Componentwise accurate Brownian motion computations using Cyclic Reduction

Giang T. Nguyen, Federico Poloni

Abstract

Markov-modulated Brownian motion is a popular tool to model continuous-time phenomena in a stochastic context. The main quantity of interest is the invariant density, which satisfies a differential equation associated with the quadratic matrix polynomial $P(z) = Vz^2 - Dz + Q$, where the matrices V and D are diagonal and Q is the transition matrix of a discrete-time Markov chain. Its solution is typically constructed by computing an invariant pair of $P(z)$ associated with its eigenvalues in the left half-plane, or by solving the matrix equation $X^2V - XD + Q = 0$. We show that these tasks can be solved using a componentwise accurate algorithm based on Cyclic Reduction, generalizing the recently appeared algorithms for the linear case ($V = 0$). We give a proof of the numerical stability of our algorithm in the componentwise sense; the same proof applies to Cyclic Reduction in a more general M-matrix setting which appears in other applications such as the modelling of QBD processes.

1 Introduction

Markov-modulated Brownian motion [Asm95, KK95] is a popular tool to model continuous-time phenomena in a stochastic context. An MMBM can be described as the pair $\{Y(t), \phi(t)\}_{t \geq 0}$, where $\phi(t)$ is a continuous-time Markov chain on a state space $\mathcal{S} = \{1, 2, \ldots, n\}$ with rate matrix $Q \in \mathbb{R}^{n \times n}$ ($Q\mathbf{1} = \mathbf{0}$, where $\mathbf{1}$ and $\mathbf{0}$ are the vectors of all ones and zeros, respectively). Whenever $\phi(t) = i \in \mathcal{S}$, $Y(t)$ evolves according to a Brownian motion process with drift d_i and variance $\frac{1}{2}v_i \geq 0$.

The main quantity of interest to determine its steady-state behaviour is the invariant density $\boldsymbol{p}(x) : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}^{1 \times n}$, which satisfies, with suitable boundary conditions, the differential equation

$$\boldsymbol{p}''(x)V - \boldsymbol{p}'(x)D + \boldsymbol{p}(x)Q = \mathbf{0}, \tag{1}$$

where $V = \operatorname{diag}(v_i)_{i \in \mathcal{S}}$ and $D = \operatorname{diag}(d_i)_{i \in \mathcal{S}}$. The solutions of this ODE are related to the eigenvalues and left eigenvectors of the matrix polynomial

$$P(z) := Vz^2 - Dz + Q. \tag{2}$$

The solution of probabilistic interest is asymptotically stable, that is, $\boldsymbol{p}(x) \to \mathbf{0}$ when $x \to \infty$. Hence, we are interested in particular in the eigenvalues λ_i with $\Re \lambda_i < 0$.

Several methods have been suggested in literature to compute this solution. Some are iterative [NL] based on Cyclic Reduction; some depend on the eigendecomposition of a linearization [KK95], or more generally a block diagonal decomposition [AS01]. A few other algorithms rely on finding a special *invariant pair* (X, U), that is, a pair of matrices $X \in \mathbb{R}^{\ell \times \ell}, U \in \mathbb{R}^{\ell \times n}$ satisfying

$$X^2 UV - XUD + Q = 0. \tag{3}$$

For instance, Ivanovs [Iva10, Sect. 3] considers a related problem—determining the steady-state behavior of a two-boundary Markov-modulated Brownian motion process—which can be solved with the same techniques. The author constructs the solution starting from an invariant pair in which

$$U = \begin{bmatrix} I & \Psi \end{bmatrix}, \tag{4}$$

where the identity block corresponds to the indices i for which $v_{ii} > 0$ or $d_{ii} > 0$. This invariant pair (X, U) has a probabilistic meaning: $\Psi \geq 0$ is the matrix recording first-return probabilities of the time-reversed process, and X is a sub-generator matrix ($X_{ij} \geq 0$ if $i \neq j$, and $X\mathbf{1} \leq \mathbf{0}$) for the downward-record process.

A special case often considered in literature is when $v_{ii} > 0$ for all $i \in \mathcal{S}$. In this case, $U = I$, and (3) reduces to

$$X^2 V - XD + Q = 0. \tag{5}$$

This matrix equation has been studied extensively, especially because of its connection to quasi-birth-death processes [BLM05, Ram99].

Another special case interesting in its own is when $V = 0$, that is, when the Markov-modulated Brownian motion $\{Y(t), \phi(t)\}$ is a *stochastic fluid model*, also known as a *fluid queue*. The papers [XXL12] and [NP15] deal with this special case, and provide quadratically convergent algorithms for the invariant density, which are *componentwise accurate*. That is, the algorithms can deliver an approximate solution $\widetilde{\boldsymbol{p}}$ such that the quantity $\max_{i \in \mathcal{S}} |\widetilde{\boldsymbol{p}}_i - \boldsymbol{p}_i|/\boldsymbol{p}_i$ is bounded, rather than $\|\boldsymbol{p} - \widetilde{\boldsymbol{p}}\|/\|\boldsymbol{p}\|$. In this informal introduction, $\boldsymbol{p}$ refers to the exact value of a vector quantity related to the solution, for instance, the value of $\boldsymbol{p}(x)$ at a determined level x, and $\widetilde{\boldsymbol{p}}$ represents its computed version in machine arithmetic.

Informally, this means that all entries of $\boldsymbol{p}$ have the same number of correct significant digits, irrespectively of their magnitude; thus, all components can be computed to a high accuracy. For example, suppose

$$\boldsymbol{p} = \begin{bmatrix} 1 - 10^{-15} & 10^{-15} \end{bmatrix}.$$

In this case, traditional linear algebra algorithms would instead ensure a high accuracy on the large component $\widetilde{\boldsymbol{p}}_1$ only, while $\widetilde{\boldsymbol{p}}_2$ could vary wildly with few theoretical guarantees: for instance, it could become negative. Componentwise error bounds are particularly meaningful for probability applications, since small

components may represent probabilities of catastrophic failure, which have to be assessed carefully.

In this paper, we focus on computing in a componentwise accurate fashion invariant pairs that solve (3) (with the additional property (4)) and solutions of the matrix equation (5), which can then be used to compute a solution $\boldsymbol{p}(x)$ of (1). This problem contains the linear models treated in [NP15, XXL12] as a special case ($V = 0$); we extend the techniques introduced there and generalize them to the more challenging second-order case.

In particular, the case in which some of the v_{ii} are zero and some are not, requires special attention. To treat it, we use a method related to both the *shift technique* [HMR02, BLM05] and the *theory of index reduction* of differential-algebraic equations [KM06]. We use these techniques in a novel way that combines the themes of these two approaches and adds componentwise accuracy and positivity preservation into the picture.

The paper is structured as follows. In Section 2, we introduce most of the concepts needed in the development of the algorithm, including invariant pairs, componentwise accurate algorithms, and Cyclic Reduction. In Section 3, we present our solution strategy and formulate a solution algorithm, first for the case $\operatorname{diag}(V) > 0$ and then in general. In Section 4, we prove the numerical stability of our algorithm in a componentwise sense. Numerical experiments in Section 5 confirm the effectiveness of this approach, and some brief conclusions follow.

To the best of our knowledge, Lemma 7 and the fully subtraction-free version of Cyclic Reduction presented in Algorithm 1, together with the proof of its componentwise stability, are new also in the context of discrete-time quasi-birth-death models [BLM05].

2   Assumptions and preliminaries

2.1   Eigenvalues and invariant pairs of matrix polynomials

For ease of analysis, we make several assumptions to make sure that the problem cannot be simplified further:

A1 The matrix Q is irreducible and aperiodic.

A2 $V \neq 0$.

A3 there is no index $i \in \mathcal{S}$ for which $v_{ii} = d_{ii} = 0$.

Assumption A1 is to eliminate the cases where our problem can be reduced to smaller disjoint cases. If Assumption A2 is not satisfied, methods for the fluid queue case like the one in [NP15] can be used. Finally, if Assumption A3 does not hold, we can replace the problem with another one, where such i is censored out.

Let $\mathbb{C}[z]$ denote the set of polynomials in the variable z. We encountered in the introduction the notion of *eigenvalues* and *left eigenvectors* of a degree-g

matrix polynomial

$$P(z) = P_0 + P_1 z + \cdots + P_g z^g \in \mathbb{C}[z]^{n \times n},$$

that is, scalars $\lambda \in \mathbb{C}$ and row vectors $\boldsymbol{u} \in \mathbb{C}^{1 \times n}$ such that $\boldsymbol{u} P(\lambda) = \boldsymbol{0}$. Under Assumptions A1–A3, $P(z)$ is a *regular* matrix polynomial, that is, the scalar polynomial $\det P(z) \in \mathbb{C}[z]$ is not identically zero, as one can see by considering its highest-degree term; hence its eigenvalues are a well-defined set of n complex numbers counted with multiplicity. When the leading term P_g is singular, $\det P(z)$ has degree strictly lower than gn. In this case we say that ∞ is an eigenvalue of $P(z)$ with algebraic multiplicity $gn - \deg P(z)$[1].

An extremely useful tool to deal with multiple eigenvalues simultaneously, both in theory and in numerical practice, is *invariant pairs* [BK11, GLR09, HK00]. For any $\ell \leq gn$, a pair $(X, U) \in \mathbb{C}^{\ell \times \ell} \times \mathbb{C}^{\ell \times n}$ is called a *left invariant pair* for the matrix polynomial $P(z) = P_0 + P_1 z + \cdots + P_g z^g \in \mathbb{C}[z]^{n \times n}$ if

$$UP_0 + XUP_1 + X^2 UP_2 + \cdots + X^g UP_g = 0 \tag{6}$$

and the matrix $\begin{bmatrix} U & XU & \cdots & X^{g-1}U \end{bmatrix}$ has full row rank. It follows from this definition that if (X, U) is a left invariant pair, then so is

$$(MXM^{-1}, MU) \quad \text{for any nonsingular } M \in \mathbb{C}^{\ell \times \ell}. \tag{7}$$

The reader not acquainted with this concept can consider a simpler case in which $P(z)$ has gn distinct finite eigenvalues; in this case, the invariant pairs for a matrix polynomial are given by

$$X = \operatorname{diag}(\lambda_1, \lambda_2, \ldots, \lambda_\ell), \quad U = \begin{bmatrix} \boldsymbol{u}_1 \\ \boldsymbol{u}_2 \\ \vdots \\ \boldsymbol{u}_\ell \end{bmatrix},$$

where $\ell \leq gn$ is arbitrary and for each $i = 1, 2, \ldots, \ell$ the row vector $\boldsymbol{u}_i$ is a left eigenvector of $P(z)$ with eigenvalue λ_i, as well as all the pairs obtained from them through the change of basis transformations in (7). Informally speaking, invariant pairs are a tool to deal with several eigenvalues and eigenvectors at the same time without resorting to a Jordan form.

Invariant pairs generalize the concept of solution of polynomial matrix equations: indeed, if X satisfies (5), then (X, I) is a left invariant pair for $P(z)$. Moreover, the following properties hold [GLR09].

Lemma 1. *Let (X, U) be a left invariant pair for a regular matrix polynomial $P(z)$. Then, the following properties hold.*

---

[1] Eigenvalues at infinity are a useful algebraic abstraction that makes several counting and transformation arguments work with little or no modification in a more general setting. We do not venture here in the theory of *geometric multiplicity* and *Jordan structure* for matrix polynomials, which is not a trivial task [GLR09].

1. *The eigenvalues of X are a subset of the finite eigenvalues of $P(z)$ (both counted with their multiplicity);*

2. *if*

$$X = \begin{bmatrix} X_{11} & 0 \\ X_{21} & X_{22} \end{bmatrix}, \quad U = \begin{bmatrix} U_1 \\ U_2 \end{bmatrix},$$

   with X_{11} and X_{22} square and U partitioned conformably, then (X_{11}, U_1) is another invariant pair for $P(z)$;

3. *(X, U) is an invariant pair also for $P(z)S(z)$, where $S(z) \in \mathbb{C}[z]^{n \times n}$ is any other regular matrix polynomial.*

Remark 1. A feature that distinguishes linear eigenvalue problems (i.e., the case in which $g = 1$, or equivalently $V = 0$ in the case of (2)) from polynomial ones is the fact that in a non-linear eigenproblem *eigenvectors do not uniquely determine their associated eigenvalues.* For instance, both $(1, e_1)$ and $(2, e_1)$, with $e_1 = \begin{bmatrix} 1 & 0 \end{bmatrix}$, are left eigenpairs of the matrix polynomial

$$P(z) = z^2 I - z \begin{bmatrix} 3 & 0 \\ 0 & 7 \end{bmatrix} + \begin{bmatrix} 2 & 0 \\ 0 & 12 \end{bmatrix}.$$

Hence, we have to deal explicitly with both elements U and X of the invariant pair, and compute them both at the same time. Instead, in the algorithms for the case $V = 0$ [NP15, XXL12], it is common to deal with the matrix Ψ in (4) as the only unknown, and then compute X from it afterwards. This is possible in the first-order case, but not in the second-order one. This point will prove crucial in Section 3.2, where the need to compute X as well will impose a nontrivial restriction not present in the linear case.

