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 "φ 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 method is based on best rational approximations to f on the negative real axis computed via the Carathéodory-Fejér procedure, and we conjecture that the accuracy scales as $(9.28903\ldots)^{-2n}$, where n is the number of complex matrix solves. In particular, three matrix solves suffice to evaluate $f(A)$ to approximately six digits of accuracy. 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 [29], 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 φ 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

$$u_{n+1} = e^{\Delta t A} u_n + \Delta t\, \varphi_1(\Delta t A) g(u_n, t_n), \tag{1.2}$$

where $\varphi_1(z) = (e^z - 1)/z$. There are many other exponential schemes and some of these ideas have been reinvented several times [29]. In recent years the interest in exponential integrators has heightened. Hochbruck and Lubich [15] introduced Krylov methods to compute the φ functions, Cox and Matthews [7] and Krogstad [21] introduced one-step methods with 4th order accuracy in many circumstances, and Hochbruck and Ostermann [17] 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 [3] and Cox and Matthews after being introduced by Nørsett [32] in 1969. Kassam and Trefethen applied the one-step method

[†]Computing Laboratory, Oxford University, United Kingdom, thoms@comlab.ox.ac.uk
[‡]Computing Laboratory, Oxford University, United Kingdom, lnt@comlab.ox.ac.uk
[§]Thomas Schmelzer is supported by a Rhodes Scholarship.

of Cox and Matthews to stiff PDEs such as the KdV, Kuramoto–Sivashinsky, Allen–Cahn and Grey–Scott equations [19, 20] and compared them with more standard schemes.

We follow a recent convention [2, 17, 18] 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

$$\varphi_l(z) = \frac{\varphi_{l-1}(z) - \varphi_{l-1}(0)}{z}, \qquad l \geqslant 1. \tag{1.3}$$

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

$$\varphi_l(z) = \sum_{k=l}^{\infty} \frac{1}{k!} z^{k-l}. \tag{1.4}$$

The functions φ_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 [13, 20]. To address this problem Cox and Matthews [7] made use of the Taylor series. This technique works for scalars and diagonal matrices. An alternative stable evaluation is based on (1.3) and a Cauchy integral representation on a circle Γ of radius 1 centered at z, for $|z| < 1/2$. Following Kassam and Trefethen [20] this is

$$\varphi_l(z) = \frac{1}{2\pi i}\int_\Gamma \frac{\varphi_l(s)}{s - z} ds \approx \frac{1}{M}\sum_{k=1}^{M} \varphi_l\left(z + e^{it_k}\right) \tag{1.5}$$

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 Γ has to enclose the spectrum of A:

$$\varphi_l(A) = \frac{1}{2\pi i}\int_\Gamma \varphi_l(s)\left(sI - A\right)^{-1} ds. \tag{1.6}$$

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 Γ. In particular this idea is successful for the matrix exponential and it is closely related to rational approximations as shown in [40]. Instead of evaluating (1.6) for $l > 0$ we give in §5 an alternative integral representation of φ_l not involving the function φ_l in the integrand.

Another method for evaluating the φ functions that is widely used nowadays is Padé approximation, as suggested by Beylkin et al. [3], Hochbruck et al. [16], and Minchev and Wright [29]. The idea is based on scaling and squaring, which is a popular method for computing the matrix exponential [14, 30]. This approach is restricted to matrices of moderate dimension as the cvaluation of $f(A)b$ requires the explicit computation of $f(A)$.

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 [23]. Lu claims that the coefficients are hard to compute and gives the rational approximation of type $(14, 14)$ for φ_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 [39] and [40].

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 behaviour. Methods based on rational approximations get their power from the fast exponential decay of the error introduced by the approximant. In the case of φ 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 [26].

Cody, Meinardus and Varga [6] showed that $E_{nn}(\exp(x), \mathbb{R}^-)$ decreases geometrically as $n \to \infty$. In 1986 Gonchar and Rakhmanov [11] 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 υ, *Halphen's constant*, is the unique positive root of the equation

$$\sum_{n=1}^{\infty} \frac{n\upsilon^n}{1-(-\upsilon)^n} = \frac{1}{8}.$$

This constant was studied by Halphen as early as 1886 [12]. The proof by Gonchar and Rakhmanov confirmed a conjecture by Magnus [24] and previous numerical computations by Trefethen and Gutknecht [39] and Carpenter, Ruttan and Varga [5].

A sharper result than (2.1) was conjectured by Magnus [25] and subsequently proved by Aptekarev [1]:

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

We define

$$E_l = \lim_{n\to\infty} E_{nn}\left(\varphi_l, \mathbb{R}^-\right)^{1/n},$$

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 [11], 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.

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 [37, 38]. By utilizing a conformal map from the unit circle to a real interval or the negative real line it was shown in [39] that the method is very efficient for real approximation, too. The idea can also be generalized for other domains in the complex plane [9].

