

TALBOT QUADRATURES AND RATIONAL APPROXIMATIONS *

LLOYD N. TREFETHEN¹, J. A. C. WEIDEMAN², and THOMAS SCHMELZER^{3†}

¹*Oxford University Computing Laboratory, Wolfson Bldg., Parks Road,
Oxford OX1 3QD, UK. email: LNT@comlab.ox.ac.uk*

²*Dept. of Applied Mathematics, University of Stellenbosch, Private Bag X1,
Matieland 7602, South Africa. email: weideman@sun.ac.za*

³*Oxford University Computing Laboratory, Wolfson Bldg., Parks Road,
Oxford OX1 3QD, UK. email: thoms@comlab.ox.ac.uk*

Abstract.

Many computational problems can be solved with the aid of contour integrals containing e^z in the integrand: examples include inverse Laplace transforms, special functions, functions of matrices and operators, parabolic PDEs, and reaction-diffusion equations. One approach to the numerical quadrature of such integrals is to apply the trapezoid rule on a Hankel contour defined by a suitable change of variables. Optimal parameters for three classes of such contours have recently been derived: (a) parabolas, (b) hyperbolas, and (c) cotangent contours, following Talbot in 1979. The convergence rates for these optimized quadrature formulas are very fast: roughly $O(3^{-N})$, where N is the number of sample points or function evaluations. On the other hand, convergence at a rate apparently about twice as fast, $O(9.28903^{-N})$, can be achieved by using a different approach: best supremum-norm rational approximants to e^z for $z \in (-\infty, 0]$, following Cody, Meinardus and Varga in 1969. (All these rates are doubled in the case of self-adjoint operators or real integrands.) It is shown that the quadrature formulas can be interpreted as rational approximations and the rational approximations as quadrature formulas, and the strengths and weaknesses of the different approaches are discussed in the light of these connections. A MATLAB function is provided for computing Cody–Meinardus–Varga approximants by the method of Carathéodory–Fejér approximation.

AMS subject classification (2000): 65D30, 41A20

Key words: rational approximation, Carathéodory–Fejér approximation, quadrature, Hankel contour, inverse Laplace transform, special functions, trapezoid rule, Talbot contour

1 Introduction.

This article is concerned with the numerical computation of contour integrals of the form

$$(1.1) \quad I = \frac{1}{2\pi i} \int_{\Gamma} e^z f(z) dz,$$

*Received. Communicated.

†This work was partially supported by the National Research Foundation in South Africa under grant NRF5289.

where f is analytic in a neighborhood of the negative real axis $\mathbb{R}^- = (-\infty, 0]$ except on \mathbb{R}^- itself, and Γ denotes a “Hankel contour” that winds from $-\infty - 0i$ in the lower half-plane, around 0, and back to $-\infty + 0i$ in the upper half-plane. We assume that Γ lies in the region of analyticity of f , in which case I is independent of Γ under mild assumptions, even if Γ is deformed to widen out to a shape such as a parabola or a hyperbola for large $|\operatorname{Re} z|$. For example, it is enough if f satisfies $|e^z f(z)| = o(z^{-1})$ as $\operatorname{Re} z \rightarrow -\infty$. Here and in what follows we mostly speak of f as a scalar, for simplicity, but the same methods apply if it is a matrix or linear operator depending analytically on z .

On the face of it (1.1) might seem quite a specialized problem, but in fact, it arises throughout scientific computing. The integral is that of the inverse Laplace transform, which is more familiarly written in the form

$$(1.2) \quad G(t) = \frac{1}{2\pi i} \int_{\Gamma} e^{st} g(s) ds;$$

the substitutions $s = z/t$, $g(s) = f(z/t)/t$ yield $I = G(t)$. Integrals of the form (1.1) or (1.2) arise in many areas of applied mathematics and are the basis of fast methods for computing special functions [11, 22, 35, 42]. In numerical linear algebra, they can lead to fast methods for computing matrix exponentials e^A and associated vectors $e^A v$ for matrices A with eigenvalues in \mathbb{R}^- [20, 39]; in such applications (1.1) takes the form

$$(1.3) \quad e^A v = \frac{1}{2\pi i} \int_{\Gamma} e^z (z - A)^{-1} v dz.$$

In partial differential equations they offer powerful tools for solving parabolic PDE [2, 4, 6, 7, 8, 9, 48]. The evaluation of integrals related to (1.1) is also a key step in some fast algorithms for solving nonlinear PDE such as reaction-diffusion equations [14, 15]. These methods generalize in natural ways to large-scale problems involving fine discretizations and large matrices, where they are readily utilized in the context of Krylov subspace iterations and multiple processors [7, 8]. And they are readily generalized to problems with singularities off \mathbb{R}^- , including evolution equations for sectorial operators such as fractional differential and integral equations associated with anomalous diffusion [25, 34].