2.2   Triplet representations and accurate matrix exponentials

As stated earlier, we are interested in performing numerical computations in a way to guarantee the componentwise accuracy of the computed quantities. To this purpose, the main resource are so-called *subtraction-free* algorithms in linear algebra: when the matrices and vectors involved have a prescribed sign structure, it is possible to carry out linear algebraic operations on a computer in terms of sums only, without ever subtracting two floating-point numbers with the same sign. In this case, there is no cancellation, and the results are provably accurate. The most famous algorithm in this class is the *GTH algorithm* [GTH85, O'C93, AXY02] and its generalizations. To introduce it, we need a few preliminary concepts.

   Here and in the following, inequalities between matrices and vectors are used in the componentwise sense: for instance, $A \leq B$ means $A_{ij} \leq B_{ij}$ for each i, j.

   For a matrix $M \in \mathbb{R}^{n \times n}$, we use the notation offdiag(M) to denote a vector $\boldsymbol{m} \in \mathbb{R}^{n^2 - n}$ which contains the elements $\{M_{ij} : i \neq j\}$, i.e., those which do not

belong to the main diagonal. The exact ordering of these elements in $\boldsymbol{m}$ is not important here. A matrix M is called *M-matrix* if it can be expressed as

$$M = sI - P, \quad P \in \mathbb{R}_{\geq 0}^{n \times n}, s \geq 0. \tag{8}$$

where $s \in \mathbb{R}$ is greater or equal than the spectral radius $\rho(P)$. It is well known that if an M-matrix M is invertible, then $M^{-1} \geq 0$ [BP94].

A *triplet representation* for an M-matrix M is a triple $(\boldsymbol{m}, \boldsymbol{v}, \boldsymbol{w}) \in \mathbb{R}_{\leq 0}^{n^2-n} \times \mathbb{R}_{>0}^n \times \mathbb{R}_{\geq 0}^n$ such that $\boldsymbol{m} = \text{offdiag}(M)$, and $\boldsymbol{v} > \boldsymbol{0}$, $\boldsymbol{w} \geq \boldsymbol{0}$ are two vectors such that $M\boldsymbol{v} = \boldsymbol{w}$. The diagonal elements of M do not appear explicitly in the triplet representation, but they are determined uniquely from the relation $M\boldsymbol{v} = \boldsymbol{w}$. Not all M-matrices admit triplet representations [Guo13, Section 1]; a counterexample is $M = \left[\begin{smallmatrix} 0 & 0 \\ -1 & 0 \end{smallmatrix} \right]$. M-matrices that admit a triplet representation are called *regular M-matrices*. Non-regular M-matrices must necessarily be singular and reducible [Guo13], so most M-matrices appearing in applications (and, in particular, all those appearing in the rest of this paper) are indeed regular.

The following result shows that one can solve linear systems with a regular M-matrix with almost perfect componentwise accuracy, given a triplet representation as input.

Theorem 2 ([O'C93, AXY02])**.** *Let* $(\boldsymbol{m}, \boldsymbol{v}, \boldsymbol{w})$*, be a triplet representation for an invertible regular M-matrix* $M \in \mathbb{R}^{n \times n}$*, and* $\boldsymbol{u} \in \mathbb{R}_{\geq 0}^n$*. Then, there is an algorithm to compute in* $O(n^3)$ *floating-point arithmetical operations (starting from the floating-point numbers* $\boldsymbol{m}, \boldsymbol{v}, \boldsymbol{w}, \boldsymbol{u}$*) an approximation* $\widetilde{\boldsymbol{x}}$ *of* $\boldsymbol{x} = M^{-1}\boldsymbol{u} \geq \boldsymbol{0}$ *such that*

$$|\widetilde{\boldsymbol{x}} - \boldsymbol{x}| \leq \big(\psi(n)\mathrm{u} + \mathcal{O}(\mathrm{u}^2)\big)\, \boldsymbol{x}, \tag{9}$$

with $\psi(n) = \frac{2}{3}(2n+5)(n+2)(n+3)$ *and* u *being the machine precision.*

Notice the remarkable absence of the condition number of M, which would be necessary in an error bound for an algorithm that uses the matrix entries rather than a triplet representation. The use of a triplet representation as an input makes it possible to solve a linear system with perfect accuracy (up to a polynomial function of the dimension), regardless of its condition number.

The algorithm basically works by computing a LU decomposition of M, in which both L and U are M-matrices. Using variants of the same algorithm, one can also perform other related tasks, again starting from a triplet representation $(\boldsymbol{m}, \boldsymbol{v}, \boldsymbol{w})$ of a regular M-matrix M:

- computing M^{-1};

- solving linear systems of the form $M^\top \boldsymbol{x} = \boldsymbol{b}$, with $\boldsymbol{b} \geq \boldsymbol{0}$;

- finding the left and right kernel of a singular irreducible M.

We shall also need the following result.

Lemma 3. *Let*

$$\left(\text{offdiag}(M), \begin{bmatrix} \boldsymbol{v}_1 \\ \boldsymbol{v}_2 \end{bmatrix}, \begin{bmatrix} \boldsymbol{w}_1 \\ \boldsymbol{w}_2 \end{bmatrix}\right), \quad M = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix}$$

(where all the matrices are partitioned conformably) be a triplet representation for the regular M-matrix M. Then,

$$(\text{offdiag}(M_{22}), \boldsymbol{v}_2, \boldsymbol{w}_2 - M_{21}\boldsymbol{v}_1), \tag{10}$$

$$(\text{offdiag}(S), \boldsymbol{v}_1, \boldsymbol{w}_1 - M_{12}M_{22}^{-1}\boldsymbol{w}_2) \tag{11}$$

are subtraction-free expressions for triplet representations of a principal submatrix M_{22} and its Schur complement $S = M_{11} - M_{12}M_{22}^{-1}M_{21}$.

Proof. The relation $M_{22}\boldsymbol{v}_2 = \boldsymbol{w}_2 - M_{21}\boldsymbol{v}_1$ which defines the first triplet representation comes from expanding the second block row of $M\boldsymbol{v} = \boldsymbol{w}$. The second relation comes from premultiplying both sides of $M\boldsymbol{v} = \boldsymbol{w}$ by

$$\begin{bmatrix} I & -M_{12}M_{22}^{-1} \\ 0 & I \end{bmatrix}.$$

Additionally, note that $M_{12}, M_{21} \leq 0$ and $M_{22}^{-1} \geq 0$ (which can be obtained in a subtraction-free way using the GTH algorithm and the triplet representation (10)), so computing the last terms in (10) and (11) does not involve subtractions. The computation of S via the formula $S = M_{11} - M_{12}M_{22}^{-1}M_{21}$ involves subtractions only for the diagonal entries, but conveniently in the triplet representation we only need offdiag(S). $\qquad \square$

Corollary 4. Given a triplet representation for M, diag(M) can be reconstructed using subtraction-free formulas only.

Proof. It is sufficient to consider (10) in the case in which M_{22} is 1×1. Then, $M_{22} = (\boldsymbol{w}_2 - M_{21}\boldsymbol{v}_1)/\boldsymbol{v}_2$, where the numerator and denominator are scalar quantities, too. $\qquad \square$

We comment briefly also on the computation of the matrix exponential, which we shall need only in the final step of our algorithm. For an M-matrix M, it holds that $\exp(-M) \geq 0$. As studied in [XY08], it is impossible to find an unconditionally accurate algorithm for this computation in the style of the GTH algorithm; we can only compute approximations $\widetilde{E}$ of $E = \exp(-M)$ satisfying a bound of the form

$$|\widetilde{E} - E| = c(M)\mathcal{O}(\mathrm{u})E,$$

where $c(M)$ is a condition number which depends explicitly on M. Algorithms for the componentwise accurate computation of matrix exponentials were discussed in [SGX14]; one of the first steps in these methods is decomposing $M = sI - P$ as in (8) and using the identity $\exp(-M) = e^{-s}\exp(P)$. Hence it is appealing to look for explicit accurate decompositions of the form (8) for the matrices whose exponentials we are going to compute.

2.3   Quadratic matrix equations and their properties

We now discuss the properties of the solutions of matrix equations of the form $A - BX + CX^2 = 0$. We focus here on the most common setting in its probabilistic applications; namely, we assume that

A4 $A \geq 0$, $C \geq 0$ and B is an M-matrix;

A5 $(A - B + C)\mathbf{1} = \mathbf{0}$;

A6 the bi-infinite matrix

$$
\begin{bmatrix}
\ddots & \ddots & & & \\
\ddots & -B & C & & \\
& A & -B & C & \\
& & A & -B & \ddots \\
& & & \ddots & \ddots
\end{bmatrix}
\tag{12}
$$

is irreducible and aperiodic.

The case most frequently appearing in the probability applications [BLM05] is the one in which $I - B \geq 0$ and $A, I - B, C$ are the transition matrices of a quasi-birth-death (QBD) process, with A being the transition to a lower level.

Under Assumptions A4, A5, A6, one can prove that $A - B + C$ is an irreducible singular M-matrix; we call its left Perron vector $\boldsymbol{u} \in \mathbb{R}^{1 \times n}$; then we have $\boldsymbol{u} > \mathbf{0}$, $\boldsymbol{u}(A - B + C) = \mathbf{0}$. Moreover, one can prove the following results [BLM05].

Theorem 5. *Let $A, B, C \in \mathbb{R}^{n \times n}$ satisfying Assumptions A4, A5, A6. Then, the following matrices exist:*

$$G \in \mathbb{R}^{n \times n} \text{such that } A - BG + CG^2 = 0, \ G \geq 0 \text{ and } G\mathbf{1} \leq \mathbf{1}; \tag{13a}$$

$$R \in \mathbb{R}^{n \times n} \text{such that } R^2 A - RB + C = 0, \ R \geq 0 \text{ and } \boldsymbol{u}R \leq \boldsymbol{u}. \tag{13b}$$

Moreover, the location in the complex plane of the eigenvalues of $F(y) = Ay^2 - By + C$, and of G and R, is related to the sign of the mean drift $d = \boldsymbol{u}(C - A)\mathbf{1}$ *as described in Table 1.*

The letters $\mathcal{S}$ and $\mathcal{U}$ in the table stand for 'stable' and 'unstable', respectively, while d stands for 'discrete time' and c in the following will stand for 'continuous time'.

Assumption A6 can be relaxed to the less stringent one where tridiag$(A, -B, C)$ has only one final class [BLM05, Section 4.7], with only some minor technical complications.

Case	Name	$\|\mathcal{S}_d\|$	$\|\mathcal{U}_d\|$	Other eigvls.	Eigvls. of G	Eigvls. of R
$d < 0$	Positive recurrent	n	$n-1$	1 (mult. 1)	$\{\lambda^{-1} : \lambda \in \mathcal{U}_d\} \cup \{1\}$	$\mathcal{S}_d$
$d = 0$	Null recurrent	$n-1$	$n-1$	1 (mult. 2)	$\{\lambda^{-1} : \lambda \in \mathcal{U}_d\} \cup \{1\}$	$\mathcal{S}_d \cup \{1\}$
$d > 0$	Transient	$n-1$	n	1 (mult. 1)	$\{\lambda^{-1} : \lambda \in \mathcal{U}_d\}$	$\mathcal{S}_d \cup \{1\}$

Tab. 1: Cardinality of the multisets $\mathcal{S}_d = \{\lambda : \lambda$ is an eigenvalue of $F(y)$ and $|\lambda| < 1$ $\}$ and $\mathcal{U}_d = \{\lambda : \lambda$ is an eigenvalue of $F(y)$ and $|\lambda| > 1\}$, and eigenvalues of the four solution matrices in three possible cases for a triple A, B, C satisfying A4, A5, A6.