These approximations are so close to optimal that the method can often be regarded as exact in practice. Magnus [25] 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 [40] a MATLAB code is presented that computes rational approximations of $\exp(x)$ on the negative real line with an error as small as 2×10^{-14}. We have made some minor modifications to adapt the code for φ 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 behaviour 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 φ_l by its CF approximation is indeed smaller than 10^{-6} for $l = 0, 1, \ldots$. The error gets smaller as l increases. For all exponential integrators we have used it is sufficient to compute φ_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.

```
function [zi,ci,r_inf] = cf_rat_phi(n,l);

  K = 75;                                  % no of Cheb coeffs
  nf = 1024;                               % no of pts for FFT
  w = exp(2i*pi*(0:nf-1)/nf);              % roots of unity
  t = real(w);                             % Cheb pts (twice over)
  scl = 9;                                 % scale factor for stability
  F = zeros(size(t));
  g = (t~=-1);
  F(g) = phiS(scl*(t(g)-1)./(t(g)+1),l);   % evaluate Phi_l on neg. axis
  c = real(fft(real(F)))/nf;               % Cheb coeffs of F
  f = polyval(c(K+1:-1:1),w);              % analytic part f of F
  [U,S,V] = svd(hankel(c(2:K+1)));         % SVD of Hankel matrix
  s = S(n+1,n+1);                          % singular value
  u = U(K:-1:1,n+1)'; v = V(:,n+1)';       % singular vectors
  zz = zeros(1,nf-K);                      % zeros for padding
  b = fft([u zz])./fft([v zz]);            % finite Blaschke product
  rt = f-s*w.^K.*b;                        % extended function r-tilde
  rtc = real(fft(rt))/nf;                  % its Laurent coeffs
  zr = roots(v); qj = zr(abs(zr)>1);       % poles
  qc = poly(qj);                           % coeffs of denominator
  pt = rt.*polyval(qc,w);                  % numerator
  ptc = real(fft(pt)/nf);                  % coeffs of numerator
  ptc = ptc(n+1:-1:1); ci = 0*qj;
  for k = 1:n                              % calculate residues
    q = qj(k); q2 = poly(qj(qj~=q));
    ci(k) = polyval(ptc,q)/polyval(q2,q);
  end
  zi = scl*(qj-1).^2./(qj+1).^2;           % poles in z-plane
  ci = 4*ci.*zi./(qj.^2-1);                % residues in z-plane
  r_inf = 0.5*(phiS(0,l)+sum(ci./zi));     % r at infinity
  [m, order]=sort(imag(zi));               % sort poles
  zi=zi(order);                            % reorder poles
  ci=ci(order);                            % reorder residues
```

FIG. 3.1. *A modification of the MATLAB function given in* [40] *to compute poles* $\{z_i\}$*, residues* $\{c_i\}$ *and the constant* r_∞ *by the Carathéodory-Fejér method for the type* (n,n) *best approximation* r^* *to* φ_l *on* $\mathbb{R}^-$.

4. Approximation in common poles. When implementing exponential integrators it is an attractive option to use a set of common poles for all φ functions. 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 [34].

PROPOSITION 4.1. *Let*

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

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

$$\varphi_l(B_z) = \begin{pmatrix} \varphi_l(z) & \varphi_{l+1}(z) \\ 0 & \varphi_l(0) \end{pmatrix} \qquad \text{for } l = 0, 1, \ldots. \tag{4.1}$$

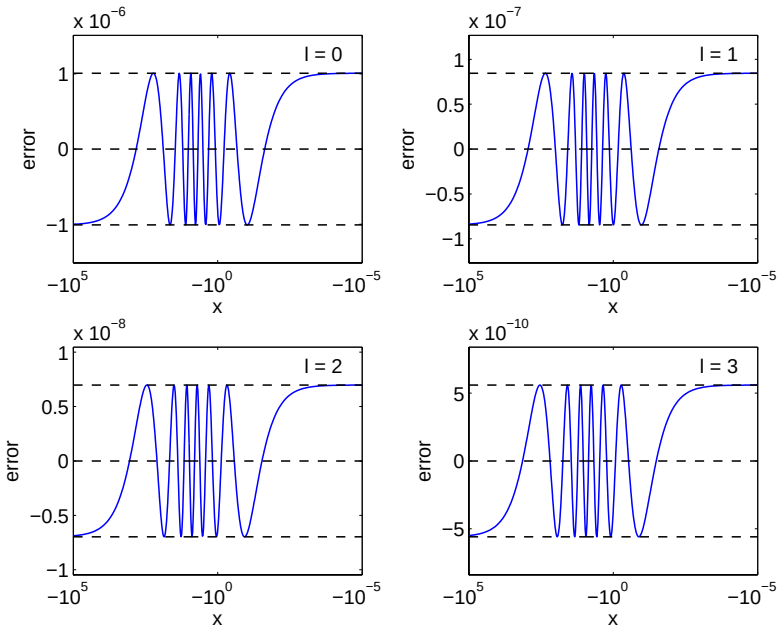

FIG. 3.2. *Error curves $\varphi_l(x) - r^*(x)$ for type $(6,6)$ best approximation to φ_l, $l = 0,1,2,3$ on $\mathbb{R}^-$, as computed by CF approximation by the program of Fig. 3.1. Note that the abscissa, the negative real axis in the z-plane, is displayed on a log scale. The dashed lines mark the minimax errors.*

Saad established this identity for $l = 0$. He was not concerned with φ 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

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

□

An alternative proof is provided by the observation that

$$f(B_z) = \begin{pmatrix} f(z) & \frac{f(z)-f(0)}{z} \\ 0 & f(0) \end{pmatrix}$$

for all z if f is an entire function, which is exactly the recursion given in (1.3).

