

Orthogonal Rational Functions on a Semi-infinite Interval

JOHN P. BOYD

*Department of Atmospheric and Oceanic Science,
University of Michigan, Ann Arbor, Michigan 48109*

Received April 11, 1986; revised September 12, 1986

By applying a mapping to the Chebyshev polynomials, we define a new spectral basis: the "rational Chebyshev functions on the semi-infinite interval," denoted by $TL_n(y)$. Continuing earlier work by the author and by Grosch and Orszag, we show that these rational functions inherit most of the good numerical characteristics of the Chebyshev polynomials: orthogonality, completeness, exponential or "infinite order" convergence, matrix sparsity for equations with polynomial coefficients, and simplicity. Seven numerical examples illustrate their versatility. The "Charney" stability problem of meteorology, for example, is solved to show the feasibility of applying spectral methods to a semi-infinite atmosphere. For functions that are singular at both endpoints, such as $K_1(y)$, one may combine rational Chebyshev functions with a preliminary mapping to obtain a single, exponentially convergent expansion on $y \in [0, \infty]$. Finally, we successfully generalize the WKB method to obtain, for the J_0 Bessel function, an amplitude-phase approximation which is convergent rather than asymptotic and is accurate not merely for large y but for all y , even the origin. © 1987 Academic Press, Inc.

1. INTRODUCTION

In an earlier paper, Boyd [1] combined an algebraic mapping from $[-1, 1]$ to $[-\infty, \infty]$ with Chebyshev polynomials to provide a new set of basis functions for solving differential equations on an infinite interval. In this follow-up, a similar set of mapped Chebyshev polynomials are defined for the semi-infinite interval, $y \in [0, \infty]$.

The earliest appearance of Chebyshev polynomials in tandem with an algebraic change-of-coordinate was Grosch and Orszag [2]. Boyd [3] used the method of steepest descent to provide guidelines for optimizing the mapping. Boyd [4-6] discusses alternatives to mappings and also the use of a change-of-coordinate to deal with functions with singularities on or near the expansion interval.

Despite this prior work, there are still many issues left for the present article. Programming mechanics, boundary conditions at both the origin and at infinity, functions that decay algebraically rather than exponentially as $y \rightarrow \infty$, and numerical examples will be discussed in turn. The new basis functions, denoted by $TL_n(y)$, are defined by

$$TL_n(y) \equiv T_n(x) \equiv \cos(nt), \quad (1.1)$$

where L is a constant map parameter and the three coordinates are related by

$$y \equiv L(1+x)/(1-x), \quad x \equiv (y-L)/(y+L) \quad (1.2)$$

$$y \equiv L \cot^2(t/2), \quad t = 2 \operatorname{arccot}([y/L]^{1/2}). \quad (1.3)$$

It may seem a little pretentious to introduce a new symbol for functions which are merely the Chebyshev polynomials in disguise. However, the Chebyshev polynomials in x are themselves the images of the terms of a Fourier cosine series under the map $x = \cos(t)$. In both cases, the change-of-coordinates alters the shape and properties of the basis set so profoundly that it is probably best to regard the result of the mapping as a new species, deserving its own symbol.

Even so, the connection with the Fourier cosine functions is nonetheless very important both for analyzing convergence, as done in the Appendix, and also for writing practical computer programs. To avoid confusion as we leap from one coordinate to another, we shall adopt the convention that $y \in [0, \infty]$ is the argument of the $TL_n(y)$, $x \in [-1, 1]$ is the argument of the ordinary Chebyshev polynomials, and $t \in [0, \pi]$ ("t" for "trigonometric") is the argument of the cosines. We are free to calculate in whichever of these three coordinates is most convenient.

We shall refer to the $TL_n(y)$ as the "rational Chebyshev functions on a semi-infinite interval"; Table I lists the explicit form of the first nine functions for $L = 1$. The graphs shown in Fig. 1 apply for any L if we replace y by (y/L) ; varying L alters the *width* of the functions without changing their *shape*. Because the dependence on L is so simple and also for notational simplicity, we shall often suppress the map parameter by writing $TL_n(y)$ [instead of the more correct $TL_n(y; L)$] and by setting $L = 1$ in some figures and tables.

By merely changing the variable in the usual orthogonality integral for the cosines, one can show that the rational Chebyshev functions are orthogonal:

$$\int_0^\infty TL_m(y; L) TL_n(y; L) (L^{1/2}/[(y+L)y^{1/2}]) dy = \begin{cases} \pi, & m = n = 0 \\ 0, & m \neq n \\ \pi/2, & m = n > 0. \end{cases}$$

TABLE I

n	$TL_n(y)$
0	1
1	$(y-1)/(y+1)$
2	$(y^2-6y+1)/(y+1)^2$
3	$(y^3-15y^2+15y-1)/(y+1)^3$
4	$(y^4-28y^3+70y^2-28y+1)/(y+1)^4$
5	$(y^5-45y^4+210y^3-210y^2+45y-1)/(y+1)^5$
6	$(y^6-66y^5+495y^4-924y^3+495y^2-66y+1)/(y+1)^6$
7	$(y^7-91y^6+1001y^5-3003y^4+3003y^3-1001y^2+91y-1)/(y+1)^7$
8	$(y^8-120y^7+1820y^6-8008y^5+12870y^4-8008y^3+1820y^2-120y+1)/(y+1)^8$

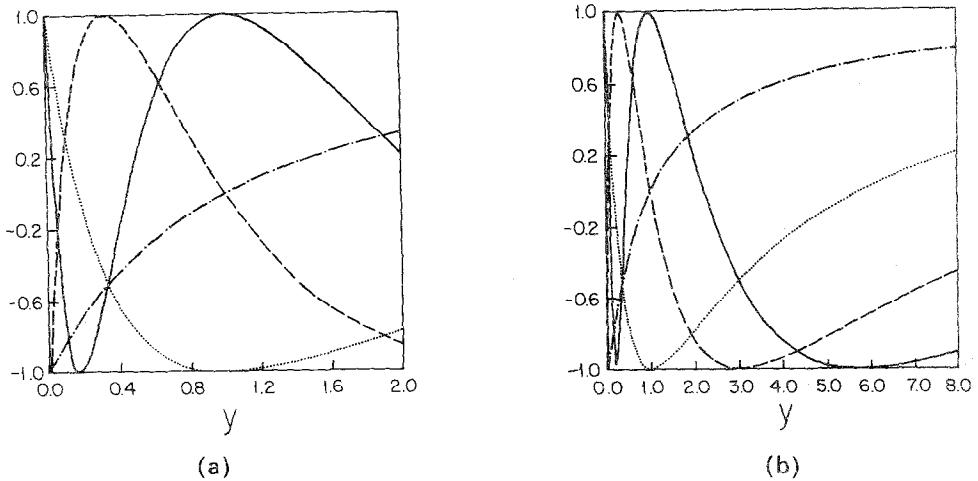


FIG. 1. Graphs of $TL_1(y)$ [dot-dash], $TL_2(y)$ [dotted], $TL_3(y)$ [dashed], and $TL_4(y)$ [solid] on (a) $y \in [0, 2]$ and (b) $y \in [0, 8]$. (The map parameter $L = 1$.)

However, this orthogonality relationship has been listed principally to explain that one never needs it! To apply Galerkin's method, it is simpler to convert to the trigonometric argument t and then evaluate the needed integrals as in Boyd [17].

In the next section, it will be shown that the rational Chebyshev functions on the semi-infinite interval, like their cousins on the infinite interval described in Boyd [1], give *banded* Galerkin matrices when the coefficients of the differential equation are *polynomials* or *rational* functions of y . For general problems, however, the pseudospectral method (alias "collocation" algorithm) is just as efficient and considerably easier to program. In the next section, the mechanics of this and of converting the problem from y to t will be described.