2.4 Cyclic Reduction

Cyclic Reduction (CR) is the following matrix iteration. Set

$$A_0 = A, \quad B_0 = \hat{B}_0 = B, \quad C_0 = C,$$

and compute for each $k = 0, 1, 2, \ldots$

$$A_{k+1} = A_k B_k^{-1} A_k, \tag{14a}$$

$$B_{k+1} = B_k - A_k B_k^{-1} C_k - C_k B_k^{-1} A_k, \tag{14b}$$

$$C_{k+1} = C_k B_k^{-1} C_k, \tag{14c}$$

$$\hat{B}_{k+1} = \hat{B}_k - C_k B_k^{-1} A_k. \tag{14d}$$

The following applicability and convergence results hold for Cyclic Reduction.

Theorem 6. *Let $A, B, C \in \mathbb{R}^{n \times n}$ satisfying Assumptions A4, A5, A6. Then,*

1. *B_k is nonsingular for $k \geq 0$; hence, CR can be applied with no breakdown.*

2. *A_k, C_k are nonnegative, and B_k and $\hat{B}_k$ are M-matrices for $k \geq 0$.*

3. *B_k and $\hat{B}_k$ converge monotonically to matrices that we shall call B_∞ and $\hat{B}_\infty$, respectively. The matrix $\hat{B}_\infty$ is invertible.*

4. *We have*

$$G = \hat{B}_\infty^{-1} A_0, \tag{15a}$$

$$R = C_0 \hat{B}_\infty^{-1}. \tag{15b}$$

5. *The convergence speed is linear with factor $1/2$ in the null recurrent case, quadratic with factor $\rho(R) < 1$ in the positive recurrent case, and quadratic with factor $\rho(G) < 1$ in the transient case.*

6. *$(A_k - B_k + C_k)\mathbf{1} = \mathbf{0}$ for each $k \geq 0$, hence $(\text{offdiag}(B_k), \mathbf{1}, (A_k + C_k)\mathbf{1})$ is a triplet representation for B_k.*

The last item in particular is useful because it allows one to perform the iteration using the GTH algorithm for the inversions required in (14). Hence, $A_k, \mathrm{offdiag}(B_k), C_k, \mathrm{offdiag}(\hat{B}_k)$ can be computed in a subtraction-free fashion. This is how Cyclic Reduction is currently implemented in software packages such as SMCSolver [BMS09]. However, to implement the final step, (15), in a subtraction-free way, we need to find a triplet representation for $\hat{B}_\infty$. To this purpose, we give the following result.

Lemma 7. *Under Assumptions A4, A5, A6, the following results hold for the iterates of Cyclic Reduction.*

- $(A_0 - \hat{B}_k + C_k)\mathbf{1} = \mathbf{0}$ *for each* $k \geq 0$, *hence*

$$\lim_{k\to\infty} C_k \mathbf{1} =: \hat{w} \quad \textit{exists},$$

  and $(\mathrm{offdiag}(\hat{B}_k), \mathbf{1}, A_0\mathbf{1} + \hat{w})$ *is a triplet representation for* $\hat{B}_\infty$.

- $\boldsymbol{u}(A_k - B_k + C_k) = \mathbf{0}$ *for each* $k \geq 0$.

- $\boldsymbol{u}(A_k - \hat{B}_k + C_0) = \mathbf{0}$ *for each* $k \geq 0$, *hence*

$$\lim_{k\to\infty} \boldsymbol{u} A_k =: \hat{v} \textit{ exists},$$

  and $(\mathrm{offdiag}(\hat{B}^\top), \boldsymbol{u}^\top, (\boldsymbol{u}C_0 + \hat{v})^\top)$ *is a triplet representation for* $\hat{B}_\infty^\top$.

Proof. We prove only the first equality, the others are analogous. The proof is by induction and similar to the one of item 6 of Theorem 6. For $k = 0$, the result holds by Assumption A5. The inductive step is

$$\begin{aligned}
(A_0 - \hat{B}_{k+1} + C_{k+1})\mathbf{1} &= (A_0 - \hat{B}_k + C_k B_k^{-1} A_k + C_k B_k^{-1} C_k)\mathbf{1} \\
&= (A_0 - \hat{B}_k)\mathbf{1} + C_k B_k^{-1}(A_k + C_k)\mathbf{1} \\
&= (A_0 - \hat{B}_k)\mathbf{1} + C_k B_k^{-1} B_k \mathbf{1} \\
&= (A_0 - \hat{B}_k + C_k)\mathbf{1} = \mathbf{0},
\end{aligned}$$

where we have used the fact that $(A_k + C_k)\mathbf{1} = B_k \mathbf{1}$ (item 6 of Theorem 6). $\square$

Armed with these triplet representations, we can formulate a fully subtraction-free version of Cyclic Reduction, Algorithm 1.

3 Derivation of the algorithm

In this section, we focus on the problem of finding a left invariant pair (X, U) for the matrix polynomial $P(z)$ in (2) associated to its eigenvalues in the left half-plane. We shall see in Section (3.8) that a solution to (1) can be constructed from this pair.

Algorithm 1: A subtraction-free version of Cyclic Reduction using triplet representations.

    Input: $A, B, C \in \mathbb{R}^{n \times n}$ satisfying A4, A5, A6
    Output: The matrices G, R defined in (13).
1 Set $A_0 = A$, $B_0 = \hat{B}_0 = B$, $C_0 = C$, and $k = 0$;
2 **repeat**
3   Compute $A_{k+1}, \text{offdiag}(B_{k+1}), C_{k+1}, \text{offdiag}(\hat{B}_{k+1})$ using (14),
        performing inversions using the triplet representation in Item 6 of
        Theorem 6;
4 $k \to k + 1$;
5 **until** $\text{offdiag}(\hat{B}_k)$ *has converged*;
6 Compute G, R using (15), performing inversions using one of the triplet representations in Lemma 7;

3.1   The spectrum of $P(z)$

We start with a theoretical result on the location of the eigenvalues of $P(z)$. To formulate it, we subdivide the indices $i \in \{1, 2, \dots, n\}$ into three disjoint subsets, according to the values of v_{ii} and d_{ii}, as shown in Table 2. Moreover, we set $n_i = |E_i|$ for $i = 1, 2, 3$, so that $n = n_1 + n_2 + n_3$, and we call $\boldsymbol{u} > \boldsymbol{0}$ the left Perron vector of Q.

Name	v_{ii}	d_{ii}
E_1	> 0	any
E_2	$= 0$	> 0
E_3	$= 0$	< 0

Tab. 2: Subdivision of each state $i \in \{1, 2, \dots, n\}$ into three different sets. Recall that we assume that there is no state with $v_{ii} = d_{ii} = 0$.

Theorem 8. *The location in the complex plane of the eigenvalues of $P(z)$ is related to the sign of the* mean drift $d = \boldsymbol{u}D\boldsymbol{1}$ *as described in Table 3.*

Proof. The case $d < 0$ appears in [KK95]; the case $d > 0$ can be proved by replacing D with $-D$ (which has the effect of changing the sign of all eigenvalues).

    For the case $d = 0$, the proof is not immediate; we give only a sketch, since the result is not necessary for the rest of the paper. A limit argument from both sides shows that $|\mathcal{S}_c| \leq n_1 + n_2 - 1$ and $|\mathcal{S}_u| \leq n_1 + n_3 - 1$. Since we assume that v_{ii} and d_{ii} are not both zero, $D_{E_2 \cup E_3, E_2 \cup E_3}$ is nonsingular, and hence the $n_2 + n_3$ Jordan chains for $\lambda = \infty$ (as defined in [GLR09, Section 1.4]) have length 1. Hence the only thing left to prove is that there are no more than 2 eigenvalues on the imaginary axis. The Gerschgorin argument in [KK95] shows that the only possible eigenvalue on the imaginary axis is zero. The multiplicity of the eigenvalue 0 is at most 2, because for h small enough the matrix polynomial

$I + hP(z)$ satisfies the hypotheses of Theorem 5, as we show in more detail in the following. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Case	Name	$\lvert\mathcal{S}_c\rvert$	$\lvert\mathcal{U}_c\rvert$	Other eigenvalues
$d < 0$	Pos. rec.	$n_1 + n_2$	$n_1 + n_3 - 1$	0 (mult. 1), ∞ (mult. $n_2 + n_3$)
$d = 0$	Null rec.	$n_1 + n_2 - 1$	$n_1 + n_3 - 1$	0 (mult. 2), ∞ (mult. $n_2 + n_3$)
$d > 0$	Transient	$n_1 + n_2 - 1$	$n_1 + n_3$	0 (mult. 1), ∞ (mult. $n_2 + n_3$)

Tab. 3: Cardinality of $\mathcal{S}_c = \{\lambda : \lambda \text{ is an eigenvalue of } P(z) \text{ and } \Re\lambda < 0\}$ and $\mathcal{U}_c = \{\lambda : \lambda \text{ is an eigenvalue of } P(z) \text{ and } \Re\lambda > 0\}$.

We now have all we need to define precisely which invariant pair we are looking for. We call a left invariant pair of $P(z)$ *c-stable*, if its associated eigenvalues are:

- the $n_1 + n_2$ eigenvalues in $\mathcal{S}_c$, if $P(z)$ is positive recurrent; or

- the $n_1 + n_2 - 1$ eigenvalues in $\mathcal{S}_c$ and the eigenvalue 0 with multiplicity 1, if $P(z)$ is null recurrent or transient.

Similarly, with respect to Table 1, we call a left invariant pair of $F(y)$ *d-stable*, if its associated eigenvalues are:

- the n eigenvalues in $\mathcal{S}_d$, if $F(y)$ is positive recurrent; or

- the $n - 1$ eigenvalues in $\mathcal{S}_d$ and the eigenvalue 1 with multiplicity 1, if $F(y)$ is null recurrent or transient.

Notice that (R, I), where R is the matrix in (13), is a d-stable invariant pair for $F(y)$.

3.2 The case $\operatorname{diag}(V) > 0$

We start by treating the simpler case in which $\operatorname{diag}(V) > 0$ (or, in probabilistic terms, the dynamic in all states has a Brownian motion component). We have $E_2 = E_3 = \varnothing$ and $n_2 = n_3 = 0$, all $2n$ eigenvalues of $P(z)$ are finite and we are looking for exactly n of them. The formulation in (4) has $U = I$, hence the task of finding an invariant pair becomes the one of finding a solution of the matrix equation (5).

We have seen that Cyclic Reduction can be applied to matrix polynomials $F(y)$ with a specific sign structure, which is associated with a specific spectral structure as shown in Table 1. The sign structure and spectral structure of $P(z)$ in (3) do not match these requirements, so we need some preprocessing to convert one case into the other. Even if the sign structure is a stricter requirement, it is useful for our analysis to focus first on the spectral structure, and describe methods of altering the position of the eigenvalues.

We start from a general lemma on rational transformations of matrix polynomials.

Lemma 9 ([MMMM15, Nof12]). *Let*

$$y = f(z) = \frac{\alpha z + \beta}{\gamma z + \delta}$$

be a degree-1 (scalar) rational function, with $\alpha, \beta, \gamma, \delta \in \mathbb{C}$ and $\alpha\delta \neq \beta\gamma$, $z = f^{-1}(y) = (\delta y - \beta)/(\alpha - \gamma y)$ its inverse, and $P(z) \in \mathbb{R}[z]^{n \times n}$ be a degree-g regular matrix polynomial with eigenvalues $\lambda_1, \dots, \lambda_{gn}$ (counted with multiplicity, and possibly including ∞). Then, the following properties hold.

1. *The matrix polynomial*

   $$F(y) = (\alpha - \gamma y)^g P(\delta y - \beta/\alpha - \gamma y)$$

   is regular and has eigenvalues $f(\lambda_i)$, for each $i = 1, 2, \dots, gn$.

2. *If (Z, U) is a left invariant pair for $P(z)$ and $\gamma Z + \delta I$ is nonsingular, then $(f(Z), U)$ is an invariant pair for $F(y)$. Conversely, if (Y, U) is an invariant pair for $F(y)$ and $\alpha I - \gamma Y$ is nonsingular, then $(f^{-1}(Y), U)$ is an invariant pair for $P(z)$.*

Note that $f(Z) = (\gamma Z + \delta I)^{-1}(\alpha Z + \beta I) = (\alpha Z + \beta I)(\gamma Z + \delta I)^{-1}$ is well-defined for a matrix argument Z since the two factors commute, and similarly for $f^{-1}(Y)$. If $U = I$, the last item gives a relation between the solutions to the unilateral matrix equations associated to $P(z)$ and $F(y)$.

Lemma 9 suggests a general strategy to approach the problem:

1. Choose a function f such that $f(0) = 1$ and the images of $\mathcal{S}_c, \mathcal{U}_c$ lie inside and outside the unit circle, respectively.

