

Efficient approximation of the exponential operator for discrete 2D advection-diffusion problems

Luca Bergamaschi^{1,*}, Marco Caliari², and Marco Vianello³

¹ *Dipartimento di Metodi e Modelli Matematici per le Scienze Applicate, Università di Padova, via
Belzoni 7, 35131 Padova, Italy,*

² *Dottorato di Ricerca in Matematica Computazionale, Università di Padova, via Belzoni 7, 35131 Padova,
Italy*

³ *Dipartimento di Matematica Pura e Applicata, Università di Padova, via Belzoni 7, 35131 Padova, Italy*

SUMMARY

In this paper we compare Krylov subspace methods with Faber series expansion for approximating the matrix exponential operator on large, sparse, nonsymmetric matrices. We consider in particular the case of Chebyshev series, corresponding to an initial estimate of the spectrum of the matrix by a suitable ellipse. Experimental results upon matrices with large size, arising from space discretization of 2D advection-diffusion problems, demonstrate that the Chebyshev method can be an effective alternative to Krylov techniques. Copyright © 2003 John Wiley & Sons, Ltd.

KEY WORDS: Exponential operator, Chebyshev series, advection-diffusion problems, sparse matrices

1. Introduction

Spatial discretization of linear advection-diffusion equations like

$$\begin{cases} \frac{\partial u}{\partial t} = \Delta u - \theta_1 \frac{\partial u}{\partial x} - \theta_2 \frac{\partial u}{\partial y}, & (x, y) \in \Omega, t > 0 \\ u(x, y, 0) = u_0(x, y), & (x, y) \in \Omega \\ u(x, y, t) \equiv 0, & (x, y) \in \partial\Omega, t > 0 \end{cases} \quad (1)$$

where $\Omega \in \mathbb{R}^2$, yields naturally large systems of ordinary differential equations of the form

$$\begin{cases} \dot{\mathbf{y}}(t) = B\mathbf{y}(t), t > 0 \\ \mathbf{y}(0) = \mathbf{y}_0, \end{cases} \quad (2)$$

where B is a large sparse *nonsymmetric* $n \times n$ matrix. In several papers [5, 8, 9, 11, 12, 13] explicit time integration schemes have been proposed for the solution of (2), which rest on the

*Correspondence to: Luca Bergamaschi, Dipartimento di Metodi e Modelli Matematici per le Scienze Applicate
Università di Padova, via Belzoni 7, 35131 Padova, Italy

Contract/grant sponsor: University of Padova, GNCS-INdAM; contract/grant number: MURST 2000

possibility of efficiently approximating the exponential operator $\exp(A)\mathbf{v}$, where $\mathbf{v} \in \mathbb{R}^n$ is a given vector, $A = \tau B$ and $\tau > 0$ is a scaling factor related to the step-size Δt .

In this framework, the Krylov subspace approach is the most popular for computing $\exp(\tau B)\mathbf{v}$; see, e. g., [4, 12, 14, 23, 25]. As known, the general idea of Krylov-based methods in computing $f(A)\mathbf{v}$ for (analytic) matrix functions, consists of approximately projecting $f(A)$ onto a “small” Krylov subspace $\mathbf{K}_m = \text{Span}\{\mathbf{v}, A\mathbf{v}, \dots, A^{m-1}\mathbf{v}\}$. This reduces the problem to compute $f(H_m)$ in view of the approximation

$$f(A)\mathbf{v} \approx \|\mathbf{v}\|_2 V_m f(H_m) \mathbf{e}_1, \quad (3)$$

$V_m = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m]$, $\{\mathbf{v}_i\}$ being the orthonormal basis of \mathbf{K}_m and H_m the Hessenberg matrix produced by the Arnoldi process with $\mathbf{v}/\|\mathbf{v}\|_2$ as initial vector. When A is symmetric, the Krylov process reduces to the Lanczos method which is based on a *three-term recurrence*. In [1], we analyzed the Chebyshev series approximation of $f(A) = \exp(A)$ for A symmetric negative-definite, which is in turn based on the three-term recurrence for Chebyshev polynomials on the spectral interval of A (and whose implementation does not require any scalar product). Despite of the preprocessing cost due to extremal eigenvalue approximation and of the generally higher convergence rate of the Lanczos method, the Chebyshev series approach proves an effective alternative to Krylov techniques, especially when memory bounds do not allow the storage of the whole subspace \mathbf{K}_m .

In the nonsymmetric case the Ritz vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_m\}$ satisfy a *long-term recurrence* which leads to a $O(m^2)$ complexity. Actual implementations of the exponential operator [23, 25] try to keep m reasonably small in order to control the quadratic cost, and possibly the storage occupancy when the problem is extremely large. This *a priori* choice of m entails that the time step Δt is dynamically subdivided into a sequence of small substeps $\tau_i, i = 1, \dots, p$, where

$$\exp(\Delta t B)\mathbf{v} = \exp(\tau_p B)(\exp(\tau_{p-1} B)(\dots(\exp(\tau_1 B)\mathbf{v})\dots)), \quad (4)$$

in order to attain a prescribed accuracy.

The natural extension of Chebyshev series to function of nonsymmetric matrices is given by Faber series. Given a compact set $\Omega \subset \mathbb{C}$ including the spectrum $\sigma(A)$ of A , and assuming f analytic on a neighborhood of Ω , we can approximate $f(A)\mathbf{v}$ with

$$f(A)\mathbf{v} \approx \mathbf{s}_m := S_m(A)\mathbf{v} := \sum_{j=0}^m a_j(f) F_j(A)\mathbf{v}, \quad (5)$$

where F_j and the a_j are respectively the j th Faber polynomial and the j th Faber coefficient of the scalar function $f(z)$, $z \in \Omega$. It is important to notice that, when f is an entire function (as the exponential), convergence of (4) is *superlinear*. As in the Krylov approach, a *long-term recurrence* is now involved. A possibility to overcome this drawback, once the spectrum has been estimated by the convex hull of some marginal eigenvalues, consists in truncating the related Faber recurrence at a fixed length [19, 21]. Alternatively, we propose here to estimate the spectrum of A by an *ellipse* in which case, as it is well known, the recurrence becomes three-term and Faber polynomials become scaled and translated Chebyshev polynomials. This choice is particularly well-justified in the present framework, since spectra of “advection-diffusion” matrices, arising from stable FD discretizations, turn out to be very similar to ellipses (possibly degenerating); see the figures in §4. Our estimating ellipse is obtained cheaply via a rough

approximation of the extremal eigenvalues, using few Arnoldi iterations (within the ARPACK routines [15]).

Starting from the pioneering work by Manteuffel [16, 17] in the '70s, it is worth recalling here the use of Chebyshev and Faber polynomials in the construction of efficient semi-iterative methods for the solution of nonsymmetric linear systems. Also in these cases, a cheap estimate of the spectrum by Arnoldi iterations represents a key ingredient; see, e. g. [6, 28]. Our numerical results show that, despite of the preprocessing stage of extremal eigenvalues evaluation, the Chebyshev approximation of the exponential matrix method is almost always faster than Krylov, on advection-diffusion discretization matrices. This is due to the need of step fractioning in Krylov implementations, and happens even when the Krylov threshold dimension m is optimally chosen, with speedup ratios up to about 3. On the other hand, the Chebyshev preprocessing stage can become negligible in practice when time discretization of (2) is considered and several matrix exponential have to be computed, since the extremal eigenvalues of B can be computed once and for all. Moreover, this stage could benefit of the growing advances in the computation of eigenvalues of nonsymmetric matrices (we quote e. g. the recent Jacobi-Davidson method [26]).