Section 3 explains that, again as true for the $TB_n(y)$, the boundary condition at infinity is usually "natural" in the sense that it need not be explicitly imposed on the basis functions. However, the Galerkin's or pseudospectral matrix must be modified by adding a row to force the numerical solution to satisfy the boundary condition at $y=0$ unless the differential equation is singular at that point also—in which case the boundary conditions may be "natural" at both endpoints.

Section 4 discusses functions which decay algebraically (as opposed to exponentially) as $y \rightarrow \infty$. It is still possible to obtain exponentially fast convergence for such $f(y)$ provided that the function either does not oscillate or has the frequency of oscillations tend to zero as $y \rightarrow \infty$.

The fifth section is the heart of the paper: seven numerical examples that illustrate each of the theoretical ideas in the article. Section 6 is a summary and prospectus. The details of the convergence theory, including the calculation of the contours of "equiconvergence" in the complex plane, have been banished to the Appendix because these concepts are only occasionally useful in applications. It is

nonetheless satisfying that thanks to the connection between the orthogonal rational Chebyshev functions and ordinary Fourier series expressed by (1.1), it is possible to give a complete discussion of the theory of $TL_n(y)$ expansions.

2. MECHANICS: COORDINATE CONVERSION, GRID POINTS, AND BANDEDNESS

In the author's experience, the simplest way to set up the Galerkin's or pseudospectral matrix for solving a boundary or eigenvalue problem is to exploit the connection with a Fourier cosine series. Table II lists the formulas for transforming derivatives, which were computed using the algebraic manipulation language REDUCE by iterating the elementary rule

$$d/dy = -\{\sin^2(t/2) \tan(t/2)\} d/dt. \quad (2.1)$$

TABLE II

Transformations of Derivatives for the Mapping $y = L \cot^2(t/2)$ which Converts a Rational-Chebyshev Series in $TL_n(y)$ into a Fourier Cosine Series in $\cos(nt)$. L Is a Constant, the "Map Parameter," $u_y = \{-\sin^3(t/2)/[L \cos(t/2)]\} u_t$, $u_{yy} = \{\sin^2(t/2)/(2L^2 \cos^3(t/2))\} \{2 \cos(t/2) \sin(t/2) u_{tt} + [3 - 2 \sin^2(t/2)] u_t\}$

u_t	u_{tt}	u_{ttt}	u_{4t}	u_{5t}	u_{6t}
u_{yyy}					
$-8 \sin^4$	$12 \cos \sin^3$	$4 \sin^4$			
$+20 \sin^2$	$-18 \cos \sin$	$-4 \sin^2$			
-15		$[\times \sin^7(t/2)/\{4L^3 \cos^5(t/2)\}]$			
u_{4y}					
$-48 \sin^6$	$88 \cos \sin^5$	$48 \sin^6$	$-8 \cos \sin^5$		
$+168 \sin^4$	$-232 \cos \sin^3$	$-120 \sin^4$	$+8 \cos \sin^3$		
$-210 \sin^2$	$+174 \cos \sin$	$+72 \sin^2$			
$+105$		$[\times \sin^9/\{8L^4 \cos^7\}]$			
u_{5y}					
$-384 \sin^8$	$800 \cos \sin^7$	$560 \sin^8$	$-160 \cos \sin^7$	$-16 \sin^8$	
$+1728 \sin^6$	$-2960 \cos \sin^5$	$-2080 \sin^6$	$+400 \cos \sin^5$	$+32 \sin^6$	
$-3024 \sin^4$	$+3900 \cos \sin^3$	$+2660 \sin^4$	$-240 \cos \sin^3$	$-16 \sin^4$	
$+2520 \sin^2$	$-1950 \cos \sin$	$-1140 \sin^2$			
-945		$[\times \sin^{11}/\{16L^5 \cos^9\}]$			
u_{6y}					
$3840 \sin^{10}$	$-8768 \cos \sin^9$	$-7200 \sin^{10}$	$+2720 \cos \sin^9$	$+480 \sin^{10}$	$-32 \cos^5 \sin^5$
$-21120 \sin^8$	$+41536 \cos \sin^7$	$+34800 \sin^8$	$-10240 \cos \sin^7$	$-1680 \sin^8$	
$+47520 \sin^6$	$-76808 \cos \sin^5$	$-65040 \sin^6$	$+13160 \cos \sin^5$	$+1920 \sin^6$	
$-55440 \sin^4$	$+67440 \cos \sin^3$	$+56160 \sin^4$	$-5640 \cos \sin^3$	$-720 \sin^4$	
$+34650 \sin^2$	$-25290 \cos \sin$	$-18720 \sin^2$			
-10395		$[\times (-1) \sin^{13}/\{32L^6 \cos^{11}\}]$			

Note. The term $[\times \sin^q/\{2^n L^p \cos^r\}]$ denotes that all entries in the table for the previous derivative must be multiplied by this common factor. Note also that all sines and cosines in the table have arguments of $(t/2)$, not t .

The main program may still use the physical coordinate y , however: Table III is a FORTRAN fragment that shows the trigonometric argument t can be entirely confined to a single subroutine with just eight executable statements (which evaluates the $TL_n(y)$ and its derivatives at a given y).

The pseudospectral grid in y is simply the image under the mapping of an evenly spaced Fourier grid

$$t_i = \pi(2i + 1)/(2N + 2), \quad i = 0, \dots, N \quad (2.2)$$

$$y_i = L \cot^2(t_i/2). \quad (2.3)$$

A differential equation of the form

$$Hu = f, \quad (2.4)$$

where H is a linear operator and $u(y)$ and $f(y)$ are functions, is pseudospectrally discretized to the matrix problem

$$\mathbf{H}\mathbf{a} = \mathbf{f}, \quad (2.5)$$

where \mathbf{a} is the column vector whose elements are the coefficients of the $TL_n(y)$ series for $u(y)$ and where

$$\mathbf{H}_{ij} \equiv (H TL_j[y])|_{y=y_i}, \quad i, j = 0, \dots, N \quad [\text{pseudospectral}] \quad (2.6)$$

$$\mathbf{f}_i \equiv f(y_i), \quad i = 0, \dots, N. \quad (2.7)$$

TABLE III

	SUBROUTINE BASIS(<i>N</i> , <i>Y</i> , <i>L</i> , <i>TL</i> , <i>TLY</i> , <i>TLYY</i>)
C	A FORTRAN subroutine for computing the rational Chebyshev
C	functions.
C	INPUT: <i>N</i> (>= 0, the degree of the basis function)
C	<i>Y</i> (on interval [0, Infinity])
C	<i>L</i> (map parameter)
C	OUTPUT: <i>TL</i> , <i>TLY</i> , <i>TLYY</i> are the <i>N</i> th Chebyshev functions
C	and their first two derivatives.
C	<i>T</i> is the trigonometric argument (on interval [0, pi]).
	REAL <i>L</i>
	<i>T</i> = 2. * ACOT(SQR(<i>Y</i> / <i>L</i>))
	<i>TL</i> = COS(FLOAT(<i>N</i>)* <i>T</i>)
C	Derivatives of cos (<i>Nt</i>) with respect to <i>t</i> :
	<i>PT</i> = -FLOAT(<i>N</i>) * SIN(FLOAT(<i>N</i>)* <i>T</i>)
	<i>PTT</i> = -FLOAT(<i>N</i>)*FLOAT(<i>N</i>) * <i>TL</i>
C	Conversion of <i>t</i> -derivatives into <i>x</i> -derivatives
	<i>C</i> = COS(<i>T</i> /2.)
	<i>S</i> = SIN(<i>T</i> /2.)
	<i>TLY</i> = - <i>PT</i> * <i>S</i> * <i>S</i> * <i>S</i> / (<i>L</i> * <i>C</i>)
	<i>TLYY</i> = (2.* <i>C</i> * <i>S</i> * <i>PTT</i> + (3. - 2.* <i>S</i> * <i>S</i>)* <i>PT</i>) * (<i>S</i> **5/(2.* <i>L</i> * <i>L</i> * <i>C</i> **3))
	RETURN
	END