2. Construct $F(y) = Ay^2 - By + C$ as in Lemma 9.

3. Apply Cyclic Reduction to find the solution R to

   $$R^2 A - RB + C = 0;$$

   then, (R, I) is a d-stable invariant pair of $F(y)$.

4. Compute $X = f^{-1}(R)$.

This general framework of relocating eigenvalues via rational transformations is quite common in literature; see for instance [BMP10] for a discussion of it in the case of fluid queues ($V = 0$). Frequent choices for f are

$$y = 1 + hz, \quad \text{and} \quad y = \frac{1 + hz}{1 - hz},$$

where $h > 0$ is a parameter. However, some care is needed here to allow for componentwise accurate computations within the framework. The first important restriction comes from the last step: once we have obtained $R \geq 0$, we need to be able to compute $f^{-1}(R)$. If one chooses $y = 1 + hz/1 - hz$, the computation

becomes $X = h^{-1}(I - R)(I + R)^{-1}$. This is problematic, because the matrix $I + R$, which we need to invert, is a nonnegative matrix; hence Theorem 2 does not apply, and we do not know of another componentwise algorithm to invert matrices with this sign pattern, even if triplet representations are available.

Things are easier if one chooses the function $y = 1 + hz$. In this case, the last step becomes $X = h^{-1}(R - I)$; subtractions are needed only to compute its diagonal, and we can avoid them completely using Corollary 4 if we manage to obtain a triplet representation for $-X$. Moreover, the matrix is now naturally expressed in the form (8). For this reason, we set $y = f(z) = 1 + hz$ in the following.

With this choice, we get $z = (y - 1)/h$, and

$$P(z) = V\frac{(y-1)^2}{h^2} - D\frac{y-1}{h} + Q$$
$$= \frac{1}{h^2}Vy^2 - \left(2\frac{1}{h^2}V + \frac{1}{h}D\right)y + \left(\frac{1}{h^2}V + \frac{1}{h}D + Q\right),$$

hence

$$A := \frac{1}{h^2}V, \quad B := 2\frac{1}{h^2}V + \frac{1}{h}D, \quad C := \frac{1}{h^2}V + \frac{1}{h}D + Q. \qquad (16)$$

Once we have decided to use (16) to convert the setting into that of a discrete-time quadratic matrix equations, we have to choose the value of the parameter h. A first requirement is that Assumption A4 is satisfied; it is easy to see that it holds provided that $v_{ii} - d_{ii}h + q_{ii}h^2 \geq 0$ for each i. Since we assume $v_{ii} > 0$ for each $i \in \mathcal{S}$ for now, this holds for sufficiently small values of h. Moreover, we have to ensure that the computed diagonal of C is componentwise accurate. For this, we follow the strategy used in [XXL12]: we choose h small enough so that all the required subtractions are of the form $b - a$ with $b \geq 2a \geq 0$. In this case, there cannot be catastrophic cancellation in the subtraction in machine arithmetic. This requirement translates to the following constraints on h:

$$v_{ii} \geq -2(d_{ii}h + q_{ii}h^2), \text{ for each } i \text{ with } v_{ii} > 0 \text{ and } d_{ii} < 0; \qquad (17a)$$
$$v_{ii} + hd_{ii} \geq -2q_{ii}h^2, \text{ for each } i \text{ with } d_{ii} > 0. \qquad (17b)$$

All these inequalities are satisfied for a sufficiently small value of h, which is easy to compute explicitly. Another possibility is performing these subtractions using machine arithmetic with a higher precision; since there are only $O(n)$ of them, this safeguard will not impact the final cost of the algorithm.

It is easy to check that Assumption A5 is always satisfied. We prove below that A6 is satisfied as well.

Theorem 10. *Suppose the matrix Q is irreducible and aperiodic, and $V \neq 0$. Then, Assumption A6 holds for the matrices $A, B,$ and C defined in* (16).

Proof. We identify each element in the index set of the infinite matrix (12) with a pair $(i, \ell) \in \{1, 2, \ldots, n\} \times \mathbb{Z}$, where the second entry denotes the block (level)

and the first denotes the position in the block. We shall prove that there is a walk in the graph associated to (12) between any two states $(i, \ell), (j, m) \in \{1, 2, \ldots, n\} \times \mathbb{Z}$.

Let $k \in \{1, 2, \ldots, n\}$ be such that $v_{kk} > 0$. By the irreducibility assumption, we can find in the graph associated to Q a walk from i to j which passes through k and has length at least $m - \ell$. Since C has the same offdiagonal nonzero structure as Q, the same walk can be used in the matrix (12), and after each step the second element of the pair goes up by one. Hence the path goes from (i, ℓ) to $(j, m + p)$, for some $p > 0$. We modify this walk by inserting p transitions using the nonzero entry A_{kk} when we first reach k as the first element of the pair. The resulting graph goes from (i, ℓ) to (j, m), as requested.

If Q is aperiodic, the same construction can be made with different values of p which are coprime; hence (16) is aperiodic, too. □

Moreover, we can prove that our transformation (16) preserves the sign of the mean drift.

Lemma 11. *Let d_c be the mean drift of the Markov-modulated Brownian motion process with parameters V, D, Q, and d_d be the mean drift of the QBD process associated to A, B, C as in (16). Then, $d_d = h^{-1} d_c$.*

Proof. First note that $A - B + C = Q$, so the left Perron vector $\boldsymbol{u}$ of $A - B + C$ coincides with the one of Q. Then it is easy to compute

$$d_d = \boldsymbol{u}(C - A)\mathbf{1} = \boldsymbol{u}\left(\frac{1}{h^2}V + \frac{1}{h}D + Q - \frac{1}{h^2}V\right)\mathbf{1} = \frac{1}{h}\boldsymbol{u}D\mathbf{1} = 1hd_c.$$

□

Finally, we note that as a byproduct of Cyclic Reduction (Algorithm (1)) we can obtain explicitly a triplet representation for $-X^\top$.

Theorem 12. *The triplet*

$$\left(\mathrm{offdiag}(-X^\top), \boldsymbol{u}^\top, \left(h^{-1}\hat{\boldsymbol{v}}\hat{B}_\infty^{-1}\right)^\top\right) \tag{18}$$

is a triplet representation for the matrix $-X^\top$, where $X = h^{-1}(R - I)$.

Proof. By Lemma 7, we have $\hat{\boldsymbol{v}} - \boldsymbol{u}\hat{B}_\infty + \boldsymbol{u}C_0 = \mathbf{0}$. Hence,

$$-\boldsymbol{u}X = \boldsymbol{u}\frac{1}{h}(I - C_0\hat{B}_\infty^{-1}) = \frac{1}{h}\hat{\boldsymbol{v}}\hat{B}_\infty^{-1}.$$

□

Summing up everything, our algorithm for the case $\mathrm{diag}(V) > 0$ is described in Algorithm 2.

Algorithm 2: Computing a c-stable invariant pair (X, I) of $P(z)$, in the case $\operatorname{diag}(V) > 0$

    Input: $V, D \in \mathbb{R}^{n \times n}$ diagonal matrices with $\operatorname{diag}(V) > \mathbf{0}$, $Q \in \mathbb{R}^{n \times n}$ a generator matrix ($Q\mathbf{1} = \mathbf{0}$, $\operatorname{offdiag}(Q) \geq \mathbf{0}$), satisfying A1, A2, A3.

    Output: the matrix X (or a decomposition (8) for it), and the triplet representation (18) for $-X^\top$.

1 compute the left Perron vector $\boldsymbol{u}$ of Q using the triplet representation ($\operatorname{offdiag}(-Q), \mathbf{1}, \mathbf{0}$);

2 choose h small enough so that $h^{-2}v_{ii} + h^{-1}d_{ii} + q_{ii} > 0$, and it can be computed without catastrophic cancellation;

3 compute A, B, C as in (16);

4 compute $R \geq 0$ via Algorithm 1;

5 using the last iterate A_k computed by Algorithm 1 and the triplet representation for $\hat{B}_\infty$, compute $\hat{\boldsymbol{v}} = \boldsymbol{u}A_k$ and $h^{-1}\hat{\boldsymbol{v}}\hat{B}_\infty^{-1}$;

6 compute $X = h^{-1}(R - I)$ (or $P = h^{-1}R$ and $s = h^{-1}$);

3.3 Shifting infinite eigenvalues in $P(z)$

The method outlined in the previous section uses the assumption that $v_{ii} > 0$ for each i. When this is not the case, it is not true in general that we can choose h small enough to have $h^{-2}v_{ii} + h^{-1}d_{ii} + q_{ii} \geq 0$. This is possible for $i \in E_2$, since $d_{ii} > 0$, but if E_3 is not empty the algorithm cannot be applied.

Moreover, if V is singular, then the matrix polynomial $P(z)$ has infinite eigenvalues, and a c-stable invariant pair (X, I), with X of size $n \times n$, cannot be constructed since even in the positive recurrent case $P(z)$ does not have n eigenvalues in the left half-plane.

Finally, the discretization methods outlined in the previous section all break down in some way: if we use the map $y = 1 + hz$, then we cannot enforce the requirement that the eigenvalue $z = \infty$ is mapped inside the unit circle by choosing a small enough h; if we use a variant of the Cayley transform, then $f(\infty) = -1$, and we are left with an eigenvalue of $F(y)$ at -1, possibly with high multiplicity; this eigenvalue often prevents the convergence of Cyclic Reduction (note indeed that we are not in the hypotheses of Theorem 6). All these issues are related, and indeed we can solve all of them with the same modification to the algorithm.

We subdivide the parameter matrices into blocks corresponding to E_1, E_2, E_3 as

$$
V = \begin{bmatrix} V_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad D = \begin{bmatrix} D_1 & 0 & 0 \\ 0 & D_2 & 0 \\ 0 & 0 & D_3 \end{bmatrix}, \quad Q = \begin{bmatrix} Q_{11} & Q_{12} & Q_{13} \\ Q_{21} & Q_{22} & Q_{23} \\ Q_{31} & Q_{32} & Q_{33} \end{bmatrix}, \quad (19)
$$

where $V_1 > 0$, $D_2 > 0$ and $D_3 \leq 0$ are diagonal matrices.

We define

$$\widetilde{P}(z) := P(z)S(z), \quad S(z) := \begin{bmatrix} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & (1+hz)I \end{bmatrix}.$$

The resulting matrix polynomial $\widetilde{P}(z) = \widetilde{V}z^2 - \widetilde{D}z + \widetilde{Q}$ has coefficients

$$\widetilde{V} := \begin{bmatrix} V_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -hD_3 \end{bmatrix}, \quad \widetilde{D} := \begin{bmatrix} D_1 & 0 & -hQ_{13} \\ 0 & D_2 & -hQ_{23} \\ 0 & 0 & D_3 - hQ_{33} \end{bmatrix}, \quad \widetilde{Q} := Q.$$

Every finite eigenvalue $\lambda \neq \infty$ of $P(z)$ is also an eigenvalue of $\widetilde{P}(z)$ (with the same left eigenvector), while n_3 infinite eigenvalues are replaced by eigenvalues $-h^{-1}$. This can be readily proved by considering the determinants $\det \widetilde{P}(z) = \det P(z) \det S(z)$ and their degrees. This formulation of shifting as multiplication by a suitable matrix polynomial has been suggested recently in [BM15].

Remark 2. We can interpret this transformation as a manipulation of the differential equation (1). Indeed, if we subdivide $\boldsymbol{p}(x) = \begin{bmatrix} \boldsymbol{p}_1(x) & \boldsymbol{p}_2(x) & \boldsymbol{p}_3(x) \end{bmatrix}$ conformably, then the third block equation reads

$$-D_3\boldsymbol{p}_3'(x) + Q_{31}\boldsymbol{p}_1(x) + Q_{32}\boldsymbol{p}_2(x) + Q_{33}\boldsymbol{p}_3(x) = \boldsymbol{0}; \tag{20}$$

differentiating this equation gives

$$-D_3\boldsymbol{p}_3''(x) + Q_{31}\boldsymbol{p}_1'(x) + Q_{32}\boldsymbol{p}_2'(x) + Q_{33}\boldsymbol{p}_3'(x) = \boldsymbol{0}; \tag{21}$$

then the equation $\boldsymbol{p}''(x)\widetilde{V} - \boldsymbol{p}'(x)\widetilde{D} + \boldsymbol{p}(x)Q = \boldsymbol{0}$ is obtained from (1) by replacing the third block equation (20) with (20) + h(21). This kind of manipulations is commonly used in the context of index reduction techniques [KM06].

