E are already computed, this algorithm may perform better when the blocks A_i^{-} are dense. We have found instances of random QBDs with dense blocks where

this last algorithm outperforms the one based on the equation $AXB + X = E$.

REFERENCES

- [1] V. Wallace, "The solution of quasi birth and death processes arising from multiple access computer systems," Ph.D. dissertation, Systems Engineering Laboratory, University of Michigan, 1969.
- [2] M. F. Neuts, *Matrix-Geometric Solutions in Stochastic Models*. Baltimore: The Johns Hopkins University Press, 1981.
- [3] G. Latouche and V. Ramaswami, "A logarithmic reduction algorithm for quasi-birth-and-death processes," *Journal of Applied Probability*, vol. 30, pp. 650–674, 1993.
- [4] D. A. Bini and B. Meini, "On the solution of a nonlinear matrix equation arising in queueing problems," *SIAM Journal of Matrix Analysis and Applications*, vol. 17, pp. 906–926, 1996.
- [5] W. K. Grassmann and J. Tavakoli, "Solving QBD processes when levels can increase only in certain phases," manuscript in preparation, presented at the MAM6 conference, Beijing (China), June 2008.
- [6] G. Latouche, "Algorithms for infinite Markov chains with repeating columns," in *Linear Algebra, Markov chains and Queueing Models*, C. D. Meyer and R. J. Plemmons, Eds. Springer-Verlag, 1993, pp. 231–265.
- [7] K. S. Meier-Hellstern, "The analysis of a queue arising in overflow models," *IEEE Transactions on Communications*, vol. 37, pp. 367–372, 1989.
- [8] J. W. Cohen, *The Single Server Queue*. North-Holland, 1969.
- [9] N. K. Jaiswal, *Priority Queues*. Academic Press, 1968.
- [10] B. Van Houdt and C. Blondia, "The waiting time distribution of a type k customer in a discrete-time MMAP[K]/PH[K]/c (c = 1, 2) queue using QBDs," *Stochastic Models*, vol. 20, pp. 55–69, 2004.
- [11] D. A. Bini, B. Meini, S. Steffé, and B. Van Houdt, "Structured Markov chains solver: software tools," in *SMCtools Workshop*. Pisa, Italy: ACM Press, 2006.
- [12] G. Latouche and V. Ramaswami, *Introduction to Matrix Analytic Methods in Stochastic Modeling*, ser. ASA-SIAM Series on Statistics and Applied Probability. Philadelphia, PA: SIAM, 1999.
- [13] G. Ciardo and E. Smirni, "ETAQA: an efficient technique for the analysis of QBD-processes by aggregation," *Performance Evaluation*, vol. 36-37, pp. 71–93, 1999.
- [14] D. Bini, G. Latouche, and B. Meini, *Numerical Methods for Structured Markov Chains*. Oxford University Press, 2005.
- [15] M. F. Neuts, *Structured stochastic matrices of M/G/1 type and their applications*. Marcel Dekker Inc., 1989.
- [16] V. Ramaswami, "A stable recursion for the steady state vector in Markov chains of M/G/1 type," *Stochastic Models*, vol. 4, pp. 183–188, 1988.
- [17] G. H. Golub, S. Nash, and C. Van Loan, "A Hessenberg-Schur method for the problem $AX+XB=C$," *IEEE Transactions on Automatic Control*, vol. 24, pp. 909–913, 1979.
- [18] J. D. Gardiner, A. J. Laub, J. J. Amato, and C. B. Moler, "Solution of the Sylvester matrix equation $AXB + CX = E$," *ACM Transactions on Mathematical Software*, vol. 18, pp. 223–231, 1992.
- [19] W. Whitt, "Approximating a point process by a renewal process, I: Two basic methods," *Operations Research*, vol. 30, pp. 125–147, 1982.
- [20] A. Heindl, "Inverse characterization of hyperexponential MAP(2)s," in *Proc. 11th Int. Conference on Analytical and Stochastic Modelling Techniques and Applications (ASMTA)*, 2004.
- [21] J. E. Diamond and A. S. Alfa, "On approximating higher order MAPs with MAPs of order two," *Queueing Systems*, vol. 34, pp. 269–288, 2000.
- [22] P. G. Taylor and B. Van Houdt, "On the dual relationship between Markov chains of GI/M/1 and M/G/1 type," *submitted to Advances in Applied Probability*, 2009.
- [23] S. Asmussen and V. Ramaswami, "Probabilistic interpretations of some duality results for the matrix paradigms in queueing theory," *Communications in Statistics Stochastic Models*, vol. 6, pp. 715–733, 1990.
- [24] L. Bright, "Matrix-analytic methods in applied probability," Ph.D. dissertation, Department of Applied Mathematics, University of Adelaide, 1996.
- [25] V. Ramaswami, "A duality theorem for the matrix paradigms in queueing theory," *Communications in Statistics Stochastic Models*, vol. 6, pp. 151–161, 1990.
- [26] G. H. Golub and C. Van Loan, *Matrix Computations*. The Johns Hopkins University Press, 1996.