Numerical methods for the quadrature of (1.1) have recently been derived by Weideman that are in certain senses optimal or close to optimal. These are based on applying an N -point trapezoid or midpoint rule to (1.1) after a transformation of z to a real variable $\theta \in \mathbb{R}$. If Γ is a parabola, parameters have been derived [50] that achieve the convergence rate

$$(1.4) \quad \text{Trapezoid rule on parabola: } O(2.85^{-N});$$

related work with parabolic contours is due to Gavriluk and Makarov [9]. If Γ is a hyperbola, a choice particularly attractive for generalizations to sectorial operators, parameters have been derived [50] that achieve the rate

$$(1.5) \quad \text{Trapezoid rule on hyperbola: } O(3.20^{-N}),$$

which relates to work with hyperbolic contours by López-Fernández, Lubich, McLean, Palencia, Schädle, Sheen, Sloan, and Thomée [17, 18, 25, 38]. The longest-established class of quadrature contours, involving the cotangent function, were introduced by Talbot in 1979 [41]. Parameters for cotangent contours have been derived [51] that achieve the rate

$$(1.6) \quad \text{Trapezoid rule on cotangent contour: } O(3.89^{-N}).$$

The convergence rates (1.4)–(1.6) are all exceedingly fast, and they can all be doubled in the common case where $f(\bar{z}) = \overline{f(z)}$ for each z , such as arises with self-adjoint matrices or operators. At the same time, the fact that the rates are different is an indication that not all of them can be optimal in an ultimate sense. One might ask, rather than optimizing within these three particular classes of integration contours, what could be achieved by a more general optimization in which the shape of the contour is arbitrary? The purpose of this article is to explore one approach to this question. In §2 we suggest that any quadrature rule such as those of (1.4)–(1.6) can be interpreted as implicitly defining a *rational approximation* of e^z on $(-\infty, 0]$. The particular rational approximations associated with (1.4)–(1.6) are examined in §3. In §4 we note that one kind of optimality can be achieved by putting contours and quadrature formulas aside and simply considering best rational approximations of e^z on $(-\infty, 0]$, a topic made famous by Cody, Meinardus, and Varga in 1969 [5]. Such approximations lead to numerical methods for (1.1) that converge at the rate

$$(1.7) \quad \text{Best rational approximations: } O(9.28903^{-N}).$$

This rate is about twice as fast as (1.4)–(1.6), and thus might appear to be distinctly superior to the other three methods. Moreover, in practice, the required approximations can be computed by the Carathéodory–Fejér method, based on the singular value analysis of a Hankel matrix of Chebyshev coefficients of $e^{9(1-x)/(1+x)}$.

Are methods based on best rational approximations superior in practice? In §6 we show that the actual situation is not so simple. Such methods are indeed the fastest in certain circumstances, but they are also somewhat delicate, losing some of their power when the problem is changed slightly.

2 Quadratures as rational approximations.

The basis of the quadrature methods we shall consider is the trapezoid rule (or equivalently for our purposes the midpoint rule), which is exponentially accurate when applied to an analytic integrand on a periodic or unbounded domain. See [47] for a review of this phenomenon of high accuracy, which was first identified by Poisson in the 1820s and analyzed further by Turing (1939), Goodwin (1949), Davies (1959), Martensen (1968), and Stenger (1981), among others. The connection we shall make to rational approximation, however, is valid for any quadrature rule, whether or not it is derived from the trapezoid formula.




Figure 2.1: If $f(z) = o(1)$ as $|z| \rightarrow \infty$, then the contours Γ' and C of (2.4) and (2.6) are equivalent. This is the essential point in the derivation of Theorem 2.1.

Let $\phi(\theta)$ be an analytic function that maps the real line \mathbb{R} onto the contour Γ . Then (1.1) can be written as

$$(2.1) \quad I = \frac{1}{2\pi i} \int_{-\infty}^{\infty} e^{\phi(\theta)} f(\phi(\theta)) \phi'(\theta) d\theta.$$

Because of the term $e^{\phi(\theta)}$, the integrand will normally decrease exponentially as $|\theta| \rightarrow \infty$, in which case one commits an exponentially small error by truncating \mathbb{R} to a finite interval. For simplicity 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 (1.1) becomes

$$(2.2) \quad I_N = -iN^{-1} \sum_{k=1}^N e^{z_k} f(z_k) w_k,$$

where $z_k = \phi(\theta_k)$ and $w_k = \phi'(\theta_k)$. To achieve geometric convergence rates as in (1.4)–(1.6), it will be necessary to have the function ϕ depend on N , and in particular, ϕ will normally be chosen so that the error introduced by the truncation of \mathbb{R} to $[-\pi, \pi]$ is of the same order as that introduced by finite step size $2\pi/N$ of the trapezoid rule.

Now (2.2) consists of a linear combination of the values taken by f at the points z_k ,

$$(2.3) \quad I_N = - \sum_{k=1}^N c_k f(z_k), \quad c_k = iN^{-1} e^{z_k} w_k.$$

According to residue calculus, this sum is equal to the integral

$$(2.4) \quad I_N = \frac{1}{2\pi i} \int_C r(z) f(z) dz,$$

where $r(z)$ is the type $(N - 1, N)$ rational function

$$(2.5) \quad r(z) = \sum_{k=1}^N \frac{c_k}{z - z_k}$$

and C is a closed contour in the region of analyticity of f that winds in the negative sense once around each point z_k . By the negative sense we mean clockwise, a direction consistent with the appearance of the minus sign in (2.3). The reason for this choice of signs is illustrated in Figure 2.1. Let Γ' be a contour for (1.1) like Γ , except lying between \mathbb{R}^- and the points $\{z_k\}$. Then Γ' is equivalent to Γ for the integral (1.1) of $e^z f(z)$. Moreover, if f is analytic in $\mathbb{C} \setminus \mathbb{R}^-$ with $f(z) = o(1)$ as $|z| \rightarrow \infty$, then $r(z)f(z) = o(|z|^{-1})$ as $|z| \rightarrow \infty$. Thus if we deform C to the contour consisting of the union of Γ' with a large circular arc of radius R , the contribution from the latter must decrease to 0 as $R \rightarrow \infty$, implying that Γ' is equivalent to C for the integral (2.4) involving $r(z)f(z)$.