We can set up the discretization scheme described in Section 3 starting from $\widetilde{P}(z)$ rather than $P(z)$. The resulting polynomial $\widetilde{F}(y)$ has coefficients

$$\widetilde{A} := \begin{bmatrix} h^{-2}V_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -h^{-1}D_3 \end{bmatrix}, \tag{22a}$$

$$\widetilde{B} := \begin{bmatrix} 2h^{-2}V_1 + h^{-1}D_1 & 0 & -Q_{13} \\ 0 & h^{-1}D_2 & -Q_{23} \\ 0 & 0 & -h^{-1}D_3 - Q_{33} \end{bmatrix}, \tag{22b}$$

$$\widetilde{C} := \begin{bmatrix} h^{-2}V_1 + h^{-1}D_1 + Q_{11} & Q_{12} & 0 \\ Q_{21} & h^{-1}D_2 + Q_{22} & 0 \\ Q_{31} & Q_{32} & 0 \end{bmatrix}. \tag{22c}$$

Notice the nontrivial simplification that zeroes out the last block column of $\widetilde{C}$. Its appearance is due to the fact that the eigenvalues $-h^{-1}$ introduced in the previous step get mapped to $f(-h^{-1}) = 0$, hence $\widetilde{F}(y)$ has n_3 zero eigenvalues.

If one chooses a sufficiently small h, the diagonals of $h^{-2}V_1 + h^{-1}D_1 + Q_{11}$ and $h^{-1}D_2 + Q_{22}$ are nonnegative (and can be computed accurately): it is sufficient to impose (17) on $i \in E_1 \cup E_2$. Hence, the matrices $\widetilde{A}$ and $\widetilde{C}$ are nonnegative, and $\widetilde{B}$ is an M-matrix, which is the correct sign structure to implement subtraction-free Cyclic Reduction (Algorithm 1).

3.4   Deflating zero eigenvalues in R

In the case v_{ii} may be zero, the solution R produced by Cyclic Reduction does not give immediately the invariant pair we need. Indeed, in view of our previous analysis of the eigenvalues of $F(y)$ and $\widetilde{F}(y)$, in the positive recurrent case the eigenvalues of R comprise of

- n_3 zero eigenvalues;

- $f(\lambda)$, for each eigenvalue λ of $P(z)$ with $\Re\lambda < 0$, counted with multiplicity.

The eigenvalues of the form $f(\lambda)$ are precisely the ones we need in our invariant pair, but there are spurious zero eigenvalues. If R were in the form

$$\begin{bmatrix} * & 0 \\ * & 0 \end{bmatrix}, \tag{23}$$

with the bottom-right block $n_3 \times n_3$, we could remove them by applying the result in point 2 of Theorem 1 to the invariant pair (R, I).

Unfortunately, we have $R = C_0 \hat{B}_\infty^{-1}$, where C_0 is in the form (23) and $\hat{B}_\infty$ is a regular M-matrix (for which we know a triplet representation). When one carries out the product, the zero block is lost. To recover it, we have to switch to a different invariant pair.

Theorem 13. *Let the matrix C in (22) and the matrix $\hat{B}_\infty$ produced by Cyclic Reduction on (22) be partitioned as*

$$C_0 = \begin{bmatrix} C_{11} & 0 \\ C_{21} & 0 \end{bmatrix}, \quad \hat{B}_\infty = \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix}, \tag{24}$$

where the bottom-right block of dimension $n_3 \times n_3$ corresponds to the indices in E_3. Then, $(Y, \begin{bmatrix} I & \Psi \end{bmatrix})$, with

$$\Psi = -B_{12}B_{22}^{-1} \geq 0,$$
$$Y = (C_{11} + \Psi C_{21})S^{-1} \geq 0, \quad with \ S := B_{11} + \Psi B_{21}$$

is a subtraction-free expression for a d-stable left invariant pair of $\widetilde{F}(y)$, and

$$(X, \begin{bmatrix} I & \Psi \end{bmatrix}), \quad with \ X = h^{-1}(Y - I) \tag{26}$$

is a subtraction-free expression for a c-stable left invariant pair of $P(z)$.

Proof. We apply to the d-stable left invariant pair (R, I) of $\widetilde{P}(z)$ a transformation of the form (7) with

$$M = \begin{bmatrix} I & \Psi \\ 0 & I \end{bmatrix},$$

obtaining

$$
\begin{aligned}
(MRM^{-1}, M) &= (MC(M\hat{B}_\infty)^{-1}, M) \\
&= \left(\begin{bmatrix} C_{11} + \Psi C_{21} & 0 \\ C_{21} & 0 \end{bmatrix} \begin{bmatrix} B_{11} + \Psi B_{21} & 0 \\ B_{21} & B_{22} \end{bmatrix}^{-1}, \begin{bmatrix} I & \Psi \\ 0 & I \end{bmatrix} \right) \\
&= \left(\begin{bmatrix} (C_{11} + \Psi C_{21})S^{-1} & 0 \\ C_{21}S^{-1} & 0 \end{bmatrix}, \begin{bmatrix} I & \Psi \\ 0 & I \end{bmatrix} \right).
\end{aligned}
$$

Notice that B_{22} and S are respectively a submatrix and a Schur complement of the regular M-matrix $\hat{B}_\infty$, so triplet representations to invert both are available by Lemma 3. We can now apply point 2 of Theorem 1 to obtain that (25) is an invariant pair associated to the d-stable eigenvalues of $F(y)$. Transforming this invariant pair with Lemma 9, we obtain (26). $\qquad\square$

3.5 A triplet representation for $-X^\top$

In this section, we obtain a triplet representation for the M-matrix $-X^\top$ using subtraction-free expressions only.

Theorem 14. *The triplet*

$$
\left(\text{offdiag}(-X^\top), \boldsymbol{u}_1^\top, \frac{1}{h} \left(\left(\hat{\boldsymbol{v}}_1 + \hat{\boldsymbol{v}}_2 B_{22}^{-1}(C_{21} - B_{21}) \right) S^{-1} \right)^\top \right) \tag{27}
$$

is a subtraction-free expression for a triplet representation of $-X^\top$, where X is defined by (26).

Proof. By introducing the partitioning (24) in Lemma 7, we get $\boldsymbol{u}_1 B_{12} + \boldsymbol{u}_2 B_{22} = \hat{\boldsymbol{v}}_2$. Hence, $\boldsymbol{u}_2 = \hat{\boldsymbol{v}}_2 B_{22}^{-1} + \boldsymbol{u}_1 \Psi$, and

$$
\boldsymbol{u} M^{-1} = \begin{bmatrix} \boldsymbol{u}_1 & \boldsymbol{u}_2 \end{bmatrix} \begin{bmatrix} I & -\Psi \\ 0 & I \end{bmatrix} = \begin{bmatrix} \boldsymbol{u}_1 & \hat{\boldsymbol{v}}_2 B_{22}^{-1} \end{bmatrix}. \tag{28}
$$

Again, from Lemma 7, we get

$$
\begin{aligned}
\hat{\boldsymbol{v}} &= \boldsymbol{u}(\hat{B}_\infty - C_0) = \boldsymbol{u} M^{-1}(M\hat{B}_\infty - MC_0) \\
&= \begin{bmatrix} \boldsymbol{u}_1 & \hat{\boldsymbol{v}}_2 B_{22}^{-1} \end{bmatrix} \left(\begin{bmatrix} S & 0 \\ B_{21} & B_{22} \end{bmatrix} - \begin{bmatrix} C_{11} + \Psi C_{21} & 0 \\ C_{21} & 0 \end{bmatrix} \right),
\end{aligned}
$$

and the first block column of this expression gives

$$
\hat{\boldsymbol{v}}_1 + \hat{\boldsymbol{v}}_2 B_{22}^{-1}(C_{21} - B_{21}) = \boldsymbol{u}_1(S - (C_{11} + \Psi C_{21}))
$$

or

$$\left(\hat{\boldsymbol{v}}_1 + \hat{\boldsymbol{v}}_2 B_{22}^{-1}(C_{21} - B_{21})\right) S^{-1} = \boldsymbol{u}_1(I - (C_{11} + \Psi C_{21})S^{-1}) = -h\boldsymbol{u}_1 X,$$

from which (27) follows. Note that $B_{21} \leq 0$ and $C_{21}, B_{22}^{-1}, S^{-1} \geq 0$, so no subtractions are needed in (27). $\qquad\square$

3.6   The algorithm

Putting everything together, we obtain Algorithm 3 for the computation of the c-stable invariant pair of a matrix polynomial $P(z)$, which generalizes Algorithm 2 by removing the assumption that $\operatorname{diag}(V) > 0$.

Algorithm 3: Computing a c-stable invariant pair of $P(z)$.

Input: $V, D \in \mathbb{R}^{n \times n}$ diagonal matrices with $\operatorname{diag}(V) \geq \boldsymbol{0}$, $Q \in \mathbb{R}^{n \times n}$ a generator matrix ($Q\boldsymbol{1} = \boldsymbol{0}$, $\operatorname{offdiag}(Q) \geq \boldsymbol{0}$), satisfying A1, A2, A3.

Output: a c-stable invariant pair $(X, \begin{bmatrix} I & \Psi \end{bmatrix})$ of $P(z) = Vz^2 - Dz + Q$ (or a decomposition (8) for it) and a triplet representation for $-X^\top$.

1 compute the left Perron vector $\boldsymbol{u}$ of Q using the triplet representation $(\operatorname{offdiag}(-Q), \boldsymbol{1}, \boldsymbol{0})$;
2 choose $h > 0$ small enough to satisfy (17);
3 compute $\widetilde{A}, \widetilde{B}, \widetilde{C}$ using the formulas (22);
4 apply Algorithm 1 to $\widetilde{A}, \widetilde{B}, \widetilde{C}$ (only the last iterate $\operatorname{offdiag}(\hat{B}_k)$ and $\hat{\boldsymbol{v}} = \boldsymbol{u}A_k$ are needed);
5 compute X (or $P = h^{-1}Y$, $s = h^{-1}$) and Ψ from (26), using the triplet representations derived from Lemma 3 to invert B_{22} and S;
6 compute the triplet representation (27);

Remark 3. In the case $V = 0$, our construction reduces to the method to transform a fluid queue into a QBD introduced by Ramaswami [Ram99], up to a diagonal scaling. Indeed, the transition matrices A_0, A_1, A_2 appearing in [Ram99, Equation 4.5] satisfy $A_0 = K^{-1}\widetilde{C}$, $A_1 = I - K^{-1}\widetilde{B}$, $A_2 = K^{-1}\widetilde{A}$, where K is the diagonal matrix with entries

$$K_{ii} = \begin{cases} d_{ii} & d_{ii} > 0, \\ 2|d_{ii}| & d_{ii} < 0. \end{cases}$$

(the case $d_{ii} = 0$ is not treated in [Ram99]).

3.7   An SDA-like variant

In the linear case, a popular algorithm for this problem is the *structured doubling algorithm* [GLX06] (SDA) and its variants [BMP10, WWL12]. It is a slightly

different iteration, which has a lower computational cost because it uses the block structure in a more effective way. Merging the derivation in [BMP10] with ours, we can obtain a SDA-lookalike variant for second-order problem. The following algorithm indeed reduces to SDA-ss [BMP10] if $n_1 = 0$.

