Complex Coordinate Methods for Hydrodynamic Instabilities and Sturm–Liouville Eigenproblems with an Interior Singularity

JOHN P. BOYD

*Department of Atmospheric and Oceanic Science,
University of Michigan, 2455 Hayward Avenue, Ann Arbor, Michigan 48109*

Calculations of inviscid, linearized waves in fluids are very difficult when a mean wind or current $U(y)$ is included because the differential equation is *singular* wherever $U(y) = c$, the phase speed. These "critical latitude," "critical level," or "critical point" singularities are particularly severe for Chebyshev methods since these global expansion algorithms are very sensitive to the analytic properties of the solution. A simple remedy is described: by making a change of coordinates $y = f(x)$ where y is the original variable and x is the new coordinate with $f(x)$ a *complex* function, one can solve the problem on an *arc* in the *complex plane* that makes a wide *detour* around the singularity. Specific guidelines for choosing $f(x)$ for different problems are given in the text. Results are impressive: for an eigenvalue problem with a pole in the middle of the original real interval (a "Sturm–Liouville problem of the fourth kind"), just six basis functions suffice to calculate the real and imaginary parts of the lowest eigenvalue to within 1.4 %. For strong instability, i.e., modes whose phase speeds have large imaginary parts, the complex mapping is unnecessary because the critical latitudes are complex and distant from the real axis. Even so, the mapping is useful for instability problems because it can be used to make calculations for very slowly growing modes to follow the changes in c right up to the "neutral curve" where the imaginary part of $c = 0$. Although especially valuable for spectral algorithms, the same trick can be applied with finite difference methods also. The main disadvantage of the algorithm is that the eigenfunction must be calculated in a second, separate step, but this is usually a minor flaw in comparison to the complex mapping's virtues for coping with singular eigenvalue problems. © 1985 Academic Press, Inc.

1. INTRODUCTION

Hydrodynamic stability and wave problems often require solving differential equations with singularities on or near the computational domain. These singularities, usually called "critical latitudes" or "critical points," create severe numerical difficulties. The purpose of this work is to show how one can obtain good results for a subset of such problems by making a transformation of the original variable y to a new variable x such that the problem is effectively solved on a curve in the complex y-plane rather than on the original interval on the real y-axis. With proper choice of map parameter, one can loop the curve of integration away from the singularity so that it does not degrade the numerical accuracy.

454

We shall discuss only eigenvalue problems because the method is not useful for ordinary boundary value problems. The difficulty is that by solving the differential equation on a complex curve, we sacrifice the solution $u(y)$ along the real axis, so the complex mapping method is useful only when something other than $u(y)$, i.e., the eigenvalue, is the principal goal of the calculation.

The great risk in a linear eigenvalue computation is that of missing modes, so Chebyshev polynomial methods in combination with the QR algorithm for solving the resulting matrix eigenvalue equation have become very popular for hydrodynamic instability problems following Orszag [1]. For self-adjoint problems, there are any number of foolproof methods, but instability problems are not self-adjoint and have complex eigenvalues. A local iterative method will certainly find some eigenvalues, but there is a real danger of missing other important modes of instability. The QR algorithm is globally convergent and does not need to be supplied with first guesses for the eigenvalues, so Orszag's Chebyshev–QR method is very robust: it will fail to find a mode using N Chebyshev polynomials only if that mode is oscillating so rapidly that N polynomials cannot resolve it.

The motivation for this paper is that Chebyshev methods are extremely sensitive to singularities on and near the real axis because the Chebyshev series is a global expansion, i.e., a single representation of $u(y)$ that must fail or succeed over the whole computational interval. Finite difference methods are much less sensitive to singularities because they are based on local Taylor expansions, and so the critical latitude will be irrelevant as long as it is more than one grid length away from the computational interval. Unfortunately, finite difference methods are much less accurate than spectral methods and thus require a much larger value of N, the number of degrees of freedom, to resolve the same wave. Even more unfortunately, the operation count of the QR algorithm grows as $O(N^3)$, whether the matrix is sparse or full. Thus, finite difference–QR methods are usually too expensive, Chebyshev–QR schemes won't work at all, and non-QR algorithms may miss eigenvalues. The complex mapping method is a way of escaping from these grim alternatives.

This technique is not a panacea or cure-all; it is unnecessary for *strong* instability because the critical latitudes are at complex values of y with large imaginary parts, and it may be ineffective for the Orr–Sommerfeld equation because the critical latitudes are so close to the boundaries that even a complex path of integration cannot get away from them. These "counter-cases" are discussed in Section 2. On the other hand, the mapping procedure is useful for finite difference methods, too, for the purpose of determining the "neutral curve" in parameter space where the flow makes the transition from stability to instability. (See, for example, Boyd and Christidis [2, 3].)

Section 2 describes in detail several classes of problems for which mapping is very useful and a few for which it is not. Section 3 will explain how to transform the differential equation and offer several simple choices of mapping. Section 4 will discuss the numerical analysis of mapping methods and some numerical illustrations. The final section is a summary and prospectus.