The paper is organized as follows. In §2 we give an outline about the fundamental properties of Faber and Chebyshev polynomials and series in the complex domain, and in §3 we present some results on error bounds for the approximation of the exponential operator. In §4 we present and discuss our implementation of the Chebyshev algorithm. Finally, in §5 we compare Krylov and Chebyshev methods on a set of numerical tests, concerning computation of $\exp(\Delta t B)v$, where B comes from Finite Difference (FD) discretization of the advection-diffusion equation (1). These results show that, even more clearly than in the symmetric case [1], the Chebyshev approach is an effective alternative to Krylov methods in computing the exponential operator.

2. Background on Faber series

In this section we recall some basic classical results, concerning approximation of analytic functions on compact sets of the complex plane based on Faber series. For a throughout discussion of this topic we refer the reader e. g. to [3, 18, 27].

Let $B[0, r] := \{w \in \mathbb{C} : |w| \leq r\}$ and $\overline{\mathbb{C}} := \mathbb{C} \cup \{\infty\}$; if $K \subset \mathbb{C}$ is a compact set with more than one point, then there is a function $w = \phi(z)$ which maps $\overline{\mathbb{C}} \setminus K$ conformally onto $\overline{\mathbb{C}} \setminus B[0, 1]$ and satisfies the conditions

$$\phi(\infty) = \infty, \quad \lim_{z \rightarrow \infty} \frac{\phi(z)}{z} = \frac{1}{\gamma},$$

where γ is called *capacity* of K ; moreover, given any integer $j > 0$, the function $[\phi(z)]^j$ has a Laurent expansion at infinity of the form

$$[\phi(z)]^j = \frac{1}{\gamma^j} \left(z^j + d_0^{(j)} z^{j-1} + \dots + d_{j-1}^{(j)} + \frac{d_j^{(j)}}{z} + \dots \right),$$

(cf. [18, Theorem 3.14 and Corollary]). The polynomial

$$F_j(z) := \frac{1}{\gamma^j} (z^j + d_0^{(j)} z^{j-1} + \dots + d_{j-1}^{(j)})$$

is called j th *Faber polynomial* generated by K . Now, let $z = \psi(w)$ the inverse function of $\phi(z)$; its Laurent expansion at infinity takes the form

$$\psi(w) = \gamma \left(w + c_0 + \frac{c_1}{w} + \dots \right). \quad (6)$$

Then, if we set $F_0(z) := 1$, we get the following recurrence for the Faber polynomials (cf. [3])

$$\begin{aligned} F_0(z) &= 1, & F_1(z) &= \frac{z}{\gamma} - c_0, \\ F_j(z) &= F_1(z) \cdot F_{j-1}(z) - (c_1 F_{j-2}(z) + c_2 F_{j-3}(z) + \dots + \\ &+ c_{j-1} F_0(z)) - (j-1)c_{j-1}, & j &\geq 2. \end{aligned} \quad (7)$$

For any $R \geq \gamma$, let $\Gamma_R := \{z : |\phi(z)| = R/\gamma\}$ and K_R the bounded domain with boundary Γ_R ($K_R \supseteq K_\gamma \equiv K$): from [18, Theorem 3.17] we know that every analytic function f on $\text{int}_{\mathbb{C}} K_{R_1}$ ($R_1 > \gamma$) can be expanded into the *Faber series*

$$f(z) = \sum_{j=0}^{\infty} a_j(f) F_j(z), \quad (8)$$

where

$$a_j(f) := \frac{1}{2\pi i} \int_{|w|=\frac{R_0}{\gamma}} \frac{f(\psi(w))}{w^{j+1}} dw, \quad j \geq 0, \quad \gamma < R_0 < R_1 \quad (9)$$

are the *Faber coefficients*.

When $f(\cdot) \equiv \exp(\cdot)$, due to its analyticity on the whole complex plane, the right-hand side of (8) converges uniformly on every compact subset of \mathbb{C} (cf. [27, §2.1.3]). Moreover, setting $S_m(z) := \sum_{j=0}^m a_j(\exp) F_j(z)$, we know from [7] that $\{S_m\}_m$ is *maximally convergent* to the exponential function in K , that is

$$\limsup_{m \rightarrow \infty} \|\exp(\cdot) - S_m(\cdot)\|_K^{1/m} = \limsup_{m \rightarrow \infty} \|\exp(\cdot) - p_m^*(\cdot)\|_K^{1/m},$$

where $p_m^*(\cdot)$ is the best uniform approximation polynomial of degree m for $\exp(\cdot)$ on K and $\|\cdot\|_K$ denotes the maximum norm on K .

From these properties, a polynomial approximation method for the *matrix exponential operator* can be derived. Indeed, given K compact set, $K \supseteq \sigma(A)$, considering the matrix polynomial

$$S_m(A) = \sum_{j=0}^m a_j(\exp) F_j(A), \quad A \in \mathcal{M}_{n \times n}(\mathbb{C}) \quad (10)$$

it is well known that

$$\lim_{m \rightarrow \infty} S_m(A) \mathbf{v} = \exp(A) \mathbf{v}, \quad \mathbf{v} \in \mathbb{C}^n \quad (11)$$

even if A is not diagonalizable, and this convergence is *asymptotically optimal* with respect to K (cf. [21] and references therein).

Remark 2.1. When K coincides with the closure of the internal part of an ellipse $\mathcal{E}(d, c, a)$ (symmetric with respect to the real axis, with center d , foci $d - c$ and $d + c$, intersections with real axis $d - a$ and $d + a$), it is known that ψ in (6) takes the form

$$\psi(w) = \gamma w + d + \frac{a-b}{2} \cdot \frac{1}{w} = \gamma \left(w + \frac{d}{\gamma} + \frac{c^2}{4\gamma^2} \frac{1}{w} \right), \quad \gamma = \frac{a+b}{2},$$

where $b = \sqrt{a^2 - c^2}$. Hence, the m -term recurrence of Faber polynomials (7) becomes a three-term recurrence, similar to that of Chebyshev polynomials:

$$\begin{aligned} F_0(z) &= 1, & F_1(z) &= \frac{z}{\gamma} - c_0, \\ F_2(z) &= F_1(z) \cdot F_1(z) - 2c_1, \\ F_j(z) &= F_1(z) \cdot F_{j-1}(z) - c_1 F_{j-2}(z), \quad j \geq 3. \end{aligned} \quad (12)$$

Indeed, in the particular case when $c \neq 0$ and $a \neq b$ (wherefrom $c^2/(4\gamma^2) \neq 0$), it can be proved that the following relation holds

$$F_j(z) = 2 \left(\frac{c}{2\gamma} \right)^j T_j \left(\frac{z-d}{c} \right), \quad j > 0, \quad (13)$$

where T_j is the j th Chebyshev polynomial (cf. [28]). In the sequel, Faber polynomials and Faber coefficients on an ellipse will be termed (scaled and translated) Chebyshev polynomials and Chebyshev coefficients, respectively.