We start from the matrix polynomial $P(z)$ in the three-blocks form (19), but this time we apply the discretization map $y = 1 + hz$ first, and then we modify the location of the infinite eigenvalues. We have $F(y) = Ay^2 + By + C$, with coefficients as in (16), that is,

$$A = \begin{bmatrix} h^{-2}V_1 & & \\ & 0 & \\ & & 0 \end{bmatrix}, \quad B = \begin{bmatrix} h^{-1}D_1 + 2h^{-2}V_1 & & \\ & h^{-1}D_2 & \\ & & h^{-1}D_3 \end{bmatrix}$$

$$C = \begin{bmatrix} Q_{11} + h^{-1}D_1 + h^{-2}V_1 & Q_{12} & Q_{13} \\ Q_{21} & Q_{22} + h^{-1}D_2 & Q_{23} \\ Q_{31} & Q_{32} & Q_{33} + h^{-1}D_3 \end{bmatrix}.$$

We postmultiply these coefficients by the inverse of the M-matrix

$$M = \begin{bmatrix} h^{-2}V_1 & 0 & 0 \\ 0 & h^{-1}D_2 & 0 \\ -Q_{31} & -Q_{32} & -Q_{33} - h^{-1}D_3 \end{bmatrix}, \tag{29}$$

an operation which does not change eigenvalues and left invariant pairs, obtaining $\widehat{F}(y) = \widehat{A}y^2 + \widehat{B}y + \widehat{C}$, with

$$\widehat{A} = \begin{bmatrix} I & & \\ & 0 & \\ & & 0 \end{bmatrix}, \quad \widehat{B} = \begin{bmatrix} B_{11} & & \\ & I & \\ -B_{31} & -B_{32} & -B_{33} \end{bmatrix}, \quad \widehat{C} = \begin{bmatrix} C_{11} & C_{12} & C_{13} \\ C_{21} & C_{22} & C_{23} \\ 0 & 0 & -I \end{bmatrix},$$

where the block coefficients are given by

$$\begin{bmatrix} C_{11} & C_{12} & C_{13} \\ C_{21} & C_{22} & C_{23} \end{bmatrix} = \begin{bmatrix} Q_{11} + h^{-1}D_1 + h^{-2}V_1 & Q_{12} & Q_{13} \\ Q_{21} & Q_{22} + h^{-1}D_2 & Q_{23} \end{bmatrix} \times$$
$$\begin{bmatrix} h^{-2}V_1 & 0 & 0 \\ 0 & h^{-1}D_2 & 0 \\ -Q_{31} & -Q_{32} & -Q_{33} - h^{-1}D_3 \end{bmatrix}^{-1}$$

and

$$\begin{bmatrix} B_{11} & 0 & 0 \\ B_{31} & B_{32} & B_{33} \end{bmatrix} = \begin{bmatrix} h^{-1}D_1 + 2h^{-2}V_1 & 0 & 0 \\ 0 & 0 & -h^{-1}D_3 \end{bmatrix} \times$$
$$\begin{bmatrix} h^{-2}V_1 & 0 & 0 \\ 0 & h^{-1}D_2 & 0 \\ -Q_{31} & -Q_{32} & -Q_{33} - h^{-1}D_3 \end{bmatrix}^{-1}.$$

To obtain these blocks with a subtraction-free expression, we can make use of the triplet representation $\left(\text{offdiag}(M), \mathbf{1}, \left[h^{-2}\,\text{diag}(V_1), h^{-1}D_2, -h^{-1}D_3 \right]^\top \right)$ for M.

Finally, we postmultiply by

$$\widehat{S}(y) = \begin{bmatrix} I & & \\ & I & \\ & & yI \end{bmatrix},$$

which has the effect of shuffling around some blocks and moving n_3 of the infinite eigenvalues to zero; the final result is $\check{F}(y) = \check{A}y^2 + \check{B}y + \check{C}$, with

$$\check{A} = \begin{bmatrix} I & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & B_{33} \end{bmatrix}, \check{B} = \begin{bmatrix} B_{11} & 0 & -C_{13} \\ 0 & I & -C_{23} \\ -B_{31} & -B_{32} & I \end{bmatrix}, \check{C} = \begin{bmatrix} C_{11} & C_{12} & 0 \\ C_{21} & C_{22} & 0 \\ 0 & 0 & 0 \end{bmatrix}. \quad (30)$$

The triple $\check{A}, \check{B}, \check{C}$ has the right signs for us to apply Cyclic Reduction, producing the same solution matrix R as the above approach, since the final location of the eigenvalues is the same. Moreover, some of the pattern in the matrices $\check{A}, \check{B}, \check{C}$ is preserved under CR iterations; namely, at each step k, the pattern is

$$A_k = \begin{bmatrix} * & * & 0 \\ * & * & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad B_k = \begin{bmatrix} * & * & * \\ * & I & * \\ * & * & I \end{bmatrix}, \quad \widehat{B}_k = \begin{bmatrix} * & * & * \\ * & I & * \\ * & * & I \end{bmatrix}, \quad C_k = \begin{bmatrix} * & 0 & * \\ 0 & 0 & 0 \\ * & 0 & * \end{bmatrix}.$$

A slightly more efficient version of Cyclic Reduction can be obtained by exploiting the knowledge of these zero and identity blocks. While in the linear case the formulas simplify notably, in our quadratic case it is dubious whether it is worthwhile dealing with the additional complication of these formulas in the implementation, despite the slight computational advantage.

3.8 Solving the ODE

The reference [GLR09, Sections 1.4, 2.4 and 2.5] contains a complete theory of the relations between invariant pairs and solution of matrix linear differential equations. Let (X, U), with $X \in \mathbb{R}^{\ell \times \ell}$ and $U \in \mathbb{R}^{\ell \times n}$ in the form (4), be a c-stable invariant pair of (2). We assume positive recurrence, since otherwise there is no invariant density to compute. Then, the eigenvalues of X coincide with the eigenvalues of $P(z)$ in the open left half-plane, and any solution $\boldsymbol{p}(x)$ of (1) such that $\lim_{x \to \infty} \boldsymbol{p}(x) = \boldsymbol{0}$ can be written as

$$\boldsymbol{p}(x) = \boldsymbol{v} \exp(Xx) U, \quad \text{for } \boldsymbol{v} \in \mathbb{R}^{1 \times \ell}.$$

Simple probability considerations show that the invariant measure of the Markov-modulated Brownian motion with coefficients V, D, Q is the sum of a mass $\begin{bmatrix} \boldsymbol{0} & \boldsymbol{p}_0 \end{bmatrix}$ at $x = 0$ (where the matrix partitioning is consistent with (4)), and the density $\boldsymbol{p}(x) = \boldsymbol{v} \exp(Xx) U$. If the computed invariant pair satisfies (4), the unknown coefficients $\boldsymbol{p}_0$ and $\boldsymbol{v}$ can be determined from the condition

$$\boldsymbol{u} = \begin{bmatrix} \boldsymbol{u}_1 & \boldsymbol{u_2} \end{bmatrix} = \begin{bmatrix} 0 & \boldsymbol{p}_0 \end{bmatrix} + \int_0^\infty \boldsymbol{p}(x)\mathrm{d}x = \begin{bmatrix} -\boldsymbol{v}X^{-1} & \boldsymbol{p}_0 - \boldsymbol{v}X^{-1}\Psi \end{bmatrix}.$$

Using the relation already derived in (28), we get

$$\boldsymbol{p}_0 = \boldsymbol{u}_2 + \boldsymbol{v}X^{-1}\Psi = \boldsymbol{u}_2 - \boldsymbol{u}_1\Psi = \hat{\boldsymbol{v}}_2 B_{22}^{-1}.$$

Moreover, $\boldsymbol{v}$ satisfies $-\boldsymbol{u}_1 X = \boldsymbol{v}$, hence it follows from (27) that

$$\boldsymbol{v} = \frac{1}{h}\left(\hat{\boldsymbol{v}}_1 + \hat{\boldsymbol{v}}_2 B_{22}^{-1}(C_{21} - B_{21})\right) S^{-1},$$

the vector that we already have computed when obtaining a triplet representation for $-X^\top$.

Hence we have all the quantities that are needed to compute the invariant density $\boldsymbol{p}(x) = \boldsymbol{v}\exp(Xx)U$.

4   Componentwise stability

In this section, we adapt the theory in [NP15] to prove that the computation of invariant pairs and matrix equation solutions with Cyclic Reduction (in the non-null-recurrent case) is componentwise stable, provided that one uses triplet representations and the GTH trick as described in Algorithm 1.

We define for each $k = 0, 1, 2, \ldots$ the 4-tuple of nonnegative matrices and vectors

$$S_k := (A_k, -\operatorname{offdiag}(B_k), -\operatorname{offdiag}(\hat{B}_k), C_k), \tag{31}$$

and we call $\mathcal{F}$ the map such that $S_{k+1} = \mathcal{F}(S_k)$, corresponding to one step of Cyclic Reduction computed with (14).

When S_k and $\widetilde{S}_k$ are two different 4-tuples in the form (31) and α is a real number, we write for short $|\widetilde{S}_k - S_k| \leq \alpha S_k$ to mean that the relation holds when we replace S_k with each of the matrices and vectors in the 4-tuple, i.e.,

$$|\widetilde{A}_k - A_k| \leq \alpha A_k,$$
$$|\operatorname{offdiag}(\widetilde{B}_k) - \operatorname{offdiag}(B_k)| \leq \alpha|\operatorname{offdiag}(B_k)|,$$
$$|\operatorname{offdiag}(\widetilde{\hat{B}}_k) - \operatorname{offdiag}(\hat{B}_k)| \leq \alpha|\operatorname{offdiag}(\hat{B}_k)|,$$
$$|\widetilde{C}_k - C_k| \leq \alpha C_k.$$

4.1   Componentwise perturbation bounds

We start from assessing the error incurred when starting from inaccurate initial values. We focus on first-order results, and adopt the notation $M \lesssim N$ to mean $M \leq N + O(\varepsilon^2)$.

The key to this result is interpreting the iterates of Cyclic Reduction as the result of a censoring operation. The connection between Cyclic Reduction and censoring is a well-established result (see, for example, [BLM05, Section 7.3]). The following lemma is one of the possible ways to formalize this connection.

Lemma 15. *Consider the sequences obtained by Cyclic Reduction* (14)*, and in addition the sequence $\breve{B}$ defined by $\breve{B}_0 = B_0$ and $\breve{B}_{k+1} = \breve{B}_k - A_k B_k^{-1} C_k$. The matrix*

$$\begin{bmatrix} A_0 + I - \hat{B}_k & C_k \\ A_k & I - \breve{B}_k + C_0 \end{bmatrix} \tag{32}$$

is the result of censoring all blocks apart from the first and last from the $n(2^k + 1) \times n(2^k + 1)$ matrix

$$\begin{bmatrix} A_0 + I - B_0 & C_0 & & & \\ A_0 & I - B_0 & C_0 & & \\ & \ddots & \ddots & \ddots & \\ & & A_0 & I - B_0 & C_0 \\ & & & A_0 & I - B_0 + C_0 \end{bmatrix} \tag{33}$$

Proof. We first censor the even-numbered blocks, obtaining

$$\begin{bmatrix} A_0 + I - B_0 & & & & \\ & I - B_0 & & & \\ & & \ddots & & \\ & & & I - B_0 & \\ & & & & I - B_0 + C_0 \end{bmatrix}$$

$$+ \begin{bmatrix} C_0 & & & & \\ A_0 & C_0 & & & \\ & A_0 & \ddots & & \\ & & \ddots & C_0 & \\ & & & A_0 \end{bmatrix} \begin{bmatrix} B_0 & & & & \\ & B_0 & & & \\ & & \ddots & & \\ & & & B_0 \end{bmatrix}^{-1} \begin{bmatrix} A_0 & C_0 & & & \\ & A_0 & C_0 & & \\ & & \ddots & \ddots & \\ & & & A_0 & C_0 \end{bmatrix}$$

$$= \begin{bmatrix} A_0 + I - \hat{B}_1 & C_1 & & & \\ A_1 & I - B_1 & C_1 & & \\ & \ddots & \ddots & \ddots & \\ & & A_1 & I - B_1 & C_1 \\ & & & A_1 & I - \breve{B}_1 + C_0 \end{bmatrix}.$$

We reiterate the same process k times in total, each time censoring the even-numbered blocks in the new matrix; after each step, we obtain a matrix with the same structure, smaller size, and the indices increased by 1. $\qquad\square$

Remark 4. If the elements in $\mathrm{diag}(B_0)$ are small enough that $I - B_0 \geq 0$, then the matrix in (33) is stochastic, and so is its censoring (32). This gives an alternative proof of the relations

$$(A_k - B_k + C_k)\mathbf{1} = (A_0 - \hat{B}_k + C_k)\mathbf{1} = (A_k - \breve{B}_k + C_0)\mathbf{1} = \mathbf{0},$$

which appeared in Theorem 6 and Lemma 7.

Once this lemma is set up, it is simple to prove the following perturbation bound.

Lemma 16. *Let $A, B, C \in \mathbb{R}^{n \times n}$ and $\widetilde{A}, \widetilde{B}, \widetilde{C} \in \mathbb{R}^{n \times n}$ be two different triples of matrices satisfying A4, A5, A6, such that*

$$|\widetilde{A} - A| \lesssim \varepsilon A, \tag{34a}$$

$$|\mathrm{offdiag}(\widetilde{B}) - \mathrm{offdiag}(B)| \lesssim \varepsilon |\mathrm{offdiag}(B)|, \tag{34b}$$

$$|\widetilde{C} - C| \lesssim \varepsilon C. \tag{34c}$$

Let S_k and $\widetilde{S}_k$ be the 4-tuples resulting from applying k steps of Cyclic Reduction (14) starting from A, B, C and $\widetilde{A}, \widetilde{B}, \widetilde{C}$, respectively. Then,

$$|\widetilde{S}_k - S_k| \lesssim n 2^k \varepsilon S_k.$$

Proof. Up to a common scaling factor (which does not alter the statement of the theorem), we can assume that $I - B \geq 0$. Then, the matrices in (32) and (33) are stochastic, and we can apply [NP15, Lemma 7.3] to this censoring operation.