If z is replaced by a matrix A this idea has been suggested amongst others by Hochbruck et al. and Saad [16, 34] 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 φ_l, and noting that

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

the entry of $r_{(l)}(B_z)$ approximating φ_{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

$$r_{(l+k)}(z) = \sum_{j=1}^{n} \frac{c_j z_j^{-k}}{z - z_j}, \qquad k \in \mathbb{Z} \tag{4.2}$$

for rational approximations using common poles of φ functions. This generalizes a result by Minchev [27]. 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 and might ultimately serve as a step towards a proof of Conjecture 2.1.

The matrices B_z are diagonalisable 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

$$\|f(B_z) - \varphi_l(B_z)\|_2 \leqslant \kappa_2(V_z) \|f(\Lambda_z) - \varphi_l(\Lambda_z)\|_2 \tag{4.3}$$

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 [10]. □

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

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

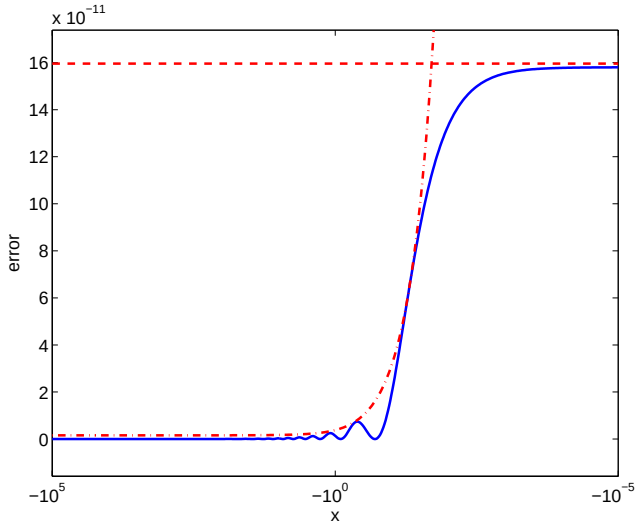

FIG. 4.1. *The poles of a CF approximation of e^x of type $(12, 12)$ have been used to evaluate the φ_1 function on the negative real line. The blue solid curve is the error. The dashed-dotted red curve is the bound induced by* (4.3). *The dashed line represents the bound of Theorem* 4.2.

If we knew more about the decay of the second term for increasing degree of the rational approximation we might be able to prove Conjecture 2.1. Nevertheless it comes as no surprise that the derivative of an analytic function approximating an analytic function is a good approximation for the derivative of the approximant.

Rather than focusing on asymptotic results we give in Table 4.1 some numerical results for small degrees of the rational approximations. These results are very promising for applications.

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 [23] 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 (4.2) and (4.4) to φ functions of higher order. We summarize our results in Table 4.2.

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

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

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 Γ is a Hankel contour, that is, a deformed Bromwich contour that winds around the

	φ_0	φ_1	φ_2	φ_3	
	$\mathbf{1.0e-06}$	$9.3e-05$	$2.2-03$	$3.0-02$	φ_0
$n=$	$5.3e-05$	$\mathbf{8.5e-08}$	$9.7e-06$	$2.7e-04$	φ_1
6	$4.6e-04$	$4.0e-06$	$\mathbf{7.0e-09}$	$9.5e-07$	φ_2
	$1.6e-03$	$3.1e-05$	$2.9e-07$	$\mathbf{5.6e-10}$	φ_3
	$\mathbf{1.2e-08}$	$1.7e-06$	$6.2e-05$	$1.2e-03$	φ_0
$n=$	$8.0e-07$	$\mathbf{7.5e-10}$	$1.3e-07$	$5.5e-06$	φ_1
8	$9.1e-06$	$4.7e-08$	$\mathbf{4.8e-11}$	$9.9e-09$	φ_2
	$4.2e-05$	$4.9e-07$	$2.8e-09$	$\mathbf{3.0e-12}$	φ_3
	$\mathbf{1.4e-10}$	$2.9e-08$	$1.5e-06$	$3.8e-05$	φ_0
$n=$	$1.1e-08$	$\mathbf{7.1e-12}$	$1.8e-09$	$1.0e-07$	φ_1
10	$1.6e-07$	$5.6e-10$	$\mathbf{3.7e-13}$	$1.1e-10$	φ_2
	$9.1e-07$	$7.3e-09$	$2.7e-11$	$\mathbf{1.9e-14}$	φ_3
	$\mathbf{1.6e-12}$	$4.7e-10$	$3.1e-08$	$1.0e-06$	φ_0
$n=$	$1.6e-10$	$\mathbf{6.8e-14}$	$2.7e-11$	$1.7e-09$	φ_1
12	$2.6e-09$	$6.5e-12$	$\mathbf{4.3e-15}$	$1.2e-12$	φ_2
	$1.8e-08$	$1.0e-10$	$2.7e-13$	$\mathbf{5.6e-16}$	φ_3

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 computed by the program of Fig. 3.1 for fixed l. An example: Given the CF approximation of degree 10 for φ_1 the error committed by approximating φ_3 with the same poles using identity (4.2) is $7.3e-09$.