Notice that we restrict our attention to ellipses symmetric with respect to the real axis since in our application to advection-diffusion equation we deal with real matrices, whose spectrum has the same property.

3. Error bounds for the Faber series approximation to the matrix exponential

In this section we give, on the basis of [19, 21], the convergence estimates for the Faber (Chebyshev) series approximation to the exponential operator. Let A be diagonalizable with diagonalization matrix P , K a compact set with capacity γ . If $\sigma(A) \subseteq K$, then

$$\|\exp(A)\mathbf{v} - S_{m-1}(A)\mathbf{v}\|_2 \leq k_2(P) \|\exp(\cdot) - S_{m-1}(\cdot)\|_K \|\mathbf{v}\|_2,$$

where $k_2(P) := \|P\|_2 \cdot \|P^{-1}\|_2$. If $\sigma(A) \not\subseteq K$, there exists $R > \gamma$ so that $\sigma(A) \subseteq K_R$ (see the previous section for the definition of K_R) and so

$$\|\exp(A)\mathbf{v} - S_{m-1}(A)\mathbf{v}\|_2 \leq k_2(P) \|\exp(\cdot) - S_{m-1}(\cdot)\|_{K_R} \|\mathbf{v}\|_2. \quad (14)$$

When K is convex, it can be shown (cf. [2, 21]) that, for $R \geq \gamma$,

$$\|\exp(\cdot) - S_{m-1}(\cdot)\|_{K_R} \leq \begin{cases} \frac{8R}{m} \exp \left(\psi \left(\frac{4R^2}{\gamma(4R-m)} \right) - \frac{m^2}{4R} \right) & \text{if } m \leq 2R \\ 4 \exp \left(\psi_0 \left(\frac{m}{\gamma} \right) \right) \left(\frac{e \cdot R}{m} \right) & \text{if } m > 2R, \end{cases} \quad (15)$$

where

$$\psi_0(w) := \psi(w) - \gamma w, \quad |w| \geq 1.$$

Clearly, the estimate is valid for $R = \gamma$ when $\sigma(A) \subseteq K = K_\gamma$. Since $\psi_0 \left(\frac{m}{\gamma} \right) = d + O \left(\frac{1}{m} \right)$ for $m \geq 2\gamma$, Faber series convergence for the (matrix) exponential is *superlinear*.

Remark 3.1. *It can be proved that the polynomial S_m does not depend on R , while it is manifest that the error bound (15) does. This means that the convergence rate is invariant in the family $\{K_R\}_R$. Therefore, a tight bound for this rate is obtained by minimizing the estimate (15) on all R such that $\sigma(A) \subseteq K_R$.*

Let us now restrict to the case when the compact set K is $K \equiv \mathcal{E}(d, c, a)$ (cf. Remark 2.1). If $\sigma(A) \subset K = \mathcal{E}(d, c, a)$, it can be easily proved that the first of (15) for $R = \gamma$ is dominated by an increasing function with respect to γ . However, the first of (15) is never of practical interest since, when dealing with advection-diffusion equations with significant advection (with eigenvalues not all real), asymptotic convergence is attained for $m > 2\gamma$. In the case with “small” advection, convergence can be reached when $m < 2\gamma$ but, in this case, the first of (15) is many orders of magnitude larger than the actual error. Bounds given in [29], valid when the spectrum is contained in a real interval, are tighter and can be used instead.

On the other hand, the second of (15), for $R = \gamma$, is itself an increasing function of $\gamma \leq \frac{m}{2}$. Therefore, among ellipses containing $\sigma(A)$ it is convenient to choose that of *smallest capacity*, which in the sequel will be termed *optimal*. If $\sigma(A) \not\subset K_\gamma$, as just said, convergence still holds, and the tighter error bound correspond to (15) where $R := \inf\{R : \sigma(A) \subset K_R\}$ that is to the ellipse with smallest capacity, among the *confocal* family $\{K_R\}_R$ containing the spectrum.

4. Numerical implementation of the Chebyshev method

The first step consists in approximating the optimal ellipse which contains all the eigenvalues of A . To this aim, only the marginal eigenvalues of A are needed. These can be computed by ARPACK (ARnoldi PACKage †), a collection of FORTRAN77 subroutines based upon an algorithmic variant of the Arnoldi process, called “*Implicitly Restarted Arnoldi Method*” (cf. [15]). This software is designed to compute a few, say k , eigenvalues with user specified features such as those of largest real part or largest magnitude using $n \cdot \mathcal{O}(k) + \mathcal{O}(k^2)$ storage, where n is the dimension of A .

In practice, for advection-diffusion matrices arising from stable centered FD discretizations (grid-Péclet numbers sufficiently smaller than 1, cf. [22]), the convex hull of the spectrum turns out to be very similar to a (possibly degenerating) ellipse; see Figure 1 for an example (here spectra have been computed by the subroutine `dgeev` of LAPACK).

In view of the shape of the spectrum, we approximate its convex hull by the rectangle constructed using the extremal eigenvalues (i. e. with extremal real and imaginary part) and we select, among the ellipses circumscribing this rectangle, that of smallest capacity (smallest sum of semi-axes). As observed in the previous section, the convergence rate of the corresponding Chebyshev series is driven by the confocal ellipse with smallest capacity containing the spectrum. The latter ellipse is not, in general, the optimal one, but it is expected to have a capacity close to the optimal, as confirmed by our numerical experiments.

To evaluate the extremal eigenvalues (smallest and largest real part, largest imaginary part) we used driver `dndrv1` of ARPACK (which does not solve linear systems). Concerning the construction of the ellipse above, denote with x , X and Y the smallest real part, the largest