In detail, we call P the matrix in (33), and $\widetilde{P}$ its equivalent built starting with the initial values with a tilde. We have for $i \neq j$

$$|(\widetilde{A}_0 + I - \widetilde{B}_0)_{ij} - (A_0 + I - B_0)_{ij}| \lesssim |(\widetilde{A}_0 - A_0)_{ij}| + |(\widetilde{B}_0 - B_0)_{ij}| \lesssim \varepsilon(A_0 + I - B_0)_{ij},$$

and similarly for all other entries, so $|\text{offdiag}(\widetilde{P}) - \text{offdiag}(P)| \leq \varepsilon\,\text{offdiag}(P)$. Thus, the first part of [NP15, Lemma 7.3] holds with $m = n(2^k - 1) \leq n2^k$. This proves that

$$|\widetilde{A}_k - A_k| \overset{.}{\leq} n2^k \varepsilon A_k,$$
$$|\widetilde{C}_k - C_k| \overset{.}{\leq} n2^k \varepsilon C_k,$$

and

$$|\widetilde{D}_k - D_k| \overset{.}{\leq} n2^k \varepsilon D_k,$$
$$|\widetilde{E}_k - E_k| \overset{.}{\leq} n2^k \varepsilon E_k,$$

where

$$D_k = B_0 - \hat{B}_k = \sum_{j=0}^{k-1} C_j B_j^{-1} A_j, \quad E_k = B_0 - \breve{B}_k = \sum_{j=0}^{k-1} A_j B_j^{-1} C_j,$$

and equivalent definitions with the tilde symbols. The bounds

$$|\text{offdiag}(\widetilde{B}_k) - \text{offdiag}(B_k)| \overset{.}{\leq} n2^k \varepsilon |\text{offdiag}(B_k)|,$$
$$|\text{offdiag}(\widetilde{\hat{B}}_k) - \text{offdiag}(\hat{B}_k)| \overset{.}{\leq} n2^k \varepsilon |\text{offdiag}(\hat{B}_k)|$$

follow by noting that $B_k = B_0 - D_k - E_k$, $\hat{B}_k = B_0 - D_k$ and using [NP15, Lemma 7.2 (i)]. □

4.2 Stability of a CR step

Our next point is investigating the stability of a step of Cyclic Reduction when performed in machine arithmetic. We rely once again on the lemmas on basic operations in [NP15, Section 7], and we hide in $M \overset{.}{\leq} N$ terms which are second-order in u.

Lemma 17. *Let the 4-tuple $S_k = (A_k, -\text{offdiag}(B_k), -\text{offdiag}(\hat{B}_k), C_k)$ be exactly-represented machine numbers. We denote by $S_{k+1} = \mathcal{F}(S_k)$ the result of performing one step of Cyclic Reduction on them, and by $\widetilde{S}_{k+1} = \widetilde{\mathcal{F}}(S_k)$ the result of performing one step of Cyclic Reduction computed in inexact machine arithmetic, starting from the same matrices.*
Then,
$$|\widetilde{S}_{k+1} - S_{k+1}| \overset{.}{\leq} (\psi(n) + n + 2)\text{u}S_{k+1},$$
where n is the size of the involved matrices and $\psi(n) = \frac{2}{3}(2n+5)(n+2)(n+3)$ is the accuracy bound for the solution of a linear system with the GTH algorithm (as in [NP15, Theorem 4.1]).

Proof. We use, with a slight abuse of notation, the notation $c(X)$ to denote the computed approximation of a quantity X along one step of the algorithm (even though it is not, strictly speaking, a function of X only).

Using [NP15, Lemma 7.9] with $a = b = 0$, we obtain that the computed value $c(B_k^{-1}A_k)$ of $B_k^{-1}A_k$ satisfies

$$|c(B_k^{-1}A_k) - B_k^{-1}A_k| \overset{.}{\leq} \psi(n)\mathrm{u}B_k^{-1}A_k.$$

Hence the computed values of $A_{k+1} = A_k B_k^{-1} A_k$ and $C_k B_k^{-1} A_k$ satisfy (by [NP15, Lemma 7.8])

$$|c(A_{k+1}) - A_{k+1}| \overset{.}{\leq} (\psi(n) + n)\mathrm{u}A_{k+1},$$
$$|c(C_k B_k^{-1} A_k) - C_k B_k^{-1} A_k| \overset{.}{\leq} (\psi(n) + n)\mathrm{u}C_k B_k^{-1} A_k,$$

Analogously we have

$$|c(C_{k+1}) - C_{k+1}| \overset{.}{\leq} (\psi(n) + n)\mathrm{u}C_{k+1},$$
$$|c(A_k B_k^{-1} C_k) - A_k B_k^{-1} C_k| \overset{.}{\leq} (\psi(n) + n)\mathrm{u}A_k B_k^{-1} C_k,$$

Using again [NP15, Lemma 7.8] for the additions, we have then

$$|c(\mathrm{offdiag}(B_{k+1})) - \mathrm{offdiag}(B_{k+1})| \overset{.}{\leq} (\psi(n) + n + 2)\mathrm{u}|\mathrm{offdiag}(B_{k+1})|,$$
$$|c(\mathrm{offdiag}(\hat{B}_{k+1})) - \mathrm{offdiag}(\hat{B}_{k+1})| \overset{.}{\leq} (\psi(n) + n + 1)\mathrm{u}|\mathrm{offdiag}(\hat{B}_{k+1})|.$$

$\square$

4.3 Stability of multiple CR steps

We can now address multiple steps of Cyclic Reduction. The proof here follows [NP15, Theorem 7.12].

Lemma 18. *Let $A, B, C \in \mathbb{R}^{n \times n}$ be three matrices satisfying Assumptions A4, A5, A6, and such that A, C and $\mathrm{offdiag}(B)$ are exactly-represented machine numbers. Denote by $S_k = \mathcal{F}^k(S_0)$ the result of performing k steps of Cyclic Reduction starting from $S_0 = (A, -\mathrm{offdiag}(B), -\mathrm{offdiag}(B), -\mathrm{offdiag}(B), C)$, and by $\widetilde{S}_k = \widetilde{\mathcal{F}}^k(S_0)$ the result of k steps of Cyclic Reduction performed in inexact machine arithmetic. Then,*

$$|\widetilde{S}_k - S_k| \leq n2^k(\psi(n) + n + 2)\mathrm{u}S_k. \tag{35}$$

Proof. We prove the result by induction on k; the base case ($k = 1$) is Lemma 17.

The following manipulation is a formal version of the statement that when considering first-order error bounds we can add up the local errors at the different steps of the algorithm. Consider the telescopic sum

$$|\widetilde{S}_k - S_k| \leq \sum_{h=1}^{k} |\mathcal{F}^{h-1}\widetilde{\mathcal{F}}(\widetilde{S}_{k-h}) - \mathcal{F}^{h-1}\mathcal{F}(\widetilde{S}_{k-h})|. \tag{36}$$

By Lemma 17, we have

$$|\widetilde{\mathcal{F}}(\widetilde{S}_{k-h}) - \mathcal{F}(\widetilde{S}_{k-h})| \dot{\leq} (\psi(n) + n + 2)\mathrm{u}\mathcal{F}(\widetilde{S}_{k-h}).$$

Then by Lemma 15 used with $\varepsilon = (\psi(n) + n + 2)\mathrm{u}$, we get

$$\begin{aligned}
|\mathcal{F}^{h-1}\widetilde{\mathcal{F}}(\widetilde{S}_{k-h}) - \mathcal{F}^{h-1}\mathcal{F}(\widetilde{S}_{k-h})| &\dot{\leq} n2^{h-1}(\psi(n) + n + 2)\mathrm{u}\mathcal{F}^{h-1}\mathcal{F}(\widetilde{S}_{k-h}) \\
&\dot{\leq} n2^{h-1}(\psi(n) + n + 2)\mathrm{u}\mathcal{F}^{h}(S_{k-h}) \\
&= n2^{h-1}(\psi(n) + n + 2)\mathrm{u}S_k.
\end{aligned}$$

Passing from the first to the second row we have replaced $\widetilde{S}_{k-h}$ with S_{k-h}; this is possible because they differ by a term of order $O(\mathrm{u})$ by inductive hypothesis.

Insert this inequality into (36) to get

$$|\widetilde{S}_k - S_k| \dot{\leq} \sum_{h=1}^{k} n2^{h-1}(\psi(n) + n + 2)\mathrm{u}S_k < n2^k(\psi(n) + n + 2)\mathrm{u}S_k.$$

$\square$

4.4 Putting everything together

The previous sections shows that the CR iteration (Algorithm 1) is componentwise stable. The computation of its initial values starting from V, D, Q can be performed with (16), (22), or (30); in all three cases, if (17) holds for each $i \in E_1 \cup E_2$, then we obtain an approximation $\widetilde{S}_0$ of the CR initial values satisfying $|\widetilde{S}_0 - S_0| \leq \alpha\mathrm{u}S_0$ for a moderate multiple α of the machine precision, and thus by Lemma 16 the computed iterates are also componentwise accurate.

Once a sufficient number k of steps is performed to achieve convergence, we compute the invariant pair (X, U) as described in Section 3.4. The computed iterates satisfy (35), and similarly the computed approximation $\widetilde{\hat{\boldsymbol{v}}}_k$ of $\hat{\boldsymbol{v}} = \hat{\boldsymbol{u}}A_k$ satisfies

$$|\widetilde{\hat{\boldsymbol{v}}}_k - \hat{\boldsymbol{v}}_k| = |\boldsymbol{u}\widetilde{A}_k - \boldsymbol{u}A_k| \dot{\leq} n2^k(\psi(n) + n + 2)\mathrm{u}\hat{\boldsymbol{v}}_k.$$

The rest of the computation only involves subtraction-free formulas: we have described in Sections 3.4, 3.5, and 3.8 how to get from the matrix R computed by CR (and $\hat{B}_\infty$ and $\hat{\boldsymbol{v}}$) to the stable invariant pair of $P(z)$, its triplet representation, and the quantities needed to compute $\boldsymbol{p}(x)$.

5 Numerical experiments

We compare the following methods.

KK The algorithm in [KK95], based on explicit computation of eigenvalues and eigenvectors of a linearizing matrix which is obtained (essentially) by deflating the infinite eigenvalues from the linearizing matrix polynomial

$$\mathcal{A} - z\mathcal{E} = \begin{bmatrix} D & -T^\top \\ I_n & 0 \end{bmatrix} - z\begin{bmatrix} V & 0 \\ 0 & I_n \end{bmatrix}. \tag{37}$$

The main drawback of this method, is that by computing explicitly an eigenvalue decomposition we expect error amplification by the condition number of the eigenvector matrix.

AS The algorithm in [AS01], based on computing the sign function of the pencil (37) using the Newton-like iteration $\mathcal{A}_{k+1} = \frac{1}{2}\left(\mathcal{A}_k + \mathcal{E}\mathcal{A}_k^{-1}\mathcal{E}\right)$, and using it to separate the infinite, stable and unstable eigenvalues into different blocks. In principle, this algorithm goes in the right direction to get better numerical properties; in practice, unfortunately, our implementation of this algorithm was affected negatively by convergence issues in this iteration. In particular, it seems that the pencil $\mathcal{A}_k - z\mathcal{E}$ converges to a singular pencil whenever $\mathrm{diag}(V)$ has zero entries, so the inversion $\mathcal{A}_k^{-1}$ becomes increasingly ill-conditioned. This complicates the choice of a stopping criterion.

QZ An algorithm similar to AS, but in which the stable subspace is computed using a permuted QZ decomposition [Kåg93] of (37) (MATLAB's `ordqz`). While we could not find an explicit reference in the applied probability literature for the use of this method in the context of Markov-modulated Brownian motion, it is the method of choice for problems of this kind in the numerical linear algebra community [BK11]. The QZ decomposition is normwise backward stable, so we expect excellent normwise stability properties.

LN The algorithm in [NL], which is based on Cyclic Reduction without the use of triplet representations, or of any particular method to preserve positiveness. The discretizing transformation is the Cayley transform with $h = 1$, $y = (z+1)/(z-1)$. This algorithm can solve only problems with $\mathrm{diag}(V) > 0$.

NP Algorithm (3) as described here, from which we expect componentwise accuracy.

We apply these algorithms to several test problems.

NP15 A modification of [NP15, Example 5.1], a problem in which there is an imbalance of several orders of magnitude between the components of the solution. We take T and D as in that problem, and add a Brownian motion component with $V = I$.