	$s=1/2$	$s=1$	$s=2$	$s=5$	
	$1.6e-06$	$2.7e-06$	$7.5e-06$	$1.5e-04$	φ_0
$n=$	$1.0e-05$	$1.1e-05$	$2.3e-05$	$2.4e-04$	φ_1
6	$2.2e-05$	$2.4e-05$	$1.8e-05$	$1.3e-04$	φ_2
	$9.7e-05$	$4.4e-05$	$4.2e-05$	$9.4e-05$	φ_3
	$1.9e-08$	$3.2e-08$	$8.7e-08$	$1.7e-06$	φ_0
$n=$	$1.5e-07$	$1.5e-07$	$2.5e-07$	$2.8e-06$	φ_1
8	$4.3e-07$	$3.8e-07$	$5.7e-07$	$3.0e-06$	φ_2
	$1.3e-06$	$6.6e-07$	$5.6e-07$	$1.6e-06$	φ_3
	$2.4e-10$	$3.7e-10$	$1.0e-09$	$2.0e-08$	φ_0
$n=$	$1.1e-09$	$1.7e-09$	$3.4e-09$	$3.9e-08$	φ_1
10	$9.0e-09$	$6.9e-09$	$7.5e-09$	$4.8e-09$	φ_2
	$1.2e-08$	$1.0e-08$	$8.8e-09$	$3.2e-08$	φ_3
	$2.6e-12$	$4.3e-12$	$1.2e-11$	$2.4e-10$	φ_0
$n=$	$2.1e-11$	$3.0e-11$	$4.9e-11$	$6.1e-10$	φ_1
12	$1.0e-10$	$5.3e-11$	$8.7e-11$	$6.0e-10$	φ_2
	$3.4e-10$	$2.3e-10$	$1.8e-10$	$7.1e-10$	φ_3

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.

negative real axis in the anti-clockwise sense, see Fig. 5.1. In particular it encloses all eigenvalues of the negative semidefinite matrix A.

In [41] 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.

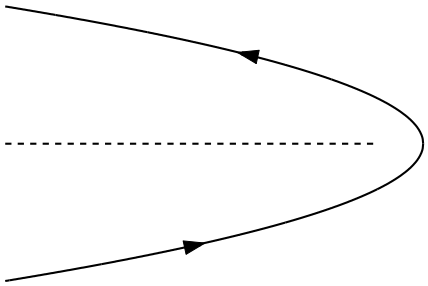

FIG. 5.1. *A typical Hankel contour, winding around the negative real axis (dashed) in the anti-clockwise sense.*

The contour Γ is represented as the image of the real line $\mathbb{R}$ under an analytic function ϕ. 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 [40] we shall arbitrarily fix this interval as $[-\pi, \pi]$. In $[-\pi, \pi]$ we take n points θ_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 $O(3.89^{-n})$ [40, 41]. 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 φ functions of higher order. The term φ_l is only algebraically decaying, which is too slow for most applications in practice [22]. 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 [31].

But given the rational approximation (5.3) of e^z the approximation for $\varphi_1(z)$ induced by (4.2) is

$$\varphi_1(z) \approx in^{-1} \sum_{k=1}^{n} \frac{e^{s_k} s_k^{-1}}{z - s_k} w_k. \tag{5.4}$$

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

$$I_1 = \frac{1}{2\pi i} \int_{\Gamma} \frac{e^s}{s} \frac{1}{s - z} ds, \tag{5.5}$$

which is indeed an alternative integral representation of φ_1.

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

$$\varphi_l(z) = \frac{1}{2\pi i} \int_C \frac{e^s}{s^l} \frac{1}{s-z} ds. \tag{5.6}$$

PROOF: Let

$$f(s) = \frac{e^s}{s^l (s-z)};$$

then if $z = 0$

$$\frac{1}{2\pi i} \int_C \frac{e^s}{s^{l+1}} ds = \text{Res}\left(\frac{e^s}{s^{l+1}}, 0\right) = \frac{1}{(l-1)!} = \varphi_l(0).$$

For $z \neq 0$ the proof is a bit more involved and the contour integral is the sum of two residues as the paths encloses two distinct poles

$$\frac{1}{2\pi i} \int_C \frac{e^s}{s^l (s-z)} ds = \text{Res}\,(f, 0) + \text{Res}\,(f, z).$$

The pole at $s = z$ is simple and therefore $\text{Res}\,(f, z) = e^z / z^l$. The pole at $s = 0$ is of order l. Hence

$$\text{Res}\,(f, 0) = \frac{1}{(l-1)!} \frac{d^{l-1}}{ds^{l-1}} \frac{e^s}{s-z}\bigg|_{s=0}.$$

It may seem odd at this stage but to avoid a technical mess we introduce the family of g functions as

$$g_l(s) = e^s \sum_{j=1}^{l} \frac{(l-1)!}{(l-j)!} (z-s)^{-j}.$$

An immediate consequence of this definition is the identity

$$\text{Res}\,(f, 0) = -\frac{1}{(l-1)!} \frac{d^{l-1}}{ds^{l-1}} g_1(s)\bigg|_{s=0}.$$

The quirky definition of the g functions was made in order to enforce

$$\frac{d}{ds} g_l(s) = g_{l+1}(s).$$

We prove this identity in a moment but note first that a consequence is the helpful identity

$$\text{Res}\,(f, 0) = -\frac{1}{(l-1)!} g_l(0).$$

The proof is a small technical exercise. Applying the product rule and shifting the index j results in

$$\begin{aligned} g_l'(s) &= g_l(s) + e^s \sum_{j=1}^{l} j \frac{(l-1)!}{(l-j)!} (z-s)^{-j-1}, \\ &= e^s \left(\sum_{j=1}^{l} \frac{(l-1)!}{(l-j)!} (z-s)^{-j} + \sum_{j=2}^{l+1} (j-1) \frac{(l-1)!}{(l-j+1)!} (z-s)^{-j} \right) \\ &= e^s \left((z-s)^{-1} + \sum_{j=2}^{l} \left[\frac{(l-1)!}{(l-j)!} + (j-1) \frac{(l-1)!}{(l-j+1)!} \right] (z-s)^{-j} + l!(z-s)^{-l-1} \right). \end{aligned}$$

