Electronic Transactions on Numerical Analysis.
Volume 29, pp. 1-18, 2007.

ISSN 1068-9613.

ETNA
Kent State University
etna@mcs.kent.edu

# EVALUATING MATRIX FUNCTIONS FOR EXPONENTIAL INTEGRATORS VIA CARATHÉODORY-FEJÉR APPROXIMATION AND CONTOUR INTEGRALS*

THOMAS SCHMELZER† AND LLOYD N. TREFETHEN‡

**Abstract.** Among the fastest methods for solving stiff PDE are exponential integrators, which require the evaluation of $f(A)$, where $A$ is a negative semidefinite matrix and $f$ is the exponential function or one of the related "$\varphi$ functions" such as $\varphi_1(z) = (e^z - 1)/z$. Building on previous work by Trefethen and Gutknecht, Minchev, and Lu, we propose two methods for the fast evaluation of $f(A)$ that are especially useful when shifted systems $(A + zI)x = b$ can be solved efficiently, e.g. by a sparse direct solver. The first method is based on best rational approximations to $f$ on the negative real axis computed via the Carathéodory-Fejér procedure. Rather than using optimal poles we approximate the functions in a set of common poles, which speeds up typical computations by a factor of 2 to 3.5. The second method is an application of the trapezoid rule on a Talbot-type contour.

**Key words.** matrix exponential, exponential integrators, stiff semilinear parabolic PDEs, rational uniform approximation, Hankel contour, numerical quadrature

**AMS subject classification.** 65L05, 41A20, 30E20

## 1. Introduction.

According to Minchev and Wright [33], the main computational challenge in the implementation of any exponential integrator is the need for fast and computationally stable evaluations of the exponential and related $\varphi$ functions. We are interested in solving problems

$$\dot{u} = Au + g(u,t) \tag{1.1}$$

where the matrix $A$ represents the spatial discretization of a linear elliptic differential operator such as the Laplacian and $g$ is a nonlinear function in $u$ and $t$. In many problems $A$ is negative semidefinite. Exponential integrators are time-stepping formulas for (1.1) that separate the linear term involving $A$, which is solved exactly by a matrix exponential, from the nonlinear term. The simplest example is the exponential forward Euler method, given by

$$u_{n+1} = e^{\Delta t A} u_n + \Delta t\, \varphi_1(\Delta t A) g(u_n, t_n),$$

where $\varphi_1(z) = (e^z - 1)/z$. There are many other exponential schemes and some of these ideas have been reinvented several times [33]. In recent years the interest in exponential integrators has heightened. Krylov methods to compute $\varphi$ functions were introduced by Saad [38] and Hochbruck and Lubich [19], Cox and Matthews [8] and Krogstad [25] introduced one-step methods with 4th order accuracy in many circumstances, and Hochbruck and Ostermann [21] showed how 4th order could be achieved for *all* problems. Exponential multi-step formulas require fewer matrix function evaluations per step than one-step methods and are computationally very promising. Recently they have been rediscovered by Beylkin [4] and Cox and Matthews after being introduced by Nørsett [36] in 1969. Kassam and Trefethen applied the one-step method of Cox and Matthews to stiff PDEs such as the Korteweg–de Vries, Kuramoto–Sivashinsky, Allen–Cahn and Grey–Scott equations [23, 24] and compared them with more standard schemes.

* Received October 31, 2006. Accepted for publication September 13, 2007. Published online on December 13, 2007. Recommended by M. Hochbruck.

† Computing Laboratory, Oxford University, United Kingdom (thoms@comlab.ox.ac.uk). Thomas Schmelzer is supported by a Rhodes Scholarship.

‡ Computing Laboratory, Oxford University, United Kingdom (lnt@comlab.ox.ac.uk).

ETNA
Kent State University
etna@mcs.kent.edu



We follow a recent convention [3, 21, 22] and introduce

$$\varphi_l(z) = \frac{1}{(l-1)!}\int_0^1 e^{(1-\theta)z}\theta^{l-1}d\theta, \qquad l = 1, 2, \ldots.$$

In addition we define $\varphi_0(z) = e^z$, which enables us to utilize the recurrence relation

$$\varphi_l(z) = \frac{\varphi_{l-1}(z) - \varphi_{l-1}(0)}{z}, \qquad l \geqslant 1. \tag{1.2}$$

For the first few values of $l$ we have

$$\varphi_1(z) = \frac{e^z - 1}{z}, \qquad \varphi_2(z) = \frac{e^z - z - 1}{z^2}, \qquad \varphi_3(z) = \frac{e^z - z^2/2 - z - 1}{z^3}.$$

The Taylor series representation of these functions is given by

$$\varphi_l(z) = \sum_{k=l}^{\infty} \frac{1}{k!} z^{k-l}. \tag{1.3}$$

The functions $\varphi_l$ are entire. Nevertheless, a numerical challenge one encounters in utilizing them is that a direct computation based on these identities suffers from cancellation errors for $z$ close to the origin [17, 24]. To address this problem Cox and Matthews [8] made use of the Taylor series. This technique works for scalars and diagonal matrices. An alternative stable evaluation is based on (1.2) and a Cauchy integral representation on a circle $\Gamma$ of radius 1 centered at $z$, for $|z| < 1/2$. Following Kassam and Trefethen [24] this is

$$\varphi_l(z) = \frac{1}{2\pi i}\int_\Gamma \frac{\varphi_l(s)}{s-z}ds \approx \frac{1}{M}\sum_{k=1}^{M} \varphi_l\left(z + e^{it_k}\right)$$

where $t_k = 2\pi k/M$. To avoid cancellation, the circles should not come close to the origin. This idea can be generalized to non-diagonal matrices $A$, where the contour $\Gamma$ has to enclose the spectrum of $A$:

$$\varphi_l(A) = \frac{1}{2\pi i}\int_\Gamma \varphi_l(s)\left(sI - A\right)^{-1} ds. \tag{1.4}$$

In practice we are interested in $f(A)b$ rather than $f(A)$ and herein lies a further strength of this idea. We can evaluate

$$\varphi_l(A)b = \frac{1}{2\pi i}\int_\Gamma \varphi_l(s)\left(sI - A\right)^{-1} b\, ds$$

by a numerical quadrature scheme that solves a linear system at each node $s_k$ on the contour $\Gamma$. In particular this idea is successful for the matrix exponential and it is closely related to rational approximations as shown in [45]. Instead of evaluating (1.4) for $l > 0$ we give in §5 an alternative integral representation of $\varphi_l$ not involving the function $\varphi_l$ in the integrand.

Another method for evaluating the $\varphi$ functions that is widely used nowadays is Padé approximation, as suggested by Beylkin et al. [4], Hochbruck et al. [20], and Minchev and Wright [33]. The idea is based on scaling and squaring, which is a popular method for computing the matrix exponential [18, 34]. This approach is restricted to matrices of moderate dimension as the evaluation of $f(A)b$ requires the explicit computation of $f(A)$.

**ETNA**
Kent State University
etna@mcs.kent.edu



We suggest instead the use of uniform rational Chebyshev approximations or the application of the trapezoid rule on Talbot-type contours. The first idea was put forward previously by Lu [27]. Lu claims that the coefficients are hard to compute and gives the rational approximation of type $(14, 14)$ for $\varphi_1$ in a partial fraction decomposition. He used the Remes algorithm and a multiple precision environment. Here we shall show that in fact, the required approximations can be computed readily in standard precision by the Carathéodory-Fejér method as in [44] and [45].

We also discuss in §4 the approximation of these functions in a common set of poles, which is advantageous in some applications.

Throughout this work we use direct methods to solve linear systems although our ideas are by no means restricted to them.

