

Componentwise accurate fluid queue computations using doubling algorithms

Giang T. Nguyen^{*} and Federico Poloni[†]

Abstract

Markov-modulated fluid queues are popular stochastic processes frequently used for modelling real-life applications. An important performance measure to evaluate in these applications is their steady-state behaviour, which is determined by the stationary density. Computing it requires solving a (nonsymmetric) M-matrix algebraic Riccati equation, and indeed computing the stationary density is the most important application of this class of equations.

Xue *et al.* (2012) provided a componentwise first-order perturbation analysis of this equation, proving that the solution can be computed to high relative accuracy even in the smallest entries, and suggested several algorithms for computing it. An important step in all proposed algorithms is using so-called *triplet representations*, which are special representations for M-matrices that allow for a high-accuracy variant of Gaussian elimination, the GTH-like algorithm. However, triplet representations for all the M-matrices needed in the algorithm were not found explicitly. This can lead to an accuracy loss that prevents the algorithms to converge in the componentwise sense.

In this paper, we focus on the *structured doubling algorithm*, the most efficient among the proposed methods in Xue *et al.* (2012), and build upon their results, providing (i) explicit and cancellation-free expressions for the needed triplet representations, allowing the algorithm to be performed in a really cancellation-free fashion; (ii) an algorithm to evaluate the final part of the computation to obtain the stationary density; and (iii) a componentwise error analysis for the resulting algorithm, the first explicit one for this class of algorithms. We also present numerical results to illustrate the accuracy advantage of our method over standard (normwise-accurate) algorithms.

1 Introduction

Markov-modulated fluid queues $\{X(t), \varphi(t) : t \geq 0\}$ are two-dimensional Markov processes frequently used for modeling real-life applications, such as risk theory, actuarial science, environmental systems and telecommunication networks; see for instance [3, 18, 19, 25]. In a fluid queue, the variable $X(t)$, called *level*, varies linearly at rate c_i whenever the Markov chain $\varphi(t)$ is in state i .

The key component for obtaining the stationary distribution of a Markov-modulated fluid queue, which determines its steady-state behaviour, is finding the minimal nonnegative solution to a *nonsymmetric algebraic Riccati equation*

^{*}The University of Adelaide, School of Mathematical Sciences, giang.nguyen@adelaide.edu.au

[†]The University of Pisa, Computer Science Department, fpoloni@di.unipi.it

(NARE) of the form

$$B - AX - XD + XCX = 0, \quad (1)$$

Several algorithms have been suggested for solving equations of this kind. We refer the reader to the review article [7] and the book [6] for a list of the possibilities. In our problem, the matrix

$$M := \begin{bmatrix} D & -C \\ -B & A \end{bmatrix}$$

is a singular irreducible M-matrix, satisfying $M\mathbf{1} = \mathbf{0}$, where $\mathbf{1}$ is the vector with all components 1 and $\mathbf{0}$ the one with all components 0. The linear algebra literature has focused on a larger class of equations, so-called *M-matrix algebraic Riccati equations*, in which M is either nonsingular, or singular and irreducible M-matrix. Nevertheless, fluid queues are the main application of this class of equations, having a broad applicability and considerable interest in applied probability. Hence, we focus in this paper on the case of NAREs appearing in fluid queues, adopting the notation and terminology closer to the probabilistic meaning of these equations.

We consider here the *structured doubling algorithm* (SDA), introduced in Guo *et al.* [13], and two variants known as SDA-ss [9] and ADDA [26]. These algorithms are the most efficient ones known for this class of equations [7]. We develop a version of these algorithms that computes componentwise accurate results with provable bounds. An approximation $\tilde{Q} \geq 0$ of a matrix $Q \geq 0$ is said to be within a componentwise error threshold $\varepsilon > 0$ if it satisfies the inequality

$$|\tilde{Q}_{ij} - Q_{ij}| \leq \varepsilon Q_{ij} \quad \text{for all } i \text{ and } j. \quad (2)$$

The threshold ε may include a condition number for the problem, or a moderately-growing function of the matrix size. Normally, linear algebra algorithms focus on normwise stability instead, that is, on error bounds of the kind $\|\tilde{Q} - Q\| \leq \varepsilon \|Q\|$, with $\|\cdot\|$ being, for instance, the Euclidean norm. The main difference between the two, is that a componentwise accurate algorithm ensures not only the dominating components but also the smallest entries in the matrices are computed with a high number of significant digits. This additional accuracy is especially useful for problems whose solutions are badly scaled. In particular, having componentwise error bounds is important in probabilistic applications: some of the states of the Markov chain $\varphi(t)$ may represent component failures or other exceptional behaviours; thus, very small entries are expected and it is crucial that they are computed with high accuracy.

A componentwise perturbation theory for equation (1) has recently appeared in [27]; its authors in addition described a first attempt to obtain a componentwise accurate algorithm. Our main improvement to their algorithm is showing how to update special representations of the M-matrices appearing in the algorithm, known as *triplet representations*. We refer the reader to Section 4 for definitions and more detail.

Moreover, we provide a componentwise accurate implementation for the part that follows the solution of (1), that is, the computation of the stationary density of the queue using this solution. This part is usually neglected in the linear algebra literature, but it is nevertheless required to complete the computation of interest in queuing theory.

We present numerical experiments to demonstrate that updating correctly the triplet representations is crucial to attain componentwise accuracy. In addition, we prove in the Appendix an explicit componentwise error bound for the computed solution Ψ , the first one to our knowledge for this class of algorithms.

We organize the paper as follows. In Section 2, we formally describe the Markov-modulated fluid queue and present a formula for its stationary distribution. In Section 3, we introduce the structure-preserving doubling algorithm. We describe our componentwise-accurate variant in Section 4. Numerical experiments are presented in Section 5, and a stability proof is included in the Appendix.

2 The stationary density of a fluid queue

Let T be the generator matrix of a continuous-time Markov chain (CTMC) $\varphi(t)$ on the finite state space \mathcal{S} ; that is, $T\mathbf{1} = \mathbf{0}$, $T_{ij} \geq 0$ for $i \neq j$ and $i, j \in \mathcal{S}$. We assume that T is irreducible. A Markov-modulated fluid queue with a reflecting boundary at 0 [20, 23, 10] is a two-dimensional Markov process $\{X(t), \varphi(t)\}$, where the *level process* $X(t) \geq 0$ varies linearly at rate c_i whenever the *phase process* $\varphi(t)$ is in state $i \in \mathcal{S}$. In other words,

$$\frac{d}{dt}X(t) = \begin{cases} c_i & \text{for } \varphi(t) = i \in \mathcal{S} \text{ and } X(t) > 0, \\ \max\{0, c_i\} & \text{for } \varphi(t) = i \in \mathcal{S} \text{ and } X(t) = 0. \end{cases}$$

We define $C = \text{diag}(c_i)_{i \in \mathcal{S}}$ to be the diagonal rate matrix of the fluid level, and assume that $c_i \neq 0$ for all $i \in \mathcal{S}$. This assumption is without loss of generality, as one can analyze models with zero rates by converting them into models without zero rates using censoring techniques [10, Section 1.7] which are componentwise-accurate without modification. Moreover, we assume that the fluid queue is *positive recurrent*; this is equivalent to requiring $\xi C \mathbf{1} < 0$, where $\xi \geq \mathbf{0}$ is the stationary distribution row vector of the Markov chain $\varphi(t)$ ($\xi T = \mathbf{0}$ and $\xi \mathbf{1} = 1$). This is a necessary condition for the existence of our quantity of interest, the stationary density.

We decompose the state space \mathcal{S} into two exhaustive and disjoint subsets: $\mathcal{S}_+ := \{i \in \mathcal{S} : c_i > 0\}$ and $\mathcal{S}_- := \{i \in \mathcal{S} : c_i < 0\}$, and define $n := |\mathcal{S}|$, $n_+ := |\mathcal{S}_+|$, and $n_- := |\mathcal{S}_-|$. Assume without loss of generality (up to reordering) that $\mathcal{S}_+ = \{1, 2, \dots, n_+\}$; then, the matrices T and C can be written as

$$T := \begin{bmatrix} T_{++} & T_{+-} \\ T_{-+} & T_{--} \end{bmatrix}, \quad C := \begin{bmatrix} C_+ & 0 \\ 0 & C_- \end{bmatrix},$$

where $T_{\ell m}$ contains entries of T_{ij} for $i \in \mathcal{S}_\ell$ and $j \in \mathcal{S}_m$, and $C_\ell := \text{diag}(c_i)_{i \in \mathcal{S}_\ell}$. Next, we denote by $\mathbf{F}(x) := (F_i(x))_{i \in \mathcal{S}}$ and $\mathbf{p} := (p_i)_{i \in \mathcal{S}}$, respectively, the stationary distribution row vector and the probability mass row vector at level zero of $\{X(t), \varphi(t)\}$, that is,

$$F_i(x) := \lim_{t \rightarrow \infty} \mathbb{P}[X(t) \leq x, \varphi(t) = i] \quad \text{for } x > 0 \text{ and } i \in \mathcal{S},$$

$$p_i := \lim_{t \rightarrow \infty} \mathbb{P}[X(t) = 0, \varphi(t) = i] \quad \text{for } i \in \mathcal{S},$$

and denote by $\mathbf{f}(x) := (f_i(x))_{i \in \mathcal{S}}$ the associated density row vector, with

$$f_i(x) := \frac{d}{dx} F_i(x) \quad \text{for } x > 0 \text{ and } i \in \mathcal{S}.$$

Based on the physical interpretations of the rates c_i , we immediately obtain $\mathbf{p} = (\mathbf{0}, \mathbf{p}_-)$, where $\mathbf{p}_- = (p_i)_{i \in \mathcal{S}_-}$. Furthermore, it is well-established [22] that the stationary density vector $\mathbf{f}(x)$ satisfies the following differential equation

$$\frac{d}{dx} \mathbf{f}(x)C = \mathbf{f}(x)T, \quad (3)$$

and is given by

$$\mathbf{f}(x) = \mathbf{p}_- T_{-+} e^{Kx} V \quad \text{for } x > 0, \quad (4)$$

with

$$e^{Kx} = \sum_{i=0}^{\infty} (Kx)^i / i!, \quad (5)$$

$$K = C_+^{-1} T_{++} + \Psi |C_-|^{-1} T_{-+}, \quad (6)$$

$$V = [C_+^{-1} \quad \Psi |C_-|^{-1}], \quad (6)$$

the vector \mathbf{p}_- being the unique solution to the system of equations

$$\mathbf{p}_- (T_{--} + T_{-+} \Psi) = \mathbf{0}, \quad (7)$$

$$\mathbf{p}_- (\mathbf{1} - T_{-+} K^{-1} V \mathbf{1}) = 1. \quad (8)$$

and Ψ the minimal nonnegative solution to the nonsymmetric algebraic Riccati equation

$$\Psi |C_-|^{-1} T_{-+} \Psi + C_+^{-1} T_{++} \Psi + \Psi |C_-|^{-1} T_{--} + C_+^{-1} T_{+-} = 0. \quad (9)$$

The matrix Ψ is stochastic and has a probabilistic meaning: each entry Ψ_{ij} is the probability of the fluid returning, from above, to the initial level $x \geq 0$ in phase $j \in \mathcal{S}_-$, after starting in phase $i \in \mathcal{S}_+$ and avoiding all levels below x [12, Theorem 2.1].