The term in squared brackets is

$$\frac{(l-1)!}{(l-j)!} + (j-1) \frac{(l-1)!}{(l-j+1)!} = \frac{(l-1)!}{(l-j+1)!} (l-j+1+j-1) = \frac{l!}{(l-j+1)!}.$$

Hence the desired identity holds

$$\frac{d}{ds} g_l(s) = e^s \sum_{j=1}^{l+1} \frac{l!}{(l+1-j)!} (z-s)^{-j} = g_{l+1}(s).$$

At this stage we have shown that for $z \neq 0$

$$\frac{1}{2\pi i} \int_C \frac{e^s}{s^l (s-z)} ds = \frac{e^z - \frac{z^l}{(l-1)!} g_l(0)}{z^l}.$$

But

$$\frac{z^l}{(l-1)!} g_l(0) = \sum_{j=1}^{l} \frac{1}{(l-j)!} z^{l-j}$$

which are the first l terms of the Taylor series for the exponential. Hence

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

Comparing this with the Taylor series of φ_l in (1.4) finishes the proof. □

In this integral representation the integrand is exponentially decaying along a Hankel contour and is therefore of greater practical use than (1.6). Although contour integrals need more points than optimal rational approximations they are somewhat more flexible when adaptive integrators are implemented as

$$\varphi_l(hA) = \frac{1}{2\pi i} \int_\Gamma \frac{e^{hs}}{(hs)^l} (sI - A)^{-1} ds, \tag{5.7}$$

where h is a time step. Hence we can work in theory with the same resolvent matrices independently of h. This behavior is typical for contours vs. best rational approximations [40].

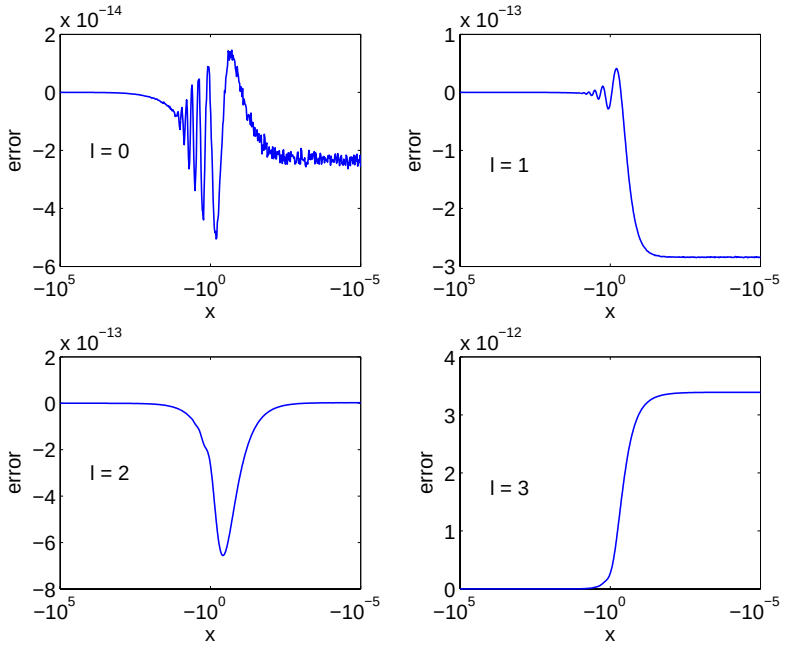

FIG. 5.2. *Error curves* $\varphi_l(x) - r_{(l)}(x)$. *The rational approximation* $r_{(0)}$ *is constructed by an application of the trapezoid rule to* (5.2) *with* $\phi(\theta) = 24\,(0.5017\theta\cot(0.6407\theta) - 0.6122 + 0.2645i\theta)$; *see* [40, 41]. *In addition we have used a shift* $s = 1$. *The functions* $r_{(l)}$ *with* $l > 0$ *are induced by* (4.2).

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 [29].

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 [21]. Like the method of Cox and Matthews [7] 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 [17]. 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

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

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

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

In order to simplify the notation we use the abbrevations

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

and

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

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 could be done "columnwise". Starting with G_{n1} it is possible to evaluate all matrix-vector products of the first column. Note that it is necessary to evaluate φ_1 and φ_2 twice with different scaling parameters c. In such situations contours offer some flexibility through equation (5.7). Using common poles, each entry in the column corresponds to a linear combination of the solutions of the 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 [32], but they have been rediscovered recently by Beylkin [3] and also by Cox and Matthews [7]. Livermore has applied them to problems from magnetohydrodynamics [22]. Ostermann et al. [33] have analysed their stability and Calvo and Palencia [4] gave details about the starting process for abstract Cauchy problems. A complete derivation is given in the thesis of Minchev [28].