**2. The asymptotic behavior.** Methods based on rational approximations get their power from the fast exponential decay of the error introduced by the approximant. In the case of $\varphi$ functions we can give precise statements about the convergence. Let $\mathcal{P}_n$ denote the set of all polynomials of degree at most $n \in \mathbb{N}$ with real coefficients. Let $\mathcal{R}_{mn}$ denote the set $\{p/q \,|\, p \in \mathcal{P}_m, q \in \mathcal{P}_n, q \not\equiv 0\}$, $m, n \in \mathbb{N}$ of rational functions. The *best rational approximant* $r^*_{mn} = r^*_{mn}(f, \mathbb{R}^-; \cdot) \in \mathcal{R}_{mn}$ to the function $f$ on $\mathbb{R}^- = (-\infty, 0]$ and the *minimal approximation error* $E_{mn} = E_{mn}(f, \mathbb{R}^-)$ are defined by

$$E_{mn}(f, \mathbb{R}^-) := \|f - r^*_{mn}\|_{\mathbb{R}^-} = \inf_{r \in \mathcal{R}_{mn}} \|f - r\|_{\mathbb{R}^-},$$

where $\|\cdot\|_{\mathbb{R}^-}$ denotes the sup-norm on $\mathbb{R}^-$. The best approximant $r^*_{mn}(\varphi_l, \mathbb{R}^-)$ exists and is unique [30].

Cody, Meinardus and Varga [7] showed that $E_{nn}(\exp(x), \mathbb{R}^-)$ decreases geometrically as $n \to \infty$. In 1986 Gonchar and Rakhmanov [14] proved that the rate of convergence is given by

$$\lim_{n\to\infty} E_{nn}\left(\exp(x), \mathbb{R}^-\right)^{1/n} = \upsilon = \frac{1}{9.28903\ldots}, \tag{2.1}$$

where $\upsilon$, *Halphen's constant*, is the unique positive root of $\sum_{n=1}^{\infty} \frac{n\upsilon^n}{1-(-\upsilon)^n} = \frac{1}{8}$. This constant was studied by Halphen as early as 1886 [16]. The proof by Gonchar and Rakhmanov confirmed a conjecture by Magnus [28] and previous numerical computations by Trefethen and Gutknecht [44] and Carpenter, Ruttan and Varga [6].

A sharper result than (2.1) was conjectured by Magnus [29] and subsequently proved by Aptekarev [2]:

$$E_{nn}\left(\exp(x), \mathbb{R}^-\right) = 2\upsilon^{n+\frac{1}{2}}(1 + o(1)) \qquad \text{as } n \to \infty.$$

We define $E_l = \lim_{n\to\infty} E_{nn}(\varphi_l, \mathbb{R}^-)^{1/n}$, if this limit exists, and conjecture that the limit does indeed exist with $E_l = E_0 = \upsilon$ for all $l$.

CONJECTURE 2.1. *For all $l \in \mathbb{N}$ the asymptotic decay of the error is $E_l = \upsilon$.*

Results about asymptotic convergence for best rational approximations are notoriously difficult to prove. A successful proof might follow the footsteps of Gonchar and Rakhmanov [14], or might be based on induction utilizing recurrence relations for non-normal matrices which we are going to introduce in §4. Numerical experiments give compelling indications that the conjecture is valid.

For practical purposes we are interested in the number of poles necessary to achieve a desired accuracy. The asymptotic convergence rate is of limited use here, although it would serve as a first indicator.

ETNA
Kent State University
etna@mcs.kent.edu



**3. Carathéodory-Fejér approximation on the negative real line.** An efficient method for constructing near-best rational approximations is the Carathéodory-Fejér (CF) method, which was introduced first for the problem of constructing approximations on the unit disc [42, 43]. By utilizing a conformal map from the unit circle to a real interval or the negative real line it was shown in [44] that the method is very efficient for real approximation, too. The idea can also be generalized for other domains in the complex plane [11].

These approximations are so close to optimal that the method can often be regarded as exact in practice. Magnus [29] has argued that the approximations of $\exp(x)$ produced by the CF method differ from the true best approximations by only about $O(56^{-n})$. The absolute difference between the best CF and best approximations is below standard machine precision for $n \geqslant 9$.

These ideas have not been exploited much over the last two decades. Today we can compute CF approximations on the fly in fractions of a second. In [45] a MATLAB code is presented that computes rational approximations of $\exp(x)$ on the negative real line with an error as small as $2 \times 10^{-14}$. We have made some minor modifications to adapt the code for $\varphi$ functions.

CF approximation enables us to estimate the error for all $n$ at once as they appear as singular values of a certain matrix within the construction process. For the applications we have in mind, a rational uniform approximation with an error of $10^{-6}$ is often appropriate. The asymptotic behavior discussed in the last section suggests that $n = 6$ poles may be sufficient to achieve this accuracy. Computations confirm that for $n = 6$, the error we commit by replacing $\varphi_l$ by its CF approximation is indeed smaller than $10^{-6}$ for $l = 0, 1, \ldots$; see Fig. 3.1. For all exponential integrators we have used it is sufficient to compute $\varphi_l$ up to $l = 4$.

We use a partial fraction expansion of the rational approximations,

$$r_n(z) = \frac{p_n(z)}{q_n(z)} = r_\infty + \sum_{j=1}^{n} \frac{c_j}{z - z_j},$$

where $c_j$ is the residue of the pole $z_j$ and $r_\infty = r(\infty)$. As the denominator $q_n(z)$ is a polynomial with real coefficients, the poles come in conjugate pairs.

**4. Approximation in common poles.** When implementing exponential integrators it is an attractive option to use a set of common poles for all $\varphi$ functions. Using this strategy one can evaluate $\varphi_l(hA)v$ for several different $l$ by linear combination of the solutions $x_j$ of a fixed set of systems $(A - z_j I)x_j = v$ with only the coefficients of the linearcombination depending on $l$. This situation is very frequent when dealing with exponential integrators. Typically the use of common poles makes exponential integrators faster by a factor of 2 to 3.5. The precise cost savings depend ultimately on the type of integrator being used and whether it is possible to store $LU$ decompositions, etc. Rather than devising an optimization problem that imposes a set of common poles as a constraint and choosing them so that associated error functions are as small as possible, we generalize an identity by Saad [38].

PROPOSITION 4.1. *Let*

$$B_z = \begin{pmatrix} z & 1 \\ 0 & 0 \end{pmatrix},$$

*where* $z \in \mathbb{C}$. *Then*

$$\varphi_l(B_z) = \begin{pmatrix} \varphi_l(z) & \varphi_{l+1}(z) \\ 0 & \varphi_l(0) \end{pmatrix} \qquad \text{for } l = 0, 1, \ldots.$$

**ETNA**
Kent State University
etna@mcs.kent.edu



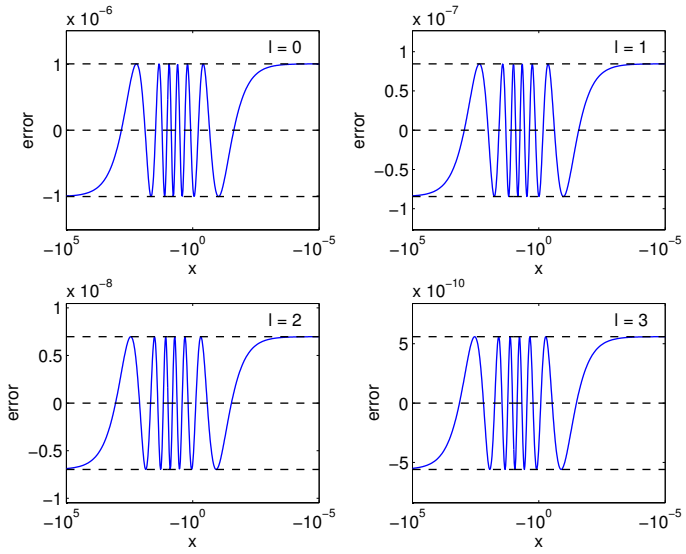


FIG. 3.1. *Error curves $\varphi_l(x) - r^*(x)$ for type $(6,6)$ best approximation to $\varphi_l$, $l = 0, 1, 2, 3$ on $\mathbb{R}^-$. Note that the abscissa, the negative real axis in the $z$-plane, is displayed on a log scale. The dashed lines mark the minimax errors.*

Saad established this identity for $l = 0$. He was not concerned with $\varphi$ functions of higher order. The proof for $l > 0$ carries over.