NP15s The same problem as NP, but with $V(n,n) = 0$, to obtain a problem with singular V.

rand(n) Random-generated problems of different sizes $n = 8, 20, 50$. The matrices are generated with the MATLAB commands

```
V = diag(abs(randn(n, 1)));
D = diag(randn(n, 1));
T = abs(randn(n)); T = T - diag(T * ones(n, 1));
```

rand(n)s The same as rand(n), but with a matrix V containing four zero diagonal entries: `V = blkdiag(diag(abs(randn(n-4, 1))), zeros(4))`.

imb(n),imb(n)s Defined as rand(n) and rand(n)s, but all the calls of the form `randn(h,k)` are replaced by a different procedure that generates numbers spanning different orders of magnitude: `randn(h, k) .* exp(5 * randn(h, k))`.

To improve reproducibility without generating the same numbers repeatedly, we have reset the random number seed once before the complete set of experiments.

As a first error measure, we have considered the residual in the Euclidean norm

$$\frac{\|X^2 UV - XUD + UQ\|}{\|U\|(\|V\| + \|D\| + \|Q\|)} \tag{38}$$

of the left stable invariant pair (X, U) as produced by the algorithms. The values of this residual are in Table 4.

Moreover, we have normalized each invariant pair to be in the form (4) with a similarity transformation (7), and checked the forward errors

$$\frac{\|X - X_{\text{exact}}\|}{\|X_{\text{exact}}\|}, \quad \frac{\|\Psi - \Psi_{\text{exact}}\|}{\|\Psi_{\text{exact}}\|}, \tag{39}$$

where the reference values X_{exact} and Ψ_{exact} are computed applying method KK with higher precision arithmetic (32 digits, using Matlab's `vpa` command). The results are in Tables (5) and (6). Note that in the problems with $V > 0$, we have $E_3 = \varnothing$, and hence the matrix Ψ is empty and computing the error does not make sense. For this reason, Table (6) does not contain all the experiments.

As one can see, the results obtained by the new algorithm are very satisfying, especially in terms of accuracy of the computed X and Ψ (which are often the quantities of interest in view of their physical interpretation). The relative residual of the computed invariant pair, however, is sometimes slightly higher than the one obtained with the QZ method.

6   Conclusions

We have described a subtraction-free algorithm to compute the quantities needed to determine the steady-state behavior of Markov-modulated Brownian motion models in a componentwise accurate fashion. The algorithm extends the one described in [NP15] for the linear case $V = 0$, and is based on a componentwise accurate variant of Cyclic Reduction. A componentwise error analysis of this CR algorithm is provided. Our analysis highlights the role of the spectral transformation which converts continuous-time to discrete-time stability. Another interesting result is the use of a transformation related to index reduction for differential-algebraic equation and to the shift technique in this novel context.

Problem	KK	AS	LN	QZ	NP
NP15	1.5e-15	5.7e-08	3.7e-15	9.9e-16	3.8e-16
NP15s	5.0e-16	5.0e-08	-	7.7e-16	2.3e-16
rand8	1.5e-15	9.4e-16	3.6e-15	9.9e-16	1.1e-15
rand8s	6.6e-15	2.8e-11	-	5.5e-15	2.6e-15
rand20	1.9e-15	5.6e-14	1.8e-14	1.8e-15	7.3e-16
rand20s	2.8e-15	5.8e-12	-	2.0e-14	1.3e-14
rand50	2.3e-15	3.0e-14	3.4e-13	1.5e-14	5.9e-15
rand50s	7.1e-14	1.2e-08	-	1.3e-14	1.7e-14
imb8	1.2e-08	8.6e-05	4.5e+04	4.2e-11	7.4e-09
imb8s	1.5e-13	4.2e-10	-	2.3e-14	2.3e-13
imb20	2.2e-15	4.8e-06	9.7e-01	2.7e-14	4.9e-13
imb20s	1.2e-09	8.0e+05	-	2.6e-13	1.9e-13
imb50	8.7e-14	7.5e-06	3.3e+01	3.3e-13	1.3e-10
imb50s	3.1e-04	3.2e+11	-	2.5e-05	2.0e-08

Tab. 4: Relative residual (38).

Problem	KK	AS	LN	QZ	NP
NP15	2.7e-12	2.5e-07	2.9e-13	1.8e-12	1.7e-16
NP15s	1.3e-12	2.2e-07	-	6.2e-13	1.8e-16
rand8	2.8e-15	1.5e-15	1.6e-15	2.4e-15	2.7e-16
rand8s	2.9e-15	1.8e-13	-	2.3e-15	3.1e-16
rand20	4.4e-15	9.6e-14	5.6e-15	4.8e-15	3.0e-16
rand20s	3.2e-15	3.0e-12	-	4.1e-14	1.1e-15
rand50	5.9e-15	4.0e-14	4.0e-14	5.6e-14	6.9e-16
rand50s	5.6e-14	1.2e-10	-	3.5e-14	5.2e-16
imb8	9.7e-12	1.9e-09	1.1e+00	7.1e-13	9.0e-13
imb8s	2.6e-14	1.3e-08	-	1.3e-12	1.1e-15
imb20	4.6e-11	2.1e-07	3.2e-04	1.1e-09	9.1e-12
imb20s	4.4e-12	6.9e-06	-	5.9e-12	4.0e-13
imb50	2.0e-10	9.8e-06	7.2e-01	1.0e-08	8.3e-10
imb50s	2.0e-10	3.3e-05	-	1.0e+00	2.6e-13

Tab. 5: Forward error (39) on X.

Problem	KK	AS	LN	QZ	NP
NP15s	2.3e-15	1.8e-11	-	2.8e-15	1.3e-16
rand8s	1.2e-14	3.7e-13	-	2.4e-15	2.5e-15
rand20s	7.1e-15	7.7e-11	-	6.7e-14	2.1e-15
rand50s	3.4e-14	3.5e-09	-	5.3e-14	4.7e-16
imb8s	8.3e-15	5.2e-09	-	1.1e-11	5.2e-15
imb20s	1.4e-10	1.9e-08	-	2.8e-11	4.0e-11
imb50s	6.9e-11	9.0e-09	-	1.0e-04	6.1e-08

Tab. 6: Forward error (39) on Ψ.

References

[AS01]   Mircea Agapie and Khosrow Sohraby. Algorithmic solution to second-order fluid flow. In *Proceedings IEEE INFOCOM 2001, The Conference on Computer Communications, Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies, Twenty years into the communications odyssey, Anchorage, Alaska, USA, April 22-26, 2001*, pages 1261–1270, 2001.

[Asm95]   S. Asmussen. Stationary distributions for fluid flow models with or without Brownian noise. *Communications in Statistics: Stochastic Models*, 11(1):21–49, 1995.

[AXY02]   Attahiru Sule Alfa, Jungong Xue, and Qiang Ye. Accurate computation of the smallest eigenvalue of a diagonally dominant *M*-matrix. *Math. Comp.*, 71(237):217–236, 2002.

[BK11]   Timo Betcke and Daniel Kressner. Perturbation, extraction and refinement of invariant pairs for matrix polynomials. *Linear Algebra Appl.*, 435(3):574–536, 2011.

[BLM05]   Dario A. Bini, Guy Latouche, and Beatrice Meini. *Numerical Methods for Structured Markov Chains*. Numerical Mathematics and Scientific Computation. Oxford University Press, Oxford, 2005.

[BM15]   Dario A. Bini and Beatrice Meini. Generalization of the brauer theorem to matrix polynomials and matrix laurent series, 2015.

[BMP10]   Dario A. Bini, Beatrice Meini, and Federico Poloni. Transforming algebraic Riccati equations into unilateral quadratic matrix equations. *Numer. Math.*, 116(4):553–578, 2010.

[BMS09]   Dario Bini, Beatrice Meini, and Sergio Steffè. SMCSolver (structured Markov chain solver) 2.1, 2009. Available at http://bezout.dm.unipi.it/SMCSolver/.

[BP94]   Abraham Berman and Robert J. Plemmons. *Nonnegative matrices in the mathematical sciences*, volume 9 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1994. Revised reprint of the 1979 original.

[GLR09]   I. Gohberg, P. Lancaster, and L. Rodman. *Matrix polynomials*, volume 58 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2009. Reprint of the 1982 original [MR0662418].

[GLX06]   X.-X. Guo, W.-W. Lin, and S.-F. Xu. A structure-preserving doubling algorithm for nonsymmetric algebraic Riccati equation. *Numer. Math.*, 103:393–412, 2006.

[GTH85]   Winfried K. Grassmann, Michael I. Taksar, and Daniel P. Heyman. Regenerative analysis and steady state distributions for Markov chains. *Oper. Res.*, 33(5):1107–1116, 1985.

[Guo13]   Chun-Hua Guo. On algebraic Riccati equations associated with M-matrices. *Linear Algebra Appl.*, 439(10):2800–2814, 2013.

[HK00]   Nicholas J. Higham and Hyun-Min Kim. Numerical analysis of a quadratic matrix equation. *IMA J. Numer. Anal.*, 20(4):499–519, 2000.

[HMR02]   C. He, B. Meini, and N. H. Rhee. A shifted cyclic reduction algorithm for quasi-birth-death problems. *SIAM J. Matrix Anal. Appl.*, 23(3):673–691 (electronic), 2001/02.

[Iva10]   Jevgenijs Ivanovs. Markov-modulated Brownian motion with two reflecting barriers. *J. Appl. Probab.*, 47(4):1034–1047, 2010.

[Kåg93]   B. Kågström. A direct method for reordering eigenvalues in the generalized real Schur form of a regular matrix pair (A, B). In *Linear algebra for large scale and real-time applications (Leuven, 1992)*, volume 232 of *NATO Adv. Sci. Inst. Ser. E Appl. Sci.*, pages 195–218. Kluwer Acad. Publ., Dordrecht, 1993.

[KK95]   R. L. Karandikar and V.G. Kulkarni. Second-order fluid flow models: Reflected Brownian motion in a random environment. *Oper. Res*, 43:77–88, 1995.

[KM06]   Peter Kunkel and Volker Mehrmann. *Differential-algebraic equations*. EMS Textbooks in Mathematics. European Mathematical Society (EMS), Zürich, 2006. Analysis and numerical solution.

[MMMM15] D. Steven Mackey, Niloufer Mackey, Christian Mehl, and Volker Mehrmann. Möbius transformations of matrix polynomials. *Linear Algebra Appl.*, 470:120–184, 2015.

[NL]   Giang T Nguyen and Guy Latouche. The morphing of fluid queues
            into Markov-modulated Brownian motion. *Stochastic systems*. To
            appear. Available online at `http://arxiv.org/abs/1311.3359`.

[Nof12]   Vanni Noferini. The behaviour of the complete eigenstructure of a
            polynomial matrix under a generic rational transformation. *Elec-
            tron. J. Linear Algebra*, 23:607–624, 2012.

[NP15]   Giang T. Nguyen and Federico Poloni.   Componentwise accu-
            rate fluid queue computations using doubling algorithms. *Numer.
            Math.*, 130(4):763–792, 2015.

[O'C93]   Colm Art O'Cinneide. Entrywise perturbation theory and error
            analysis for Markov chains. *Numer. Math.*, 65(1):109–120, 1993.

[Ram99]   V. Ramaswami. Matrix analytic methods for stochastic fluid flows.
            In D. Smith and P. Hey, editors, *Teletraffic Engineering in a Com-
            petitive World (Proceedings of the 16th International Teletraffic
            Congress)*, pages 1019–1030. Elsevier Science B.V., Edinburgh,
            UK, 1999.

[SGX14]   Meiyue Shao, Weiguo Gao, and Jungong Xue. Aggressively trun-
            cated Taylor series method for accurate computation of exponen-
            tials of essentially nonnegative matrices. *SIAM J. Matrix Anal.
            Appl.*, 35(2):317–338, 2014.

[WWL12]   Wei-Guo Wang, Wei-Chao Wang, and Ren-Cang Li. Alternating-
            directional doubling algorithm for M-matrix algebraic Riccati
            equations. *SIAM J. Matrix Anal. Appl.*, 33(1):170–194, 2012.

[XXL12]   J. Xue, S. Xu, and R.-C. Li. Accurate solutions of M-matrix alge-
            braic Riccati equations. *Numer. Math.*, 120:671–700, 2012.

[XY08]   Jungong Xue and Qiang Ye.   Entrywise relative perturbation
            bounds for exponentials of essentially non-negative matrices. *Nu-
            mer. Math.*, 110(3):393–403, 2008.