Galerkin's method replaces (2.6) and (2.7) by

$$\mathbf{H}_{ij} \equiv (TL_i, H TL_j), \quad i, j = 0, \dots, N \quad [\text{Galerkin's}] \quad (2.8)$$

$$\mathbf{f}_i \equiv (TL_i, f), \quad i = 0, \dots, N, \quad (2.9)$$

where the inner product is defined by [for arbitrary $g(y)$, $h(y)$]

$$(g, h) \equiv \int_0^\pi g(L \cot^2[t/2]) h(L \cot^2[t/2]) dt. \quad (2.10)$$

Since nonlinear solutions can be calculated by the Newton–Kantorovich method (among others), solving a linear differential equation at each iteration [7, 8], we discuss only linear problems in this article.

As explained in Gottlieb and Orszag [9], it is rarely worth the bother to replace *evaluation* (pseudospectral) by *integration* (Galerkin's) because the two methods give almost equal accuracy for the same N , but collocation is simpler. There is one exception: when the coefficients of the differential equation are *polynomial* or *rational* in y , the coefficients will be rational functions of $\sin(t)$ and $\cos(t)$. After the denominators have been multiplied out, the Galerkin matrix elements (2.8) and (2.9) will be integrals of trigonometric basis functions multiplied by trigonometric polynomials.

By using elementary trigonometric identities, one can show, as done in [1], that if $P_k(t)$ is a trigonometric polynomial of degree k , then letting $\phi_n(t)$ denote either $\sin(nt)$ or $\cos(nt)$ as appropriate,

$$P_k(t) \phi_n(t) = \sum_{m=n-k}^{n+k} p_m \phi_m(t). \quad (2.11)$$

If k is the highest degree of the trigonometric polynomials that appear in the differential equation (after clearing denominators), then (2.11)—and the orthogonality of the sines and coefficients—imply that the Galerkin's matrix will be *banded* in the sense that

$$\mathbf{H}_{ij} = 0 \quad \text{if} \quad |i - j| > k. \quad (2.12)$$

This is an extremely important property because banded matrices can be solved in a small fraction of the expense of inverting full matrices. Thus in spite of its greater complexity, Galerkin's method is still a key ingredient of many semi-implicit time-stepping codes. Cain *et al.* [10] seem to have been the first to appreciate this point; they also mapped Fourier series to a semi-infinite or infinite domain, but did not make the connection with Chebyshev polynomials that is central to [1] and [2].

The present paper will not dwell on the issue of bandedness because the precise bandwidth varies from problem to problem. Many older spectral papers went to great lengths to apply recurrence relations and elaborate algebra to analytically evaluate the Galerkin's matrix. This is reasonable only when the end result is a

sparse matrix and the programmer intends to exploit the bandedness. Otherwise, the pseudospectral algorithm is simpler to program and easier to adapt from one differential equation to another. All the numerical examples given in Section 5 were solved pseudospectrally.

3. BOUNDARY CONDITIONS: IMPLICIT VERSUS EXPLICIT, NATURAL VERSUS ESSENTIAL

When applying Chebyshev polynomials to solve a differential equation on $[-1, 1]$, it is usually necessary to modify the matrix H defined above by reserving two rows to impose the boundary conditions. The only exception is when the differential equation is *singular* at one or both endpoints. One might suppose that such singularities would cause enormous difficulties, but they often make the problem simpler: the desired solution is the only linearly independent solution which is analytic at the singularity. One can therefore apply (2.6) without modification; the Chebyshev series will converge to the bounded solution without imposition of explicit constraints.

In the language of finite element methods, boundary conditions which must be *explicitly* imposed on the numerical solution are called "essential." When the *unconstrained* sum of basis functions will converge to the correct boundary behavior—although not exactly satisfying the boundary conditions for any finite N —the boundary conditions are said to be "natural."

It is possible to construct differential equations which are analytic at infinity, but the *usual* case is that all solutions of the differential equation except the desired one blow up as $y \rightarrow \infty$. The boundary condition at infinity is then "natural."

At $y=0$, it may be necessary to reserve one row of the pseudospectral matrix to impose an "essential" condition such as $u(0)=1$. We shall show how this is done in Section 5 in the context of the two numerical examples for which it is necessary (the BVP whose solution is $K_1(y)$ and the "Charney" problem of meteorology). On the other hand, the differential equation may be singular at the origin as well as at infinity; this is true for the Laguerre and Whittaker eigenproblems of Section 5. Thus, the normal cases are "natural-natural" or "essential-natural" for the boundary conditions at the origin and at infinity, respectively.

When both boundary conditions are natural, (2.6) applies without modification. The only constraints on the boundary behavior are the *implicit* constraints imposed by the differential equation.

4. SOLUTIONS WHICH DECAY ALGEBRAICALLY WITH y OR ASYMPTOTE TO A CONSTANT

Laguerre functions, which are an alternative to the rational Chebyshev functions, also have series coefficients that decay exponentially with n if $f(y)$, the solution

being expanded, decays *exponentially* with y as $y \rightarrow \infty$. However, the Laguerre functions are the product of $\exp(-y/2)$ with the n th Laguerre polynomial, so all these basis functions decay exponentially with y for large y . If $f(y)$ does not decay exponentially, too—if instead, $f(y)$ decreases as some inverse power of y , or perhaps asymptotes to a constant—then the mismatch in asymptotic behavior between the function and the Laguerre basis will cause the Laguerre coefficients to decrease *algebraically* with n .

Because the $TL_n(y)$ and its derivatives all decrease as inverse powers of y (or tend to a constant), they retain the property of exponentially rapid convergence even for such slowly decaying $f(y)$, and are therefore the method of choice for this situation. In Section 5, we shall discuss a couple of examples where we exploit this strength of the rational Chebyshev functions to compute very accurate approximations for the ground state eigenvalue of the quantum anharmonic oscillator and for the Bessel function $J_0(y)$.

There is, alas, a restriction. If $f(y)$ oscillates, then the frequency of these oscillations must tend to 0 as $y \rightarrow \infty$ unless $f(y)$ decreases exponentially fast with y . For example,

$$f(y) = \exp(-y) \cos(y) \quad (4.1)$$

has an infinite number of sign changes on $[0, \infty]$. To obtain an approximation with a uniformly small *absolute* error ε , however, it is only necessary to resolve the oscillations on $y \in [0, -\log(\varepsilon)]$ because $|f(y)| < \varepsilon$ for all larger y . The $TL_n(y)$ series for (4.1) converges exponentially fast with n .

Unfortunately,

$$f(y) \equiv (1+y)^{1/2} J_0(y) \sim a(y) \cos(y - \pi/4 + \phi[y]), \quad y \rightarrow \infty \quad (4.2)$$

is a counterexample, and its rational Chebyshev series is useless for all practical purposes. The problem is that the amplitude of the oscillations in the Bessel function does *not* decay. If the spectral series cannot resolve a region of sinusoidal variation, even one for very, very large y , then this will give an $O(1)$ absolute error. A *uniform* approximation on $y \in [0, \infty]$ through a direct expansion is impossible.