We have proved:

THEOREM 2.1. *Let f be analytic in $\mathbb{C} \setminus \mathbb{R}^-$ with $f(z) \rightarrow 0$ uniformly as $|z| \rightarrow \infty$, and define $r(z)$ by (2.5). Let Γ' be a contour for (1.1) lying between \mathbb{R}^- and the points $\{z_k\}$, as illustrated in Figure 2.1. Then the sum I_N defined by (2.2) is equal to*

$$(2.6) \quad I_N = \frac{1}{2\pi i} \int_{\Gamma'} r(z) f(z) dz,$$

and together with (1.1), this gives the quadrature error estimate

$$(2.7) \quad I - I_N = \frac{1}{2\pi i} \int_{\Gamma'} (e^z - r(z)) f(z) dz.$$

The last equation suggests that $r(z)$ is likely to be a good approximation to e^z near \mathbb{R}^- . Thus any quadrature formula can be viewed as a rational approximation; we shall look at three examples in the next section. Conversely, any rational approximation can be interpreted as a quadrature formula—see §4.

3 Parabolic, hyperbolic, and cotangent contours.

We now consider the three particular quadrature contours Γ mentioned in §1. The simplest is a parabola, a choice proposed by Gavriluk and Makarov [9]. Several parameters can be varied in defining such a curve, and by balancing various error terms Weideman has devised the particular combination

$$(3.1) \quad z(\theta) = N[0.1309 - 0.1194\theta^2 + 0.2500i\theta],$$

which achieves the convergence rate $O(2.85^{-N})$ mentioned in (1.4). According to Theorem 2.1, for each N , this corresponds to a rational function $r(z)$ of type $(N - 1, N)$ that we may consider as an approximation to e^z for values of z near \mathbb{R}^- . Figure 3.1 shows the quality of the approximation for the case $N = 32$. Since $2.85^{-32} \approx 3 \times 10^{-15}$, we expect 14 or more digits of accuracy, and this is




Figure 3.1: Contour plot in the complex z -plane of the error $|e^z - r(z)|$, where $r(z)$ is the rational function (2.5) associated with the trapezoid rule quadrature formula with $N = 32$ on the parabolic contour (3.1), marked by the grey curve. The solid dots mark the quadrature points z_k , which are the poles of r . From right to left, the contours correspond to $10^0, 10^{-1}, \dots, 10^{-14}$.




Figure 3.2: Like Fig. 3.1 but for the hyperbolic contour (3.2). The innermost contour is affected by rounding errors.




Figure 3.3: Like Fig. 3.1 but for the cotangent contour (3.3).

just what the figure shows. The high accuracy is achieved not only on \mathbb{R}^- but in a wide region in the left half-plane.

Another simple choice of contour is a hyperbola, as has been considered by López-Fernández, Lubich, McLean, Palencia, Schädle, Sheen, Sloan, and Thomée [17, 18, 25, 38]. Here Weideman has proposed the parameters

$$(3.2) \quad z(\theta) = 2.246N[1 - \sin(1.1721 - 0.3443i\theta)],$$

with convergence rate $O(3.20^{-N})$. (That this formula describes a hyperbola follows from standard properties of the sine function in the complex plane.) The slightly improved convergence rate would permit N to be reduced from 32 to, say, 28, but to facilitate comparisons, we take $N = 32$ again in Figure 3.2. Once more, the prescribed accuracy is achieved in a wide region of the left half-plane.

Talbot's original contours were defined by a formula involving the cotangent function. Again Weideman has developed parameters for such contours that are in certain senses optimal [51], and one of his choices amounts to

$$(3.3) \quad z(\theta) = N[0.5017\theta \cot(0.6407\theta) - 0.6122 + 0.2645i\theta],$$

with convergence rate $O(3.89^{-N})$. In this case we might reduce N to 24, but again we keep it at 32 for Figure 3.3, which looks much like the other two. Indeed, we have been quite surprised to find how similar these images all are. This suggests that there will be little difference between parabolas, hyperbolas, and cotangent contours in practice.

Figures 3.1–3.3 confirm the expectation suggested by Theorem 2.1: by applying the trapezoid rule over a Talbot contour, we implicitly construct excellent rational approximations to e^z .

```

J = 100; h = 2/J; s = (-1+h:h:1-h)';           % 1D grid
[xx,yy] = meshgrid(s,s);                       % 2D grid
x = xx(:); y = yy(:);                         % 2D grid stretched to 1D
L = -gallery('poisson',J-1)/h^2;               % 2D Laplacian
I = speye((J-1)^2);                           % identity
N = 32; theta = pi*(1:2:N-1)/N;               % quad pts in (0,pi)
z = N*(.1309-.1194*theta.^2+.2500i*theta);     % quad pts on contour
w = N*(-.1194*2*theta+.2500i);                 % derivatives
c = (1i/N)*exp(z).*w;                         % quadrature weights
u = (1-x.^2).*(1-y.^2).*exp(x);               % initial condition
v = zeros(size(u));
for k = 1:N/2                                  % quadrature via
    v = v - c(k)*(z(k)*I-0.02*L)\u;           % sparse linear solves
end
v = 2*real(v);                                % exploit symmetry
IN = v(J^2/2-J+1)                            % value at origin

```

Figure 3.4: MATLAB code fragment to compute the exponential of a matrix of dimension 9801 for the numerical solution of the heat equation in three dimensions. On a 2004 laptop, this code prints the correct result 0.93864312688253 in about 10 seconds.