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 β_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 [8, Chapter 7.4]. Comparing the predicted and corrected result serves as an error control in implementations of the classic multipstep methods. In order to predict u_n it is necessary to evaluate $e^{hA}u_{n-1}$ and $\varphi_l(hA)g_{n-1}$, $l \geqslant 1$, which can be done using a set of common poles with u_{n-1} and g_{n-1} as right-hand sides. Next one has to evaluate $g(u_n, t_n)$ in order to compute β_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. The reaction-diffusion equation

$$u_t = \varepsilon \triangle u + \underbrace{u - u^3 - u^3 \cos^2 4t}_{g(u,t)}. \tag{7.1}$$

serves here as a test example. We have implemented the two exponential integrators introduced above and compare them with MATLAB's ODE15s function [36], which is a standard (and very impressive) routine for solving large stiff systems of ODEs.

We solve this equation for $0 \leqslant t \leqslant 5$ with $\varepsilon = 0.0025$ on the square $[0, 1]^2$ with homogeneous Neumann boundary conditions. The initial condition is a trigonometric polynomial:

$$u_0 = c \sum_{i=1}^{8} \sum_{j=1}^{8} r_{ij} \cos(i\pi x)\, \cos(j\pi y). \tag{7.2}$$

The coefficient matrix R is constructed by taking the first 64 digits of π:

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

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 $N^2 \times N^2$ and symmetric.

The equation above is a close relative of the well understood Allen-Cahn equation

$$u_t = \varepsilon \triangle u + u - u^3. \tag{7.3}$$

The solution of this equation has stable equilibria at $u = \pm 1$. Relatively sharp and slowly moving interfaces separate areas where the solution is rather flat and near ± 1. This interesting phenomenon is called *metastability*. In a situation like this adaptive solvers are very effective. To avoid such a "steady state" we have introduced the additional periodic forcing term making the cubic nonlinearity time-dependent.

For all experiments we have computed an "exact" solution by using ODE15s with very tight error tolerances. The accuracy is the relative error measured in the 2-norm at $t_f = 5$. Initial experimental results are given in Table 7.1–7.3. The numerical experiments were performed with MATLAB 7.0.4 on a HP workstation xw4200 with a 3.2 GHz Pentium 4 processor and 1.00 GByte of RAM running Windows XP.

The exponential multistep method (Table 7.1) gives good results for as few as 6 poles. It seems high accuracy in evaluating the φ functions is of very limited

Exponential Multistep, 4th order				
N	Time for iteration	Time for LU decomp.	Accuracy	n = no. of poles
100	9.4	1.8	0.025	4
100	13.4	2.6	0.0049	6
100	17.2	3.5	0.0049	8

TABLE 7.1

Experimental results using an exponential multistep method with time step $h = 0.25$. It is enough to compute as few as $n/2$ LU decompositions. At each step one has to solve $n/2 \times 2$ and $n/2$ linear systems. This nonstandard notation means we are solving first $n/2$ linear systems with two right-hand sides and again $n/2$ linear systems in the corrector step.

ODE15s						
N	Time	Accuracy	Steps	LU d.	lin. sys.	eval. g
100	29.4	0.0023	66	23	152	154

TABLE 7.2

Experimental results using MATLAB's ODE15s function. Details are specified in Fig. 7.2

impact. Only when using as few as 4 poles do we observe weaker results. MATLAB's ODE15s function [36] (Table 7.2) needs far more steps and LU decompositions and evaluations of the nonlinear function g to achieve a similar accuracy. Krogstad's method (Table 7.3) is slower than the exponential multistep method but of better accuracy.

8. Conclusions and outlook. We have shown that φ 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 (4.2) 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 [35].

9. Acknowledgments. Parts of this work were done when T.S. stayed for six weeks at the University of Düsseldorf, 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.

REFERENCES

[1] A. I. Aptekarev, *Sharp constants for rational approximations of analytic functions*, Mat. Sb. **193** (2002), 1–72.

[2] Håvard Berland, Bård Skaflestad, and Will Wright, *EXPINT — A MATLAB package for exponential integrators*, ACM Trans. Math. Software, to appear.

[3] Gregory Beylkin, James M. Keiser, and Lev Vozovoi, *A new class of time discretization schemes for the solution of nonlinear PDEs*, J. Comput. Phys. **147** (1998), 362–387.

[4] M. P. Calvo and C. Palencia, *A class of explicit multistep exponential integrators for semilinear problems*, Numer. Math. **102** (2006), 367–381.

[5] A. J. Carpenter, A. Ruttan, and R. S. Varga, *Extended numerical computations on the "1/9"*