*Proof.* We observe that $B_z^0 = I$ and

$$B_z^n = \begin{pmatrix} z^n & z^{n-1} \\ 0 & 0 \end{pmatrix} \qquad n \geqslant 1.$$

As $\varphi_l(z) = \sum_{k=l}^{\infty} \frac{1}{k!} z^{k-l}$ we get

$$\varphi_l(B_z) = \sum_{k=l}^{\infty} \frac{1}{k!} B_z^{k-l} = \begin{pmatrix} \sum_{k=l}^{\infty} \frac{1}{k!} z^{k-l} & \sum_{k=l+1}^{\infty} \frac{1}{k!} z^{k-l-1} \\ 0 & \frac{1}{l!} \end{pmatrix} = \begin{pmatrix} \varphi_l(z) & \varphi_{l+1}(z) \\ 0 & \varphi_l(0) \end{pmatrix}. \quad \square$$

If $z$ is replaced by a matrix $A$ this idea has been suggested amongst others by Hochbruck et al. [20] and Saad [38] for computing $\varphi_1(A)$. However, this direct approach has two disadvantages: the resulting matrix $B_A$ has twice the size of $A$, and it does not inherit properties such as symmetry from $A$.

Given a rational approximation

$$r_{(l)}(z) = r_\infty + \sum_{j=1}^{n} \frac{c_j}{z - z_j}$$

of $\varphi_l$, and noting that

$$(B_z - z_j I)^{-1} = \begin{pmatrix} (z - z_j)^{-1} & (z - z_j)^{-1} z_j^{-1} \\ 0 & -z_j^{-1} \end{pmatrix},$$

ETNA
Kent State University
etna@mcs.kent.edu



the entry of $r_{(l)}(B_z)$ approximating $\varphi_{l+1}$ is

$$r_{(l)}(B_z)_{1,2} = \sum_{j=1}^{n} \frac{c_j z_j^{-1}}{z - z_j}.$$

This identity implies the recurrence relation

$$r_{(l+k)}(z) = \sum_{j=1}^{n} \frac{c_j z_j^{-k}}{z - z_j}, \qquad k \in \mathbb{Z} \tag{4.1}$$

for rational approximations using common poles of $\varphi$ functions. This generalizes a result by Minchev [31]. These approximations are far from optimal and yet they provide reasonable quality on the negative real axis. For $k > 0$ we can give error bounds that provide some insight.

The matrices $B_z$ are diagonalizable for $z \neq 0$ with eigenvalues 0 and $z$, hence $B_z = V_z \Lambda_z V_z^{-1}$. And yet it is not appropriate to estimate the error by

$$\|f(B_z) - \varphi_l(B_z)\|_2 \leqslant \kappa_2(V_z) \|f(\Lambda_z) - \varphi_l(\Lambda_z)\|_2 \tag{4.2}$$

where $\kappa_2(V_z) = \|V_z\|_2 \|V_z^{-1}\|_2$ is the 2-norm condition number of $V_z$, as the condition number of $V_z$ tends to infinity for $z$ approaching the origin, see Fig. 4.1. If this were not the case we would be able to give a simple proof based on induction for Conjecture 2.1. A more useful error bound for small $z$ is given by the following result.

THEOREM 4.2. *If $f$ and $g$ are analytic on the negative real axis $\mathbb{R}^-$ then*

$$\|f(B_z) - g(B_z)\|_F \leqslant \sqrt{2} \sup_{s \in \mathbb{R}^-} |f(s) - g(s)| + \sup_{s \in \mathbb{R}^-} |f'(s) - g'(s)|.$$

*Proof.* This is a special case of Theorem 11.2.2 in [13]. □

If $r_{(l+1)}$ is the rational approximation induced by $r_{(l)}$ using the recurrence relation above we can bound the error by

$$\begin{aligned} \sup_{s \in \mathbb{R}^-} |r_{(l+1)}(s) - \varphi_{l+1}(s)| &\leqslant \sup_{s \in \mathbb{R}^-} \|r_{(l)}(B_s) - \varphi_l(B_s)\|_F \\ &\leqslant \sqrt{2} \sup_{s \in \mathbb{R}^-} |r_{(l)}(s) - \varphi_l(s)| + \sup_{s \in \mathbb{R}^-} |r'_{(l)}(s) - \varphi'_l(s)|. \end{aligned} \tag{4.3}$$

Using approximations in common poles the accuracy of best rational approximations can not be achieved when using the same number of poles for both ideas. Nevertheless, this approach is rather robust as the derivatives of the error term in (4.3) are small, too. In the context of exponential integrators it is important to work with approximations providing an accuracy beyond the level of the truncation error of the integrator.

We give in Table 4.1 some numerical results for small degrees of the rational approximations.

The entries in the first column of Table 4.1 can be further improved by introducing a positive shift $s$. As $e^z = e^s e^{z-s}$, we can approximate the second factor by a CF approximation,

$$e^z \approx e^s \left( r_\infty + \sum_{j=1}^{n} \frac{c_j}{z - (s + z_j)} \right). \tag{4.4}$$

Lu [27] introduced this shift to deal with matrices that have negative and in addition small positive eigenvalues. We have observed that a shift of $O(1)$ gives only slightly weaker results for $e^z$ but significantly better rational approximations induced by (4.1) and (4.4) to $\varphi$ functions of higher order. We summarize our results in Table 4.2.

**ETNA**
Kent State University
etna@mcs.kent.edu



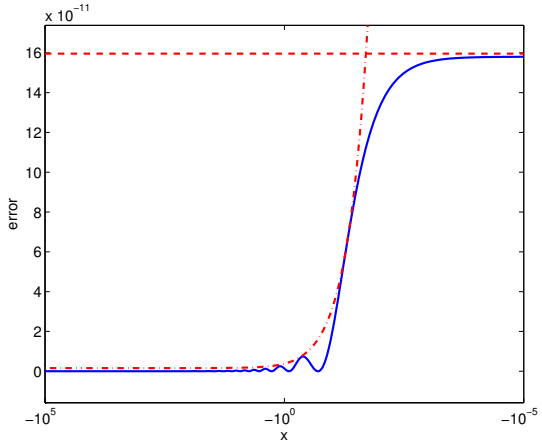


FIG. 4.1. *The poles of a CF approximation of $e^x$ of type $(12, 12)$ have been used to evaluate the $\varphi_1$ function on the negative real line. The blue solid curve is the error. The dashed-dotted red curve is the bound induced by (4.2). The dashed line represents the bound of Theorem 4.2.*

TABLE 4.1
*Using degree $n$ rational approximations with common poles the maximal error $\varphi_l - r$ is given for $l = 0, 1, 2, 3$ and $n = 6, 8, 10, 12$. A column represents CF approximations for fixed $l$. An example: Given the CF approximation of degree $10$ for $\varphi_1$ the error committed by approximating $\varphi_3$ with the same poles using identity (4.1) is $7.3e-09$. To determine an approximation for the maximal error we have computed the difference in $500$ points distributed in $-10^5$ to $-10^{-5}$.*

| | $\varphi_0$ | $\varphi_1$ | $\varphi_2$ | $\varphi_3$ | |
|---|---|---|---|---|---|
| | $\mathbf{1.0e-06}$ | $9.3e-05$ | $2.2-03$ | $3.0-02$ | $\varphi_0$ |
| $n=$ | $5.3e-05$ | $\mathbf{8.5e-08}$ | $9.7e-06$ | $2.7e-04$ | $\varphi_1$ |
| 6 | $4.6e-04$ | $4.0e-06$ | $\mathbf{7.0e-09}$ | $9.5e-07$ | $\varphi_2$ |
| | $1.6e-03$ | $3.1e-05$ | $2.9e-07$ | $\mathbf{5.6e-10}$ | $\varphi_3$ |
| | $\mathbf{1.2e-08}$ | $1.7e-06$ | $6.2e-05$ | $1.2e-03$ | $\varphi_0$ |
| $n=$ | $8.0e-07$ | $\mathbf{7.5e-10}$ | $1.3e-07$ | $5.5e-06$ | $\varphi_1$ |
| 8 | $9.1e-06$ | $4.7e-08$ | $\mathbf{4.8e-11}$ | $9.9e-09$ | $\varphi_2$ |
| | $4.2e-05$ | $4.9e-07$ | $2.8e-09$ | $\mathbf{3.0e-12}$ | $\varphi_3$ |
| | $\mathbf{1.4e-10}$ | $2.9e-08$ | $1.5e-06$ | $3.8e-05$ | $\varphi_0$ |
| $n=$ | $1.1e-08$ | $\mathbf{7.1e-12}$ | $1.8e-09$ | $1.0e-07$ | $\varphi_1$ |
| 10 | $1.6e-07$ | $5.6e-10$ | $\mathbf{3.7e-13}$ | $1.1e-10$ | $\varphi_2$ |
| | $9.1e-07$ | $7.3e-09$ | $2.7e-11$ | $\mathbf{1.9e-14}$ | $\varphi_3$ |
| | $\mathbf{1.6e-12}$ | $4.7e-10$ | $3.1e-08$ | $1.0e-06$ | $\varphi_0$ |
| $n=$ | $1.6e-10$ | $\mathbf{6.8e-14}$ | $2.7e-11$ | $1.7e-09$ | $\varphi_1$ |
| 12 | $2.6e-09$ | $6.5e-12$ | $\mathbf{4.3e-15}$ | $1.2e-12$ | $\varphi_2$ |
| | $1.8e-08$ | $1.0e-10$ | $2.7e-13$ | $\mathbf{5.6e-16}$ | $\varphi_3$ |