Here is an illustration of how powerful these quadrature formulas can be in scientific computing. Suppose the heat equation $u_t = \varepsilon \Delta u$ in the square $[-1, 1]^2$, with zero boundary conditions and $\varepsilon = 0.02$, is semidiscretized by standard finite differences with a 5-point stencil on a regular grid with $\Delta x = \Delta y = 1/100$, with the variable v representing the discretization of u . The discretization introduces a matrix A of dimension $99^2 = 9801$. Suppose we take $u(x, y, 0) = e^x(1-x^2)(1-y^2)$ and seek the value

$$(3.4) \quad v(1)_{0,0} = \frac{1}{2\pi i} \int_{\Gamma} e^z (z - A)^{-1} v(0) dz$$

at the grid point $x = y = 0$, where $v(0)$ is the initial condition sampled on the grid. The MATLAB code fragment shown in Figure 3.4 evaluates (3.4) by the quadrature method (3.1) with $N = 32$. It runs in about 10 seconds on a 2004 laptop,¹ whereas MATLAB's `expm` function, which computes the whole matrix e^A rather than just a vector, would require about 6 hours.

4 Best approximations.

An alternative idea is to bypass the use of quadrature formulas and work directly with rational functions designed as good approximations to e^z on \mathbb{R}^- . There are various ways in which such approximations could be defined, and we shall follow the route that is the most familiar. Given m and n , let r^* be the real

¹All tests reported in this article were run in Matlab 7 on an IBM ThinkPad T42 with 512 MB of RAM and a 1.70 GHz clock speed.

rational function of type (m, n) (i.e., numerator of degree $\leq m$ and denominator of degree $\leq n$) such that

$$\|e^z - r^*(z)\| = \inf_r \|e^z - r(z)\| = E_{mn},$$

where the infimum is over all real rational functions of type (m, n) and $\|\cdot\|$ denotes the supremum norm on \mathbb{R}^- . Standard results of approximation theory ensure that r^* exists and is unique [26]. A great deal has been learned about these functions in the years since the publication of an influential 1969 paper by Cody, Meinardus and Varga [5], which followed earlier work by Varga [48]. Cody et al. proved that E_{nn} decreases geometrically as $n \rightarrow \infty$, a conclusion that we now see also could follow from a suitably rigorous treatment of any of the quadrature estimates (1.4)–(1.6). In 1973 Schönhage [36] proved that E_{0n} converges at the rate

$$\lim_{n \rightarrow \infty} E_{0n}^{1/n} = \frac{1}{3}.$$

In 1983 Trefethen and Gutknecht observed numerically [46] that the rate for E_{nn} is²

$$(4.1) \quad \lim_{n \rightarrow \infty} E_{nn}^{1/n} = \frac{1}{9.28903\dots}.$$

Carpenter et al. confirmed this result by high-precision computations by the Remes algorithm [3]. A few years later, Magnus and independently Gonchar and Rakhmanov found exact expressions for this number, with Gonchar and Rakhmanov proving that the limit is indeed (4.1) [12, 13]. Magnus also found that $H = 1/9.28903\dots$ is a quantity associated with elliptic integrals first studied and even evaluated to 6-digit precision by Halphen in 1886 [13], and accordingly, it is now known as *Halphen's constant*. A much sharper result than (4.1) was conjectured by Magnus [24] and subsequently proved by Aptekarev [1]:

$$E_{nn} \sim 2H^{n+1/2} \quad (n \rightarrow \infty).$$

This “1/9 problem” has been the starting point of a great deal of fundamental work in approximation theory, building on the powerful connections with potential theory introduced by Gonchar and Rakhmanov [24, 32, 40]. As far as we can determine, it has been used only occasionally for applications in scientific computing [4, 16, 20, 39].

A comparison of (4.1) with (1.4)–(1.6) suggests that it should be possible to roughly double the speed of the algorithms of the last section:

$$(4.2) \quad \text{Best rational approximations: } O(9.28903^{-N}).$$

To put this idea into effect, we proceed in the reverse of the order of §3. There, we started with quadrature points and weights and interpreted these as poles

²LNT remembers his excitement when this number emerged from the computer one morning when he was a graduate student at Stanford (May 23, 1981). He sent Gutknecht a telegram in Switzerland: “9.28903?”.

and residues for a rational function. Now, we start with a rational function and interpret its poles and residues as points and weights for a quadrature formula. Specifically, suppose r is a rational function of type $(N-1, N)$ that is a good approximation to e^z on \mathbb{R}^- . Then we may expect the quantity I_N defined by (2.6) to be a good approximation to the integral I of (1.1). Let the poles and residues of r be z_1, \dots, z_N and c_1, \dots, c_N , as in (2.5). Then just as in Theorem 2.1, I_N will be equal to the integral (2.4) and the sum (2.3), and it is the sum we will use to evaluate it in practice. In effect we are applying a quadrature formula with weights defined by the second half of (2.3), that is, $w_k = -iNe^{-z_k}c_k$. The error formula (2.7) will hold as before.