Another way of (i) reaching the same conclusion and (ii) testing a function $f(y)$ for convergence is to apply ordinary *Fourier* convergence theory to $f(L \cot^2(t/2))$. (The Fourier coefficients of this function of t are *identical* with the TL_n coefficients of $f(y)$.) For (4.2), for example, $f(L \cot^2(t/2)) \sim [\text{constant}] \cos(4L/t^2 + \text{constant})$ as $t \rightarrow 0$. Since $\cos(4L/t^2)$ lacks bounded derivatives of any order at $t = 0$, it follows that the Fourier (and TL_n) coefficients of the Bessel function cannot decrease even as fast as $O(1/n)$.

Fortunately, however, there are, at least sometimes, alternatives: for (4.2), the amplitude $a(y)$ and the derivative of the phase $\phi_y(y)$ both asymptote to constants, and these two functions do have rapidly convergent $TL_n(y)$ series. An explanation of how to extract *two* expansions from *one* $f(y)$ will be given in Section 5—but it requires a special trick. A single expansion in terms of any standard basis set will fail for (4.2).

5. NUMERICAL EXAMPLES

In this section, we discuss seven problems, which subdivide into a simple boundary value problem, three eigenproblems, and three expansions of known functions. If numerical analysis is an art and not merely a science, it is hoped that this art show will illustrate the little tricks that are important in applications.

Example One: K_1 Bessel Boundary Value Problem

$$y^2 u_{yy} + y u_y - (1 + y^2) u = 0, \quad y \in [1, \infty] \quad (5.1)$$

$$u(1) = K_1(1) \quad (5.2)$$

with the exact solution

$$u(y) = K_1(y), \quad (5.3)$$

where $K_1(y)$ is the usual "imaginary" Bessel function. This boundary value problem requires two slight modifications to the standard procedures explained earlier.

First, if the problem is posed on the interval $y \in [y_0, \infty]$ instead of $[0, \infty]$, then we merely generalize (1.3) to

$$y = y_0 + L \cot^2(t/2), \quad (5.4)$$

where in the present case $y_0 = 1$. Since this is merely a translation by a constant, the derivative formulas given in Table II are unchanged.

The second modification is that when the basis includes $\{TL_0, \dots, TL_n\}$, the collocation grid is changed to

$$y_i = y_0 + L \cot^2(t_i/2), \quad i = 0, \dots, N-1 \quad (5.5a)$$

$$t_i = \pi(2i+1)/(2N), \quad i = 0, \dots, N-1, \quad (5.5b)$$

which is the same form as (2.2) except that there are only N points on the interior of $[0, \pi]$. N rows of the matrices \mathbf{H} and \mathbf{f} are given by (2.6) and (2.7), but the last row is altered to

$$\mathbf{H}_{Nj} = (-1)^j, \quad j = 0, \dots, N, \quad \mathbf{f}_N = u(y_0) \quad (5.6)$$

to impose the boundary condition. Note that we have used

$$TL_n(y_0) = \cos(n\pi) = (-1)^n \quad (5.7)$$

in the left half of (5.6).

Boyd [3] has argued that one of the virtues of using mapped Chebyshev polynomials as the basis is that the efficiency is insensitive to the map parameter L . Figure 2, which compares three solutions computed for three values of L that differ




FIG. 2. A comparison of numerical solutions with $N=6$ (six interior collocation points plus the boundary condition) for Example One, whose exact solution is $K_1(y)$. Exact solution [solid], $L=0.25$ [dotted], $L=1.25$ [solid], and $L=6.25$ [dashed] are shown. Note that the exact solution and the approximation with map parameter = 1.25 are graphically indistinguishable, so both are shown as the solid curve.

by a factor of 25, certainly confirms his assertion. The numerical solution for $L=1.25$ is indistinguishable from the exact solution to within the thickness of the curve, but the approximations for L five times smaller and five times larger also have small absolute errors: no worse than 0.01 for the former and 0.015 for $L=6.25$.

No comparison is made with Boyd's formulas [3] for predicting the optimum L because these are asymptotic formulas for large N —and we only need $N=6$ for this problem.

Example Two: Whittaker's Equation Eigenproblem

The differential equation

$$u_{yy} + [-1/4 + 1/y + \lambda/y] u = 0, \quad y \in [0, \infty], \quad (5.8)$$

where λ is the eigenvalue is a special case of Whittaker's equation. The exact solution is $u \equiv \exp(-y/2)$ for $\lambda = -1$ plus

$$u \equiv e^{-0.5y} y L_n^1(y) \quad (5.9a)$$

$$\lambda \equiv n, \quad n \geq 0 \quad \text{is an integer,} \quad (5.9b)$$

where L_n^1 is the associated Laguerre polynomial of first order and degree n . Because the differential equation is singular at both endpoints, an unmodified $TL_n(y)$ basis, as is appropriate for "natural" boundary conditions at both $y=0$ and $y=\infty$, gave exceptionally accurate results.

One special difficulty with eigenvalue problems (whether on an unbounded interval or not) is that N basis functions will never give enough resolution to even

crudely represent all of the first N eigenmodes. It follows that of the N eigenvalues of the pseudospectral matrix, only the lowest few will be good approximations to those of the differential equation while the larger matrix eigenvalues are numerical artifacts that must be discarded. The only effective way to determine which eigenvalues are good and which are inappropriate is to repeat the calculation with different N and accept only those eigenvalues that are approximately the same in both computations.

Parenthetically, we note that for this example, some of these spurious eigenvalues are complex even though all eigenvalues of the differential equation are real. Galerkin's method, which gives a symmetric matrix \mathbf{H} if applied to a self-adjoint operator, is better in this respect because all its numerical eigenvalues are real. Even with Galerkin's, however, the large eigenvalues of the matrix are still hopelessly wrong, so the fact that some pseudospectral eigenvalues are complex is irrelevant unless the eigenproblem is solved using an iterative method that requires real eigenvalues.

The discretized matrix problem is a "generalized" eigenproblem in which the eigenvalue λ is multiplied by a full matrix rather than the identity matrix. However, the QZ algorithm from the EISPACK library will compute all the eigenvalues and eigenfunctions—real or complex—without requiring any input except the two square matrices themselves.

In Figs. 3 and 4, we have arbitrarily defined a "good" eigenvalue as one which is within 0.05 of the corresponding exact eigenvalue and then plotted the number of "good" eigenvalues as a function of the map parameter L for $N=10$ and $N=40$, respectively. The graphs show several things: (i) as for the first example, accuracy is rather insensitive to L . Indeed, for $N=40$ one can vary L by three orders of magnitude and still obtain the first two eigenvalues correctly! (ii) The optimum





FIG. 3. The number of "good" eigenvalues for the Whittaker eigenvalue equation for $N=10$ as a function of map parameter L . A "good" eigenvalue is arbitrarily defined to be one which is within 0.05 of the exact eigenvalues.


FIG. 4. Same as Fig. 2 but for $N=40$.

eigenvalue increases roughly linearly with N as predicted by the steepest descent theory of Boyd [3] for solutions which are entire functions. (iii) Boyd's prediction that

$$L_{\text{optimum}} = 1.414N \quad (5.10)$$

is an overestimate—but just barely; the accuracy is so insensitive to L in the neighborhood of the optimum L that it is difficult to judge. (iv) The fact that the differential equation is singular at both endpoints and $u(y)$ is singular at infinity seems quite irrelevant; the prediction (5.10) depends only on the knowledge that the solutions are without singularities for any finite y .

Boyd [11] offered the rule-of-thumb that in a pseudospectral eigenvalue calculation, the number of “good” eigenvalues is roughly $N/2$. Figure 5 tests this

FIG. 5. The number of “good” eigenvalues as a function of N with $L=0.8N$ for the Whittaker eigenvalue problem.

conjecture by plotting the number of "good" eigenvalues against N where $L = 0.8N$. The graph clearly shows that the number of accurate eigenvalues does rise linearly with N , but is equal to $0.3N$ to $0.4N$ (with the larger ratio applying for larger N), which is a little smaller than Boyd's rule-of-thumb. Clearly, however, the penalty extracted by the unboundedness of the interval and the singularities of the differential equation at both endpoints is quite modest.