In an alternative linear algebra view, the structure-preserving doubling algorithm determines a basis $[I_{n_+} \quad -\Psi]$ for the antistable left invariant subspace of the matrix $C^{-1}T$; that is, there exists a matrix—which turns out to be K as defined in (5)—of order $n_+ \times n_+$ with all its eigenvalues in the right half-plane such that

$$[I \quad -\Psi] C^{-1}T = K [I \quad -\Psi]. \quad (10)$$

The formulas (4)–(6) show that solving the Riccati equation (9) for Ψ and calculating the matrix-exponential e^{Kx} are two main tasks in the evaluation of the stationary density $\mathbf{f}(x)$.

3 Structure-preserving doubling algorithms

Guo *et al.* [13] developed a structure-preserving doubling algorithm for solving nonsymmetric algebraic Riccati equation, and proved that convergence to the minimal nonnegative solution is monotonically increasing and quadratic. Later, Bini *et al.* [9] and Wang *et al.* [26] presented two variants, SDA-ss and ADDA

respectively, which differ only in the initialization step and have faster convergence speeds. We describe below the algorithms, starting from the most recent and general one, the ADDA [26].

We first define a map $\mathcal{F} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ as follows. Given a matrix $P \in \mathbb{R}^{n \times n}$, consider the following partition

$$P = \begin{bmatrix} E & G \\ H & F \end{bmatrix}, \quad \text{where } E \in \mathbb{R}^{n_+ \times n_+} \text{ and } F \in \mathbb{R}^{n_- \times n_-}, \quad (11)$$

and define

$$\widehat{E} := E(I - GH)^{-1}E, \quad (12a)$$

$$\widehat{F} := F(I - HG)^{-1}F, \quad (12b)$$

$$\widehat{G} := G + E(I - GH)^{-1}GF, \quad (12c)$$

$$\widehat{H} := H + F(I - HG)^{-1}HE. \quad (12d)$$

Then,

$$\mathcal{F}(P) := \begin{bmatrix} \widehat{E} & \widehat{G} \\ \widehat{H} & \widehat{F} \end{bmatrix}, \quad \text{where } \widehat{E} \in \mathbb{R}^{n_+ \times n_+} \text{ and } \widehat{F} \in \mathbb{R}^{n_- \times n_-}. \quad (13)$$

The map \mathcal{F} is known as the *doubling* map, which is well-defined whenever the matrices $I - GH$ and $I - HG$ are nonsingular. Moreover, define

$$\alpha_{\text{opt}} := \min_{i \in \mathcal{S}_-} \left| \frac{C_{ii}}{T_{ii}} \right|, \quad \beta_{\text{opt}} := \min_{i \in \mathcal{S}_+} \left| \frac{C_{ii}}{T_{ii}} \right|, \quad (14)$$

and choose two nonnegative reals $\alpha \leq \alpha_{\text{opt}}$ and $\beta \leq \beta_{\text{opt}}$, not both being zero, and let

$$P_0 := Q^{-1}R \quad (15)$$

where

$$Q := \begin{bmatrix} C_+ - \alpha T_{++} & -\beta T_{+-} \\ -\alpha T_{-+} & |C_-| - \beta T_{--} \end{bmatrix}, \quad R := \begin{bmatrix} C_+ + \beta T_{++} & \alpha T_{+-} \\ \beta T_{-+} & |C_-| + \alpha T_{--} \end{bmatrix} \quad (16)$$

Note that the conditions $\alpha \leq \alpha_{\text{opt}}$ and $\beta \leq \beta_{\text{opt}}$ are equivalent to requiring $R_{ii} \geq 0$ for each i .

Applying the doubling map to P_0 iteratively defines a sequence as follows

$$P_k := \mathcal{F}^k(P_0), \quad P_k := \begin{bmatrix} E_k & G_k \\ H_k & F_k \end{bmatrix}, \quad k \geq 0. \quad (17)$$

where \mathcal{F}^k denotes the composition of \mathcal{F} with itself k times. The following theorem on the limiting behavior of the four blocks of P_k follows from [26, Theorem 3.3] when $\alpha \neq 0$ and $\beta \neq 0$ and from [9, Theorem 8] otherwise. Here and in the following, inequalities among matrices are intended in the componentwise sense.

Theorem 3.1. *Consider the iteration (17) applied with initial values (15), where (T, C) are the irreducible generator and the rate matrix of a positive recurrent fluid queue $\{X(t), \varphi(t)\}$. For each $k \geq 0$, the matrices $I - G_k H_k$ and $I - H_k G_k$ are nonsingular, and the matrices E_k, F_k, G_k , and H_k satisfy the following properties:*

- (i) $E_k > 0$ and $\lim_{k \rightarrow \infty} E_k = 0$,
- (ii) $F_k > 0$ and $\lim_{k \rightarrow \infty} F_k =: F_\infty$, for some matrix F_∞ .
- (iii) $0 < G_k < G_{k+1}$ and $\lim_{k \rightarrow \infty} G_k = \Psi$, where Ψ is the minimal nonnegative solution of the NARE (9).
- (iv) $0 < H_k < H_{k+1}$ and $\lim_{k \rightarrow \infty} H_k = \widehat{\Psi}$, where $\widehat{\Psi}$ is the minimal nonnegative solution to the NARE

$$|C_-|^{-1}T_{-+} + \widehat{\Psi}C_+^{-1}T_{++} + |C_-|^{-1}T_{--}\widehat{\Psi} + \widehat{\Psi}C_+^{-1}T_{+-}\widehat{\Psi} = 0. \quad (18)$$

- (v) The convergence rate in the above limits is quadratic, that is,

$$|P_k - \lim_{k \rightarrow \infty} P_k| \leq K_0 \delta^{2^k} \mathbf{1}^\top, \quad (19)$$

for a suitable $K_0 > 0$ and $0 < \delta < 1$.

The algorithm described above is called ADDA [26]; older algorithms appeared in literature are SDA-ss [9], which corresponds to choosing $\alpha = 0$, and SDA [13], which corresponds to choosing $\alpha = \beta = \min(\alpha_{\text{opt}}, \beta_{\text{opt}})$. In general, we call them *doubling algorithms*. Clearly, Theorem 3.1 also holds for both SDA-ss and SDA, but the values of K_0 and δ for the three algorithms are different: the smallest possible δ that satisfies the bound (19) is

$$\delta = \max \left\{ \left| \frac{1 + \beta\lambda}{1 - \alpha\lambda} \right| : \lambda \text{ is an eigenvalue of } C^{-1}T \text{ and } \Re\lambda < 0 \right\}, \quad (20)$$

where $\Re\lambda$ denotes the real part of λ .

Remark 3.2. Here, there are two significant changes in the presentation of these algorithms with respect to the papers in which they originally appeared. First, the parameters α and β here are the inverses of the parameters α' and β' in the original papers. That is, $\alpha := 1/\alpha'$ and $\beta := 1/\beta'$. This makes more apparent the relation with SDA-ss, since we can set $\alpha = 0$. Second, Equation (15), which can be derived using [21, Theorem 5.1], is a simpler formula for the initial values than those given in the original papers.

4 A cancellation-free version of doubling

4.1 M-matrices and the GTH-like algorithm

We begin with some definitions. Let us define the *off-diagonal* of an $m \times m$ matrix $A = (A_{i,j})$ as the vector $\text{offdiag}(A) \in \mathbb{R}^{m^2-m}$ which contains all the non-diagonal entries of A ordered one column after the other, that is,

$$\text{offdiag}(A) = (A_{2,1}, A_{3,1}, \dots, A_{m,1}, A_{2,1}, A_{2,3}, \dots, A_{2,m}, \dots, A_{m,m-1})^\top.$$

(In fact, the actual order of these elements does not matter in our paper.) A square matrix A is called a *Z-matrix* if $\text{offdiag}(A) \leq \mathbf{0}$, and an $m \times m$ *Z-matrix* A is called an *M-matrix* if there exists a vector $\mathbf{v} \in \mathbb{R}^m$ such that $\mathbf{v} > \mathbf{0}$ and

$A\mathbf{v} \geq \mathbf{0}$. It is well-known [4] that a nonsingular Z -matrix is an M -matrix if and only if $Z^{-1} \geq \mathbf{0}$, and that a Z -matrix A is an M -matrix if and only if all its eigenvalues have nonnegative real parts. In particular, the transpose of an M -matrix is also an M -matrix.

A triple $(\mathbf{u}, \mathbf{v}, \mathbf{w}) \in \mathbb{R}^{m^2-m} \times \mathbb{R}^m \times \mathbb{R}^m$ with $\mathbf{u} \leq \mathbf{0}$, $\mathbf{v} > \mathbf{0}$, $\mathbf{w} \geq \mathbf{0}$ is called a *right triplet representation* for the Z -matrix A if $\text{offdiag}(A) = \mathbf{u}$ and $A\mathbf{v} = \mathbf{w}$. In a sense, providing a triplet representation serves as a ‘‘certificate’’ that A is an M -matrix, because it corresponds to verifying our definition explicitly. It is remarkable that the same certificate can be used to perform several matrix-algebraic operations with high accuracy. Indeed, we recall here the GTH-like algorithm, introduced in [1], which is an algorithm for solving linear systems $A\mathbf{x} = \mathbf{b}$, with A being a nonsingular M -matrix and $\mathbf{b} \geq \mathbf{0}$, using a triplet representation of A and \mathbf{b} as the only input. We present the algorithm as Algorithm 1 below.

Algorithm 1: GTH-like algorithm for solving $A\mathbf{x} = \mathbf{b}$, where A is a nonsingular M -matrix

Input: a triplet representation $(\text{offdiag}(A), \mathbf{v}, \mathbf{w})$, where

$$\begin{aligned} \text{offdiag}(A) &\in \mathbb{R}^{m^2-m}, \mathbf{v} = (v_1, \dots, v_m), \mathbf{w} = (w_1, \dots, w_m), \mathbf{b} \in \mathbb{R}^m, \\ \text{offdiag}(A) &\leq \mathbf{0}, \mathbf{v} > \mathbf{0}, \mathbf{w}, \mathbf{b} \geq \mathbf{0} \end{aligned}$$

Output: $\mathbf{x} = A^{-1}\mathbf{b}$

```

1  $L \leftarrow I_m$ ;
2  $U \leftarrow 0_{m \times m}$ ;  $\text{offdiag}(U) \leftarrow \text{offdiag}(A)$ ;
3 for  $k = 1$  to  $m$  do
4    $U_{k,k} \leftarrow (w_k - U_{k,k+1:m}v_{k+1:m})/v_k$ ;
5    $L_{k+1:m,k} \leftarrow U_{k+1:m,k}/U_{k,k}$ ;
6    $w_{k+1:m} \leftarrow w_{k+1:m} - L_{k+1:m,k}w_k$ ;
7    $U_{k+1:m,k} \leftarrow 0$ ;
8    $\text{offdiag}(U_{k+1:m,k+1:m}) \leftarrow \text{offdiag}(U_{k+1:m,k+1:m} - L_{k+1:m,k}U_{k,k+1:m})$ ;
9 end
10  $\mathbf{y} \leftarrow L^{-1}\mathbf{b}$  (computed by forward-substitution);
11  $\mathbf{x} \leftarrow U^{-1}\mathbf{y}$  (computed by back-substitution).
```

The following result demonstrates the high level of stability of Algorithm 1 when implemented on a computer. Here and in the following, we use the common floating point arithmetic model in which $\text{fl}(x \text{ op } y) = (x \text{ op } y)(1 + \varepsilon)$, $|\varepsilon| \leq u$, for $\text{op} \in \{+, -, \times, \div\}$ and machine precision u [15].