†available from <ftp://ftp.caam.rice.edu/pub/software/ARPACK>

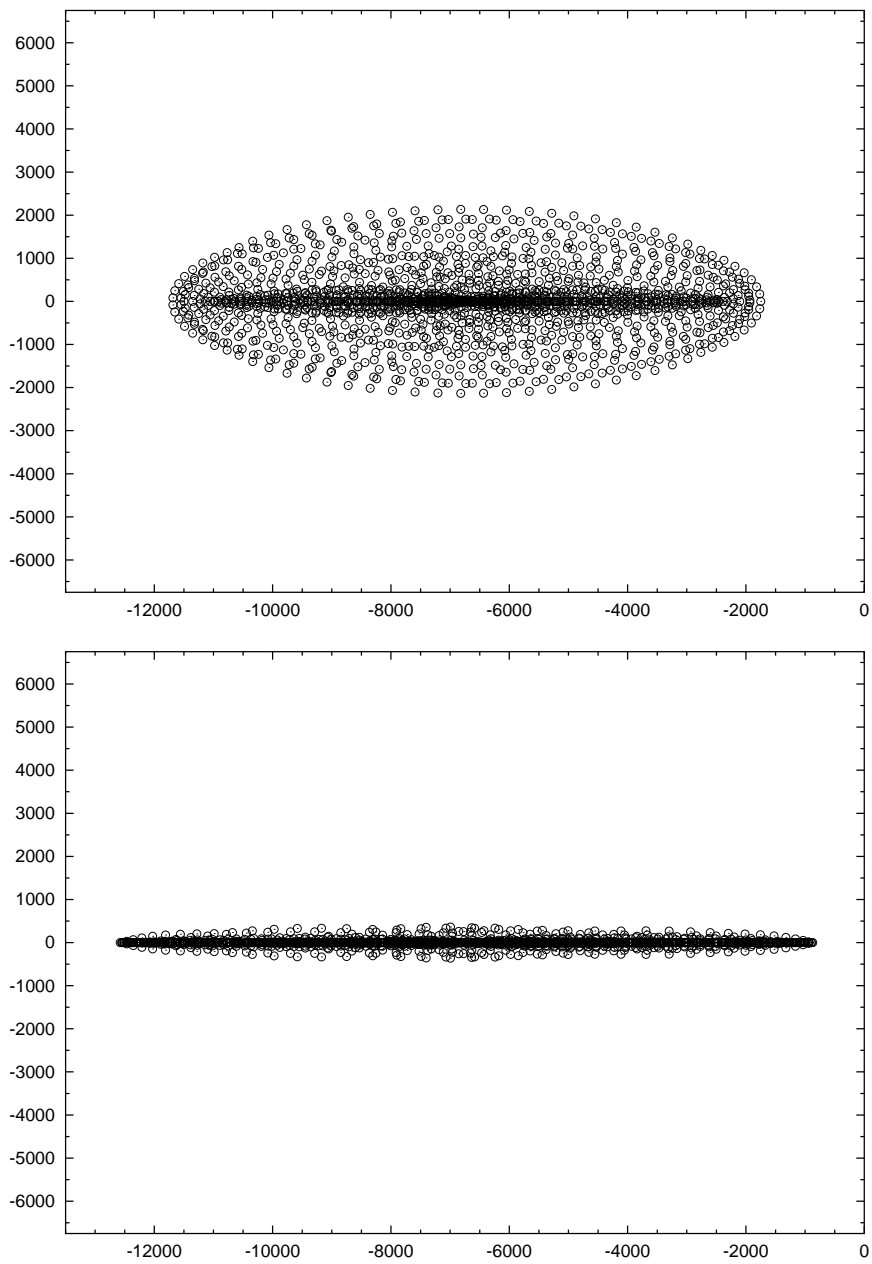


Figure 1. Eigenvalues distribution of advection-diffusion matrices with $\theta_1 = 60$, $\theta_2 = 60$ and $n = 1600$ (top) and with $\theta_1 = 40$, $\theta_2 = 40$ and $n = 1600$ (bottom).

real part and the largest imaginary part of $\sigma(A)$, respectively. A rectangle \mathbf{R} of vertices $x + iY$, $X + iY$ and the symmetric to real axis is defined: it is easy to verify that the ellipse $\mathcal{E}_{\mathbf{R}}^+$ with center $d = (X + x)/2$ and

$$c^2 = (C^{\frac{2}{3}} + Y^{\frac{2}{3}})(C^{\frac{4}{3}} - Y^{\frac{4}{3}}), \quad \gamma = \frac{C^{\frac{2}{3}}\sqrt{C^{\frac{2}{3}} + Y^{\frac{2}{3}}} + \sqrt{(CY^2)^{\frac{2}{3}} + Y^2}}{2}, \quad \left(C := \frac{X - x}{2}\right) \quad (16)$$

is that circumscribed to \mathbf{R} with smallest capacity.

Once the ellipse has been determined, Chebyshev coefficients can be evaluated. To this aim, we used a numerical integration scheme based on trapezoidal rule and the fast Fourier transform (*FFT*) algorithm (cf. [10]). If K is an ellipse, we can choose $R = \gamma$ in (9), whereas if K is a segment of capacity γ we can choose an ellipse of capacity $\gamma' = 1.1 \times \gamma$ and $R = \gamma'$ (this is possible because an ellipse is a Jordan curve, cf. [7]) and applying the transformation of variables $e^{i2\pi\theta} = w$, we have

$$a_j(\text{exp}) = \int_0^1 \exp\left(\left(\gamma + \frac{c^2}{4\gamma}\right) \cos 2\pi\theta + d + i\left(\gamma - \frac{c^2}{4\gamma}\right) \sin 2\pi\theta\right) e^{-ij2\pi\theta} d\theta;$$

using a $(M + 1)$ -point trapezoidal discretization ($M = 2^N$, $N > 1$) we have, for $0 \leq j \leq M - 1$,

$$a_j^{(M)}(\text{exp}) = \frac{1}{M} \sum_{k=0}^{M-1} \exp\left[\left(\gamma + \frac{c^2}{4\gamma}\right) \cos 2\pi \frac{k}{M} + d + i\left(\gamma - \frac{c^2}{4\gamma}\right) \sin 2\pi \frac{k}{M}\right] e^{-ij2\pi \frac{k}{M}} \quad (17)$$

and so we can use *FFT*. From error bounds (15), since 2γ iterations are necessary to reach superlinear convergence, we calculate a priori a safety number of Chebyshev coefficients, say 2^{N^*} , where N^* is the smallest integer such that $2^{N^*} > \lfloor 4\gamma \rfloor$.

At this point, using (10–12) we could compute in principle the Chebyshev series approximating $\exp(A)\mathbf{v}$ stopping the iteration at a suitable index m , since theoretical convergence is guaranteed by estimate (15). In practice, however, convergence may not take place when the capacity of the smallest ellipse confocal to that numerically computed, and containing the spectrum, is large. Indeed, as already observed in [24], in this case one is forced to use Chebyshev coefficients affected by serious cancellation errors. There are two main strategies to overcome this problem, both adopted in our implementation: first, computing the Chebyshev coefficients in higher precision and, when this does not suffice, reducing the step-size (fractionizing the capacity), i. e. computing $\exp(A)\mathbf{v}$ as

$$\exp(A)\mathbf{v} = \exp(\eta_1 A)(\exp(\eta_2 A) \dots (\exp(\eta_p A) \dots))\mathbf{v}, \quad \sum_{i=1}^p \eta_i = 1. \quad (18)$$

We now give and make some comments on our implementation of Chebyshev series approximation of $\exp(A)\mathbf{v}$.

Algorithm 4.1

1. INPUT: $A, v, tol, maxitr, x, X, Y$
2. Compute c^2, γ by (16)
3. IF $Y/C < 10^{-3}$ OR $C/Y < 10^{-3}$ THEN $\gamma := \gamma \times 1.1$
4. $E := \frac{1}{\gamma}A - \frac{d}{\gamma}I$
5. IF $(4\gamma \geq maxitr)$ THEN
6. $nstep := \left\lceil \frac{4\gamma}{maxitr} \right\rceil$
7. $\gamma := \frac{\gamma}{nstep}, \quad c := \frac{c}{nstep}$
8. ELSE
9. $nstep := 1$
10. END IF
11. $istep := 0, \quad \beta := \|v\|_2, \quad conv := \text{FALSE}$
12. DO WHILE (NOT $conv$)
13. $m_{max} := \lceil 4\gamma \rceil, \quad M := \max\{128, 2^{\lceil \log_2(4\gamma) \rceil}\}$
14. Compute the first M Chebyshev coefficients by (17)
15. $conv := \text{TRUE}$
16. DO WHILE ($(istep < nstep)$ AND ($conv$))
17. $f_0 := v, \quad s_0 := a_0 f_0$
18. $f_1 := Ev, \quad s_1 := s_0 + a_1 f_1,$
19. $r_1 := s_1 - s_0, \quad f_0 := 2f_0, \quad k := 1$
20. DO WHILE ($(\|r_k\|_2 > tol \cdot \beta)$ OR $(k < \gamma)$) AND ($conv$))
21. $k := k + 1$
22. $f_k := Ef_{k-1} - \frac{c^2}{4\gamma^2} f_{k-2}, \quad s_k := s_{k-1} + a_k f_k,$
23. $r_k := s_k - s_{k-1}$
24. IF ($(k > 2\gamma$ AND $\|r_k\|_2 > \|r_{k-1}\|_2$) OR $(k = m_{max})$) $conv := \text{FALSE}$
25. END DO
26. IF ($conv$) THEN
27. $v := s_k, \quad istep := istep + 1$
28. ELSE
29. $nstep := 2 \times (nstep - istep), \quad istep := 0$
30. $\gamma := \frac{\gamma}{2}, \quad c := \frac{c}{2}$
31. END IF
32. END DO
33. END DO
34. OUTPUT: $s_k \approx e^A v$