Example Three: Laguerre Eigenproblem

The differential equation

$$yu_{yy} + (y + 1)u_y + \lambda u = 0 \quad (5.11)$$

has the exact eigensolutions

$$u \equiv 1, \quad \lambda = 0 \quad (5.12a)$$

$$u = e^{-y}L_n(y), \quad \lambda \equiv n + 1, \quad n \geq 0, \text{ an integer,} \quad (5.12b)$$

where the $L_n(y)$ are the usual Laguerre polynomials. This problem is clearly a close cousin of the previous example, and the boundary conditions are "natural" at both endpoints. However, excluding the trivial eigenvalue ($\lambda = 0$), numerical solutions of (5.11) were extremely disappointing: with $L = 4$ and $N = 30$, only one eigenvalue was even close, and the approximation was the complex conjugate pair

$$\lambda_{\text{approx}} = 0.9998 \pm 0.064i \approx 1. \quad (5.13)$$

The eigenfunction is also complex, but is a good approximation (within 1%) of $\exp(-y)$. Still, this is astonishingly poor accuracy for such a simple problem with so large an N , and varying the map parameter did not produce much improvement.

The problem is that the second linearly independent solution of (5.11) has the asymptotic form

$$u \sim 1/y^2, \quad y \rightarrow \infty. \quad (5.14)$$

Thus, the form of the differential equation forces *both* solutions to vanish as $y \rightarrow \infty$. When no boundary constraints are applied, i.e. when the boundary conditions at

rather than algebraically with y —but just barely. Imposing the explicit constraint $u(\infty) = 0$ will be quite useless since both solutions satisfy this condition.

The difficulty can be cured by introducing the new unknown

$$w(y) \equiv e^{y/2}u(y) \quad (5.15)$$

which satisfies the differential equation

$$yw_{yy} + w_y + [(\lambda - 1/2) - (1/4)/y]w = 0. \quad (5.16)$$

Now, the offending solution blows up exponentially instead of a slow, inverse power decay. When (5.16) is pseudospectrally solved using $N=40$, the result is 15 eigenvalues to within 3%—and all 41 matrix eigenvalues are real whereas the “bad” version of the same problem, (5.11), yields only three “good” eigenvalues for $N=40$ even if one is willing to ignore the small imaginary parts.

The moral of the story is quite clear: it is important to analyze the asymptotic behavior of the differential equation and its solutions.

Example Four: “Charney” Eigenproblem

The physical background of the “Charney” problem is described in Pedlosky [12]. Here, it will suffice to note that Charney’s work is typical of one-dimensional hydrodynamic stability theory, motivated in his case by the need to explain the growth of the travelling storm systems (“synoptic cyclones,” alias “baroclinic instability”) that are the “Highs” and “Lows” on a daily weather map. The differential equation is

$$(y - c) u_{yy} + [r - y/4 + c/4] u = 0, \quad (5.17)$$

where y is height above the ground, r is a constant parameter (which depends on the nondimensional longitudinal wavenumber of the unstable modes) and c , the eigenvalue, is the phase speed of the waves. For the special case known as the “Boussinesq approximation,” the boundary condition is

$$cu_y + u = 0 \quad \text{at } y = 0. \quad (5.18)$$

Most of the principles that apply to the standard eigenproblems above must be discarded. The only physically relevant mode is the unstable one, that is, the mode with a complex phase speed c . The matrix has real eigenvalues, too, but [12] explains why these should be ignored.

The unstable modes have branch points (“critical levels”) for complex y which move closer and closer to the real interval $y \in [0, \infty]$ when r tends to either 0 or 1. To track the instability in these limits, one must use the special mapping tricks explained in Boyd [5]. Fortunately, the most interesting modes are those that grow most rapidly since they will leave their more slowly amplifying brethren behind and dominate the weather. For strongly unstable waves, we can obtain good accuracy with small N and no special tricks.

The boundary condition at infinity [the top of the atmosphere] is “natural,” but (5.18) must be explicitly imposed. Curiously, the fact that the eigenvalue appears in the boundary condition actually simplifies life because the result of applying a pseudospectral discretization to both the differential equation and boundary condition can still be written in the standard form

$$\mathbf{Aa} = c\mathbf{Ba}, \quad (5.19)$$

where \mathbf{A} and \mathbf{B} are square matrices such that $\mathbf{A} - c\mathbf{B} = \mathbf{H}$, where \mathbf{H} is defined by (2.6) except for the last row, where (5.18) is imposed. If the boundary condition did

not involve c , then the square matrix \mathbf{B} would have a row of zeros, so before applying the QZ algorithm, a preprocessing step to reduce the problem to standard form would be necessary [18].

The derivative in the boundary condition, however, does require special treatment because Table II gives

$$dTL_n(y)/dy = n\{\sin^3(t/2)/\cos(t/2)\} \sin(nt), \tag{5.20}$$

where $t = 2 \operatorname{arccot}([y/L]^{1/2})$. Unfortunately, both $\sin(nt)$ and $\cos(t/2)$ vanish at $t = \pi$, which corresponds to $y = 0$. Thus, to follow the strategy of transforming from y to the trigonometric argument t , we must apply l'Hopital's rule [or equivalently, Taylor expand the vanishing functions and then take the ratio] to obtain

$$dTL_n/dy|_{y=0} = -2n^2 \cos(n\pi)/L. \tag{5.21}$$

It then follows that the elements of the last rows of \mathbf{A} and \mathbf{B} are given by $A_{Nj} = TL_j(0) = \cos(j\pi)$ and $B_{Nj} = -dTL_j/dy(0)$ as given by (5.21).

For $r = 0.5$ (near the r of maximum growth rate), the exact phase speed for the one unstable mode is

$$c = 0.901016 + i0.544987. \tag{5.22}$$

Figure 6 shows the errors in the real and imaginary parts of c for L ranging from $L = 1/16$ to $L = 16$ for $N = 10$. Despite the small number of basis functions, the branch point in $u(y)$ for complex y , and the tremendous variations in the map parameter, the errors are all rather modest. For L between 0.25 and 4, the relative error in both the real and imaginary parts is never more than 2%.

Boyd [3] showed that when a solution has singularities for complex y , this will make the optimum L much smaller. Figure 6 clearly supports this. For the Whit-




FIG. 6. The errors in the real part (left half of each pair) and imaginary part (right half) of the eigenvalue c for the "Charney" problem for $r = 0.5$ with $N = 10$, plotted as a function of map parameter L .

taker eigenproblem, the optimum $L_{\text{optimum}}(N)=8$ for $N=10$ versus only $L_{\text{optimum}}(10)=1$ here even though both solutions decay as $\exp(-y/2)$. The difference is simply that the Laguerre functions are entire functions whereas the Charney waves have branch points.

Example Five: "Global" Expansion of Smallest Eigenvalue, Quantum Anharmonic Oscillator

For this and the next two examples, the standard way to approximate the function is by using (i) a Taylor series for small values of the argument and (ii) an asymptotic expansion for $|y| \gg 1$. Because the three examples have quite different behavior for small and large y , it is impossible to extend the asymptotic series (which is divergent!) to small y or vice versa. Nonetheless, by using mapping, we shall obtain rapidly convergent approximations that are "global," that is to say, are a *single* expansion accurate over the *whole* interval $y \in [0, \infty]$.