Theorem 4.1 ([1]). *Suppose that Algorithm 1 is carried out in floating-point arithmetic with machine precision u , with its input $\text{offdiag}(A), \mathbf{v}, \mathbf{w}$, and \mathbf{b} given as exact floating point numbers. Then, the computed $\tilde{\mathbf{x}}$ satisfies the following inequality*

$$|\tilde{\mathbf{x}} - \mathbf{x}| \leq (\psi(m)u + O(u^2))\mathbf{x}, \quad (21)$$

where $\psi(m) := \frac{2}{3}(2m + 5)(m + 2)(m + 3)$.

Notice in (21) the surprising absence of the condition number of A , which one would expect from classical error analysis [15, Chapters 7 and 9]. This is possible only because the input values of the algorithm are not the entries of A but a different object, its triplet representation. In addition, we point out that

the term $\psi(m)$ is usually a pessimistic estimate: in practice, the growth factor is typically $O(m)$ rather than $O(m^3)$ [1].

The GTH-like algorithm works by computing an LU factorization of the matrix A ; it differs only slightly from a standard LU-based linear system solver [11]. Indeed, the only difference here are Lines 4 and 6, which compute each diagonal entry of U with an equivalent subtraction-free expression instead of the usual update rule (that is, performing Lines 2 and 8 on the whole matrix U rather than on the off-diagonal entries only). It is easy to use the same LU factorization to solve linear systems of the form $A^\top \mathbf{x} = \mathbf{b}$ or to compute the (left or right) one-dimensional kernel of a singular M -matrix, by simple variants of Algorithm 1.

In a similar fashion, we say that $(\text{offdiag}(A), \mathbf{v}^\top, \mathbf{w}^\top)$ is a *left triplet representation* for A if $\mathbf{v}^\top A = \mathbf{w}^\top$, with $\mathbf{v}^\top > \mathbf{0}$ and $\mathbf{w}^\top \geq \mathbf{0}$ (or, equivalently, if $(\text{offdiag}(A^\top), \mathbf{v}, \mathbf{w})$ is a right triplet representation for A^\top). It is clear that we can adapt the GTH-like algorithm to work with left triplet representations.

4.2 Triplet representations in doubling

Recently, an implementation of SDA taking advantage of the GTH-like algorithm was proposed in [27], but triplet representations for the matrices to invert in the algorithm, $I - H_k G_k$ and $I - G_k H_k$, were not provided. Rather, the authors advised to compute them from the entries of the matrices using an iterative algorithm. However, when a matrix is ill-conditioned, computing its triplet representation starting from its entries is an ill-conditioned problem, hence most of the benefits of using the GTH-like algorithm are lost in this way.

We show in the following theorem that triplet representations for all the matrices to be inverted in the algorithm are available explicitly with subtraction-free formulas.

Theorem 4.2.

1. A triplet representation for Q , defined in (15), is

$$(\text{offdiag}(Q), \mathbf{1}, |C|\mathbf{1}). \quad (22)$$

2. A triplet representation for $I - G_k H_k$ is

$$(\text{offdiag}(-G_k H_k), \mathbf{1}, G_k F_k \mathbf{1} + E_k \mathbf{1}). \quad (23)$$

3. A triplet representation for $I - H_k G_k$ is

$$(\text{offdiag}(-H_k G_k), \mathbf{1}, H_k E_k \mathbf{1} + F_k \mathbf{1}). \quad (24)$$

Proof. The first representation follows easily from $T\mathbf{1} = \mathbf{0}$. For the second representation, since the matrix P_k is stochastic for $k \geq 0$ (see Corollary A.6 in the Appendix),

$$\begin{bmatrix} I & -G_k \\ -H_k & I \end{bmatrix} \mathbf{1} = \begin{bmatrix} E_k & 0 \\ 0 & F_k \end{bmatrix} \mathbf{1}. \quad (25)$$

Pre-multiplying both sides by $[H_k \ I]$ gives

$$[0 \ I - H_k G_k] \mathbf{1} = [H_k E_k \ F_k] \mathbf{1},$$

which implies that $(I - G_k H_k)\mathbf{1} = G_k F_k \mathbf{1} + E_k \mathbf{1}$. Similarly, pre-multiplying both sides of (25) by $\begin{bmatrix} I & \\ & G_k \end{bmatrix}$ we get the equality needed for the third triplet representation (24). \square

Using these triplet representations, we can perform in a subtraction-free fashion with the GTH-like algorithm all the inversions required for computing the matrix Ψ starting from T and C . However, going through the algorithm outlined in Section 3, there is still a place left where we need subtractions: when computing R using (15), we need to evaluate the diagonals of the matrices $C_+ + \beta T_{++}$ and $|C_-| + \alpha T_{--}$.

Recall that the parameters α and β are chosen so that these diagonal entries are positive, and the optimal values α_{opt} and β_{opt} are the smallest ones that achieve this goal. The same issue is discussed in [27, Section 4.3]; a simple fix is choosing $\alpha = \eta \alpha_{\text{opt}}$ and $\beta = \eta \beta_{\text{opt}}$ for some $\eta < 1$ not too close to 1 (for instance, $\eta = 0.9$ or $\eta = 0.5$). In this way, the algorithm still contains n subtractions but they are numerically safe, since we guarantee that the two terms are not too close to each other.

Remark 4.3. If one wishes to avoid subtractions completely, the following strategy can be used. Choose

$$\alpha := \left(\sum_{k \in \mathcal{S}_-} \frac{|T_{kk}|}{|c_k|} \right)^{-1} = \left(\sum_{k \in \mathcal{S}_-, j \in \mathcal{S}, j \neq k} \frac{T_{kj}}{|c_k|} \right)^{-1}.$$

Now, for each $i \in \mathcal{S}_-$ we can write

$$|c_i| + \alpha T_{ii} = \alpha |c_i| \left(\alpha^{-1} - \frac{|T_{ii}|}{|c_i|} \right) = \alpha |c_i| \left(\sum_{k \in \mathcal{S}_-, j \in \mathcal{S}, k \notin \{i, j\}} \frac{T_{kj}}{|c_k|} \right),$$

and the last expression contains only sums, products and inverses of nonnegative numbers. A similar expression holds for β , replacing \mathcal{S}_- with \mathcal{S}_+ . This choice of α satisfies $n^{-1} \alpha_{\text{opt}} \leq \alpha \leq \alpha_{\text{opt}}$, so we expect the algorithm with these values of α and β to require, in the worst case, $\log_2 n$ more steps than the one with α_{opt} and β_{opt} . \square

Larger values of α and β mean more steps required in the SDA, so this additional safeguard makes the algorithm slower. Nevertheless, thanks to the quadratic convergence speed, choosing $\eta = 0.5$ should correspond roughly to one more step needed.

To summarize, our proposed algorithm is presented as Algorithm 2.

4.3 Numerical accuracy

Relying on the results on the stability of the GTH-like algorithm, we can obtain explicit error bounds for the accuracy of the SDA algorithm. These bounds are stated in the following theorem.

Theorem 4.4. *Let \tilde{G}_k be the approximation of the matrix G_k computed by Algorithm 2, performed in floating-point arithmetic with machine precision u using $\alpha \leq 2\alpha_{\text{opt}}$ and $\beta \leq 2\beta_{\text{opt}}$. Then,*

$$|\tilde{G}_k - G_k| \leq \left(\left(k + K_1 + \frac{2(K_1 + K_2)n}{1 - \delta} \tilde{K}_0 \right) u + O(u^2) \right) \Psi, \quad (26)$$

Algorithm 2: Componentwise accurate SDA**Input:** $\text{offdiag}(T)$, C , ε , where $T \in \mathbb{R}^{n \times n}$ is the transition matrix of a continuous-time Markov chain, $C \in \mathbb{R}^{n \times n}$ is a diagonal rate matrix with positive rates in the first n_+ entries and negative ones in the remaining $n_- = (n - n_+)$ entries, ε is the convergence tolerance.**Output:** The minimal nonnegative solutions $\Psi \in \mathbb{R}^{n_+ \times n_-}$ of (9) and $\widehat{\Psi} \in \mathbb{R}^{n_- \times n_+}$ of (18)

- 1 Set $\alpha \leftarrow \eta\alpha_{\text{opt}}$, $\beta \leftarrow \eta\beta_{\text{opt}}$ using, for instance, $\eta = 0.5$;
- 2 Compute $T_{ii} = -\sum_{i \neq j} T_{ij}$;
- 3 Compute initial values E_0, F_0, G_0, H_0 according to (15), solving the required linear systems with Algorithm 1 and triplet representation (22);
- 4 $k \leftarrow 0$;
- 5 **repeat**
- 6 Compute $E_{k+1}, F_{k+1}, G_{k+1}, H_{k+1}$ according to (12), using Algorithm 1 for inversions with triplet representations (23) and (24);
- 7 $k \leftarrow k + 1$;
- 8 **until** $|G_{k+1} - G_k| \leq \varepsilon G_{k+1}$;
- 9 $\Psi \leftarrow G_k$; $\widehat{\Psi} \leftarrow H_k$.

where $\tilde{K}_0 > 0$ is a constant such that $\Psi - G_k \leq \tilde{K}_0 \delta^{2^k} \Psi$,

$$K_1 := \psi(n) + 5n + 2,$$

$$K_2 := \psi(n_{\max}) + 4n_{\max}^2 + 2n_{\max} - 1, \quad \text{with } n_{\max} = \max(n_+, n_-).$$

We recall that $\psi(n)$, defined in Theorem 4.1, is a degree-3 polynomial. The proof of Theorem 4.4 is rather technical and lengthy; therefore, we defer it to the Appendix.

4.4 Other triplet representations

We now focus on the computation of the stationary densities as described in Section 2, relying on the solutions computed with Algorithm 2. Using the Riccati solutions Ψ and $\widehat{\Psi}$, one can obtain the four matrices

$$\begin{aligned} W &= T_{--} + T_{-+}\Psi, & K &= C_+^{-1}T_{++} + \Psi|C_-^{-1}|T_{-+}, \\ \widehat{W} &= T_{++} + T_{+-}\widehat{\Psi}, & \widehat{K} &= |C_-^{-1}|T_{--} + \widehat{\Psi}C_+^{-1}T_{+-}. \end{aligned}$$

The matrix K was defined in (5), while W appeared in (7). They are all $-M$ -matrices, so the computation of their off-diagonal entries can be done directly without subtractions. In addition, triplet representations for their negatives can be computed explicitly.