2. Cases, Counter-Cases, and Alternatives

The complex mapping method is useful for

(i) hydrodynamic instability problems with complex eigenvalues of small or zero imaginary part, and

(ii) Sturm–Liouville eigenproblems of the "fourth kind" as defined by Boyd [4].

In both cases, the singularity in the differential equation arises at that point ("critical latitude," "critical level," or "critical point") where the local value of the mean flow, $U(y)$, is equal to the wave phase speed c. For the stability calculation, c is the eigenvalue, so complex c implies that the critical latitude y_c defined by $U(y_c) = c$, is complex, too. It is only when the imaginary part of c is small or zero that the singularity slows down the convergence of the Chebyshev series. However, the "neutral curve," which is that curve in parameter space defined by $\mathrm{Im}(c) = 0$, is always a primary goal because it marks the boundary between stability and instability. Thus, one is always in trouble in a hydrodynamic instability problem in the neighborhood of the neutral curve; even a less sensitive finite difference method is doomed when $\mathrm{Im}(c) = 0$.

For the Sturm–Liouville problem of the "fourth kind," c is always real and therefore the critical latitude is always on the real axis. A typical example, taken from Boyd [4] with the coordinate y shifted so that $y = 0$ is the critical latitude, is

$$u'' + [1/y - \lambda]u = 0 \qquad \text{with} \quad u(A) = u(B) = 0; \; A < 0; \; B > 0 \qquad (2.1)$$

where the double prime denotes double differentiation and where λ is the eigenvalue. If A and B, the boundaries, are both of the same sign, then (2.1) is a Sturm–Liouville eigenproblem of the first kind; that is to say, a self-adjoint problem with no singularities on $y \in [A, B]$. Nineteenth century theorems then prove that λ is always real.

When A and B are of opposite sign, however, the singularity at $y = 0$ is in the *interior* of the interval, and the eigenvalues of (2.1) are complex. There are Sturm–Liouville problems of the second and third kind (as defined in [4]) for which the singularities of the differential equation at the boundaries or interior, respectively, are only "apparent" in the sense that all the eigenfunctions are analytic at these singularities (Boyd [4]). In contrast, the eigenfunctions of (2.1) have logarithmic branch points at $y = 0$. Indeed, (2.1) is not even properly posed until one specifies whether the branch cut should be in the upper or lower complex y-plane. (The argument that shows that the upper half-plane is the correct choice in fluid mechanics is given in Boyd [4].)

For Sturm–Liouville problems of the fourth kind, the critical latitude is *always* on the real y-axis, so one is always in need of some algorithm for coping with it. In compensation, the location of the critical latitude is always known in advance because the phase speed c and the mean flow $U(y)$ are both part of the prespecified

data of the problem. Thus, Sturm–Liouville problems of the fourth kind are particularly well suited to the complex mapping technique described below. Some explicit results for (2.1) with complex mapping and the Chebyshev–QR algorithm are given in Table II.

Instability calculations, in contrast, can be one form of "counter-case": the method is useful only for *weak* instabilities for which the imaginary part of the phase speed is small. The reason is that when the imaginary part of the phase speed is large, the critical latitude y_c will have a large imaginary part, too. When the singularity of $u(y)$ is distant from the real interval $[A, B]$, the Chebyshev–QR algorithm will work just as well without the mapping as with, and perhaps better. Tables III and IV illustrate this for the most rapidly growing modes of the so-called "Bickley jet." As explained above, however, the complex mapping is still needed for this same problem to find the neutral curve.

Another counter-case is provided by the so-called "wall modes" of the Orr–Sommerfeld equation

$$u'''' - 2k^2 u'' + k^4 u = ikR[(U[y] - c)(u'' - k^2 u) - U'' u] \tag{2.2}$$

where k and R are parameters ("zonal wavenumber" and "Reynolds' number"). The "wall modes" have critical latitudes very close to the boundaries $[A, B]$. It is not possible to deform the contour of integration so that it does not pass close to the critical latitude because both the old and new contours must intersect both boundary points, $y = A$ and $y = B$. Since these boundary points are very close to the complex critical latitudes of (2.2), one is simply stuck with a very difficult numerical problem, and it may be just as efficient to solve (2.2) without mapping as with it. Table I below and Section 4 suggest, however, that the complex mapping may be surprisingly effective even for critical latitudes close to the endpoints, so it is not possible to definitively label the Orr–Sommerfeld equation as a "case" or "counter-case" at the present time.

In point of fact, one can obtain fairly good results for (2.2) without mapping—albeit at the expense of using 35–40 Chebyshev polynomials [1]—because the critical latitudes of (2.2) are technically not singularities of the differential equation. The terms on the L.H.S of (2.2) represent a viscous dissipation which smears out the logarithmic singularity into a thin, non-singular transition region. In practice, R is so large that the transition region is very narrow, so resolving this neighborhood of the critical latitudes is still the major numerical challenge. Because the critical latitudes are so close to the boundaries, however, this weak viscosity may do as much for numerical efficiency as any mapping could.