Boyd [13] gives the physical background of the anharmonic oscillator. Here, our goal is to approximate the dependence of the smallest ("ground state") eigenvalue on the coupling constant λ . Boyd [13] calculated two separate Chebyshev series approximations, one on $\lambda \in [0, 0.2]$ and the other on $[0.2, \infty]$, by solving the problem at a set of discrete λ chosen to be the Chebyshev interpolation points on each interval. To express the results in a simpler and more transparent form, these Chebyshev approximations were converted to ordinary polynomials to obtain

$$\begin{aligned} E_0(\lambda) &\approx 0.50077 + 0.66420\lambda - 0.79280\lambda^2, & \lambda \in [0, 0.2] \\ E_0(\lambda) &\approx \lambda^{1/3} \{0.66839 + 1.41059\lambda^{-2/3} - 0.005977\lambda^{-4/3}\}, & \lambda \in [.2, \infty]. \end{aligned} \quad (5.23)$$

Since the error in these two approximations is only 1 part in 700 and 1 part in 1500, respectively, it is clear that they are an efficient method for representing the energy of the ground state of this quantum oscillator.¹ Nonetheless, it is still appealing to find a single expression that works for all λ .

The form of (5.23) determines the form of our "global" expansion. First, since the asymptotic expansion for large λ involves fractional powers, we define

$$y \equiv \lambda^{1/3}. \quad (5.24)$$

Second, since $E_0(\lambda)$ is unbounded as $\lambda \rightarrow \infty$, we apply the $TL_n(y)$ expansion to

$$\varepsilon_0(y) \equiv E_0(y^3)/(1+y), \quad (5.25)$$

where the denominator causes $\varepsilon_0(y)$ to asymptote to a constant as $y \rightarrow \infty$. Figure 7 is a bar graph that dramatically shows the rapid convergence of the series.

¹ Parenthetically, we note a couple of minor errors in [13]: the first entry in column one of Table II should be 1.11059... instead of 1.1059..., and also the claim that $\lambda^{-2/3} \in [0, 0.731]$ should be replaced by $[0, 2.924]$. All other expressions in [13] are correct.




FIG. 7. The absolute values of the coefficients of the rational Chebyshev expansion of the ground state of the quantum quartic oscillator. The argument of the $TL_n(y)$ is $y = \lambda^{1/3}$, where λ is the "coupling constant" of the oscillator and the function expanded is the ground state eigenvalue divided by $(1 + \lambda^{1/3})$ so that the TL_n series is applied to an $f(y)$ which is bounded as $y \rightarrow \infty$.

Table IV lists the numerical value of the coefficients. Recalling that

$$|TL_n(y)| \leq 1, \quad y \in [0, \infty] \text{ for all } n, \quad (5.26)$$

we see that the error in truncating the series after a given number of terms is bounded by the sum of the absolute values of all the neglected coefficients. Thus, the

TABLE IV

n	a_n
0	0.497194
1	0.093215
2	0.091423
3	-0.011715
4	-0.004460
5	0.002851
6	-0.000334
7	-0.000360
8	0.000206
9	-0.000014
10	-0.000037
11	0.000020
12	-0.000001
13	-0.000004
14	0.000002

Note. This lists the $TL_n(y)$ coefficients for the ground state eigenvalue of the quartic oscillator of quantum mechanics.

three-term approximation has a relative error of no worse than 1 in 25 over the whole semi-infinite interval—poorer than either part of (5.23), but still remarkably good for such a low order approximation—and keeping the terms up to and including $TL_4(y)$ gives an absolute error of less than 0.004.

Example Six: “Global” Expansion of $K_1(Y)$

The Bessel function $K_1(y)$ has a simple pole at $y=0$, so we will compute approximations to

$$f(y) \equiv yK_1(y). \quad (5.27)$$

Unfortunately, $K_1(y)$ also has a logarithm at the origin, and this cannot be multiplied out since the coefficient of the logarithm is an infinite series [that for $I_1(y)$]. Fortunately, Stenger [14] has observed that it is possible to force the error to decrease exponentially with N by using a mapping from $y \in [0, \infty]$ to $Y \in [-\infty, \infty]$ of exponential character, that is to say, one which clusters the pseudospectral interpolation points around $y=0$ so that the density increases exponentially with N . Boyd [6] notes that once the mapping has been applied, one can use any of a variety of basis sets on $Y \in [-\infty, \infty]$ including the rational Chebyshev functions of Boyd [1]. To approximate (5.27), we therefore need to combine ideas from all three of these previous works: [14, 6, 1].

Since the $K_1(y)$ decays as $\exp(-y)$ as $y \rightarrow \infty$, we need to transform y only for small y , so we choose a map such that $Y(y)$ is proportional to y as $y \rightarrow \infty$. A simple choice from [14] is

$$y = \sinh^{-1}(e^Y), \quad Y \in [-\infty, \infty]. \quad (5.28)$$

The rational Chebyshev functions $TB_n(Y)$ are in turn the images of $\cos(nt)$ under the mapping $Y = L \cot(t)$. Thus, the coefficients of the required global expansion are those of the Fourier cosine series of

$$f(t) \equiv \{\sinh^{-1}[\exp(L \cot(t))]\} K_1\{\sinh^{-1}[\exp(L \cot(t))]\} \quad (5.29)$$

and the series approximation can be written

$$f(y) = \sum_{n=0}^{\infty} a_n \cos\{n \cot^{-1}[(1/L) \log(\sinh[y])]\}. \quad (5.30)$$

A well-known theorem asserts that a polynomial approximation to a function which is singular at one endpoint—even if $f(y)$ is bounded at that point—cannot converge faster than algebraically as the number N of terms in the series is increased. We are able to bypass this restriction because our double mappings have replaced polynomials (or cosines) by the more complicated transcendentals shown in (5.30).

Indeed, the convergence is almost too good! Boyd [1, 3] explains that when $f(y)$ is singular at an endpoint, the convergence must be “subgeometric” in the sense

that the decay, although exponential with n , is slower than that of $\exp(-pn)$ for any positive p . Figure 8 graphs the logarithm of the error bound (obtained by adding up the absolute values of all the neglected coefficients) for (5.30). For a series whose terms decrease like those of a geometric series (the usual case for a Chebyshev expansion on $x \in [-1, 1]$), the graph would asymptote to a straight line. When the series has only “subgeometric” convergence as defined in [4], the graph should flatten for large n so that the slope approaches 0 from above.

Nonetheless, Fig. 8 approximates a straight line. If one has good eyesight, one can detect that the slope is *slowly* decreasing as n increases—but for $n=24$, the approximation is already accurate to better than five decimal places and the deviation from a straight line is still small. The concept of “geometric” versus “subgeometric” convergence is an *asymptotic* notion. Figure 8 shows that sometimes the distinction between the two is detectable only for such large n as to have no practical significance. [Table VIII of Boyd [13] makes the same point.] It also implies that the mapping treatment of a singularity as weak as $y^2 \log(y)$ is unnecessary unless one wants more than five decimal places.

This series differs from the other six examples in that we do not actually use the $TL_n(y)$. However, the desired expansion interval is semi-infinite, and this example does illustrate the way in which different mappings can be combined to deal with the simultaneous problems of (i) a logarithm at an endpoint and (ii) an unbounded interval.

Example Seven: Amplitude and Phase Expansions of Bessel Function $J_0(y)$

The Bessel function $J_0(y)$ can be approximated by a power series in y for small y and by a divergent asymptotic series of the form

$$J_0(y) \sim [2/(\pi y)]^{1/2} \{ [1 - 9/(128y^2) + O(y^{-4})] \cos(y - \pi/4) + [-1/(8y) + O(y^{-3})] \sin(y - \pi/4) \}. \tag{5.31}$$




FIG. 8. A plot of the logarithm (base 10) of the error bound E_N (obtained by summing the absolute values of all coefficients with $n > N$) versus N for the function $yK_1(y)$ as computed using the arcsinh-exponential mapping with $L = 1$.