Lemma 4.5.

1. $(\text{offdiag}(-W), \mathbf{1}, \mathbf{0})$ and $(\text{offdiag}(-\widehat{W}), \mathbf{1}, T_{-+}F_\infty \mathbf{1})$ are the right triplet representations for $-W$ and $-\widehat{W}$, respectively.
2. $(\text{offdiag}(-K), \boldsymbol{\xi}_+ C_+, \boldsymbol{\xi}_- |C_-| F_\infty |C_-^{-1}| T_{-+})$ and $(\text{offdiag}(-\widehat{K}), \boldsymbol{\xi}_- |C_-|, \mathbf{0})$ are the left triplet representations for $-K$ and $-\widehat{K}$, respectively.

Proof. Analogously to (25), by Corollary A.6

$$[\xi_+ C_+ \quad \xi_- |C_-|] \begin{bmatrix} I & -G_k \\ -H_k & I \end{bmatrix} = [\xi_+ C_+ \quad \xi_- |C_-|] \begin{bmatrix} E_k & 0 \\ 0 & F_k \end{bmatrix}.$$

We take the limit as $k \rightarrow \infty$ of the second column of both sides to obtain

$$\xi_- |C_-| - \xi_+ C_+ \Psi = \xi_- |C_-| F_\infty.$$

Since $\xi T = \mathbf{0}$, we have $\xi_+ T_{++} + \xi_- T_{-+} = \mathbf{0}$ and therefore

$$-\xi_+ C_+ K = (\xi_- |C_-| - \xi_+ C_+ \Psi) |C_-^{-1}| T_{-+} = \xi_- |C_-| F_\infty |C_-^{-1}| T_{-+}.$$

The other representations are proved in a similar way, using (25) and the fact that $E_k \rightarrow 0$ as $k \rightarrow \infty$. \square

4.5 Matrix exponentials, scaling and squaring, and SDA

The problem of computing e^{Kx} in (4) with high componentwise accuracy is discussed in [24], where different algorithms are presented for computing $\exp(A)$, A being a $-Z$ -matrix. For instance, the more basic algorithm consists in

- (i) choosing $\hat{A} \geq 0$ and $z \geq 0$ such $A = \hat{A} - zI$, and
- (ii) computing e^A as

$$e^A \approx e^{-z} \left(\mathcal{T}_m(2^{-s} \hat{A}) \right)^{2^s}, \quad (27)$$

where \mathcal{T}_m is the Taylor series of the exponential, truncated after its m th term, and s is a suitable positive integer whose magnitude depends on the norm of A .

The papers [28, 29] discuss the computation of essentially nonnegative matrices to high componentwise accuracy, with the componentwise condition number obtained in [28] and the error analysis for (27) presented in [29].

Algorithms of the form (27), where \mathcal{T}_m may be a rational approximation rather than a Taylor expansion, are known as *scaling and squaring* [17]. Matlab's standard algorithm `expm`, for instance, uses an algorithm of this family [16], and we found that it often delivers componentwise accurate results on the matrix on which we are using it, although it is not explicitly designed for this purpose¹. It is interesting to see that doubling is, indirectly, a scaling-and-squaring algorithm. Indeed, at each step the factorization

$$\begin{bmatrix} I & -G_k \\ 0 & F_k \end{bmatrix}^{-1} \begin{bmatrix} E_k & 0 \\ -H_k & I \end{bmatrix} = ((I_n - \alpha C^{-1} T)^{-1} (I_n + \beta C^{-1} T))^{2^k} \quad (28)$$

holds. The right-hand side is the same expression that would be obtained by applying a scaling and squaring method for computing $\exp(2^k(\alpha + \beta)C^{-1}T)$, with the inner function

$$\mathcal{T}(x) = \frac{1 + \frac{\beta}{\alpha + \beta} x}{1 - \frac{\alpha}{\alpha + \beta} x},$$

which is indeed a rational first-order approximation of e^x .

The formula (28) is known in the SDA literature, but, as far as we know, the observation that this corresponds to a scaling and squaring method, albeit simple, is novel. \square

¹ For an analysis of the componentwise accuracy of scaling and squaring algorithms, see [2].

4.6 Density computation

Putting everything together, we can evaluate the stationary density $\mathbf{f}(x)$ and the stationary probability mass at level zero of a Markov-modulated fluid queue $\{X(t), \varphi(t)\}$ using Algorithm 3. Note that all the computations are subtraction-free; in particular, for Line 3, recall that $K^{-1} \leq 0$ while $V \geq 0$ and $T_{-+} \geq 0$.

Algorithm 3: Stationary density and probability mass of a fluid queue

Input: The irreducible generator $T \in \mathbb{R}^{n \times n}$, diagonal rate matrix $C \in \mathbb{R}^{n \times n}$, $x > 0$.

Output: The stationary density $\mathbf{f}(x)$ and the stationary probability mass \mathbf{p}_- at level zero.

- 1 Compute Ψ using Algorithm 2;
 - 2 Compute a vector \mathbf{q} such that $\mathbf{q}W = \mathbf{0}$ using (a variant of) the GTH-like algorithm and the triplet representation in Lemma 4.5;
 - 3 Compute $\mathbf{p}_- = \beta^{-1}\mathbf{q}$, where $\beta = \mathbf{q}(\mathbf{1} - T_{-+}K^{-1}V\mathbf{1})$;
 - 4 Compute K and V according to (5) and (6);
 - 5 Compute $\mathbf{f}(x)$ according to (4), evaluating the matrix exponential as described in Section (4.5) or in [24], or with Matlab's `expm`;
-

5 Numerical experiments

We present two examples to demonstrate that updating correctly the triplet representations is crucial for attaining componentwise accuracy in the algorithm. For simplicity, we focus on the SDA and choose $\eta = 0.5$; hence, $\alpha = \beta = 0.5 \min(\alpha_{\text{opt}}, \beta_{\text{opt}})$.

Example 6.1 (weakly connected queue). Consider a weakly recurrent fluid model $\{X(t), \varphi(t)\}$, where the phase process $\varphi(t)$ has the state space $\mathcal{S} = \{1, 2, 3, 4, 5, 6\}$ and the transition matrix

$$T = \begin{bmatrix} -4 & 0 & 0 & 0 & 0 & 4 \\ 0 & -(15 + 10^{-8}) & 5 & 5 & 5 & 10^{-8} \\ 0 & 5 & -15 & 5 & 5 & 0 \\ 0 & 5 & 5 & -15 & 5 & 0 \\ 0 & 5 & 5 & 5 & -15 & 0 \\ 4 & 1 & 0 & 0 & 0 & -5 \end{bmatrix}, \quad (29)$$

and $C = \text{diag}(1, 1, 1, -1.001, -1.001, -1.001)$. The transition graph of the phase process $\varphi(t)$ is depicted in Figure 1.

This is a fluid queue whose phase process is composed of two weakly connected sets of states: the states $\{1, 6\}$ are difficult to reach from the others due to the comparatively small transition rates 10^{-8} and 1. Hence, we expected the stationary density of the queue to have very small values in these two states. We computed the exact solution Ψ with a high number of significant digits using variable precision arithmetic, and compared the normwise and componentwise




Fig. 1: The modulating Markov chain of a fluid queue, with two weakly connected sets of states, $\{4, 2, 3, 5\}$ and $\{1, 6\}$.

errors defined as

$$e_{\text{norm}} := \frac{\|\tilde{\Psi} - \Psi\|}{\|\Psi\|}, \quad e_{\text{cw}} := \max \frac{\tilde{\Psi}_{ij} - \Psi_{ij}}{\Psi_{ij}}.$$

The resulting matrix

$$\Psi \approx \begin{bmatrix} 0.195 & 0.195 & 0.61 \\ 0.5 & 0.5 & 2 \cdot 10^{-9} \\ 0.5 & 0.5 & 1.7 \cdot 10^{-9} \end{bmatrix}$$

has two entries, $\Psi_{2,3}$ and $\Psi_{3,3}$, that are much smaller in magnitude than the others. Probabilistically, these correspond to returning to the initial level in the hard-to-reach state 6 after starting from states 2 and 3, respectively. We compare in Table 1 the results of three algorithms:

GLX-SDA the non-componentwise accurate “standard” version of SDA introduced in Guo *et al.* [13],

XXL-SDA the variant suggested by Xue *et al.* [27], which employs triplet representations but recomputes them from the matrix entries at each step;

comp-SDA the componentwise accurate version of SDA that we introduce in this paper (Algorithm 2).

GLX-SDA		XXL-SDA		comp-SDA	
e_{norm}	e_{cw}	e_{norm}	e_{cw}	e_{norm}	e_{cw}
$1.3 \cdot 10^{-13}$	$2.1 \cdot 10^{-09}$	$1.0 \cdot 10^{-13}$	$3.6 \cdot 10^{-13}$	$2.0 \cdot 10^{-16}$	$9.3 \cdot 10^{-16}$

Tab. 1: Accuracy of various algorithms on the queue in Figure 1

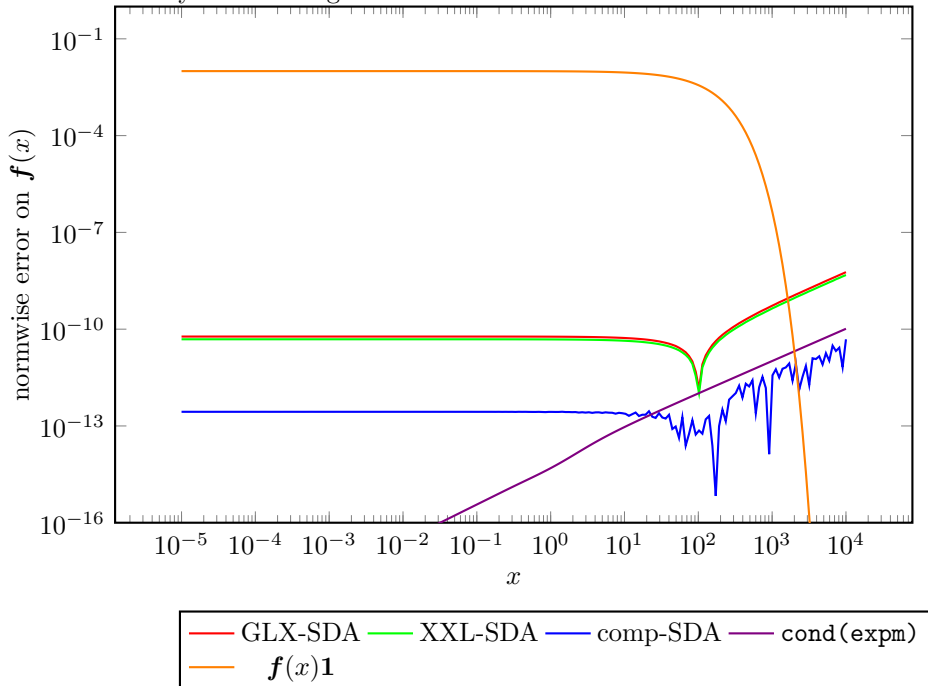
The non-componentwise accurate variant GLX-SDA computes a normwise-accurate matrix $\tilde{\Psi}$, but the two small entries are computed with a larger relative error:

$$\left[\frac{\tilde{\Psi}_{ij} - \Psi_{ij}}{\Psi_{ij}} \right] \approx \begin{bmatrix} -4 \cdot 10^{-13} & -4 \cdot 10^{-13} & 1 \cdot 10^{-16} \\ -6 \cdot 10^{-14} & -6 \cdot 10^{-14} & -2 \cdot 10^{-9} \\ -6 \cdot 10^{-14} & -6 \cdot 10^{-14} & -2 \cdot 10^{-9} \end{bmatrix}.$$

The algorithm XXL-SDA reaches a smaller error on these tiny entries, but the overall error is of the same order of magnitude as the standard algorithm, and is far from reaching the full machine precision of 16 significant digits.