4.1. Comments on Algorithm 4.1

1. We assume that a suitable approximation of the extremal eigenvalues of A , say x , X , Y , is known. This will be useful in the application to advection-diffusion equations where $A = \tau B$ for different τ -values and the spectrum of A can be estimated by that of B , which is computed only once. The parameter *maxitr* is the maximum degree of the truncated Chebyshev series, and *tol* is the exit tolerance relative to $\|\mathbf{v}\|_2$.
3. Manages the case in which the ellipse degenerates into a segment.
4. Also this matrix can be computed, as the extremal eigenvalues, only once, being independent of any scaling of A , and it is displayed for sake of clarity.
- 5.–10. Fractionizes the capacity if it is too large, to guarantee convergence using less than *maxitr* Chebyshev iterations; *nstep* is the number p of steps in (18).
11. *conv* = FALSE when convergence is not reached, and γ needs to be furtherly fractionized.
- 12.–33. Main loop: manages possible fractionizing.
13. $m_{max} = 4\gamma$ is the maximum number of allowed Chebyshev iterations taking into account that after 2γ iterations superlinearity should begin (cf. (15)) as confirmed by our numerical results. A minimum value of M must be given (heuristically set to 128) in order to guarantee accuracy of the quadrature rule.
14. Chebyshev coefficients are computed using quadruple precision arithmetics.
- 16.–32. First level loop: implements the factorization (18).
- 17.–19. Initialization of the Chebyshev recurrence and series.
- 20.–25. Second level loop: computes a single exponential operator on the matrix with the (possibly) fractionized capacity. Stopping criteria: relative “residual” (last computed term of the Chebyshev expansion) below the tolerance, provided that a sufficient number of iterations have been done to avoid erroneous convergence.
24. If the residual norm is not decreasing after 2γ iterations (expected starting of superlinearity, cf. (15) with $R = \gamma$), or the maximum number of iteration m_{max} has been reached, then convergence has not occurred and we need to fractionize γ .
- 26.–31. If *conv* = TRUE, proceeds with the stepwise computation of the exponential with unchanged γ . Otherwise, γ is halved and the number of remaining steps is consequently doubled.

5. Numerical results

In this Section we consider the spatial discretization of the linear advection-diffusion equation (1) by five point Finite Differences with constant step-size h on $\Omega = (0, 1) \times (0, 1)$, which yields the large system of ODEs (2), where

$$B = \frac{1}{h^2} \text{tridiag}(D_+, G, D_-) \in \mathcal{M}_{n \times n}(\mathbb{R}), \quad h = \frac{1}{\nu + 1}, \quad \nu^2 = n \quad (19)$$

and

$$D_+ = \left(1 + \frac{\theta_2 h}{2}\right) I_\nu, \quad G = \text{tridiag}\left(1 + \frac{\theta_1 h}{2}, -4, 1 - \frac{\theta_1 h}{2}\right), \quad D_- = \left(1 - \frac{\theta_2 h}{2}\right) I_\nu.$$

Below, we compare the Chebyshev method as implemented in the previous Section with the well-known Krylov method, as implemented by Saad in the EXPPRO routine [23, 9], in

the evaluation of $\exp(\Delta t B)\mathbf{v}$ (where $\mathbf{v} := (1, \dots, 1)^T$) for different values of the timestep Δt , at a given tolerance $tol = 10^{-8}$ (relative to $\|\mathbf{v}\|_2$). We fixed the spatial discretization corresponding to $h = 1/101$ (leading to $n = 10000$) and select three values of the pair (θ_1, θ_2) , namely $(100, 100)$, $(0, 50)$, $(50, 50)$, which lead to discretization matrices with different spectral features. In all cases the convex hull of the spectrum of B resembles an ellipse (cf. Figure 1) with $C/Y \approx 6.4, 137, 1310$ respectively.

For each test case we report the results of the two methods for five values of Δt . The largest Δt has been selected, for each test case, in order to yield a nearly steady-state solution (the norm of the solution is approximately four orders of magnitude smaller than the norm of \mathbf{v}). The performance of both methods is evaluated in terms of ‘equivalent’ scalar products (e.s.p.), that is operations costing $2n$ flops (a *mat-vect* product counts as l e.s.p. with l the average number of nonzeros per row — in our case $l = 5$) and in terms of total CPU seconds on a 600Mhz Alpha station. We observe that the complexity measure in terms of e.s.p. takes into account only the bulk of both methods, in the sense that the Chebyshev “preprocessing” stage (estimate of the spectrum and computation of the series coefficients) and the Krylov “postprocessing” stage (evaluation of the Hessenberg matrix exponential) are not considered.

As for the Krylov implementation, the maximum dimension m of the Krylov subspace is an input parameter. The input timestep Δt is automatically subdivided, in view of (4), into a number p of smallest time intervals τ_i , such that $\exp(\tau_i B)$ times a vector can be evaluated within a number of iterations less or equal than m . We employed values of m ranging from 5 to 50 since the optimal value of m is not known *a priori* and it is problem dependent. In the tables we also report the number p of substeps in which each step is subdivided, the number $m \times p$ of *mat-vect* products (total number of iterations).

Concerning the Chebyshev method, we used Algorithm 4.1 with $A = \Delta t B$ and *maxitr* = 500. The tables report the results obtained with three different approximations of the spectrum corresponding to

$\text{ncv} :=$ number of ARPACK iterations

equal to 3, 5 or 10 for the approximation of each extremal eigenvalue, and with the ‘true’ values of the extremal eigenvalues (that is, computed by ARPACK with a relative tolerance of 10^{-6}), in order to estimate the impact of possible inaccuracy at this stage. Note that, once the extremal eigenvalues of B (say x_B, X_B, Y_B) have been approximated, we can use Algorithm 4.1 for every choice of Δt after setting: $\mathbf{x} = \Delta t x_B, \mathbf{X} = \Delta t X_B$ and $\mathbf{Y} = \Delta t Y_B$. In the tables we give the number p of substeps, and the number of total *mat-vect* products (number of iterations) as a sum of those necessary to assess the value of γ which guarantees convergence, plus those needed for the actual computation of the right hand side of (4). The total CPU time in seconds and the partial times for eigenvalue approximation, Faber coefficients evaluation, series evaluation, respectively, are also reported.