Since no finite sum of basis functions can possibly hope to *uniformly* approximate an *infinite* number of oscillations on $y \in [0, \infty]$, it is obvious that for this example, the rational Chebyshev series and all other direct spectral approximations will fail miserably.

The remedy is to define a modified $f(y)$ with a multiplicative factor that eliminates the leading square root in (5.13) and then compute *two separate expansions* for the functions $P(y)$ and $Q(y)$ in

$$f(y) \equiv (1+y)^{1/2} J_0(y) \quad (5.32)$$

$$= P(y) \cos(y - \pi/4) + Q(y) \sin(y - \pi/4). \quad (5.33)$$

The form of the asymptotic expansion suggests what numerical experiments confirm: $P(y)$ and $Q(y)$ have rapidly convergent TL_n expansions whose coefficients are listed in Table V. With $L=2$, the absolute error with a combined total of seventeen coefficients is uniformly less than 2×10^{-7} on $y \in [0, \infty]$.

Nonetheless, there are some serious and unresolved issues hidden behind this glitter of success. First, (5.33) represents one function in terms of two: how do we choose $P(y)$ and $Q(y)$, out of the infinite range of possibilities, to maximize numerical efficiency?

Krylov and Bogoliubov, as reviewed by Nayfeh [15], successfully solved this uniqueness problem for *perturbation* theory by demanding that

$$u = a(y) \cos(\phi[y]) \quad (5.34a)$$

$$u_y = -a(y) \sin(\phi[y]). \quad (5.34b)$$

Their motive is that in the asymptotic regime ($y \gg 1$ for our example), the amplitude a and derivative of the phase ϕ vary slowly with y . They therefore imposed (5.34b), which would be trivially true if a and ϕ_y were constant. (Note that

TABLE V

n	P_n	Q_n
0	1.10499	0.112125
1	-0.422833	-0.158486
2	0.157614	0.060126
3	-0.055648	-0.016579
4	0.017623	0.003121
5	-0.004784	-0.000308
6	0.001132	
7	-0.000258	
8	0.000060	
9	-0.000013	
10	0.000003	

an expression like (5.34a) can always be converted into (5.33) through elementary trigonometry, so we shall refer to approximations in either form as “amplitude-phase” expansions.) This pair of equations is equivalent to

$$a(y) \equiv (u^2 + [u_y]^2)^{1/2}, \tag{5.35}$$

which gives us a unique, explicit expression for the amplitude. We can expand this as a $TL_n(y)$ series through either a matrix multiplication or a Fast Fourier Transform after evaluating $a(y)$ at the usual interpolation points (2.3).

The fatal flaw is that if we define ε to be their perturbation parameter, the Krylov-Bogoliubov amplitude function has an $O(\varepsilon)$ part which oscillates for all real y . For our present $f(y)$, one can show from (5.31) that

$$a^2(y) \sim 1 + 1/y - 1/(4y^2) + \{ \cos[2(y - \pi/4)] - 4 \sin[2(y - \pi/4)] \} / (8y^2). \tag{5.36}$$

Numerical experiments (not presented) show that this oscillation destroys the convergence of the $TL_n(y)$ series—the coefficients level off and do not become small even for large n . Worse still, the Krylov-Bogoliubov procedure is only asymptotic, so it is impossible to refine it.

We therefore tried the simple-minded alternative of expanding $P(y)$ and $Q(y)$ in (5.33) as

$$P(y) = \sum_{n=0}^M P_n TL_n(y), \quad Q(y) = \sum_{n=0}^N Q_n TL_n(y) \tag{5.37}$$

and then imposing the collocation conditions

$$P(y_i) \cos(y_i - \pi/4) + Q(y_i) \sin(y_i - \pi/4) = f(y_i), \tag{5.38}$$

$$i = 0, \dots, N + M + 1,$$

where

$$y_i = L \cot^2(t_i/2),$$

$$t_i = \pi(2i + 1)/(2M + 2N + 4), \quad i = 0, \dots, N + M + 1 \tag{5.39}$$

are the usual rational Chebyshev interpolation points. This gives a simple set of *linear* equations to determine the unknown coefficients $\{P_n\}$ and $\{Q_n\}$. (Note that an approximation in the form $u = a(y) \cos(\phi[y])$ depends *nonlinearly* on the phase ϕ , so (5.38) is preferable.) Taking $M = N = 8$ gave a maximum absolute error on $y \in [0, \infty]$ of only 0.00002.

Unfortunately, increasing M and N gave no improvement even in 16-decimal precision because the set of linear equations is ill-conditioned. The wonder

change-of-coordinates from y to t and then invoking the theory of trigonometric interpolation], there is no theory for a double expansion such as we attempt here. The interpolation (5.37) is a shot in the dark—albeit a successful one.

Some experimentation showed that taking $M > N$ reduced the ill-conditioning because the factors of $TL_n(y) \sin(y - \pi/4)$ closely resemble $TL_n(y) \cos(y - \pi/4)$ for large n , which is why Table V lists 11 coefficients for the multiplier of $\cos(y - \pi/4)$ and only 6 for that of $\sin(y - \pi/4)$. Additionally, we improved the stability of the calculation by using the interpolation points (5.3) to define a discrete inner product (an approximation to the integral on $[0, \infty]$ with the weight given in (1.4)) and then applied a Gram-Schmidt orthogonalization to define new basis functions. After applying this least-squares procedure using 30 quadrature points and 21 basis functions, we converted the coefficients back into the form of (5.37). Figure 9 shows the amplitude and phase functions when (5.33) is converted into $f(y) = a(y) \cos(\phi[y])$.

When we replace interpolation by a least-squares procedure, we are guaranteed that adding more terms will reduce the residual at least at the quadrature points. However, there is still no theory to guarantee that minimizing errors at any finite set of points will give small errors over the continuous interval $y \in [0, \infty]$. In addition, M and N are still limited because of as yet poorly understood round-off errors.

In short, this amplitude-phase approximation would make a pure mathematician recoil in horror. Nonetheless, it gives seven decimal places of accuracy. Clearly, it is possible for creative use of interpolation to give approximations on unbounded intervals that mimic asymptotic expansions and the WKB method, and yet have (i) a much wider range of usefulness and (ii) are free of the curse of asymptoticity, which limits the accuracy of WKB formulas for fixed y even when y is large.




FIG. 9. The amplitude (solid line) and phase (dashed) in the amplitude-phase approximation to $f(y) = (1+y)^{1/2} J_0(y)$.

6. SUMMARY

By making a change-of-coordinate, we have shown that Chebyshev polynomials (or equivalently, a Fourier cosine series) can be applied to a semi-infinite interval as first suggested by Grosch and Orszag [2]. By exploiting the mapping-induced connection between Fourier cosine series and the new basis functions defined here, the rational Chebyshev functions $TL_n(y)$, we are able to fill in the gaps of Boyd [3].

In the Appendix, we display the “equiconvergence” contours in the complex plane. In the body of the text, we discuss the issues of natural versus essential boundary conditions, banded Galerkin matrices, and explain how the $TL_n(y)$ —in contrast to Laguerre functions—can be applied to bounded functions with asymptotic expansions in inverse powers of y as $y \rightarrow \infty$.

The seven numerical examples illustrate the possibilities. The Charney problem shows that numerical weather prediction, which uses mixed spectral–finite difference models (spherical harmonics in the horizontal, fourth order differences in height) can instead employ pure pseudospectral codes by using the TL_n in height. The success of the uniform amplitude–phase expansion for $J_0(y)$ shows that it is possible to replace the WKB method by an approximation which is *convergent* rather than *asymptotic* and which can be applied for *all* y , not just when y is large.