There has been so much attention to the Cody–Meinaruds–Varga problem of best approximations with $m = n$ that we shall stay within this framework. Given N , let r^* be the best approximation to e^z on \mathbb{R}^- of type (N, N) , and define

$$(4.3) \quad r(z) = r^*(z) - r^*(\infty).$$

This is a rational function of type $(N-1, N)$, not quite optimal but nonetheless with deviation from e^z decreasing at the optimal rate $(9.28903\dots)^{-N}$ as $N \rightarrow \infty$. The function r^* can be computed by the Remes algorithm as in [3] and [5]; we have also been given very high accuracy computed results in a private communication from Alphonse Magnus. A simple practical alternative is to use the Carathéodory–Fejér method as proposed in the final section of [46] and analyzed at length in [24], based on the singular value decomposition of a Hankel matrix of Chebyshev coefficients of the function e^z transplanted from \mathbb{R}^- to $[-1, 1]$. For the easiest introduction to CF approximation, see [43]. Though the CF approach is in principle only approximate, Magnus has shown that the functions it produces differ from the true best approximations by about $O(56^{-N})$, which is so much smaller than 9.28^{-N} that the method can be regarded as exact in practice, apart from the effects of rounding errors; the error in exact arithmetic is below standard machine precision for $N \geq 9$, and for $N = 14$ it is about 10^{-26} [24]. We have found that by modifying codes of [44] to work with pole-residue form rather than polynomial coefficients, and carrying out the transplantation just mentioned by the formula $z = s(t-1)/(t+1)$ with $s \approx 9$, we can determine r^* by this method to high accuracy for $1 \leq N \leq 13$ and to good accuracy for $N = 14$, without the need for extended precision arithmetic. Figure 4.1 shows our MATLAB code, and Figure 4.2 shows the result it produces for $N = 14$ in about 0.2 sec. on our laptop.

Figure 4.3 is analogous to Figures 3.1–3.3, but for the near-best approximation r of (4.3). (A figure like this has previously been published by Sidje [39].) The first thing that strikes us is how similar this figure is to the earlier ones, though the methods of derivation have been entirely different. The high accuracy, however, is now achieved with the smaller value $N = 14$. It is striking that as before, the poles line up along a smooth arc curving leftward around the origin. As before, this arc scales in proportion to N , and its shape in the limit $N \rightarrow \infty$ was determined by Gonchar and Rakhmanov by methods of potential theory, though we are unaware of a simple formula; see [1, 12, 32].

```

function [zk,ck] = cf(n);
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 = exp(scl*(t-1)./(t+1e-16)); % exp(x) transpl. to [-1,1]
c = real(fft(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 vector
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); qk = zr(abs(zr)>1); % poles
qc = poly(qk); % coeffs of denominator
pt = rt.*polyval(qc,w); % numerator
ptc = real(fft(pt))/nf; % coeffs of numerator
ptc = ptc(n+1:-1:1); ck = 0*qk;
for k = 1:n % calculate residues
    q = qk(k); q2 = poly(qk(qk~=q));
    ck(k) = polyval(ptc,q)/polyval(q2,q);
end
zk = scl*(qk-1).^2./(qk+1).^2; % poles in z-plane
ck = 4*ck.*zk./(qk.^2-1); % residues in z-plane

```

Figure 4.1: MATLAB function to compute poles $\{z_k\}$ and residues $\{c_k\}$ by the Carathéodory–Fejér method for the type (n, n) best approximation r^* to e^z on \mathbb{R}^- [24, 43, 44, 46].




Figure 4.2: Error curve $e^z - r^*(z)$ against z for type $(14, 14)$ best approximation to e^z on \mathbb{R}^- , as computed by CF approximation by the program of Fig. 4.1. Note that the abscissa, the negative real axis in the z -plane, is displayed on a log scale. The dashed lines mark the ideal minimax error. The imperfect equioscillation is caused by rounding errors but will do less than one bit of damage in applications.




Figure 4.3: Error $|e^z - r^*(z)|$ in best rational approximation to e^z on $(-\infty, 0]$ with $N = 14$, with contours again as in Figs. 3.1–3.3. The dots mark the poles of r^* , and this time, the grey curve marks the limiting shape that the poles approach as $N \rightarrow \infty$, as shown by Gonchar and Rakhmanov [12]. Note that we get approximately the same accuracy as in Figs. 3.1–3.3 on \mathbb{R}^- , though N has been cut in half. The enhanced accuracy is mostly confined to a region near \mathbb{R}^- , however.

```

J = 100; h = 2/J; s = (-1+h:h:1-h)';           % 1D grid
[xx,yy] = meshgrid(s,s);                         % 2D grid
x = xx(:); y = yy(:);                           % 2D grid stretched to 1D
L = -gallery('poisson',J-1)/h^2;                 % 2D Laplacian
I = speye((J-1)^2);                             % identity
N = 14;                                           % degree of approx
[z,c] = cf(N);                                   % compute the approx
u = (1-x.^2).*(1-y.^2).*exp(x);                 % initial condition
v = zeros(size(u));
for k = 1:2:N
    v = v - c(k)*(z(k)*I-0.02*L)\u);             % quadrature via
                                                % sparse linear solves
end
v = 2*real(v);                                  % exploit symmetry
IN = v(J^2/2-J+1)                               % value at origin

```

Figure 4.4: Variant of the code fragment of Fig. 3.4 to compute the same result by means of best approximation. With $N = 14$ instead of $N = 32$, this code gets 14-digit accuracy in less than half the time.

Figure 4.4 repeats Figure 3.4, but now making use of best approximation. With $N = 14$ we now get the same accuracy obtained before with $N = 32$. The behavior of this code is remarkable: it evaluates $e^A v$ to 14-digit precision by means of just seven sparse matrix solves, each requiring only a fraction of a second.