As a second test encompassing the whole Algorithm 3, we computed the value of the probability density function $\mathbf{f}(x)$ with the three algorithms for several values of x spanning different orders of magnitude. We used Matlab's `expm` function rather than the simple one described in Section 4.5, since we found that it yields slightly better results on this example, despite having no explicit componentwise accuracy guarantees. The normwise errors on the resulting vectors $\mathbf{f}(x)$ are reported in Figure 2, while the componentwise ones in Figure 3.

Fig. 2: Normwise error for the stationary density $\mathbf{f}(x)$ computed with Ψ obtained by different algorithms.




The plots contain also an estimate of the condition number of the matrix exponential, computed using the method in [14] for the normwise condition number and [24, Equation 2.2] for the componentwise condition number, and the aggregate magnitude of the density function $\mathbf{f}(x)\mathbf{1}$. As one can see, for higher values of the fluid level x , the result cannot be computed with full machine precision anymore, and this is due to the higher condition number of the matrix exponential for large arguments, which dominates the error. This happens roughly at the same order of magnitude as the decay of the density function.

Notice also the unexpected dip of the error in correspondence of $x = 10^2$ — this seems to be a numerical artifact, possibly due to the computation of the matrix exponential; its position seems to vary when the parameter -1.001 in the matrix C is altered, but we do not have a compelling probabilistic or numeric explanation for its existence. ■

Example 6.2 (cascading process). Consider a fluid queue $\{X(t), \varphi(t)\}$ whose transition graph for $\varphi(t)$ is depicted in Figure 4, with state space $\mathcal{S} =$

Fig. 3: Componentwise error for the stationary density $f(x)$ computed with Ψ obtained by different algorithms.



$\{1, 2, \dots, 8\}$, generator T given by

$$T = \begin{bmatrix} -1 & 0 & 0 & 0 & 0 & 0 & 0 & 1.00 \\ 0 & -1.01 & 0 & 0 & 0.01 & 0 & 0 & 1.00 \\ 0 & 0 & -1.01 & 0 & 0 & 0.01 & 0 & 1.00 \\ 0 & 0 & 0 & -1.01 & 0 & 0 & 0.01 & 1.00 \\ 0.01 & 0 & 0 & 0 & -1.01 & 0 & 0 & 1.00 \\ 0 & 0.01 & 0 & 0 & 0 & -1.01 & 0 & 1.00 \\ 0 & 0 & 0.01 & 0 & 0 & 0 & -1.01 & 1.00 \\ 0 & 0 & 0 & 0.01 & 0 & 0 & 0 & -0.01 \end{bmatrix},$$

and $C = \text{diag}(\kappa, 1, 1, 1, -1, -1, -1, -1)$ for a real parameter κ . The state space \mathcal{S} consists of a “base” state, labeled 8, and a sequence of states each of which can be reached only by the previous one with a low transition rate; from each state, there is also a higher transition rate to the base state (see Figure 4).

This process models a fluid queue in which the stationary probabilities have varying orders of magnitudes not due to a single very low transition rate but to a series of moderately small ones: for instance, several unreliable backup systems failing one after the other. By varying the parameter κ across several orders of magnitude, we alter the fluid rate in the most unlikely state.

In Figure 5, we plot the magnitude of the two error measures e_{norm} and e_{cw} for the same algorithms as the previous example, for different values of M . As can be seen from the pictures, in this example XXL-SDA has essentially the same numerical behaviour as the non-componentwise accurate GLX-SDA: both have increasing errors as the fluid rate κ in the base state 8 increases. On the





Fig. 4: A fluid queue model for a cascading process

Fig. 5: Numerical results for the cascading process



other hand, our new algorithm can solve this problem with accuracy that is close to the machine precision $u \approx 2.2 \cdot 10^{-16}$. ■

Acknowledgement

The first author would like to acknowledge the financial support of the Australian Research Council through the Discovery Grant DP110101663. The second author would like to acknowledge the financial support of Istituto Nazionale di Alta Matematica, and thank N. Higham and M.-W. Shao for useful discussions on the accuracy of various algorithms for the computation of the exponential of a $-M$ -matrix; in particular, N. Higham pointed us to the paper [2]. Both authors are grateful to helpful comments of N. Bean.

A Doubling and censoring

In this section, we derive some results connecting the doubling map \mathcal{F} to the concept of *censoring* in Markov chains, which we shall need in the next section when we prove the numerical accuracy of Algorithm 2, our componentwise accurate SDA. In particular, Theorem A.1 and Lemma A.2 are standard results on censoring that we shall need.

Theorem A.1. *Consider $\mathcal{S} = \{1, \dots, m\} = S_1 \cup S_2$, where $S_1 \cap S_2 = \emptyset$, and let P be an $m \times m$ stochastic matrix partitioned as*

$$P = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix}, \quad (30)$$

where $P_{\ell m}$ contains the transition probabilities from S_ℓ to S_m . Suppose that $I - P_{11}$ is invertible. Then,

(i) $P' := P_{22} + P_{21}(I - P_{11})^{-1}P_{12}$ is stochastic.

(ii) If a row vector $\xi = [\xi_1 \quad \xi_2] \geq \mathbf{0}$ is such that $\xi P = \xi$, then $\xi_2 P' = \xi_2$.

The matrix P' is said to be the matrix obtained by *censoring* the set S_1 or the submatrix P_{11} . Censoring can alternatively be described in fully linear algebraic terms: the matrix $I - P'$ is the Schur complement of the submatrix $I - P_{11}$ in $I - P$.

Lemma A.2. *Let P be an $m \times m$ stochastic matrix, and \mathcal{T} and \mathcal{U} be two disjoint subsets of $\mathcal{S} = \{1, 2, \dots, m\}$, with $\mathcal{S} \setminus (\mathcal{T} \cup \mathcal{U}) \neq \emptyset$. Define P' to be the matrix obtained from P by censoring \mathcal{T} , and P'' the matrix obtained from P' by censoring \mathcal{U} .*

Then, P'' is also the matrix obtained from P by censoring $\mathcal{T} \cup \mathcal{U}$.

It was proved in [9] that the doubling map \mathcal{F} (13) corresponds to a step in another, similar iteration, *Cyclic Reduction*, on the matrices

$$A_+ := \begin{bmatrix} E & 0 \\ 0 & 0 \end{bmatrix}, \quad A_0 := \begin{bmatrix} 0 & G \\ H & 0 \end{bmatrix}, \quad A_- := \begin{bmatrix} 0 & 0 \\ 0 & F \end{bmatrix}. \quad (31)$$

Moreover, it was also established [5, 8] that Cyclic Reduction corresponds to Schur complementation on a suitable tridiagonal block-Toeplitz matrix; thus, it should not be surprising that there is a direct interpretation of doubling as Schur complementation/censoring, which is our next result. This result is slightly different from the existing ones, since we are interested with block-circulant matrices rather than block-tridiagonal ones here; nevertheless, the proof is similar (see, for instance, [8, Section 7.3]).

Let Z_m be the $m \times m$ *circulant generator matrix*, defined by

$$(Z_m)_{ij} = \begin{cases} 1 & j - i \equiv 1 \pmod{m}, \\ 0 & \text{otherwise,} \end{cases}$$

and denote by \otimes the Kronecker product.

Theorem A.3. *Let P be a $n \times n$ stochastic matrix, and define E , F , G , and H as in (11), and A_+ , $A_=$, and A_- as in (31). Then, for $k \geq 0$ the matrix $\mathcal{F}^k(P)$ is obtained by censoring the top $(2^k - 1)n \times (2^k - 1)n$ block of the $2^k n \times 2^k n$ matrix*

$$P^{(2^k)} := Z_{2^k}^{-1} \otimes A_- + I_{2^k} \otimes A_+ + Z_{2^k} \otimes A_+ \quad (32)$$

$$= \begin{bmatrix} A_= & A_+ & & & A_- \\ A_- & A_= & A_+ & & \\ & & A_- & A_= & \ddots \\ & & & \ddots & \ddots & A_+ \\ A_+ & & & & A_- & A_= \end{bmatrix}.$$

Here and in the following, for clarity we omit blocks of zeros from some large block matrices such as the one in (32).

Proof. We prove the theorem by induction. For $k = 1$,

$$P^{(2)} = \begin{bmatrix} A_= & A_- + A_+ \\ A_- + A_+ & A_= \end{bmatrix} = \left[\begin{array}{cc|cc} 0 & G & E & 0 \\ H & 0 & 0 & F \\ \hline E & 0 & 0 & G \\ 0 & F & H & 0 \end{array} \right].$$

Then, using the identity

$$\left(I - \begin{bmatrix} 0 & G \\ H & 0 \end{bmatrix} \right)^{-1} = \begin{bmatrix} I & G \\ H & I \end{bmatrix} \begin{bmatrix} (I - GH)^{-1} & 0 \\ 0 & (I - HG)^{-1} \end{bmatrix}$$

we can verify the result directly.

Now, we assume that the result holds for a certain k and prove it for $k + 1$. Let $\hat{P} := \mathcal{F}(P)$. We first prove that the matrix obtained by censoring out the components corresponding to the blocks with odd index numbers from $P^{(2^{k+1})}$ is exactly $\hat{P}^{(2^k)}$, that is, the matrix with the same structure (32) built starting from \hat{P} . The result then follows from Lemma A.2.