5.1. Comments on the numerical results

The experimental results, collected in the Tables I–VIII, show that the Chebyshev method with a rough approximation of extremal eigenvalues ($\text{ncv} = 10$) performs better than the Krylov one with the optimal choice m of the dimension of the subspace. In Table I we show the evolution of the solution norm for different timesteps Δt up to a nearly steady-state for the three test problems.

Tables II, III and IV refer to the test case with $\theta_1 = \theta_2 = 100$. From Table II we note

$\theta_1 = \theta_2 = 100$		$\theta_1 = 0, \theta_2 = 50$		$\theta_1 = \theta_2 = 50$	
Δt	$\ \cdot \ _2$	Δt	$\ \cdot \ _2$	Δt	$\ \cdot \ _2$
5.0e-4	92.002	5.0e-4	93.514	5.0e-4	93.280
1.0e-3	86.106	1.0e-3	90.106	1.0e-3	89.441
5.0e-3	42.468	5.0e-3	71.505	5.0e-3	65.39
1.0e-2	1.128	1.0e-2	50.903	1.0e-2	37.686
1.2e-2	0.018	3.4e-2	0.019	2.6e-2	0.016

Table I. Solution norms for different timesteps

Krylov-EXPPRO ($\theta_1 = \theta_2 = 100$), $tol = 10^{-8}$					
Δt	m	p	$mat-vect$	e.s.p.	CPU s.
5.0e-4	5	70	350	4445	1.41
	10	9	90	1584	0.41
	20	3	60	1653	0.41
	50	1	50	2876	0.81
1.0e-3	5	95	475	6033	1.98
	10	14	140	2464	0.66
	20	4	80	2204	0.56
	50	2	100	5752	1.71
5.0e-3	5	233	1165	14796	4.85
	10	43	430	7568	2.37
	20	15	300	8265	2.33
	50	7	350	20132	6.06
1.0e-2	5	360	1800	22860	7.75
	10	69	690	12144	3.94
	20	25	500	13775	3.89
	50	10	500	28760	8.64
1.2e-2	5	401	2005	25509	8.74
	10	77	770	13552	4.38
	20	28	560	15428	4.41
	50	11	550	31636	9.55

Table II. Results of the Krylov method for the test case with $\theta_1 = \theta_2 = 100$

that the Krylov method is very sensitive to the choice of m , which governs the fractionizing of the timestep and hence the number of *substeps* p . In particular, when m is too small, the complexity is affected by the extremely large number p of substeps, while for large m the algorithm is penalized by the quadratic complexity (in m) of the single steps. In practice the ‘optimal’ value of m ranges between 10 and 20 for this problem. We notice that large values of m are in general not recommended due to memory occupancy.

In Table III we present the results of the Chebyshev-Algorithm 4.1 for the various timesteps. The timestep is subdivided into p substeps if:

- the approximation of the eigenvalues is good but the capacity of the ellipse is too large to

Chebyshev-Algorithm 4.1 ($\theta_1 = \theta_2 = 100$), $tol = 10^{-8}$, $maxitr = 500$								
Δt	ncv	p	mat-vect	e.s.p.	partial CPU s.			CPU s.
					eig.	coeff.	series	
5.0e-4	3	1	45	313	0.09	0.01	0.14	0.24
	5	1	45	313	0.15	0.01	0.14	0.30
	10	1	40	278	0.27	0.01	0.12	0.40
	'true'	1	40	278		0.01	0.14	0.15
1.0e-3	3	2	120	833	0.09	0.02	0.35	0.46
	5	1	65	453	0.15	0.01	0.20	0.36
	10	1	60	418	0.27	0.01	0.18	0.46
	'true'	1	60	418		0.01	0.18	0.19
5.0e-3	3	16	875	6075	0.09	0.15	2.65	2.80
	5	2	430	3003	0.15	0.02	1.28	1.45
	10	1	180	1258	0.27	0.03	0.55	0.85
	'true'	2	215	1500		0.05	0.64	0.69
1.0e-2	3	32	1420	9850	0.09	0.25	4.14	4.48
	5	4	650	4538	0.15	0.10	1.95	2.20
	10	2	300	2095	0.27	0.05	0.89	1.21
	'true'	3	385	2688		0.05	1.15	1.20
1.2e-2	3	32	1570	10900	0.09	0.25	4.59	4.93
	5	16	1255	8737	0.15	0.20	3.71	4.06
	10	3	355	2478	0.27	0.05	1.04	1.36
	'true'	4	460	3210		0.05	1.38	1.43

Table III. Results of the Chebyshev method for the test case with $\theta_1 = \theta_2 = 100$

Chebyshev-Algorithm 4.1 ($\theta_1 = \theta_2 = 100$), $tol = 10^{-8}$, $maxitr = 1000$								
Δt	ncv	p	mat-vect	e.s.p.	partial CPU s.		CPU s.	
1.2e-2	10	2	335	2340	0.27	0.05	0.99	1.31
	'true'	2	450	3145		0.05	1.31	1.36

Table IV. As in the previous Table with $maxitr = 1000$.

guarantee convergence within the allowed maximum number of iterations (the capacity and hence the timestep is fractionized a priori, cf. steps 5–10 of Algorithm 4.1). This occurs, for example, with $\Delta t = 0.01, 0.012$ and $ncv = 10$ where $p = 2$ and $p = 3$, respectively.

- The extremely rough approximation of the spectrum leads to non-decreasing error estimate after 2γ iterations, so that the capacity is halved (step 24 of the Algorithm 4.1). This occurs, for instance, with $\Delta t = 0.01$ and $ncv = 3$, where the timestep is halved 5 times leading to $p = 32$.

The cost of the extremal eigenvalue computation is displayed as the first of the partial CPU seconds. For the two smallest timesteps a very rough approximation of the spectrum (with 3

Krylov-EXPPRO ($\theta_1 = 0, \theta_2 = 50$), $tol = 10^{-8}$					
Δt	m	p	$mat-vect$	e.s.p.	CPU s.
5.0e-4	5	52	260	3302	0.98
	10	7	70	1232	0.33
	20	2	40	1102	0.27
	50	1	50	2876	0.80
1.0e-3	5	73	365	4636	1.39
	10	11	110	1936	0.53
	20	4	80	2204	0.58
	50	1	50	2876	0.79
5.0e-3	5	176	880	11176	3.36
	10	31	310	5456	1.59
	20	11	220	6061	1.58
	50	4	200	11504	3.28
1.0e-2	5	300	1500	19050	5.69
	10	52	520	9152	2.70
	20	17	340	9367	2.49
	50	5	250	14380	4.11
3.4e-2	5	578	2890	36703	13.32
	10	122	1220	21472	6.71
	20	37	740	20387	5.79
	50	11	550	31636	9.34