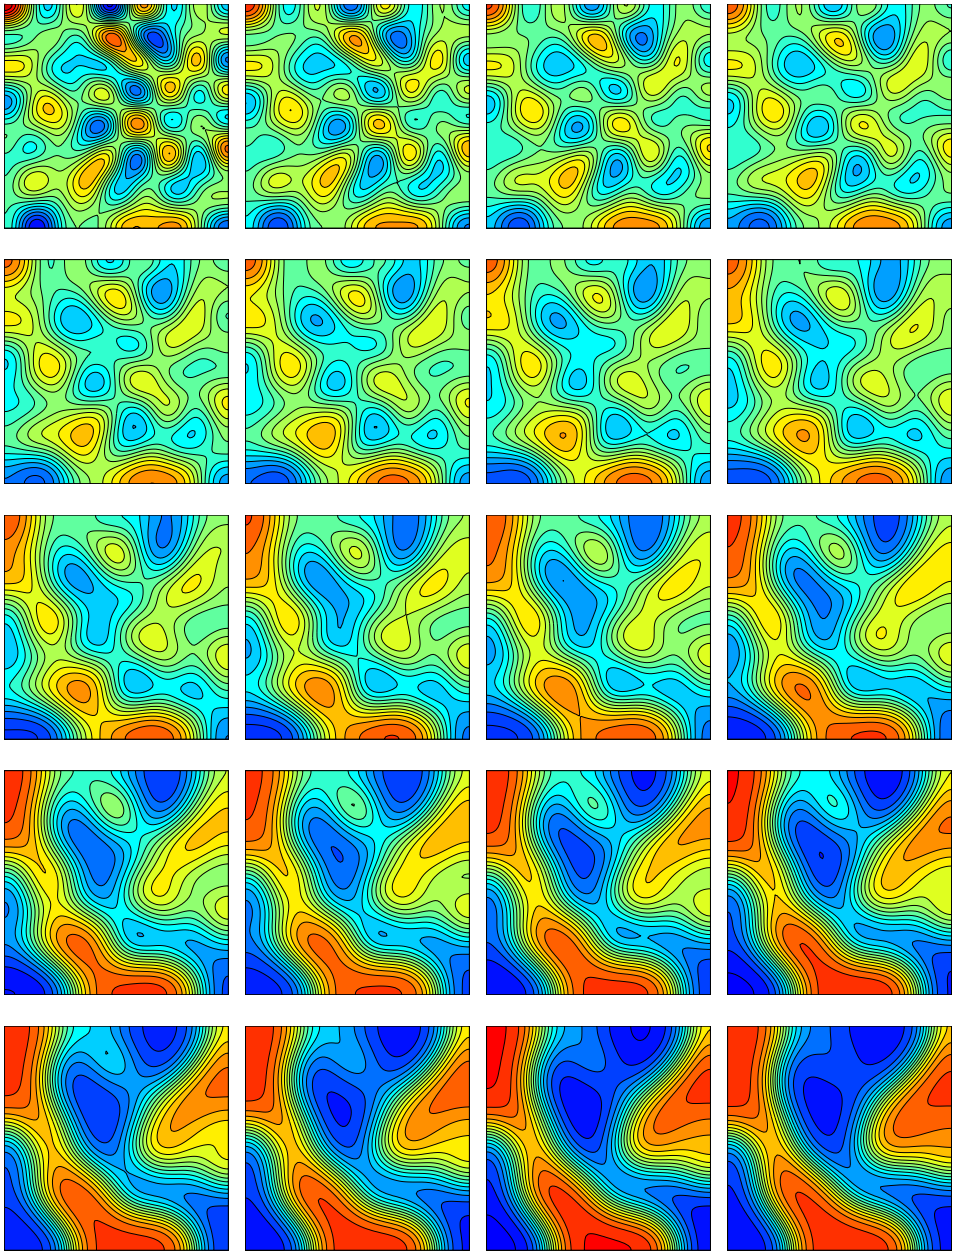

FIG. 7.1. *Evolution in time of* (7.3). *Starting in the upper left corner with the initial state* u_0 *the plots show* $u(t_i)$ *with* $t = 0, 0.25, 0.5, \ldots, 4.75$.

conjecture in rational approximation theory, Rational approximation and interpolation (P.R. Graves-Morris, E.B. Saff, and R.S. Varga, eds.), Lecture Notes in Math. 1105, Springer, Berlin, 1984, pp. 383–411.

[6] 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), 50–65.

[7] S. M. Cox and P. C. Matthews, *Exponential time differencing for stiff systems*, J. Comput. Phys. **176** (2002), 430–455.

```
tmax = 5; h = 0.25;
options = odeset('relTol',10e-4,'abstol',10e-4,'jacobian',@jac,'stats','on');
tic, [T,U] = ode15s(@evol,0:h:tmax,u0,options,L); toc

% modified Allen Cahn
function [up] = evol(t,u, L)
  up = L*u + (u-u.^3) - (cos(4*t))^2*u.^3;

% The Jacobian
function [J]=jac(t,u, L)
  n = size(L,1); J = L+spdiags(1-3*u.^2-(cos(4*t))^2*3*u.^2,0,n,n);
```

FIG. 7.2. *A code fragment for solving* (7.3) *by MATLAB's ODE15s function.*

Krogstad's method				
N	Time	Accuracy	$n = \sharp$poles	Stepsize
100	15.9	0.079	2	0.25
100	25.4	9.35×10^{-4}	4	0.25
100	38.1	5.79×10^{-4}	6	0.25
100	46.5	5.76×10^{-4}	8	0.25
100	7.8	0.097	2	0.50
100	13.8	0.010	4	0.50
100	20.0	0.0088	6	0.50
100	26.7	0.0088	8	0.50

TABLE 7.3

Experimental results using Krogstad's method. It is enough to compute n LU decompositions. Per step one has to solve $n/2, n/2, n/2 \times 3$ *and* $n/2$ *linear systems corresponding to the* 4 *stages of the process. In addition we have to evaluate the nonlinear function g once at every stage.*

[8] Peter Deuflhard and Folkmar Bornemann, *Scientific computing with ordinary differential equations*, Springer-Verlag, New York, 2002.

[9] S. W. Ellacott, *On the Faber transform and efficient numerical rational approximation*, SIAM J. Numer. Anal. **20** (1983), 989–1000.