5 Aside: Padé approximation and Gauss quadrature.

Over the years numerous connections have been made between inverse Laplace transforms, rational functions, and quadrature formulas. We shall not attempt a full survey but just point to a few ideas of this kind that are particularly close to those of the present paper. In (1.1), suppose that the integrand $e^z f(z)$ is dominated by its behavior near $z = 0$. Then instead of approximating e^z by a rational function on \mathbb{R}^- , one may choose to approximate it at the origin, i.e., by a Padé approximant. Methods like this have been investigated by Luke, Piessens, Rjabov, and Zakian, among others [23, 28, 29, 30]; the earliest reference seems to be a 1969 paper of Vlach [49]. In contrast to the functions we have considered in this paper, the Padé approximants have poles entirely in the right half-plane. Their disadvantage comes if values of z deep in the left half-plane are important, which may happen in an inverse Laplace transform (1.2) when t is close to zero.

Piessens points out in [28] that these Padé approximation methods have connections with a form of Gauss quadrature. Suppose we pose the problem of finding nodes and weights for a quadrature formula that is exact in (1.1) when f is a polynomial in z^{-1} of maximal order, with no constant term. The nodes for such a formula turn out to be the poles of certain orthogonal polynomials related to Padé approximants; see [28]. This Gauss quadrature approach to inverse Laplace transforms goes back at least to Salzer in 1959 [33], and related contributions are by Krylov, Piessens, Rjabov, Rodrigues, and Skoblya, among others [29, 31, 30].

6 Comparison of quadratures and best approximations.

We have seen that in comparison to Talbot-type quadratures, best rational approximations lead to methods for evaluation of (1.1) that are

1. Based on coefficients that are harder to compute, but
2. Up to twice as fast.

In this final section we comment further on these two important practical matters.

There is no doubt that the parabola, hyperbola, and cotangent quadrature formulas (3.1)–(3.3) are all very simple. Each can be implemented for arbitrary N in a few lines of computer code and negligible computation time. By contrast, best approximations demand special effort. One can precompute the required poles and residues, perhaps starting from the coefficients published by Carpenter et al. [3]. This gives programs that execute at top speed, but the programs are not very elegant, and they do not permit variable N . For example, if an

extended precision environment becomes available, one cannot increase the accuracy of such a code by simply adjusting N . Alternatively, the Carathéodory–Fejér method can be used to find the poles and residues, as in Figure 4.1. This is a general procedure for arbitrary N , but it is based on complex mathematical tools whose implementation on a computer requires fractions of a second, not fractions of a millisecond. For many applications, one might prefer to skip these complications and just use a Talbot contour.

Still, a factor of two in efficiency is sometimes very attractive. In Figures 3.4 and 4.4 we have seen that in some applications this speedup is realized, and there are undoubtedly many other such cases. For example, suppose one wants to compute the gamma function $\Gamma(t)$ for real or complex values of t by means of the Hankel integral

$$(6.1) \quad \frac{1}{\Gamma(t)} = \frac{1}{2\pi i} \int_{\Gamma} z^{-t} e^z dz.$$

(Of course, the contour Γ and the function Γ in this formula are different objects.) As shown in [35], this approach is eminently practical, coming within a small factor in efficiency of the best known methods for evaluating $\Gamma(t)$. Both Talbot contours and best approximations for (6.1) are effective, and the latter are twice as efficient for a wide range of values of t . An excellent gamma function code could be based on best approximations implemented with precomputed poles and residues. The Carathéodory–Fejér approach is also effective in a context where many values of t are to be treated at once.

On the other hand, the factor of two advantage of best approximations may begin to be lost when the problem is perturbed. A conspicuous example of such problems comes in any inverse Laplace transform (1.2) where one would like results for multiple values of t . For such situations, as was pointed out in §1, a change of variables offers two mathematically equivalent integrals:

$$(6.2) \quad G(t) = \frac{1}{2\pi i} \int_{\Gamma} e^{st} g(s) ds = \frac{1}{2\pi i} \int_{\Gamma} e^z g(z/t) t^{-1} dz,$$

or in the special case $u_t = Au$, as in (3.4),

$$(6.3) \quad u(t) = \frac{1}{2\pi i} \int_{\Gamma} e^{st} (s - A)^{-1} u(0) ds = \frac{1}{2\pi i} \int_{\Gamma} e^z (z/t - A)^{-1} t^{-1} u(0) dz.$$

Now in each of (6.2) and (6.3), the integral on the right, containing just e^z , is the one to which the quadrature contours and best approximations analyzed in this paper are tuned. The integral on the left with e^{st} , however, has the big advantage that $g(s)$ or $(s - A)^{-1}$ is independent of t . This raises the possibility that this quantity might be evaluated once and for all at a fixed set of quadrature points s_k , and the results combined with t -dependent coefficients to provide results for various t . For the problem of Figures 3.4 and 4.4, for example, this would offer the hope that whereas computing a result for the first value of t would take many seconds expended in large matrix solves, each subsequent value could be computed almost instantaneously by recombination of the same vectors with new coefficients. Impressive algorithms that realize this vision, based on




Figure 6.1: Comparison of errors in solution of the problem (3.4) for various t via cotangent contours (dashed) and best approximations (solid) with $N = 12$. The best approximation method is far more accurate for $t = 1$, but its advantage lessens for other values of t .

quadratures explicitly adapted to ranges of t , have been developed by López-Fernández, Palencia, and Schädle [18, 19].