**5. Talbot contours and integrals of Cauchy type.** The results of the last section give us a new perspective on contour integrals, too. If $\Gamma$ is a contour enclosing $z$ with winding number 1, then

$$e^z = \frac{1}{2\pi i}\int_\Gamma \frac{e^s}{s-z}ds. \tag{5.1}$$

**ETNA**
Kent State University
etna@mcs.kent.edu



TABLE 4.2
*Same as the first column of Table 4.1, but with a shift $s$ introduced as in (4.4). The numbers are better, and $s = 1$ seems a good choice in practice.*

| | $s = 1/2$ | $s = 1$ | $s = 2$ | $s = 5$ | |
|---|---|---|---|---|---|
| | $1.6e-06$ | $2.7e-06$ | $7.5e-06$ | $1.5e-04$ | $\varphi_0$ |
| $n =$ | $1.0e-05$ | $1.1e-05$ | $2.3e-05$ | $2.4e-04$ | $\varphi_1$ |
| 6 | $2.2e-05$ | $2.4e-05$ | $1.8e-05$ | $1.3e-04$ | $\varphi_2$ |
| | $9.7e-05$ | $4.4e-05$ | $4.2e-05$ | $9.4e-05$ | $\varphi_3$ |
| | $1.9e-08$ | $3.2e-08$ | $8.7e-08$ | $1.7e-06$ | $\varphi_0$ |
| $n =$ | $1.5e-07$ | $1.5e-07$ | $2.5e-07$ | $2.8e-06$ | $\varphi_1$ |
| 8 | $4.3e-07$ | $3.8e-07$ | $5.7e-07$ | $3.0e-06$ | $\varphi_2$ |
| | $1.3e-06$ | $6.6e-07$ | $5.6e-07$ | $1.6e-06$ | $\varphi_3$ |
| | $2.4e-10$ | $3.7e-10$ | $1.0e-09$ | $2.0e-08$ | $\varphi_0$ |
| $n =$ | $1.1e-09$ | $1.7e-09$ | $3.4e-09$ | $3.9e-08$ | $\varphi_1$ |
| 10 | $9.0e-09$ | $6.9e-09$ | $7.5e-09$ | $4.8e-09$ | $\varphi_2$ |
| | $1.2e-08$ | $1.0e-08$ | $8.8e-09$ | $3.2e-08$ | $\varphi_3$ |
| | $2.6e-12$ | $4.3e-12$ | $1.2e-11$ | $2.4e-10$ | $\varphi_0$ |
| $n =$ | $2.1e-11$ | $3.0e-11$ | $4.9e-11$ | $6.1e-10$ | $\varphi_1$ |
| 12 | $1.0e-10$ | $5.3e-11$ | $8.7e-11$ | $6.0e-10$ | $\varphi_2$ |
| | $3.4e-10$ | $2.3e-10$ | $1.8e-10$ | $7.1e-10$ | $\varphi_3$ |

When $z$ becomes a matrix $A$ instead of a scalar, the same approach works, with the term $1/(s - z)$ becoming the resolvent matrix $(sI - A)^{-1}$. In the context of this work $\Gamma$ is a Hankel contour, that is, a deformed Bromwich contour that winds around the negative real axis in the counter-clockwise sense, see Fig. 5.1. In particular it encloses all eigenvalues of

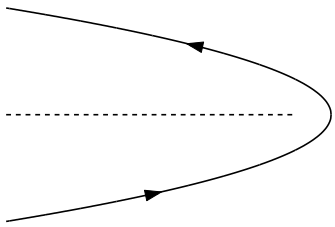

FIG. 5.1. *A typical Hankel contour, winding around the negative real axis (dashed) in the anti-clockwise sense.*

the negative semidefinite matrix $A$.

In [46] various choices for such contours are discussed. Although the integral does not depend on this choice, the convergence of the results gained by an evaluation with the trapezoid rule on the contour can be optimized a great deal.

The contour $\Gamma$ is represented as the image of the real line $\mathbb{R}$ under an analytic function $\phi$. Then (5.1) can be written as

$$e^z = \frac{1}{2\pi i} \int_{-\infty}^{\infty} \frac{e^{\phi(\theta)}}{\phi(\theta) - z} \phi'(\theta)\, d\theta. \tag{5.2}$$

The integrand is an exponentially decaying function . By truncating $\mathbb{R}$ to a finite interval one therefore commits only an exponentially small error. Following [45] we shall arbitrarily fix

ETNA
Kent State University
etna@mcs.kent.edu



this interval as $[-\pi, \pi]$. In $[-\pi, \pi]$ we take $n$ points $\theta_k$ spaced regularly at a distance $2\pi/n$, and our trapezoid approximation to (5.2) becomes

$$I_n = in^{-1} \sum_{k=1}^{n} \frac{e^{s_k}}{z - s_k} w_k, \tag{5.3}$$

where $s_k = \phi(\theta_k)$ and $w_k = \phi'(\theta_k)$.

Using an optimized version of Talbot's original contours it is possible to achieve a convergence rate of $O(3.89^{-n})$ [45, 46]. In particular it is possible to get almost down to machine precision with as few as 24 poles, which come in 12 conjugate pairs.

The exponential decay of the integrand in (5.2) is missing once we try to generalize the approach for $\varphi$ functions of higher order. The term $\varphi_l$ is only algebraically decaying, which is too slow for most applications in practice [26]. An alternative approach to enforce the exponential decay might be to introduce an additional reparametrization of the real line $\mathbb{R}$ by transformations as discussed in [35].

But given the rational approximation (5.3) of $e^z$ the approximation for $\varphi_1(z)$ induced by (4.1) is

$$\varphi_1(z) \approx in^{-1} \sum_{k=1}^{n} \frac{e^{s_k} s_k^{-1}}{z - s_k} w_k.$$

This is the trapezoidal rule applied to the integral of Cauchy type

$$I_1 = \frac{1}{2\pi i} \int_\Gamma \frac{e^s}{s} \frac{1}{s - z} ds,$$

which is indeed an alternative integral representation of $\varphi_1$.

THEOREM 5.1. *Let $C$ be a closed contour encircling the points $0$ and $z \in \mathbb{C}$ with winding number $1$. Then*

$$\varphi_l(z) = \frac{1}{2\pi i} \int_C \frac{e^s}{s^l} \frac{1}{s - z} ds.$$

*Proof.* Multiplying the Taylor series representation (1.3) of $\varphi_l(z)$ by $z^l$ reveals that $\varphi_l(z)$ is the regular part of the Laurent series for $e^z/z^l$:

$$\frac{e^z}{z^l} = \varphi_l(z) + \sum_{k=0}^{l-1} \frac{1}{k!} z^{k-l}.$$

Given the linearity of the path integral it is sufficient to show that

$$I = \frac{1}{2\pi i} \int_C \sum_{k=0}^{l-1} \frac{1}{k!} s^{k-l} \frac{1}{s - z} ds$$

vanishes. Changing the order of summation and integration yields

$$I = \sum_{k=0}^{l-1} \frac{1}{k!} \frac{1}{2\pi i} \int_C s^{k-l} \frac{1}{s - z} ds.$$

ETNA
Kent State University
etna@mcs.kent.edu



If $z = 0$ all integrals vanish. Let $z \neq 0$ and define $I_k = \frac{1}{2\pi i} \int_C s^{k-l} \frac{1}{s-z} ds$. The integrals $I_k$ are solved by residue calculus

$$I_k = \text{Res} \left( \frac{s^{k-l}}{s-z}, 0 \right) + \text{Res} \left( \frac{s^{k-l}}{s-z}, z \right).$$