Much theoretical work is needed to put such “double expansions” on a firm basis. In addition, we have not discussed condition numbers, the eigenvalue spectrum, and “pre-conditioning” methods for pseudospectral matrix problems. In one dimension, the pre-conditioned iterative procedures are unnecessary (most of the calculations in this article took less than five seconds on an IBM AT!). In higher dimensions, the best choice of pre-conditioner is still a research frontier even for ordinary Chebyshev expansions [16].

Similarly, we have avoided the important topic of time-stepping algorithms combined with $TL_n(y)$ series for the spatial coordinates. As emphasized by Trefethen and Trummer [19], the eigenvalue spectrum of the first derivative operator is poorly understood even for ordinary Chebyshev polynomials. At present, we do not know how to predict whether a pseudospectral representation of a differential operator will have complex eigenvalues (which could cause time-stepping instability) even for a self-adjoint equation; the examples above show only that the spectrum is sometimes purely real—and sometimes not.

Clearly, orthogonal rational functions offer many topics for future research. Nonetheless, the examples and analysis given here show that these mapped Chebyshev polynomials are an extremely effective tool for solving difficult problems on a semi-infinite interval.

APPENDIX: CONTOURS OF EQUICONVERGENCE

The domain of convergence of a Fourier cosine series is a strip parallel to the real t -axis, that is, the series converges within the region

$$|t_i| \leq \rho \quad \text{Fourier domain of convergence} \quad (\text{A.1})$$

where ρ is the largest possible positive constant such that the *interior* of the strip (but not its boundary) contains no singularities of the function $f(t)$. Although we are usually interested in evaluating a function only for real t , (A.1) is of considerable practical importance because the asymptotic rate at which the Fourier coefficients decrease with n is directly related to the width of the strip of the convergence via

$$a_n \sim [\] (1/\delta)^n, \quad n \rightarrow \infty, \quad (\text{A.2})$$

where the empty $[\]$ denote functions that vary more slowly than δ^{-n} and where

$$|\delta| = e^\rho. \quad (\text{A.3})$$

Thus, the wider the strip of convergence, the faster the rate at which the Fourier coefficients decrease. This motivates the following.

DEFINITION. The “contours of equiconvergence” are the curves in the complex plane where the asymptotic rate of convergence is a constant, that is, are the contours of constant δ (or ρ).

For Fourier series, the equiconvergence contours are straight lines parallel to the real t -axis. The “equiconvergence” contours for the $TL_n(y)$ series are the images of these straight lines under the mapping $t \rightarrow y$. Therefore, if $f(y)$ has a singularity at the point $y = y_r + iy_i$, then the width of the strip of convergence in the t -plane is

$$\rho = \text{Im} \{ 2 \operatorname{arccot} [([y_r + iy_i]/L)^{1/2}] \} \quad (\text{A.4})$$

$$= -1/2 \log \{ [\zeta_r^2 + (\zeta_i + 1)^2] / [\zeta_r^2 + (\zeta_i - 1)^2] \} \quad (\text{A.5})$$

via the logarithmic representation of the arccot function where

$$\zeta = (y/L)^{1/2}. \quad (\text{A.6})$$

Knowledge of $\rho(y_r, y_i)$ determines the asymptotic behavior of the TL_n series via (A.2) and (A.3). When the location of the singularities of $f(y)$ are known in advance, either because $f(y)$ itself is known or because it satisfies a linear differential equation with coefficients whose singularities are known, then one can evaluate (A.4) for each branch point or pole, and the smallest ρ so obtained determines the region of convergence. Figure 10 shows the contours of $\rho(y)$ in the complex y -plane.

One can directly trace these contours by defining

$$s \equiv \sinh(\rho) \quad (\text{A.7})$$

and then solving the equation (with $L = 1$ for simplicity)

$$\begin{aligned} y_r^4 [-s^4] + 2y_i^2 [2 - s^4 - 2s^2(s^2 + 1)y_r - s^4 y_r^2] \\ + s^2 (y_r + 1)^2 \{ -4y_r - s^2 (y_r + 1)^2 \} = 0 \end{aligned} \quad (\text{A.8})$$




FIG. 10. The equiconvergence contours in the upper half of the complex y -plane for the rational Chebyshev functions on the semi-infinite interval. The map parameter $L = 1$; only the upper half-plane is shown because the curves in the lower half-plane are the reflection with respect to the real y -axis of those for $\text{Im}(y) > 0$.

for y_r^2 as a function of y_r and ρ . Equation (A.7) is derived by decomposing the trigonometric functions in (1.3) into their real and imaginary parts, solving for $\cos(t_r)$ and $\sin(t_r)$ and then applying the identity $\cos^2(t_r) + \sin^2(t_r) = 1$ to eliminate the real part of t . Although (A.7) can be solved as a quadratic in y_r^2 , the equation is nonetheless quartic, which implies that the equiconvergence contours for the $TL_n(y)$ are more complicated than the ellipses that play the same role for the ordinary Chebyshev polynomials, $T_n(x)$.

The reason that this analysis has been banished to this appendix is that unfortunately it applies only to functions that are *analytic at infinity*, such as certain rational functions and other simple forms. If $f(y)$ has poles, branch points, or essential singularities at infinity, then the $TL_n(y)$ series will converge only on the real interval $y \in [0, \infty]$. Equation (A.2) does not describe the asymptotic form of the coefficients and the contours in Fig. 10 become irrelevant.

Fortunately, if $f(y)$ has all its derivatives bounded at infinity, the series coefficients will still decrease *exponentially* with n , but the (asymptotic!) rate of convergence will be "subgeometric" in the sense defined in Boyd [4]. Characterizing this requires the more subtle steepest descent analysis of Boyd [3, 4] rather than the equiconvergence contours described here.

ACKNOWLEDGMENT

This work was supported by the NSF under Grant OCE8509923.

REFERENCES

1. J. P. BOYD, *J. Comput. Phys.* **69**, 112–142 (1987).
2. C. E. GROSCH AND S. A. ORSZAG, *J. Comp. Phys.* **25**, 273 (1977).
3. J. P. BOYD, *J. Comput. Phys.* **45**, 43 (1982).
4. J. P. BOYD, *J. Comput. Phys.* **54**, 382 (1984).
5. J. P. BOYD, *J. Comput. Phys.* **57**, 454 (1985).
6. J. P. BOYD, *J. Comput. Phys.* **64**, 266 (1986).
7. H. J. NORTON, *Comput. J.* **7**, 76 (1964).
8. J. P. BOYD, *Physica D* **21**, 227 (1986).
9. D. GOTTLIEB AND S. A. ORSZAG, “Numerical Analysis of Spectral Methods: Theory and Applications” (SIAM Philadelphia, 1977).
10. A. B. CAIN, J. H. FERZIGER, AND W. C. REYNOLDS, *J. Comput. Phys.* **56**, 272 (1984).
11. J. P. BOYD, *Monthly Weather Rev.* **106**, 1192 (1978).
12. J. PEDLOSKY, “Geophysical Fluid Dynamics” (Springer-Verlag, New York, 1979).
13. J. P. BOYD, *J. Math. Phys.* **19**, 1445 (1978).
14. F. STENGER, *SIAM Rev.* **23**, 165 (1981).
15. A. H. NAYFEH, “Perturbation Methods” (Wiley, New York, 1973).
16. C. CANUTO AND A. QUARTERONI, *J. Comput. Phys.* **60**, 315 (1985).
17. J. P. BOYD, *J. Sci. Comput.* **1**, 183 (1986).
18. J. GARY AND R. HELGASON, *J. Comput. Phys.* **5**, 169 (1970).
19. L. N. TREFETHEN AND M. R. TRUMMER, *SIAM J. Numer. Anal.* (1986), in press.