[10] Gene H. Golub and Charles F. Van Loan, *Matrix computations*, 3rd ed., Johns Hopkins University Press, Baltimore, 1996.

[11] A. A. Gonchar and E. A. Rakhmanov, *Equilibrium distributions and degree of rational approximation of analytic functions*, Math. USSR Sbornik **62** (1989), no. 2, 305–348.

[12] G. H. Halphen, *Traité des fonctions elliptiques et de leurs applications. Premiére partie. Théorie des fonctions elliptiques et de leurs développement en séries*, http://moa.cit.cornell.edu/, Gauthier-Villars, Paris, 1886.

[13] Nicholas J. Higham, *Accuracy and stability of numerical algorithms*, 2nd ed., SIAM, Philadelphia, 2002.

[14] ———, *The scaling and squaring method for the matrix exponential revisited*, SIAM J. Matrix Anal. Appl. **26** (2005), 1179–1193.

[15] Marlis Hochbruck and Christian Lubich, *On Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal. **34** (1997), 1911–1925.

[16] Marlis Hochbruck, Christian Lubich, and Hubert Selhofer, *Exponential integrators for large systems of differential equations*, SIAM J. Sci. Comput. **19** (1998), 1552–1574.

[17] Marlis Hochbruck and Alexander Ostermann, *Explicit exponential Runge-Kutta methods for semilinear parabolic problems*, SIAM J. Numer. Anal. **43** (2005), 1069–1090.

[18] ———, *Exponential Runge-Kutta methods for parabolic problems*, Appl. Numer. Math. **53** (2005), 323–339.

[19] Aly-Khan Kassam, *Solving reaction-diffusion equations ten times faster*, Tech. Report NA 03/16, Oxford University Computing Laboratory, 2003.

[20] Aly-Khan Kassam and Lloyd N. Trefethen, *Fourth-order time-stepping for stiff PDEs*, SIAM J. Sci. Comput. **26** (2005), 1214–1233.

[21] S. Krogstad, *Generalized integrating factor methods for stiff PDEs*, J. Comput. Phys. **203** (2005), 72–88.

[22] P. Livermore, *An implementation of the exponential time differencing scheme to the magneto-hydrodynamics equations in a spherical shell*, J. Comp. Phys., to appear.
[23] Ya Yan Lu, *Computing a matrix function for exponential integrators*, J. Comput. Appl. Math. **161** (2003), 203–216.
[24] Alphonse P. Magnus, *CFGT determination of Varga's constant '1/9'*, 1985, handout of the Padé approximation meeting organized by C. Brezinski at Marseille-Luminy, 14-18 Oct. 1985.
[25] ———, *Asymptotics and super asymptotics of best rational approximation error norms for the exponential function (the '1/9' problem) by the Carathéodory–Fejér method*, Nonlinear Methods and Rational Approximation II (A. Cuyt et al., ed.), Kluwer, Dordrecht, 1994, pp. 173–185.
[26] Günter Meinardus, *Approximation of functions: Theory and numerical methods*, Springer, Berlin, 1967.
[27] Borislav V. Minchev, *Computing analytic matrix functions for a class of exponential integrators*, Reports in Informatics 278, University of Bergen, Bergen, Norway, 2004.
[28] ———, *Exponential integrators for semilinear problems*, Ph.D. thesis, University of Bergen, 2004.
[29] Borislav V. Minchev and Will 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.
[30] Cleve Moler and Charles Van Loan, *Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later*, SIAM Rev. **45** (2003), 3–49.
[31] Masatake Mori, *Discovery of the double exponential transformation and its developments*, Publ. RIMS, Kyoto University **41** (2005), 897–935.
[32] S. P. Nørsett, *An A-stable modification of the Adams–Bashforth methods*, Conf. on Numerical Solution of Differential Equations (Dundee, Scotland, 1969), Springer, Berlin, 1969, pp. 214–219.
[33] A. Ostermann, M. Thalhammer, and W. Wright, *A class of explicit exponential general linear methods*, BIT Numerical Mathematics **46** (2006), 409–431.
[34] Y. Saad, *Analysis of some Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal. **29** (1992), 209–228.
[35] Thomas Schmelzer and Lloyd N. Trefethen, *Computing the gamma function using contour integrals and rational approximations*, to appear in SIAM J. Numer. Anal.
[36] Lawrence F. Shampine and Mark W. Reichelt, *The MATLAB ODE suite*, SIAM J. Sci. Comput. **18** (1997), 1–22.
[37] Lloyd N. Trefethen, *Rational Chebyshev approximation on the unit disk*, Numer. Math. **37** (1981), 297–320.
[38] ———, *Chebyshev approximation on the unit disk*, Computational aspects of complex analysis (H. Werner et al., ed.), D. Reidel Publishing, Dordrecht, 1983, pp. 309–323.
[39] Lloyd N. Trefethen and Martin H. Gutknecht, *The Carathéodory-Fejér method for real rational approximation*, SIAM J. Numer. Anal. **20** (1983), 420–436.
[40] Lloyd N. Trefethen, J.A.C. Weideman, and Thomas Schmelzer, *Talbot quadratures and rational approximations*, BIT Numerical Mathematics **46** (2006), 653–670.
[41] J. A. C. Weideman and Lloyd N. Trefethen, *Parabolic and hyperbolic contours for computing the Bromwich integral*, Math. Comp., to appear.