The latter residue is $z^{k-l}$. The residue of $\frac{s^{-m}}{s-z}$ at the origin, where $m$ is a positive integer, is given as

$$\frac{1}{(m-1)!} \frac{d^{m-1}}{ds^{m-1}} \frac{1}{s-z}\bigg|_{s=0} = \frac{1}{(m-1)!} \frac{d^{m-1}}{da^{m-1}} \frac{1}{a}\bigg|_{a=-z} = -z^{-m},$$

which implies $I_k = 0$ for all $k = 0, \ldots, l-1$. □

In this integral representation the integrand is exponentially decaying along a Hankel contour and is therefore of greater practical use than (1.4). Although contour integrals need more points than optimal rational approximations they are somewhat more flexible when adaptive integrators are implemented as

$$\varphi_l(hA) = \frac{1}{2\pi i} \int_\Gamma \frac{e^{hs}}{(hs)^l} (sI - A)^{-1} ds, \tag{5.4}$$

where $h$ is a time step. This identity may give hope that $\varphi_l(hA)$ can be evaluated with the same resolvent matrices independently of $h$. Unfortunately, a balance of the truncation and discretization error cannot be maintained. The truncation error grows as $h$ shrinks, whereas the discretization error grows with $h$. However, accuracy can be maintained throughout certain ranges of $h$. Experiments to identify them for the exponential are reported in [45].

Exponential integrators, such as the method by Krogstad, often need $\varphi_l(hA)v$ and $\varphi_l\left(\frac{1}{2}hA\right)v$. Using best rational approximation we have to solve $n$ linear systems for both products. Achieving similar accuracy with Talbot contours it is typically enough to solve $2n$ linear systems once and to apply (5.4). Hence, we do not expect a trade-off with respect to best rational approximations in this situation.

**6. Exponential integrators.** It would go far beyond the scope of this article to introduce exponential integrators in detail. Instead we have decided to pick two typical but rather distinctive members of the huge family of exponential integrators. Various others are classified in [33].

Quite often exponential integrators have close relatives amongst the more established methods. There are exponential versions of Runge-Kutta methods and multistep methods which try to overcome the problems of their relatives for stiff problems by treating the linear term exactly.

A typical exponential Runge-Kutta method is the one-step method of Krogstad [25]. Like the method of Cox and Matthews [8] it has 4th order accuracy in many circumstances. In the worst case the order reduces to 2 for Cox and Matthews or 3 for Krogstad. Order reduction has been studied in detail by Ostermann and Hochbruck [21]. They have introduced a scheme that gives 4th order in all cases. We follow their formulation and introduce exponential Runge-Kutta methods for (1.1) as

$$u_{n+1} = u_n + h \sum_{i=1}^{s} b_i(hA) G_{ni},$$

$$U_{ni} = u_n + h \sum_{j=1}^{i-1} a_{ij}(hA) G_{nj},$$

**ETNA**
Kent State University
etna@mcs.kent.edu



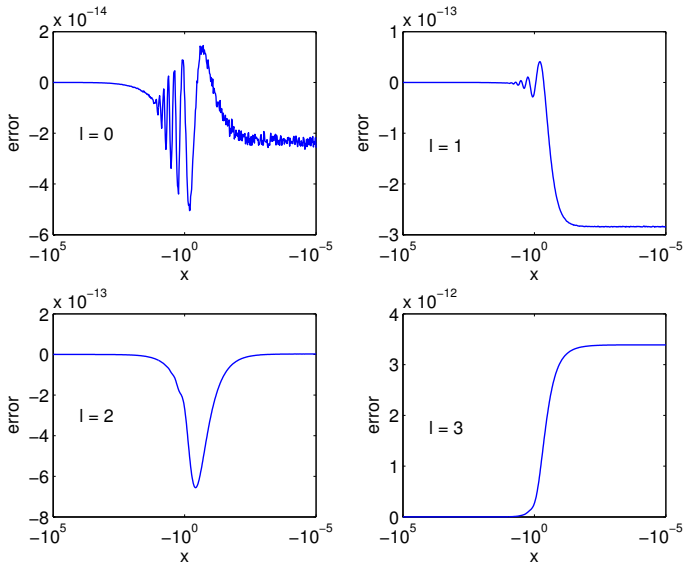


FIG. 5.2. *Error curves* $\varphi_l(x) - r_{(l)}(x)$. *The rational approximation* $r_{(0)}$ *is constructed by an application of the trapezoid rule with* 24 *nodes (symmetric with respect to the real axis) to* (5.2) *with* $\phi(\theta) = 24\,(0.5017\theta\cot(0.6407\theta) - 0.6122 + 0.2645i\theta)$*; see* [45, 46]. *In addition we have used a shift* $s = 1$. *The functions* $r_{(l)}$ *with* $l > 0$ *are induced by* (4.1).

$$G_{nj} = g\left(t_n + c_j h, U_{nj}\right) + Au_n.$$

In order to simplify the notation we use the abbreviations

$$\varphi_{i,j} = \varphi_{i,j}\left(hA\right) = \varphi_i\left(c_j hA\right), \qquad 2 \leqslant j \leqslant s$$

and

$$\varphi_i = \varphi_i\left(hA\right).$$

We only give the method of Krogstad:

| | | | | |
|---|---|---|---|---|
| $c_1 = 0$ | | | | |
| $c_2 = \frac{1}{2}$ | $a_{21} = \frac{1}{2}\varphi_{1,2}$ | | | |
| $c_3 = \frac{1}{2}$ | $a_{31} = \frac{1}{2}\varphi_{1,3} - \varphi_{2,3}$ | $a_{32} = \varphi_{2,3}$ | | |
| $c_4 = 1$ | $a_{41} = \varphi_{1,4} - 2\varphi_{2,4}$ | | $a_{43} = 2\varphi_{2,4}$ | |
| | $b_1 = \varphi_1 - 3\varphi_2 + 4\varphi_3$ | $b_2 = 2\varphi_2 - 4\varphi_3$ | $b_3 = b_2$ | $b_4 = -\varphi_2 + 4\varphi_3$ |

The implementation can be done "columnwise". Starting with $G_{n1}$ it is possible to evaluate all matrix-vector products of the first column. We calculate the savings introduced by approximating in common poles in the first column by counting the matrix-vector products. The terms $a_{21}$ and $a_{31}$ can be solved by solving one set of shifted linear systems rather than two sets. The terms $a_{41}$ and $b_1$ need an alternative set of poles, as a different scaling parameter $c$ is used. In such situations contours offer some flexibility through equation (5.4). Using optimal poles we would have to solve the systems for $\varphi_1$, $\varphi_2$ and $\varphi_3$. Thus we save the factor $5/2$ for the first column. Repeating the argument for all columns we save a factor of $12/6$ in

ETNA
Kent State University
etna@mcs.kent.edu



total. This factor does not take into account the additional initial costs of computing the $LU$ decompositions for 5 instead of 2 sets of shifted systems.

Exponential relatives of classic multistep methods are more attractive in terms of the underlying linear algebra than exponential Runge-Kutta methods. An early reference is an article by Nørsett [36], but they have been rediscovered recently by Beylkin [4] and also by Cox and Matthews [8]. Livermore has applied them to problems from magnetohydrodynamics [26]. Ostermann et al. [37] have analyzed their stability and Calvo and Palencia [5] gave details about the starting process for abstract Cauchy problems. A complete derivation is given in the thesis of Minchev [32].