Such methods appear to work better for Talbot contours than best approximations: the latter lose accuracy more quickly as t is varied. Figure 6.1 illustrates this phenomenon. This experiment repeats that of Figures 3.4 and 4.4, but now for a range of values of t , all with $N = 12$. (For speed, we also increase Δx and Δy to $1/10$, with negligible effect on the results. Here we use a cotangent contour rather than a parabola, but this change has little effect.) In each case we start with a formula tuned to solving the problem (3.4) at $t = 1$, and adjust the quadrature coefficients as indicated in the first integral of (6.3) for other values of t . The best approximation method gets 12-digit accuracy at $t = 1$, but its accuracy diminishes rapidly as t is varied. The cotangent contour gets only 7 digits, but that accuracy is maintained throughout the interval $[0.6, 1.2]$.

In summary, best approximations are more accurate than Talbot quadratures, but more fragile. We believe that this relationship is analogous to the relationship between Gauss quadrature on $[-1, 1]$ and the simpler method of Clenshaw–Curtis quadrature, and will develop this analogy in a forthcoming publication [45].

Acknowledgments.

We are grateful for advice from Martin Gutknecht, Nick Higham, and our fellow contour followers Maria López-Fernández, César Palencia, and Achim

Schädle. The ‘1/9’ wizard Alphonse Magnus also helped in many ways, and provided the program that drew the limiting curve shown in gray in Figure 4.3.

REFERENCES

1. A. I. Aptekarev, *Sharp constants for rational approximation of analytic functions*, Sb. Math. 193 (2002), pp. 1–72.
2. D. Calvetti, E. Gallopoulos, and L. Reichel, *Incomplete partial fractions for parallel evaluation of rational matrix functions*, J. Compt. Appl. Math. 59 (1995), pp. 349–380.
3. A. J. Carpenter, A. Ruttan, and R. S. Varga, *Extended computations on the ‘1/9’ conjecture in rational approximation theory*, in P. R. Graves-Morris, E. B. Saff, and R. S. Varga, eds., *Rational Approximation and Interpolation*, Lect. Notes in Math. 1105 (1984), Springer, pp. 383–411.
4. J. C. Cavendish, W. E. Culham and R. S. Varga, *A comparison of Crank–Nicolson and Chebyshev rational methods for numerically solving linear parabolic equations*, J. Comp. Phys. 10 (1972), pp. 354–368.
5. W. J. Cody, G. Meinardus and R. S. Varga, *Chebyshev rational approximations to e^{-x} in $[0, +\infty)$ and applications to heat-conduction problems*, J. Approx. Th.z 2 (1969), pp. 50–65.
6. E. Gallopoulos, *A partial fraction decomposition approach to improved efficiency of some parabolic solvers*, Tech. Rep. 874, Ctr. for Supercomputing Res. Dev., May 1989.
7. E. Gallopoulos and Y. Saad, *On the parallel solution of parabolic equations*, Proc. 1989 ACM Internat. Conf. on Supercomputing, Heraklion, Greece, ACM, 1989, pp. 17–28.
8. E. Gallopoulos and Y. Saad, *Efficient solution of parabolic equations by Krylov approximation methods*, SIAM J. Sci. Statist. Comp. 13 (1992), pp. 1236–1264.
9. I. P. Gavriluk and V. L. Makarov, *Exponentially convergent parallel discretization methods for the first order evolution equations*, Comp. Meth. Appl. Math. 1 (2001), pp. 333–355.
10. I. P. Gavriluk and V. L. Makarov, *Exponentially convergent algorithms for the operator exponential with applications to inhomogeneous problems in Banach spaces*, preprint, 2004.
11. A. Gil, J. Segura, and N. M. Temme, *Computing special functions by using quadrature rules*, Numer. Algs. 33 (2003), pp. 265–275.
12. A. A. Gonchar and E. A. Rakhmanov, *Equilibrium distributions and degree of rational approximation of analytic functions*, Mat. Sb. 134 (1987), pp. 306–352; English transl. in Math. USSR-Sb. 62 (1989).
13. G. H. Halphen, *Traité des fonctions elliptiques et de leurs applications. I Théorie des fonctions elliptiques et de leurs développement en séries*, Gauthier-villars, Paris, 1886. <http://moa.cit.cornell.edu/>
14. A.-K. Kassam, *Solving reaction-diffusion equations ten times faster*, Numer. Anal. Rep. NA 03/16, Oxford U. Computing Lab., 2003.
15. A.-K. Kassam and L. N. Trefethen, *Fourth-order time-stepping for stiff PDE*, SIAM J. Sci. Comp. 26 (2003), pp. 1214–1233.
16. J. D. Lawson and D. A. Swayne, *High-order near best uniform approximations to the solution of heat conduction problems*, Information Processing 80, North-Holland, Amsterdam, 1980, pp. 741–746.