We first reorder the blocks in $P^{(2^{k+1})}$ to put the odd-numbered ones on the top, obtaining

$$\begin{bmatrix} A_= & & & & A_+ & & & & A_- \\ & A_= & & & A_- & A_+ & & & \\ & & \ddots & & & \ddots & \ddots & & \\ & & & A_= & & & A_- & A_+ \\ \hline A_- & A_+ & & & A_= & & & & \\ & & A_- & \ddots & & & & & \\ & & & \ddots & A_+ & & & & \\ A_+ & & & & A_- & & & & A_= \end{bmatrix}$$

$$= \left[\begin{array}{cccc|cccc} I \otimes A_= & & & & Z^{-1} \otimes A_- + I \otimes A_+ & & & \\ \hline I \otimes A_- + Z \otimes A_+ & & & & I \otimes A_= & & & \end{array} \right].$$

Next, we censor the top part, $I \otimes A_{=}$, to obtain

$$\begin{aligned} & I \otimes A_{=} + (I \otimes A_{-} + Z \otimes A_{+})(I \otimes (I - A_{=}))^{-1}(Z^{-1} \otimes A_{-} + I \otimes A_{+}) \\ &= Z^{-1} \otimes A_{-}(I - A_{=})^{-1}A_{-} + I \otimes (A_{=} + A_{-}(I - A_{=})^{-1}A_{+} \\ &\quad + A_{+}(I - A_{=})^{-1}A_{-}) + Z \otimes A_{+}(I - A_{=})^{-1}A_{+} \\ &= Z^{-1} \otimes \begin{bmatrix} 0 & 0 \\ 0 & \widehat{F} \end{bmatrix} + I \otimes \begin{bmatrix} 0 & \widehat{G} \\ \widehat{H} & 0 \end{bmatrix} + Z \otimes \begin{bmatrix} \widehat{E} & 0 \\ 0 & 0 \end{bmatrix}, \end{aligned}$$

where \widehat{E} , \widehat{F} , \widehat{G} , and \widehat{H} are the blocks of \widehat{P} as in (12). The last expression is precisely $\widehat{P}^{(2^k)}$. By the inductive hypothesis, the matrix obtained by censoring the top $(2^k - 1)n \times (2^k - 1)n$ of this matrix is $\mathcal{F}^k(\widehat{P}) = \mathcal{F}^{k+1}(P)$, which completes the proof. \square

Remark A.4. One can obtain a different, probabilistic proof for this result relying on a physical interpretation of doubling algorithms in terms of the underlying fluid queue. We are currently working on an explanation for these algorithms from a probabilistic point of view, giving some insight that complements the linear-algebraic point of view presented in this paper.

Similarly, we can express the initial values of the doubling algorithms as the censoring of a suitable stochastic matrix.

Theorem A.5. *Let (T, C) be the transition and rate matrices of a fluid queue, and let $\gamma \leq (\max_i Q_{ii})^{-1}$. Then, P_0 as in (15) is obtained by censoring the first n states out of*

$$S = \begin{bmatrix} I - \gamma Q & \gamma R \\ I & 0 \end{bmatrix}. \quad (33)$$

Proof. Clearly, $0 + I(I - (I - \gamma Q))^{-1}\gamma R = Q^{-1}R$, so it suffices to prove that S is stochastic. For the nonnegativity, we need only to verify that the diagonal entries of R are nonnegative, which holds because $\alpha \leq \alpha_{\text{opt}}, \beta \leq \beta_{\text{opt}}$, and that the diagonal entries of $I - \gamma Q$ are nonnegative, which holds because of the choice of γ . Moreover, substituting the definitions of Q and R from (15), and using $T\mathbf{1} = \mathbf{0}$, one sees that $S\mathbf{1} = \mathbf{1}$. \square

Corollary A.6. *Let $\xi \geq 0$ be such that $\xi T = \mathbf{0}$. Then, for $k \geq 0$, P_k is stochastic and $\xi|C|P_k = \xi|C|$.*

Proof. By their definitions and the fact that $\xi T = \mathbf{0}$,

$$\xi Q = \xi|C| = \xi R,$$

and consequently that $[\xi \quad \gamma\xi|C|]S = [\xi \quad \gamma\xi|C|]$. Hence, by applying Theorem A.1 to P_0 as a censoring of S , the corollary holds for $k = 0$.

For $k > 0$, one can verify directly that $(\mathbf{1}^\top \otimes \xi)P^{(2^k)} = (\mathbf{1}^\top \otimes \xi)$, so the result follows again from applying Theorem A.1 to P_k as a censoring of $P^{(2^k)}$. \square

B Componentwise stability of SDA

In this section, we prove the results on componentwise stability of SDA, thus giving the proof of Theorem 4.4. We begin with a perturbation bound that tells us how the iterates are affected by a small change in the initial values.

B.1 Componentwise perturbation bounds

We focus on first-order results, assuming ε to be a small parameter and hence ignoring all terms containing ε^2 and higher powers of ε . For brevity, we use the notation $A \dot{\leq} B$ to denote $A \leq B + O(\varepsilon^2)$. Moreover, we focus at first on the case $n_+ \approx n_-$; hence, we replace n_+ and n_- liberally with $n_{\max} := \max(n_+, n_-)$ in our bound. The results of a more accurate analysis that keeps track of the differences between n_+ and n_- are reported later in Section B.3.

We state here two useful lemmas. The first is a small variant of [1, Lemma 2.2], while the second is a simple consequence of the triangle inequality.

Lemma B.1. *Let $(\text{offdiag}(A), \mathbf{1}, \mathbf{w})$ and $(\text{offdiag}(\tilde{A}), \mathbf{1}, \tilde{\mathbf{w}})$ be triplet representations for two $m \times m$ M -matrices A and \tilde{A} such that*

$$|\text{offdiag}(A) - \text{offdiag}(\tilde{A})| \leq \varepsilon |\text{offdiag}(A)| \quad \text{and} \quad |\mathbf{w} - \tilde{\mathbf{w}}| \leq \varepsilon \mathbf{w}.$$

Then,

$$|A^{-1} - \tilde{A}^{-1}| \dot{\leq} (2m - 1)\varepsilon A^{-1}. \quad (34)$$

Lemma B.2. *Let A, B, \tilde{A} , and \tilde{B} be nonnegative matrices of suitable sizes such that $|A - \tilde{A}| \dot{\leq} a\varepsilon A$ and $|B - \tilde{B}| \dot{\leq} b\varepsilon B$. Then,*

- (i) $|A + B - (\tilde{A} + \tilde{B})| \dot{\leq} \max(a, b)\varepsilon(A + B)$.
- (ii) $|AB - \tilde{A}\tilde{B}| \dot{\leq} (a + b)\varepsilon AB$.

The following lemma provides a perturbation bound for censoring.

Lemma B.3. *Consider two stochastic matrices*

$$P = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix} \quad \text{and} \quad \tilde{P} = \begin{bmatrix} \tilde{P}_{11} & \tilde{P}_{12} \\ \tilde{P}_{21} & \tilde{P}_{22} \end{bmatrix},$$

with $P_{11}, \tilde{P}_{11} \in \mathbb{R}^{m \times m}$, $|\text{offdiag}(P - \tilde{P})| \leq \varepsilon \text{offdiag}(P)$, and suppose that $I - P_{11}$ is invertible. Then,

$$|\tilde{P}_{21}(I - \tilde{P}_{11})^{-1}\tilde{P}_{12} - P_{21}(I - P_{11})^{-1}P_{12}| \dot{\leq} (2m + 1)\varepsilon. \quad (35)$$

If, in addition, $|\text{diag}(P_{22}) - \text{diag}(\tilde{P}_{22})| \leq (2m + 1)\varepsilon P_{22}$, then

$$|\tilde{P}_{22} + \tilde{P}_{21}(I - \tilde{P}_{11})^{-1}\tilde{P}_{12} - P_{22} - P_{21}(I - P_{11})^{-1}P_{12}| \dot{\leq} (2m + 1)\varepsilon. \quad (36)$$

Proof. We have

$$|(I - P_{11})\mathbf{1} - (I - \tilde{P}_{11})\mathbf{1}| = |P_{12}\mathbf{1} - \tilde{P}_{12}\mathbf{1}| \leq \varepsilon P_{12}\mathbf{1}.$$

Thus, $I - P_{11}$ and $I - \tilde{P}_{11}$ satisfy the hypotheses of Lemma B.1 and hence $|(I - \tilde{P}_{11})^{-1} - (I - P_{11})^{-1}| \leq (2m - 1)\varepsilon(I - P_{11})^{-1}$. The results now follow by propagating errors with the help of Lemma B.2. \square

We start from a componentwise perturbation bound for P_0 .

Lemma B.4. *Let (T, C) and (\tilde{T}, \tilde{C}) be two pairs of the transition matrix and the rate matrix of two fluid queues such that*

$$|\text{offdiag}(T) - \text{offdiag}(\tilde{T})| \leq \varepsilon \text{offdiag}(T) \quad \text{and} \quad |C - \tilde{C}| \leq \varepsilon |C|.$$

Consider $P_0 = Q^{-1}R$ and its analogous $\tilde{P}_0 = \tilde{Q}^{-1}\tilde{R}$ computed using the same formulas (15) but starting from \tilde{T} and \tilde{C} , and with the same values of α and β , chosen such that $\alpha \leq \alpha_{\text{opt}}, \beta \leq \beta_{\text{opt}}$. Then,

$$|P_0 - \tilde{P}_0| \leq (2n + 1)\varepsilon P_0.$$

Proof. We have

$$\begin{aligned} |\tilde{R}_{ii} - R_{ii}| &= |\tilde{C}_+ + \beta \tilde{T}_{ii} - C_+ - \beta T_{ii}| \\ &\leq |\tilde{C}_+ - C_+| + \beta |\tilde{T}_{ii} - T_{ii}| \\ &\leq \varepsilon C_+ + \beta \sum_{j \neq i} |\tilde{T}_{ij} - T_{ij}| \\ &\leq \varepsilon C_+ + \beta \varepsilon \sum_{j \neq i} T_{ij} \\ &= \varepsilon R_{ii}. \end{aligned}$$

Hence, for S from (33) and its analog \tilde{S} (where we take the same value of γ for both), we have the bound $|\text{offdiag}(\tilde{S}) - \text{offdiag}(S)| \leq \varepsilon \text{offdiag}(S)$. By Lemma B.3, we obtain the required result. \square

Next, we derive a perturbation bound for the result of k steps of SDA.

Lemma B.5. *Let P and \tilde{P} be two $n \times n$ stochastic matrices such that $|\tilde{P} - P| \leq \varepsilon P$, and \mathcal{F} be the doubling map (13). Then,*

$$|\mathcal{F}^k(\tilde{P}) - \mathcal{F}^k(P)| \leq n2^k \varepsilon \mathcal{F}^k(P).$$

Proof. Consider the matrices $P^{(2^k)}$ and its analog $\tilde{P}^{(2^k)}$ defined in the same way but starting from \tilde{P} . They are stochastic matrices, and they satisfy the hypotheses of Lemma B.3. We apply the second part of this lemma, which completes the proof. \square

By applying the first part of Lemma B.3 rather than the second, we get a corresponding bound for

$$J_k := G_k - G_{k-1} = E_k(I - G_{k-1}H_{k-1})^{-1}G_{k-1}F_{k-1} \geq 0.$$

Corollary B.6. *With the hypotheses of Lemma B.5, we have*

$$|\tilde{J}_k - J_k| \leq n2^k \varepsilon J_k.$$

The above bound gives us immediately $|\tilde{G}_k - G_k| \leq n2^k \varepsilon G_k$ for the doubling iterates starting from two nearby sets of matrices. This bound, however, is quite weak: in particular, it would allow for $\lim_{k \rightarrow \infty} |\tilde{G}_k - G_k| = \infty$, while one would expect this quantity to converge, since $G_k \rightarrow \Psi$ and $\tilde{G}_k \rightarrow \tilde{\Psi}$.