The underlying formula for (1.1) is here

$$u_n = e^{hA} u_{n-1} + h \sum_{l=0}^{k} \beta_l g_{n-l}. \tag{6.1}$$

The coefficients $\beta_l$ are linear combinations of matrix functions. The explicit ($\beta_0 \equiv 0$) exponential Adams-Bashforth method of order 4 is given by

$$\begin{pmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \end{pmatrix} = \begin{pmatrix} 1 & 11/6 & 2 & 1 \\ 0 & -3 & -5 & -3 \\ 0 & 3/2 & 4 & 3 \\ 0 & -1/3 & -1 & -1 \end{pmatrix} \begin{pmatrix} \varphi_1(hA) \\ \varphi_2(hA) \\ \varphi_3(hA) \\ \varphi_4(hA) \end{pmatrix}.$$

This method serves as the predictor. An implicit exponential Adams-Moulton method may serve as corrector,

$$\begin{pmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \end{pmatrix} = \begin{pmatrix} 0 & 1/3 & 1 & 1 \\ 1 & 1/2 & -2 & -3 \\ 0 & -1 & 1 & 3 \\ 0 & -1/6 & 0 & -1 \end{pmatrix} \begin{pmatrix} \varphi_1(hA) \\ \varphi_2(hA) \\ \varphi_3(hA) \\ \varphi_4(hA) \end{pmatrix}.$$

It is standard for classic multistep methods to perform only the first step of a fixed-point iteration to solve (6.1). This idea is often referred as "PECE" form [10, Chapter 7.4]. Comparing the predicted and corrected result serves as an error control in implementations of the classic multistep methods. In order to predict $u_n$ it is necessary to evaluate $e^{hA}u_{n-1}$ and $\varphi_l(hA)g_{n-1}$, $l \geqslant 1$. It is sufficient to solve two sets of shifted linear systems when working with a set of common poles. For a $k$-step method we therefore save the factor $(k+1)/2$ when using the same set of poles for $\varphi_1, \ldots, \varphi_k$. This factor does not take into account the possibility to reuse the poles for $\varphi_1, \ldots, \varphi_k$ to compute $e^{hA}u_{n-1}$. Using that approach we reduce the number of $LU$ decompositions by a factor of 2 and may achieve an additional speedup by solving linear systems with 2 right-hand sides rather than solving 2 linear systems with one right-hand side each. Next one has to evaluate $g(u_n, t_n)$ in order to compute $\beta_0$ by solving again a set of shifted systems. All other coefficients are linear combinations of vectors already available from previous steps.

**7. Numerical experiments.** In this section we illustrate the potential of exponential integrators relying on rational approximations by two examples. Both examples are nonlinear reaction-diffusion equations with rather large diffusion constants to make the linear term both stiff and dominant in the evolution. The equations are posed in 1D and 2D, which results in advantageous band structures in the finite difference discretizations of the diffusion term.

These particular equations could be solved more efficiently using a diagonalization by means of Fourier transforms, or by spectral collocation methods. It should therefore be emphasized that they serve solely as a convenient means to illustrate our technique for dealing

ETNA
Kent State University
etna@mcs.kent.edu



with $\varphi$ functions. Errors introduced by the spatial discretization are not taken into account here.

We compare exponential integrators with established numerical methods for stiff systems of ordinary differential equations. Although we have worked with a uniform mesh in time, we find that our implementation can be competitive even with state-of-the-art adaptive methods such as Radau collocation methods [15] or the multistep backward differentiation formulas of MATLAB [40] when applied to reaction-diffusion equations with a mild nonlinearity. The methods we have compared are:

- The explicit exponential Runge-Kutta methods of Krogstad. We used two CF approximations of $\varphi_1(s)$ and $\varphi_1(\frac{1}{2}s)$ with 12 poles each, and varied the number of poles; see Figures 7.2 and 7.3.
- The exponential multistep methods of order 4, 6 and 8. Starting values were computed using MATLAB's `ode15s` integrator. The startup calculation was not taken into account in the timings. We used a CF approximation of the exponential with 12 poles and a shift of $s = 1$.
- MATLAB's `ode15s` integrator, which is an adaptive solver based on backward differentiation formulas. Rather than reducing the stepsize to achieve higher accuracy, we reduced the tolerances for the absolute and relative error. The method is implicit and the linear systems are solved using direct methods, that is UMFPACK, which has been incorporated with effective matrix reorderings in MATLAB's `ode15s` in recent releases.
- The method `RADAU5`, which is an implementation of a fifth-order implicit Runge-Kutta method of RADAU IIA type, with 3 stages and automatic stepsize control. For this we used a MATLAB implementation adapted from the thesis of Tee [41]. The code is available online.[1]

The results reported here depend very strongly on the underlying numerical linear algebra, and in particular, an exponential integrator may perform better or worse than a competitor simply because of a switch from a direct method to an iterative method or vice versa. Therefore, for consistency, we used direct methods for all of our computations, relying on UMFPACK as the common underlying framework for solving all linear systems. The complex shifted symmetric matrices were reordered using a sparse reverse Cuthill-McKee ordering and symmetric approximate minimum degree permutations [9].

Before we present the results in detail, we shall take a moment to describe common properties of the graphs describing the performance of the methods being used. In Figures 7.2 and 7.3 the vertical axes represent the *relative error*. For all experiments we have computed an "exact" solution by using `ode15s` with very tight error tolerances. The error is then calculated as the 2-norm of the difference between the approximation and the exact solution, divided by the 2-norm of the solution. Thus the error plotted in the graphs is a relative one. The scaling of the horizontal axes varies. The *relative time-steps* displayed are scaled by the overall simulation time scale. A relative timestep $h$ implies that $1/h$ steps have been taken. The *computer time* displayed is the CPU time we have measured. As is well known, such timings should not be relied upon too precisely, since in MATLAB we have limited control of internal routines optimizing certain code fragments. The *number of evaluations of* $g$ refers to the number of evaluations of the nonlinear reaction term.

The experiments were performed with MATLAB 7.4 on a HP workstation xw4200 with a 3.2 GHz Pentium 4 processor and 1 GByte of RAM running Windows XP.

[1] See the webpage of the author: www.comlab.ox.ac.uk/thomas.schmelzer

ETNA
Kent State University
etna@mcs.kent.edu



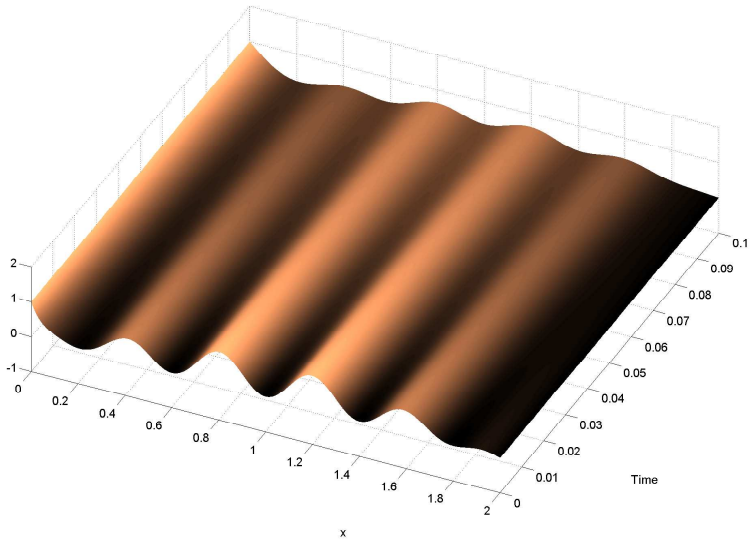


FIG. 7.1. *Solution of the Fisher equation.*

**7.1. The Fisher equation.** The Fisher equation [12] is a one-dimensional reaction-diffusion problem with a "logistic" reaction term

$$u_t = Du_{xx} + ru(1-u) \qquad x \in [0,2].$$

For the experiments reported here we chose $D = 0.05$ and $r = 0.01$. The initial function is chosen as

$$u_0(x, t=0) = \exp(-20x) - x(x-2)\cos^2(5x\pi/2).$$

The boundary values are fixed as

$$u(x=2,t) = 0, \qquad u(x=0,t) = 1.$$

We solve the equation for $0 \leqslant t \leqslant 0.1$. The equation is semidiscretized by standard finite differences with a 3-point stencil on a regular grid with $\triangle x = 1/2000$. This results in an extremely stiff system. The matrix $A$ introduced by this discretization is of dimension $N = 1999$ and tridiagonal. The left-most eigenvalue of the scaled matrix $DA$ is approximately $-2 \times 10^5$. Tridiagonal systems can be solved in $O(N)$ operations, so the matrix is very well suited for any integrator based on rational approximation.