17. M. López-Fernández, Ch. Lubich, C. Palencia and A. Schädle, *Fast Runge-Kutta approximation of inhomogeneous parabolic equations*, ZIB-Report 05-10, Zuse-Institut Berlin, 2005.
18. M. López-Fernández and C. Palencia, *On the numerical inversion of the Laplace transform in certain holomorphic mappings*, *Appl. Numer. Math.* 51 (2004), 289–303.
19. M. López-Fernández, C. Palencia, and A. Schädle, *A spectral order method for inverting sectorial Laplace transforms*, ZIB-Report 05-26, Zuse-Institut Berlin, 2005.
20. Y. Y. Lu, *Exponentials of symmetric matrices through tridiagonal reductions*, *Lin. Alg. Applics.* 279 (1998), pp. 317–324.
21. Ch. Lubich & A. Schädle, *Fast convolution for nonreflecting boundary conditions*, *SIAM J. Sci. Comp.* 24 (2002), pp. 161–182.
22. Y. L. Luke, *The Special Functions and Their Approximations*, v. 1–2, Academic Press, New York, 1969.
23. Y. L. Luke, *Error estimation in numerical inversion of Laplace transforms using Padé approximation*, *J. Franklin Inst.* 305 (1978), 259–273.
24. A. P. Magnus, *Asymptotics and super asymptotics of best rational approximation error norms for the exponential function (the ‘1/9’ problem) by the Carathéodory–Fejér method*, in A. Cuyt, et al, eds., *Nonlinear Methods and Rational Approximation. II*, Kluwer, Dordrecht, 1994. pp. 173–185.
25. W. McLean and V. Thomée, *Time discretization of an evolution equation via Laplace transforms*, *IMA J. Numer. Anal.* 24 (2004), pp. 439–463.
26. G. Meinardus, *Approximation of Functions: Theory and Numerical Methods*, Springer, 1967.
27. C. Moler and C. Van Loan, *Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later*, *SIAM Review* 45 (2003), pp. 3–49.
28. R. Piessens, *On a numerical method for the calculation of transient responses*, *J. Franklin Inst.* 292 (1971), 57–64.
29. R. Piessens, *Gaussian quadrature formulas for the numerical integration of Bromwich’s integral and the inversion of the Laplace transform*, *J. Eng. Math.* 5 (1971), pp. 1–9.
30. V. M. Rjabov, *Application of Padé approximations to Laplace transformation inversion* (Russian), *Vestnik Leningrad. Univ. Math.* 2 (1970), p. 119.
31. A. J. Rodrigues, *Properties of constants for a quadrature formula to evaluate Bromwich’s integral*, *J. Inst. Maths. Applics.* 18 (1976), pp. 49–56.
32. E. B. Saff and V. Totik, *Logarithmic Potentials with External Fields*, Springer, 1997.
33. H. E. Salzer, *Orthogonal polynomials arising in the numerical evaluation of inverse Laplace transforms*, *Math. Comp.* 9 (1955), p. 164–177.
34. A. Schädle, M. López-Fernández and Ch. Lubich, *Fast and oblivious convolution quadrature*, ZIB-Report 05-09, Zuse-Institut Berlin, 2005.
35. T. Schmelzer, *Computing the gamma function using contour integrals and rational approximations*, manuscript in preparation.
36. A. Schönhage, *Zur rationalen Approximierbarkeit von e^{-x} über $[0, \infty)$* , *J. Approx. Th.* 7 (1973), pp. 395–398.
37. D. Sheen, I. H. Sloan, and V. Thomée, *A parallel method for time-discretization of parabolic problems based on contour integral representation and quadrature*, *Math. Comp.* 69 (2000), pp. 177–195.

38. D. Sheen, I. H. Sloan, and V. Thomée, *A parallel method for time discretization of parabolic equations based on Laplace transformation and quadrature*, IMA J. Numer. Anal. 23 (2003), pp. 269–299.
39. R. B. Sidje, *Expokit: A software package for computing matrix exponentials*, ACM Trans. Math. Softw. 24 (1998), pp. 130–156.
40. H. Stahl and V. Totik, *General Orthogonal Polynomials*, Cambridge U. Press, 1992.
41. A. Talbot, *The accurate numerical inversion of Laplace transforms*, J. Inst. Maths. Applics. 23 (1979), pp. 97–120.
42. N. M. Temme, *Special Functions*, Wiley, New York, 1996.
43. L. N. Trefethen, *Chebyshev approximation on the unit disk*, in H. Werner et al., eds., *Computational Aspects of Complex Analysis*, D. Reidel Publishing, Dordrecht, 1983, pp. 309–323.
44. L. N. Trefethen, *Matlab programs for CF approximation*, in *Approximation Theory V*, Academic Press, 1986, pp. 599–602.
45. L. N. Trefethen, *Is Gauss quadrature better than Clenshaw–Curtis?*, manuscript in preparation.
46. L. N. Trefethen and M. H. Gutknecht, *The Carathéodory–Fejér method for real rational approximation*, SIAM J. Numer. Anal. 20 (1983), pp. 420–436.
47. L. N. Trefethen and J. A. C. Weideman, *The fast trapezoid rule in scientific computing*, manuscript in preparation.
48. R. S. Varga, *On higher order stable implicit methods for solving parabolic partial differential equations*, J. Math. Phys. 40 (1961), pp. 220–231.
49. J. Vlach, *Numerical method for transient responses of linear networks with lumped, distributed or mixed parameters*, J. Franklin Inst. 288 (1969), 99–113.
50. J. A. C. Weideman, *Computing special functions via inverse Laplace transforms*, ICNAAM-2005, 2005.
51. J. A. C. Weideman, *Optimizing Talbot’s contours for the inversion of the Laplace transform*, Numer. Anal. Rep. NA 05/05, Oxford U. Computing Lab., 2005.
52. V. Zakian, *Properties of I_{MN} and J_{MN} approximants and applications to numerical inversion of Laplace transforms and initial value problems*, J. Math. Anal. Applic. 50 (1975), pp. 191–222.