Here, we overestimate the error: essentially, we compute a new term of the sum $G_k = G_0 + J_1 + J_2 + \dots + J_k$ at each step, and at the same time make the following estimate:

$$\begin{aligned} |\tilde{G}_k - G_k| &\leq |\tilde{G}_0 - G_0| + |\tilde{J}_1 - J_1| + |\tilde{J}_1 - J_1| + |\tilde{J}_2 - J_2| + \dots + |\tilde{J}_k - J_k| \\ &\leq \varepsilon G_0 + n\varepsilon J_1 + 2n\varepsilon J_2 + \dots + 2^k n\varepsilon J_k \\ &\leq 2^k n(G_0 + J_1 + J_2 + \dots + J_k). \end{aligned}$$

Indeed, the last line is how errors in sums are bounded using Lemma B.2. However, the factors J_k decay exponentially like $O(\delta^{2^k})$, and this compensates for the growing powers of two. The following theorem formalizes this intuition to improve this bound.

Theorem B.7. *With the hypotheses of Lemma B.5, we have*

$$|\tilde{G}_k - G_k| \leq \left(1 + \frac{4\hat{K}_0 n}{1 - \delta}\right) \varepsilon \Psi.$$

for a suitable constant \hat{K}_0 .

Proof. We know from Theorem 3.1 that $0 \leq \Psi - G_k \leq K_0 \delta^{2^k}$, and that $\Psi - G_k \leq \Psi$. Hence we can write, for a suitable constant \hat{K}_0 (different in general from K_0), $\Psi - G_k \leq \hat{K}_0 \delta^{2^k} \Psi$.

$$\begin{aligned} |\tilde{G}_k - G_k| &\leq |\tilde{G}_0 - G_0| + \sum_{h=1}^k |\tilde{J}_h - J_h| \\ &\leq \varepsilon G_0 + n\varepsilon \sum_{h=1}^k 2^h J_h \\ &\leq \varepsilon G_0 + n\varepsilon \sum_{h=1}^k 2^h (\Psi - G_{h-1}) \\ &\leq \varepsilon G_0 + n\varepsilon \sum_{h=1}^k 2^h \hat{K}_0 \delta^{2^{h-1}} \Psi \\ &\leq \varepsilon G_0 + 2\hat{K}_0 n\varepsilon \sum_{h=1}^k 2^{h-1} \delta^{2^{h-1}} \Psi \\ &\leq \varepsilon G_0 + 4\hat{K}_0 n\varepsilon \sum_{j=1}^{2^k} \delta^j \Psi \leq \varepsilon \Psi + \frac{4\hat{K}_0 n}{1 - \delta} \varepsilon \Psi \quad \square \end{aligned}$$

Clearly, an analogous result for H_k holds.

B.2 Numerical stability

We now establish a more complex result that keeps track also of the numerical errors in an implementation of the algorithm in floating-point arithmetic. We denote by u the machine precision, and ignore second-order terms in u by hiding them again in the syntax $A \dot{\leq} B$.

We shall rely on the two following lemmas. The first one is a version of Lemma B.2 that keeps track of floating point arithmetic errors as well.

Lemma B.8. *Let A and B be nonnegative matrices of suitable sizes, and let \tilde{A} and \tilde{B} be matrices of floating point numbers such that $|A - \tilde{A}| \leq auA$, $|B - \tilde{B}| \leq buB$.*

1. *Let \tilde{S} be the result of the matrix addition $\tilde{A} + \tilde{B}$ performed in floating point arithmetic. Then,*

$$|\tilde{S} - (A + B)| \leq (\max(a, b) + 1)u(A + B).$$

2. *Let \tilde{P} be the result of the matrix product $\tilde{A}\tilde{B}$ performed in floating point arithmetic. Then,*

$$|\tilde{P} - (AB)| \leq (a + b + m)u(AB),$$

where m is the number of columns of A (and rows of B).

The second one is a slight variation of Theorem 4.1 to convert it to the same form as the bounds in Lemma B.8.

Lemma B.9. *Let $A, \tilde{A} \in \mathbb{R}^{m \times m}$ be two M -matrices with triplet representations $(\text{offdiag}(A), \mathbf{1}, \mathbf{w})$ and $(\text{offdiag}(\tilde{A}), \mathbf{1}, \tilde{\mathbf{w}})$, and $B, \tilde{B} \in \mathbb{R}^{m \times k}$ be nonnegative, such that*

$$\begin{aligned} |\text{offdiag}(A) - \text{offdiag}(\tilde{A})| &\leq au|\text{offdiag}(A)|, \\ |B - \tilde{B}| &\leq buB, \\ |\mathbf{w} - \tilde{\mathbf{w}}| &\leq au\mathbf{w}, \end{aligned}$$

and all of $\text{offdiag}(\tilde{A})$, \tilde{B} , and $\tilde{\mathbf{w}}$ are floating-point numbers.

Let \tilde{X} be the result of the operation $\tilde{A}^{-1}\tilde{B}$ performed in floating-point arithmetic with the GTH-like algorithm. Then,

$$|\tilde{X} - A^{-1}B| \leq ((2m - 1)a + b + \psi(m))uA^{-1}B.$$

Proof. By Lemma B.1, $|A^{-1} - \tilde{A}^{-1}| \leq (2m - 1)uA^{-1}$. Using this result and Lemma B.2, we get

$$|\tilde{A}^{-1}\tilde{B} - A^{-1}B| \leq ((2m - 1)a + b)uA^{-1}B.$$

By Theorem 4.1,

$$|\tilde{X} - \tilde{A}^{-1}\tilde{B}| \leq \psi(m)u\tilde{A}^{-1}\tilde{B} \leq \psi(m)uA^{-1}B.$$

The result follows using the triangle inequality. \square

Our first result towards an accuracy proof concerns the initial values.

Proposition B.10. *Let Algorithm 2 be performed in floating-point arithmetic with $\alpha \leq 2\alpha_{\text{opt}}$ and $\beta \leq 2\beta_{\text{opt}}$. Then, the computed approximation \tilde{P}_0 to P_0 satisfies the bound*

$$|\tilde{P}_0 - P_0| \leq K_1uP_0,$$

with $K_1 := \psi(n) + 5n + 2$ and $\psi(n)$ defined as in Theorem 4.1.

Proof. We assume without loss of generality that α and β are exact machine numbers (otherwise, we simply replace them with their machine representation). The machine representations of T and c_i satisfy $|\tilde{T} - T| \leq uT$ and $|\tilde{c}_i - c_i| \leq uc_i$.

Let \tilde{Q} and \tilde{R} be the computed versions of Q and R in floating-point arithmetic. They are computed by taking the machine representation of the corresponding entry of T and multiplying it by α or β , hence

$$|\text{offdiag}(\tilde{Q}) - \text{offdiag}(Q)| \leq 2u \text{offdiag}(Q)$$

and

$$|\text{offdiag}(\tilde{R}) - \text{offdiag}(R)| \leq 2u \text{offdiag}(Q).$$

The quantities $d_i := \beta \sum_{j \neq i} T_{ij}$, for $i \in \mathcal{S}_+$, are computed (with $n-1$ additions and one multiplication) up to an error $|\tilde{d}_i - d_i| \leq (n+1)ud_i$, as we get by applying repeatedly Lemma B.8. Moreover, due to our choice of β , $c_i \geq 2d_i$, which implies $d_i \leq c_i - d_i$ and $c_i \leq 2(c_i - d_i)$. Hence we have

$$\begin{aligned} |\text{fl}(\tilde{c}_i - \tilde{d}_i) - (c_i - d_i)| &= |(\tilde{c}_i - \tilde{d}_i)(1 + \varepsilon) - (c_i - d_i)| \\ &\leq |\tilde{c}_i - c_i| + |\tilde{d}_i - d_i| + u|\tilde{c}_i - \tilde{d}_i| \\ &\leq uc_i + (n+1)ud_i + u(c_i - d_i) \\ &\leq (2+n+1)u(c_i - d_i) = (n+4)u(c_i - d_i). \end{aligned}$$

A similar bound applies for $i \in \mathcal{S}_-$. Putting everything together, we have $|\tilde{R} - R| \leq (n+4)R$, $|\text{offdiag}(\tilde{Q}) - \text{offdiag}(Q)| \leq 2u \text{offdiag}(Q)$, and $|\tilde{C}\mathbf{1} - C\mathbf{1}| \leq uC\mathbf{1}$. Hence, by Lemma B.9

$$|\tilde{P}_0 - P_0| \leq (2(2n-1) + (n+4) + \psi(n))uP_0.$$

□

We are now ready to assess the numerical error a single SDA step.

Lemma B.11. *Let $\tilde{\mathcal{F}}$ be the map that computes approximately the result of one iteration of SDA implementing Equations (12) in machine arithmetic. For any $n \times n$ stochastic matrix P , we have*

$$|\tilde{\mathcal{F}}(P) - \mathcal{F}(P)| \leq K_2 u \mathcal{F}(P),$$

with $K_2 := \psi(n_{\max}) + 4n_{\max}^2 + 2n_{\max} - 1$

Proof. Let E, F, G, H be the blocks of P , $\hat{E}, \hat{F}, \hat{G}, \hat{H}$ those of $\mathcal{F}(P)$, and $\tilde{E}, \tilde{F}, \tilde{G}, \tilde{H}$ those of $\tilde{\mathcal{F}}(P)$. For simplicity, with a slight abuse of notation, for any matrix $X \geq 0$ appearing in the algorithm we denote here by $c(X)$ its computed value in machine arithmetic along the SDA step computation. Applying Lemma B.8 repeatedly, we have $|c(GH) - GH| \leq n_- u \tilde{G}\tilde{H}$ and

$$\begin{aligned} |c(F\mathbf{1}) - F\mathbf{1}| &\leq n_- u F\mathbf{1}, \\ |c(GF\mathbf{1}) - GF\mathbf{1}| &\leq 2n_- u GF\mathbf{1}, \\ |c(E\mathbf{1}) - E\mathbf{1}| &\leq n_+ u E\mathbf{1}, \\ |c(GF\mathbf{1} + E\mathbf{1}) - GF\mathbf{1} - E\mathbf{1}| &\leq (2n_{\max} + 1)u(GF\mathbf{1} - E\mathbf{1}). \end{aligned}$$

Hence

$$\begin{aligned} |c((I - GH)^{-1}G) - (I - GH)^{-1}G| \\ \leq ((2n_{\max} - 1)(2n_{\max} + 1) + \psi(n_{\max}))u(I - GH)^{-1}G, \end{aligned}$$

and carrying on the error propagation in the same fashion leads to

$$\begin{aligned} |\tilde{G} - \hat{G}| &= |c(G + E(I - GH)^{-1}GF) - G - E(I - GH)^{-1}GF| \\ &\leq (\psi(n_{\max}) + 4n_{\max}^2 + 2n_{\max} - 1)u\hat{G}. \end{aligned}$$

The error computations for the other three iterates \tilde{E} , \tilde{F} , \tilde{H} are analogous. \square

We are now ready to tackle several doubling steps together.

Theorem B.12. *Let $\tilde{P}_k = \tilde{\mathcal{F}}^k(\tilde{P}_0)$ be the approximation to $P_k = \mathcal{F}^k(P_0)$ computed by Algorithm 2 performed in floating-point arithmetic with machine precision u , using values of the parameters α, β such that $\alpha \leq 2\alpha_{\text{opt}}$ and $\beta \leq 2\beta_{\text{opt}}$. Then,*

$$|\tilde{P}_k - P_k| \leq n2^k(K_1 + K_2)uP_k.$$

Proof. We prove the result by induction on k .