The results in Figure 7.2 emphasize the power of exponential integrators. In the first plot one should note the large relative timesteps of $O(10^{-1})$. This implies that only 10 steps have been taken. In the second plot one sees that exponential integrators outperform `radau5` and `ode15s`. The Krogstad, Multistep 4 and Multistep 6 curves in this plot have benefitted by factors of approx. 2, 2.5 and 3.5 , respectively, through the use of common poles. In the third plot graphs labeled Krogstad$_n$ show the relative error of Krogstad's method when using only $n$ poles. The tremendous advantage of exponential integrators is revealed in fourth plot. They need far fewer evaluations of the nonlinear term.

**ETNA**
Kent State University
etna@mcs.kent.edu



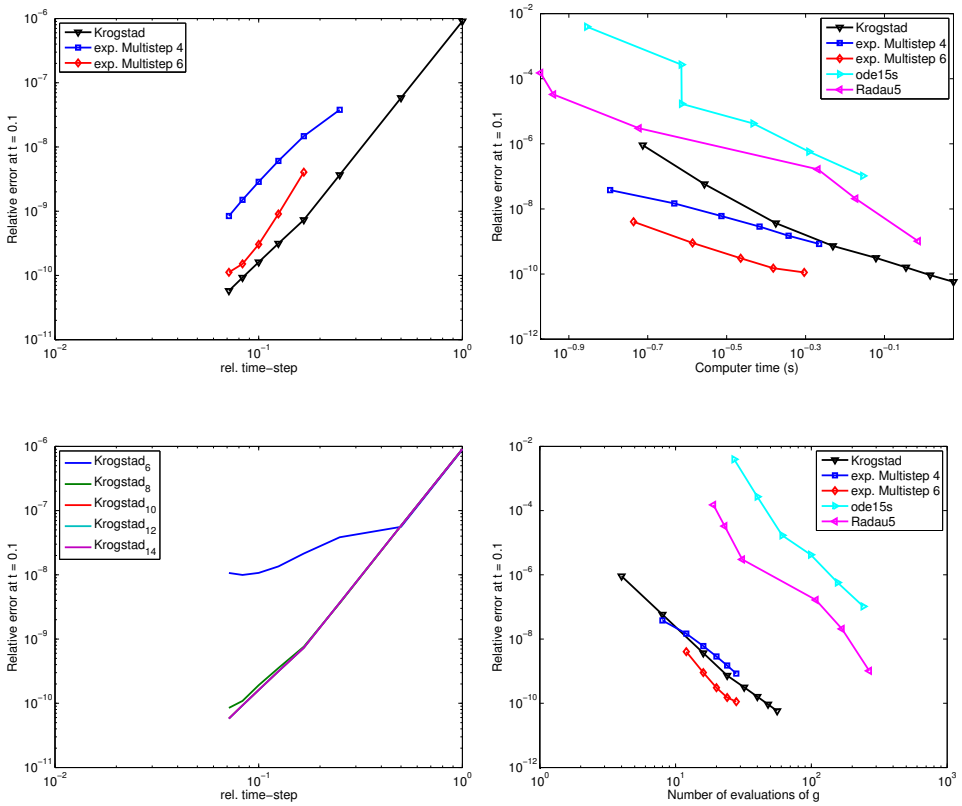


FIG. 7.2. *Results for the Fisher equation.*

**7.2. The Allen-Cahn equation in 2D.** The Allen-Cahn equation in 2D reads as [1]

$$u_t = \varepsilon \triangle u + (u - u^3).$$

We solve this equation for $0 \leqslant t \leqslant 0.1$ with $\varepsilon = 0.1$ on the square $[0,1]^2$ with homogeneous Neumann boundary conditions. The initial condition is a trigonometric polynomial:

$$u_0 = c \sum_{i=1}^{8} \sum_{j=1}^{8} r_{ij} \cos(i\pi x)\ \cos(j\pi y).$$

The coefficient matrix $R$ is constructed, so as to be arbitrary but reproducible, by taking the first 64 digits of $\pi$:

$$r_{11} = 3,\ r_{21} = 1,\ r_{31} = 4,\ \ldots,\ r_{12} = 5,\ \ldots, r_{88} = 2.$$

These numbers are then normalized by

$$r_{ij} = \frac{r_{ij}}{5} - 1.$$

The parameter $c$ is chosen such that $\max |u_0| = 1$ on the square. For the spatial discretization $A$ of the Laplacian we use standard finite differences on $[0,1]^2$. The matrix $A$ is of dimension $10^4 \times 10^4$ and symmetric.

**ETNA**
Kent State University
etna@mcs.kent.edu



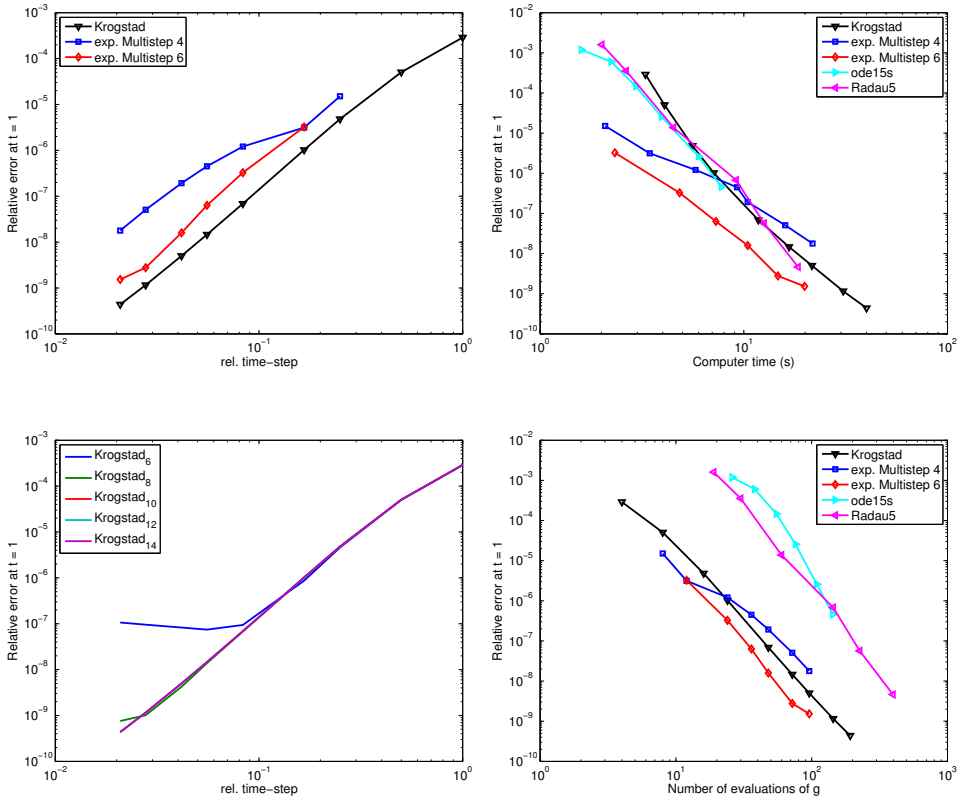


FIG. 7.3. *Results for the Allen-Cahn equation in* 2*D.*

The results of the experiment are shown in Figure 7.3. Broadly speaking the conclusions for this experiment are like those for the Fisher equation. Again the use of common poles has speeded up the Krogstad and multistep timings by factors of 2 to 4.

**8. Conclusions and outlook.** We have shown that $\varphi$ functions can be evaluated for matrix arguments efficiently using rational approximations constructed via Carathéodory-Fejér approximation or contour integrals. This enables us to implement competitive exponential integrators for large stiff systems of ODEs. The rational approximations are typically twice as fast as the contour integrals as they require half as many poles for the same accuracy.

Exponential integrators rely on the fast evaluation of the matrix-vector product $\varphi_l(A)b$ for several $l$ at once. Therefore we proposed the approximation in a set of common poles as in equation (4.1) and found that this enables us to reduce the work per step dramatically.

We should also mention that similar techniques can be used for other functions of interest [39].

**Acknowledgments.** Parts of this work were done when T.S. stayed for six weeks at the University of Dusseldorf, enjoying the hospitality of Marlis Hochbruck. She and Alexander Ostermann taught us a great deal about exponential integrators. Herbert Stahl gave us advice on best rational approximations. Christoph Ortner's insight on reaction-diffusion equations has always been helpful.