Table V. Results of the Krylov method for the test case with $\theta_1 = 0, \theta_2 = 50$

or 5 ARPACK iterations, respectively) gives convergence comparable to that obtained with the ‘true’ eigenvalues. For larger timesteps, 10 ARPACK iterations produce a satisfactory approximation of the spectrum together with a computational cost which is at most the 30% of the whole algorithm. In any case, Chebyshev with $ncv = 10$ for the extremal eigenvalues is more efficient than the Krylov method for all the timesteps. Comparison of Chebyshev with ‘optimal’ Krylov shows the following *speedups* (ratios between total CPU times): 1.0, 1.2, 2.7, 3.2 and 3.2, from the smallest to the largest Δt . It is worth noticing the apparently strange feature that the number of *mat-vect* products using the ‘true’ extremal eigenvalues is larger than using an approximation of the spectrum with $ncv = 10$. There are two reasons for this behavior: one is that the ellipse \mathcal{E}_1 , constructed from the ‘true’ extremal eigenvalues (solid line in Figure 2), has a much larger capacity than the one (\mathcal{E}_2) constructed using the approximated spectrum (dashed line), thus producing a priori a larger number of substeps for $\Delta t \geq 0.005$ (cf. Table III). The second reason is that the optimal ellipse (dot-dashed in Figure 2) among those confocal to \mathcal{E}_1 , has a larger capacity than the optimal ellipse (dotted line) among those confocal to \mathcal{E}_2 , and then convergence is slower in the first case. This behavior is manifest in Table IV where *maxitr* has been increased (= 1000) in order to force the method using the exact eigenvalues and that with $ncv=10$ to fractionize a priori (cf. step 5 of Algorithm 4.1) with the same number of substeps $p = 2$: indeed, the number of iterations per step is smaller using the approximated eigenvalues.

The computational results for the second test case are summarized in Tables V and VI.

Chebyshev-Algorithm 4.1 ($\theta_1 = 0, \theta_2 = 50$), $tol = 10^{-8}$								
Δt	ncv	p	mat-vect	e.s.p.	partial CPU s.			CPU s.
					eig.	coeff.	series	
5.0e-4	3	1	45	313	0.09	0.01	0.14	0.24
	5	1	40	278	0.15	0.01	0.11	0.27
	10	1	35	243	0.27	0.01	0.11	0.39
	'true'	1	30	208		0.01	0.10	
1.0e-3	3	2	160	1113	0.09	0.02	0.52	0.63
	5	1	60	418	0.15	0.01	0.16	0.32
	10	1	50	348	0.27	0.01	0.14	0.42
	'true'	1	45	313		0.01	0.12	
5.0e-3	3	16	985	6845	0.09	0.15	2.75	2.99
	5	1	180	1258	0.15	0.03	0.49	0.67
	10	1	140	978	0.27	0.03	0.46	0.76
	'true'	1	120	838		0.03	0.33	
1.0e-2	3	32	1545	10725	0.09	0.25	4.43	4.77
	5	2	355	2480	0.15	0.03	1.00	1.18
	10	2	275	1920	0.27	0.03	0.81	1.11
	'true'	2	230	1605		0.03	0.64	
3.4e-2	3	40	2755	19178	0.09	0.25	7.90	8.24
	5	6	1090	7615	0.15	0.05	3.61	3.81
	10	6	825	5760	0.27	0.05	2.74	3.06
	'true'	6	760	5305		0.05	2.25	

Table VI. Results of the Chebyshev method for the test case with $\theta_1 = 0, \theta_2 = 50$

From Table V we see that, as in the previous problem, the optimal m -value for the Krylov method lies in the interval $[10, 20]$. Comparing Tables V and VI, we see again that Chebyshev with $ncv = 10$ is better than 'optimal' Krylov, with the exception of the smallest Δt . The *speedups* are: 0.7, 1.3, 2.1, 2.2 and 1.9, respectively.

The results for the third test case are reported in Tables VII and VIII. Also here Chebyshev with $ncv = 10$ is faster than optimal Krylov, except for $\Delta t = 5 \cdot 10^{-4}$. The *speedups* are 0.7, 1.2, 2.3, 2.9 and 2.9, respectively.

As observed before, the cost of extremal eigenvalue approximation for the Chebyshev method can be a large percentage of the computation for small timesteps. However, if we have to compute $\exp(\Delta t B)\mathbf{v}$ several times, for different \mathbf{v} , and this is the case when an accurate timestepping is performed, the eigenvalue computation can be carried on only once and for all and becomes completely negligible.

REFERENCES

1. L. BERGAMASCHI AND M. VIANELLO, *Efficient computation of the exponential operator for large, sparse, symmetric matrices*, Numerical Linear Algebra with Applications, 7 (2000), pp. 27–45.
2. M. CALIARI, *Computation of the exponential operator for large, sparse, nonsymmetric matrices (in italian)*, Laurea thesis, Università di Padova, (1999).

Krylov-EXPPRO ($\theta_1 = \theta_2 = 50$), $tol = 10^{-8}$					
Δt	m	p	$mat\text{-}vect$	e.s.p.	CPU s.
5.0e-4	5	55	275	3493	1.18
	10	7	70	1232	0.33
	20	2	40	1102	0.27
	50	1	50	2876	0.83
1.0e-3	5	79	395	5017	1.73
	10	11	110	1936	0.53
	20	4	80	2204	0.60
	50	1	50	2876	0.82
5.0e-3	5	181	905	11494	4.60
	10	33	330	4488	2.08
	20	11	220	6061	1.82
	50	4	200	11504	3.53
1.0e-2	5	312	1560	19812	8.22
	10	56	560	9856	3.59
	20	19	380	10469	3.39
	50	6	300	17256	5.38
2.6e-2	5	546	2730	34671	14.39
	10	108	1080	19008	7.69
	20	36	720	19836	6.47
	50	11	550	31636	10.19

Table VII. Results of the Krylov method for the test case with $\theta_1 = \theta_2 = 50$

3. J. H. CURTISS, *Faber polynomials and Faber series*, The American Mathematical Monthly, 78 (1971), pp. 577–596.
4. V. L. DRUSKIN AND L. A. KNIZHERMAN, *Two polynomial methods for calculating functions of symmetric matrices*, Computational Mathematics and Mathematical Physics, 6 (1989), pp. 112–121.
5. W. S. EDWARDS, L. S. TUCKERMAN, R. A. FRIESNER, AND D. C. SORENSEN, *Krylov methods for the incompressible navier-stokes equations*, Journal of Computational Physics, 110 (1994), pp. 82–102.
6. M. EIERMANN, *On semiiterative methods generated by Faber polynomials*, Numerische Mathematik, 56 (1989), pp. 139–156.
7. S. W. ELLACOTT, *Computation of Faber series with application to numerical polynomial approximation in the complex plane*, Mathematics of Computation, 40 (1983), pp. 575–587.
8. R. A. FRIESNER, L. S. TUCKERMAN, B. C. DORNBLASER, AND T. V. RUSSO, *A method for exponential propagation of large system of stiff nonlinear differential equations*, Journal of Scientific Computing, 4 (1989), pp. 327–354.
9. E. GALLOPOULOS AND Y. SAAD, *Efficient solution of parabolic equations by Krylov subspace methods*, SIAM Journal on Scientific and Statistical Computing, 13 (1992), pp. 1236–1264.
10. K. O. GEDDES, *Near-minimax polynomial approximation in an elliptical region*, SIAM Journal on Numerical Analysis, 15 (1978), pp. 1225–1233.
11. E. HAIRER, G. BADER, AND C. LUBICH, *On the stability of semi-implicit methods for ODEs*, BIT. Numerical Mathematics, 22 (1982), pp. 211–232.
12. M. HOCHBRUCK AND J. J. LUBICH, *On Krylov subspace approximations to the matrix exponential*, SIAM Journal on Numerical Analysis, 34 (1997), pp. 1911–1925.
13. M. HOCHBRUCK, C. LUBICH, AND H. SELHOFER, *Exponential integrators for large system of differential equations*, SIAM Journal on Scientific Computing, 19 (1998), pp. 1552–1574.
14. L. A. KNIZHERMAN, *Computation of functions of unsymmetric matrices by means of arnoldi's method*, Computational Mathematics and Mathematical Physics, 31 (1991), pp. 5–16.
15. R. B. LEHOUCQ, D. C. SORENSEN, AND C. YANG, *ARPACK USERS GUIDE: Solution of Large Scale Eigenvalue Problems by Implicitly Restarted Arnoldi Methods*, 8 Oct. 1997.