The following manipulation is nothing but a formal version of the statement that when considering first-order error bounds, we can simply add up the local errors in every step of the algorithm. We start from the telescopic sum

$$\begin{aligned} |\tilde{\mathcal{F}}^k(\tilde{P}_0) - \mathcal{F}^k(\tilde{P}_0)| &\leq \sum_{h=1}^k |\mathcal{F}^{h-1}\tilde{\mathcal{F}}^{k-h+1}(\tilde{P}_0) - \mathcal{F}^h\tilde{\mathcal{F}}^{k-h}(\tilde{P}_0)| \\ &= \sum_{h=1}^k |\mathcal{F}^{h-1}\tilde{\mathcal{F}}(\tilde{P}_{k-h}) - \mathcal{F}^{h-1}\mathcal{F}(\tilde{P}_{k-h})|. \end{aligned} \quad (37)$$

By Lemma B.11,

$$|\tilde{\mathcal{F}}(\tilde{P}_{k-h}) - \mathcal{F}(\tilde{P}_{k-h})| \leq K_2u\mathcal{F}(\tilde{P}_{k-h});$$

hence we can apply Lemma B.5 with $\varepsilon = K_2u$, to obtain for each $h \geq 1$

$$\begin{aligned} |\mathcal{F}^{h-1}\tilde{\mathcal{F}}(\tilde{P}_{k-h}) - \mathcal{F}^{h-1}\mathcal{F}(\tilde{P}_{k-h})| &\leq n2^{h-1}K_2u\mathcal{F}^h(\tilde{P}_{k-h}) \\ &\leq n2^{h-1}K_2u\mathcal{F}^k(P_0). \end{aligned}$$

In the last line, we have substituted P_{k-h} for \tilde{P}_{k-h} , which is possible because the two expressions differ only for a term of magnitude $O(u)$ by inductive hypothesis.

Insert this inequality into (37) to obtain

$$\begin{aligned} |\tilde{\mathcal{F}}^k(\tilde{P}_0) - \mathcal{F}^k(\tilde{P}_0)| &\leq \sum_{h=1}^k n2^{h-1}K_2u\mathcal{F}^k(P_0) \\ &< n2^k K_2u\mathcal{F}^k(P_0). \end{aligned} \quad (38)$$

Moreover, using again Lemma B.5 we get

$$|\mathcal{F}^k(\tilde{P}_0) - \mathcal{F}^k(P_0)| \leq n2^k K_1u\mathcal{F}^k(P_0). \quad (39)$$

Combining (38) and (39) we get the required result. \square

Again, this error bound is exponential in k , which is quite unwelcome. We can improve it using the same trick as in Theorem B.7. Since computing J_k is just removing the final addition from the algorithm, the previous bound can be easily modified to yield $|\tilde{J}_k - J_k| \leq n2^k(K_1 + K_2)uJ_k$. Now we proceed as in the proof of Theorem B.7 to prove Theorem 4.4.

Proof. (of Theorem 4.4) The matrix \tilde{G}_k is the result of the addition (in machine arithmetic) of the matrices $\tilde{G}_0, \tilde{J}_1, \tilde{J}_2, \dots, \tilde{J}_k$, performed in this order at different steps along the algorithm. For this addition, we have

$$\begin{aligned}
|\tilde{G}_k - G_k| &\leq |(\tilde{G}_0 + \sum_{h=1}^k \tilde{J}_h)(1 + ku) - (G_0 + \sum_{h=1}^k J_h)| \\
&\leq kuG_k + |\tilde{G}_0 - G_0| + \sum_{h=1}^k |\tilde{J}_h - J_h| \\
&\leq kuG_k + K_1uG_0 + \sum_{h=1}^k n2^h(K_1 + K_2)uJ_h \\
&\leq kuG_k + K_1uG_0 + (K_1 + K_2)n \sum_{h=1}^k 2^h u \tilde{K}_0 \delta^{2^h} \Psi \\
&\leq ku\Psi + K_1u\Psi + \frac{2(K_1 + K_2)n}{1 - \delta} \tilde{K}_0 \Psi.
\end{aligned}$$

□

B.3 The unbalanced case

In the previous bounds we have replaced n_+ and n_- with n_{\max} for simplicity; hence, these bounds are overestimated whenever $n_- \neq n_+$. It turns out that there is a different strategy that we can adopt in this case, which leads to a slightly smaller value of K_2 . First of all, we recall from [27] that the following alternative formulas equivalent to (12) can also be used for the doubling step

$$E_{k+1} := E_k(I + G_k(I - H_k G_k)^{-1} H_k) E_k, \quad (40a)$$

$$F_{k+1} := F_k(I + H_k(I - G_k H_k)^{-1} G_k) F_k, \quad (40b)$$

$$G_{k+1} := G_k + E_k G_k (I - H_k G_k)^{-1} F_k, \quad (40c)$$

$$H_{k+1} := H_k + F_k H_k (I - G_k H_k)^{-1} E_k. \quad (40d)$$

If $n_+ = n_-$, the formulas (12) are the ones with the lower computational cost; however, if $n_+ < n_-$, one can invert the $n_+ \times n_+$ matrix $I - G_k H_k$ rather than the larger $n_- \times n_-$ matrix $I - H_k G_k$, that is, use (12a), (40b), (12c), (40d), and vice versa if $n_- < n_+$.

By using the appropriate formulas, one can derive a slightly tighter version of Lemma B.11. We report only the final result here, since the proof is basically unchanged.

Theorem B.13. *If we choose for each iterate the formula with the smallest matrix to invert among (12) and (40), then in Lemma B.11 (as well as in*

Theorems B.12 and 4.4) we may replace the factor K_2 with

$$\begin{aligned}\tilde{K}_2 &:= \psi(n_{\min}) + (2n_{\max} + 1)(2n_{\min} - 1) + 2n_{\max} + n_{\min} \\ &= \psi(n_{\min}) + 4n_{\max}n_{\min} + 3n_{\min}.\end{aligned}$$

References

- [1] A. S. Alfa, J. Xue, and Q. Ye. Accurate computation of the smallest eigenvalue of a diagonally dominant M -matrix. *Math. Comp.*, 71(237):217–236, 2002.
- [2] M. Arioli, B. Codenotti, and C. Fassino. The Padé method for computing the matrix exponential. *Linear Algebra Appl.*, 240:111–130, 1996.
- [3] N. G. Bean, M. M. O’Reilly, and J. E. Sargison. A stochastic fluid flow model of the operation and maintenance of power generation systems. *IEEE Transactions on Power Systems*, 25(3):1361–1374, 2010.
- [4] A. Berman and R. 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.
- [5] D. A. Bini and L. Gemignani. Solving quadratic matrix equations and factoring polynomials: new fixed point iterations based on Schur complements of Toeplitz matrices. *Numer. Linear Algebra Appl.*, 12(2-3):181–189, 2005.
- [6] D. A. Bini, B. Iannazzo, and B. Meini. *Numerical solution of algebraic Riccati equations*, volume 9 of *Fundamentals of Algorithms*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2012.
- [7] D. A. Bini, B. Iannazzo, B. Meini, and F. Poloni. Nonsymmetric algebraic Riccati equations associated with an M -matrix: recent advances and algorithms. In V. Olshevsky and E. Tyrtysnikov, editors, *Matrix Methods: Theory, Algorithms and Applications*, chapter 10, pages 176–209. World Scientific Publishing, 2010.
- [8] D. A. Bini, G. Latouche, and B. Meini. *Numerical Methods for Structured Markov Chains*. Numerical Mathematics and Scientific Computation. Oxford University Press, Oxford, 2005.
- [9] D. A. Bini, B. Meini, and F. Poloni. Transforming algebraic Riccati equations into unilateral quadratic matrix equations. *Numer. Math.*, 116(4):553–578, 2010.
- [10] A. da Silva Soares. *Fluid queues – Building upon the analogy with QBD processes*. PhD thesis, Université libre de Bruxelles, 2005.
- [11] G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, third edition, 1996.
- [12] M. Govorun, G. Latouche, and M.-A. Remiche. Stability for fluid queues: characteristic inequalities. *Stoch. Models*, 29(1):64–88, 2013.

-
- [13] 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.
- [14] N. J. Higham. The Matrix Function Toolbox. <http://www.ma.man.ac.uk/~higham/mfttoolbox>.
- [15] N. J. Higham. *Accuracy and stability of numerical algorithms*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, second edition, 2002.
- [16] N. J. Higham. The scaling and squaring method for the matrix exponential revisited. *SIAM J. Matrix Anal. Appl.*, 26(4):1179–1193 (electronic), 2005.
- [17] N. J. Higham. *Functions of matrices*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2008. Theory and computation.
- [18] G. Latouche and P. G. Taylor. A stochastic fluid model for an ad hoc mobile network. *Queueing Systems*, 63(1):109–129, 2009.
- [19] M. Mandjes, D. Mitra, and W. Scheinhardt. Simple models of network access, with applications to the design of joint rate and admission control. In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 1, pages 3–12, 2002.
- [20] D. Mitra. Stochastic theory of a fluid model of producers and consumers coupled by a buffer. *Adv. in Appl. Probab.*, 20(3):646–676, 1988.
- [21] F. Poloni and T. Reis. The SDA method for numerical solution of Lur’ equations. Technical Report 1101.1231, arXiv.org, 2011.
- [22] V. Ramaswami. Matrix analytic methods for stochastic fluid flows. In D. Smith and P. Hey, editors, *Teletraffic Engineering in a Competitive World (Proceedings of the 16th International Teletraffic Congress)*, pages 1019–1030. Elsevier Science B.V., Edinburgh, UK, 1999.
- [23] L. C. G. Rogers. Fluid models in queueing theory and Wiener-Hopf factorization of Markov chains. *Ann. Appl. Probab.*, 4:390–413, 1994.
- [24] M. Shao, W. Gao, and J. Xue. Componentwise high relative accuracy algorithms for the exponential of an essentially nonnegative matrix. Technical Report UMINF 12/04, Umeå University, 2012.
- [25] N. van Foreest, M. Mandjes, and W. Scheinhardt. Analysis of a feedback fluid model for heterogeneous TCP sources. *Stoch. Models*, 19(3):299–324, 2003.
- [26] W.-G. Wang, W.-C. Wang, and R.-C. Li. Alternating-directional doubling algorithm for M -matrix algebraic Riccati equations. *SIAM J. Matrix Anal. Appl.*, 33(1):170–194, 2012.
- [27] J. Xue, S. Xu, and R.-C. Li. Accurate solutions of M -matrix algebraic Riccati equations. *Numer. Math.*, 120:671–700, 2012.

-
- [28] J. Xue and Q. Ye. Entrywise relative perturbation bounds for exponentials of essentially nonnegative matrices. *Numer. Math.*, 110:393–403, 2008.
- [29] J. Xue and Q. Ye. Computing exponentials of essentially nonnegative matrices entrywise to high relative accuracy. *Math. Comp.*, 82:1577–1596, 2013.