**ETNA**
**Kent State University**
**etna@mcs.kent.edu**



REFERENCES

[1] S. M. ALLEN AND J. W. CAHN, *A microscope theory for antiphase boundary motion and its application to antiphase domain coarsening*, Acta Metall., 27 (1979), pp. 1085–1095.
[2] A. I. APTEKAREV, *Sharp constants for rational approximations of analytic functions*, Mat. Sb., 193 (2002), pp. 1–72.
[3] H. BERLAND, B. SKAFLESTAD, AND W. WRIGHT, *EXPINT—A MATLAB package for exponential integrators*, ACM Trans. Math. Softw., 33 (2007).
[4] G. BEYLKIN, J. M. KEISER, AND L. VOZOVOI, *A new class of time discretization schemes for the solution of nonlinear PDEs*, J. Comput. Phys., 147 (1998), pp. 362–387.
[5] M. P. CALVO AND C. PALENCIA, *A class of explicit multistep exponential integrators for semilinear problems*, Numer. Math., 102 (2006), pp. 367–381.
[6] A. J. CARPENTER, A. RUTTAN, AND R. S. VARGA, *Extended numerical computations on the "$1/9$" conjecture in rational approximation theory*, in Rational approximation and interpolation, P. Graves-Morris, E. Saff, and R. Varga, eds., Lecture Notes in Mathematics, 1105, Springer, Berlin, 1984, pp. 383–411.
[7] W. J. CODY, G. MEINARDUS, AND R. S. VARGA, *Chebyshev rational approximations to $\mathbf{e^{-x}}$ in $[\mathbf{0}, +\infty)$ and applications to heat-conduction problems*, J. Approx. Theory, 2 (1969), pp. 50–65.
[8] S. M. COX AND P. C. MATTHEWS, *Exponential time differencing for stiff systems*, J. Comput. Phys., 176 (2002), pp. 430–455.
[9] T. A. DAVIS, *Direct methods for sparse linear systems*, SIAM, Philadelphia, PA, 2006.
[10] P. DEUFLHARD AND F. BORNEMANN, *Scientific computing with ordinary differential equations*, Springer-Verlag, New York, 2002.
[11] S. W. ELLACOTT, *On the Faber transform and efficient numerical rational approximation*, SIAM J. Numer. Anal., 20 (1983), pp. 989–1000.
[12] R. FISHER, *The wave of advance of advantageous genes*, Annals of Eugenics, 7 (1937), pp. 355–369.
[13] G. H. GOLUB AND C. F. VAN LOAN, *Matrix computations*, 3rd ed., Johns Hopkins University Press, Baltimore, 1996.
[14] A. A. GONCHAR AND E. A. RAKHMANOV, *Equilibrium distributions and degree of rational approximation of analytic functions*, Math. USSR Sbornik, 62 (1989), pp. 305–348.
[15] E. HAIRER AND G. WANNER, *Solving ordinary differential equations. II. Stiff and differential-algebraic problems*, Springer-Verlag, Berlin, 1991.
[16] G. H. HALPHEN, *Traité des fonctions elliptiques et de leurs applications. Premiére partie. Théorie des fonctions elliptiques et de leurs développement en séries*, Gauthier-Villars, Paris, 1886. http://moa.cit.cornell.edu/
[17] N. J. HIGHAM, *Accuracy and stability of numerical algorithms*, 2nd ed., SIAM, Philadelphia, 2002.
[18] ———, *The scaling and squaring method for the matrix exponential revisited*, SIAM J. Matrix Anal. Appl., 26 (2005), pp. 1179–1193.
[19] M. HOCHBRUCK AND C. LUBICH, *On Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal., 34 (1997), pp. 1911–1925.
[20] M. HOCHBRUCK, C. LUBICH, AND H. SELHOFER, *Exponential integrators for large systems of differential equations*, SIAM J. Sci. Comput., 19 (1998), pp. 1552–1574.
[21] M. HOCHBRUCK AND A. OSTERMANN, *Explicit exponential Runge-Kutta methods for semilinear parabolic problems*, SIAM J. Numer. Anal., 43 (2005), pp. 1069–1090.
[22] ———, *Exponential Runge-Kutta methods for parabolic problems*, Appl. Numer. Math., 53 (2005), pp. 323–339.
[23] A.-K. KASSAM, *Solving reaction-diffusion equations ten times faster*, Tech. Report NA 03/16, Oxford University Computing Laboratory, 2003.
[24] A.-K. KASSAM AND L. N. TREFETHEN, *Fourth-order time-stepping for stiff PDEs*, SIAM J. Sci. Comput., 26 (2005), pp. 1214–1233.
[25] S. KROGSTAD, *Generalized integrating factor methods for stiff PDEs*, J. Comput. Phys., 203 (2005), pp. 72–88.
[26] P. LIVERMORE, *An implementation of the exponential time differencing scheme to the magnetohydrodynamic equations in a spherical shell*, J. Comput. Phys., 220 (2007), pp. 824–838.
[27] Y. Y. LU, *Computing a matrix function for exponential integrators*, J. Comput. Appl. Math., 161 (2003), pp. 203–216.
[28] A. P. MAGNUS, *CFGT determination of Varga's constant '$1/9$'*, 1985. handout of the Padé approximation meeting organized by C. Brezinski at Marseille-Luminy, October 14-18, 1985.
[29] ———, *Asymptotics and super asymptotics of best rational approximation error norms for the exponential function (the '$1/9$' problem) by the Carathéodory–Fejér method*, in Nonlinear Methods and Rational Approximation II, A. Cuyt et al., ed., Kluwer, Dordrecht, 1994, pp. 173–185.
[30] G. MEINARDUS, *Approximation of functions: Theory and numerical methods*, Springer, Berlin, 1967.
[31] B. V. MINCHEV, *Computing analytic matrix functions for a class of exponential integrators*, Reports in

**ETNA**
**Kent State University**
**etna@mcs.kent.edu**



Informatics 278, University of Bergen, Bergen, Norway, 2004.
[32] ———, *Exponential integrators for semilinear problems*, PhD thesis, University of Bergen, 2004.
[33] B. V. MINCHEV AND W. M. WRIGHT, *A review of exponential integrators for first order semi-linear problems*, Tech. Report 2/05, Norwegian University of Science and Technology, Trondheim, Norway, 2005.
[34] C. MOLER AND C. VAN LOAN, *Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later*, SIAM Rev., 45 (2003), pp. 3–49.
[35] M. MORI, *Discovery of the double exponential transformation and its developments*, Publ. RIMS, Kyoto University, 41 (2005), pp. 897–935.
[36] S. P. NØRSETT, *An A-stable modification of the Adams–Bashforth methods*, in Conf. on Numerical Solution of Differential Equations (Dundee, Scotland, 1969), Springer, Berlin, 1969, pp. 214–219.
[37] A. OSTERMANN, M. THALHAMMER, AND W. WRIGHT, *A class of explicit exponential general linear methods*, BIT, 46 (2006), pp. 409–431.
[38] Y. SAAD, *Analysis of some Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal., 29 (1992), pp. 209–228.
[39] T. SCHMELZER AND L. N. TREFETHEN, *Computing the gamma function using contour integrals and rational approximations*, SIAM J. Numer. Anal., 45 (2007), pp. 558–571.
[40] L. F. SHAMPINE AND M. W. REICHELT, *The MATLAB ODE suite*, SIAM J. Sci. Comput., 18 (1997), pp. 1–22.
[41] T.-W. TEE, *An adaptive rational spectral method for differential equations with rapidly varying solutions*, PhD thesis, Oxford University, 2006.
[42] L. N. TREFETHEN, *Rational Chebyshev approximation on the unit disk*, Numer. Math., 37 (1981), pp. 297–320.
[43] ———, *Chebyshev approximation on the unit disk*, in Computational aspects of complex analysis, H. Werner et al., ed., D. Reidel Publishing, Dordrecht, 1983, pp. 309–323.
[44] L. N. TREFETHEN AND M. H. GUTKNECHT, *The Carathéodory-Fejér method for real rational approximation*, SIAM J. Numer. Anal., 20 (1983), pp. 420–436.
[45] L. N. TREFETHEN, J. WEIDEMAN, AND T. SCHMELZER, *Talbot quadratures and rational approximations*, BIT, 46 (2006), pp. 653–670.
[46] J. A. C. WEIDEMAN AND L. N. TREFETHEN, *Parabolic and hyperbolic contours for computing the Bromwich integral*, Math. Comp., 76 (2007), pp. 1341–1356.