Chebyshev-Algorithm 4.1 ($\theta_1 = \theta_2 = 50$), $tol = 10^{-8}$								
Δt	ncv	p	mat-vect	e.s.p.	partial CPU s.			CPU s.
					eig.	coeff.	series	
5.0e-4	3	1	45	313	0.09	0.01	0.15	0.25
	5	1	40	278	0.15	0.01	0.13	0.29
	10	1	35	243	0.27	0.01	0.10	0.38
	'true'	1	30	208		0.01	0.08	
1.0e-3	3	2	155	1078	0.09	0.02	0.43	0.54
	5	1	60	418	0.15	0.01	0.20	0.36
	10	1	55	383	0.27	0.01	0.15	0.43
	'true'	1	45	313		0.01	0.15	
5.0e-3	3	16	980	6810	0.09	0.15	3.30	2.54
	5	1	180	1258	0.15	0.03	0.50	0.68
	10	1	150	1048	0.27	0.03	0.50	0.80
	'true'	1	135	943		0.03	0.37	
1.0e-2	3	32	1540	10690	0.09	0.25	4.33	4.67
	5	2	360	2515	0.15	0.03	0.99	1.17
	10	2	295	2060	0.27	0.03	0.85	1.15
	'true'	2	240	1675		0.03	0.67	
2.6e-2	3	32	2250	15663	0.09	0.25	6.25	6.59
	5	5	865	6043	0.15	0.05	2.39	2.59
	10	5	675	4713	0.27	0.05	1.92	2.24
	'true'	5	580	4048		0.05	1.61	

Table VIII. Results of the Chebyshev method for the test case with $\theta_1 = \theta_2 = 50$

16. T. A. MANTEUFFEL, *The tchebychev iteration for nonsymmetric linear systems*, Numerische Mathematik, 28 (1977), pp. 307–327.
17. ———, *Adaptive procedure for estimation of parameter for the nonsymmetric tchebychev iteration*, Numerische Mathematik, 28 (1978), pp. 187–208.
18. A. I. MARKUSHEVIC, *Theory of Functions of a Complex Variable*, Prentice-Hall, Inc., Englewood Cliffs, 1967.
19. I. MORET AND P. NOVATI, *The computation of functions of matrices by truncated faber series*, Numerical Functional Analysis and Optimization, 22 (2001), pp. 697–719.
20. ———, *An interpolatory approximation of the matrix exponential based on Faber polynomials*, Journal of Computational and Applied Mathematics, 131 (2001), pp. 361–380.
21. P. NOVATI, *Polynomial methods for the computation of functions of unsymmetric matrices*, PhD thesis, Trieste, 2000. Dip.to di Scienze Matematiche, Università di Trieste.
22. A. QUARTERONI AND A. VALLI, *Numerical Approximation of Partial Differential Equations*, Springer series in Computational Mathematics 23, Springer-Verlag, 1994.
23. Y. SAAD, *Analysis of some Krylov subspace approximations to the matrix exponential operator*, SIAM Journal on Numerical Analysis, 29 (1992), pp. 209–228.
24. M. J. SCHAEFER, *A polynomial based iterative method for linear parabolic equations*, Journal of Computational and Applied Mathematics, 29 (1990), pp. 35–50.
25. R. B. SIDJE, *EXPOKIT: Software package for computing matrix exponentials*, ACM. Association for Computing Machinery. Transactions on Mathematical Software, 24 (1998), pp. 130–156.
26. G. L. G. SLEIJPEN AND H. A. VAN DER VORST, *A Jacobi-Davidson method for linear eigenvalue problems*, SIAM Journal on Matrix Analysis and Applications, 17 (1996), pp. 401–425.
27. V. I. SMIRNOV AND N. A. LEBEDEV, *Functions of a Complex Variable, Constructive Theory*, Iliffe Books, London, 1968.
28. G. STARKE AND R. S. VARGA, *A hybrid Arnoldi-Faber iterative method for nonsymmetric systems of*

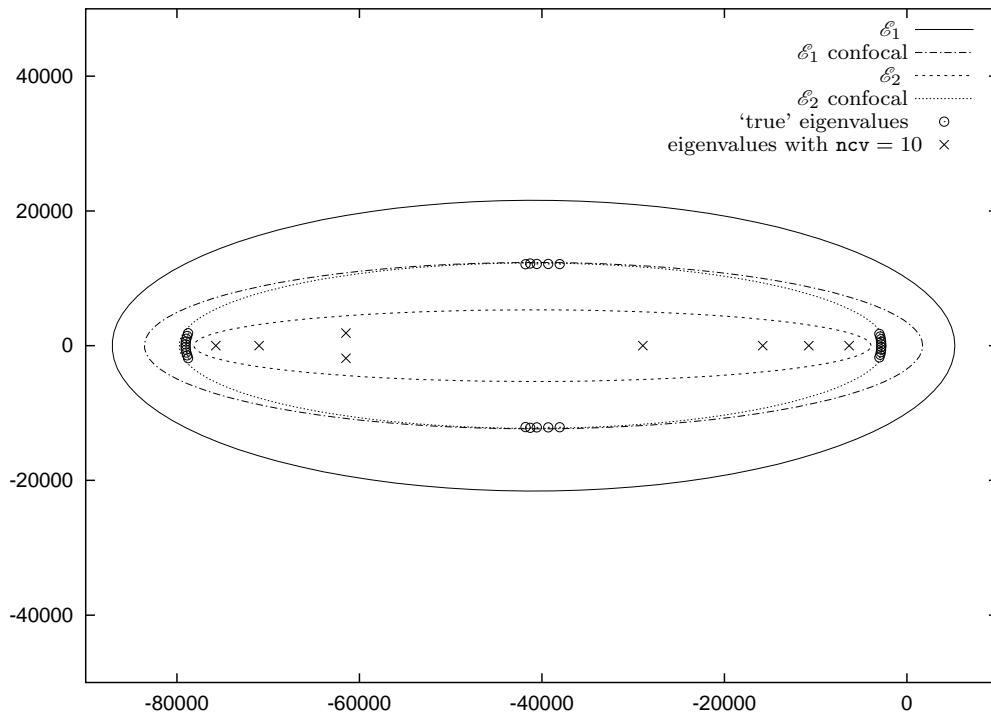


Figure 2. ‘True’ and approximated extremal eigenvalues and optimal ellipses for the test case with $\theta_1 = \theta_2 = 100$ and $n = 10000$

- linear equations*, *Numerische Mathematik*, 64 (1993), pp. 213–240.
29. D. E. STEWART AND T. S. LEYK, *Error estimates for Krylov subspace approximations of matrix exponentials*, *Journal of Computational and Applied Mathematics*, 72 (1996), pp. 359–369.