Indeed, an artificial computational damping has been the principal means for coping with critical latitudes in the past. The usual Laplacian viscosity leads to a fourth order equation like (2.2), but Simmons [5] deliberately chose a damping proportional to the biharmonic operator, which gave rise to an equation similar to (2.2), but of sixth order. The motive for the higher order diffusion is that it damps different length scales proportional to the fourth power of the wavenumber whereas

an ordinary viscosity damps as the square of the wavenumber. Since the eigen-function is smooth except very close to the singularity, the biharmonic diffusion provides very strong damping around the critical latitude but has almost no effect outside the neighborhood of the singularity.

In some cases, replacing the inviscid second order differential equation by a sixth order equation is inefficient and undesirable, so Boyd [6] used a simple, scale-independent friction in combination with a real mapping of the coordinate y that provided high resolution in the vicinity of the critical latitude. This, too, was suc-cessful, so it should not be inferred that the complex mapping described in the next couple of sections is the only way to cope with singular differential equations. However, artificial damping methods do distort the inviscid physics of the problem and require some trial-and-error and finesse to choose a damping coefficient large enough so that one can accurately solve the differential equation and at the same time small enough so that the computed eigenvalue is close to that for zero damp-ing. The complex mapping method avoids both raising the order of the differential equation and extrapolating dissipative eigenvalues to zero damping.

3. COMPLEX MAPPING

A second order differential equation of the form

$$a_2(y) u_{yy} + a_1(y) u_y + a_2 u(y) = 0 \tag{3.1}$$

where the y subscripts denote differentiation with respect to y can be transformed via the mapping

$$y = f(x) \tag{3.2}$$

into

$$\frac{a_2(f[x])}{[f'(x)]^2} u_{xx} + \left[\frac{a_1(f[x])}{f'(x)} - \frac{a_2(f[x]) f''(x)}{[f'(x)]^3} \right] u_x + a_0(f[x]) u = 0 \tag{3.3}$$

where $f'(x)$ and $f''(x)$ are the first and second derivatives of the mapping function. The map (3.2) and transformation (3.3) are completely general and also describe the purely real-valued mappings used in Boyd [6] and in Tables III and IV.

A useful complex mapping must satisfy several constraints.

First, since the standard interval for the Chebyshev polynomials is $x \in [-1, 1]$ and the new contour of integration must pass through both boundary points, we impose

$$f(-1) = A, \qquad f(1) = B. \tag{3.4}$$

Second, the transformation should introduce no additional singularities of $u(f[x])$. Third, the transformation should be as simple and smooth as possible. Four representative maps are shown in Fig. 1 and discussed in turn below.

The simplest choice is the parabolic map

$$y = (A + B)/2 + [(B - A)/2]\{x + i\Delta(x^2 - 1)\}. \tag{3.5}$$

With its single adjustable constant Δ, (3.5) is the most general parabolic mapping which is consistent with the constraints (3.4). Its simplicity makes it possible to invert (3.5) by solving a quadratic equation. This inversion does not need to be done in practice, however (except to optimize Δ), since the Chebyshev expansion converges rapidly only along that curve in the complex y-plane for which x lies on the real interval $x \in [-1, 1]$, so the eigenfunction must be computed by different means as explained below.

The one remaining question is: Given the location of the critical latitude y_c, what is the optimum value of the mapping parameter Δ? Unfortunately, a general analytic answer does not seem possible and some trial-and-error will be needed in most cases. However, when there is but a single critical latitude on or very near the real axis, it is possible to calculate and tabulate the most effective Δ as a function of y_c; the result is given in Table I and justified in the next section.

One complication is that there may be more than one critical latitude. The Bickley jet, for example, for which some results are given in Tables III and IV, is the wind profile

$$U(y) = -\text{sech}^2(y) \tag{3.6}$$

which is symmetric about $y = 0$. This implies that for any choice of phase c, there will be *two* critical latitudes such that $U(y_c) = c$, which have equal modulus but differ in complex phase by π. The line of reasoning given in Boyd [4] shows that the branch cut from each critical latitude must pass to ∞ in the upper half-plane when dU/dy is positive (which includes (2.1)) and in the lower half-plane when dU/dy is negative. This implies that the contour of integration for (3.6) should pass below the real y-axis for positive y and above the real axis for $y < 0$. The simplest mapping that does this is the cubic

$$y = (A + B)/2 + [(B - A)/2]\{x + i\,\Delta x(x^2 - 1)\}. \tag{3.7}$$

The generalization to a larger number of critical latitudes is straightforward.

Unfortunately, critical latitudes are not the only singuarities that a differential equation is likely to have in the complex plane. Equation (2.1) is atypical in this respect. The wind profile (3.6) generates an equation which is the more usual case: since $\text{sech}(y)$ has poles at $y = \pm i\pi/2$, $\pm i3\pi/2, ...$, the differential equation does also. These can cause major problems in choosing a suitable contour of integration because deforming the contour of integration *across* a branch point of the solution will effectively choose a new, unphysical branch and give a spurious answer. The research of one of the author's students was an inadvertent case study of the need for care. The student, using a single rather large value of Δ in the mapping (3.6) for the mean wind $U(y) = \exp(-y^2)$, obtained interesting unstable modes which had
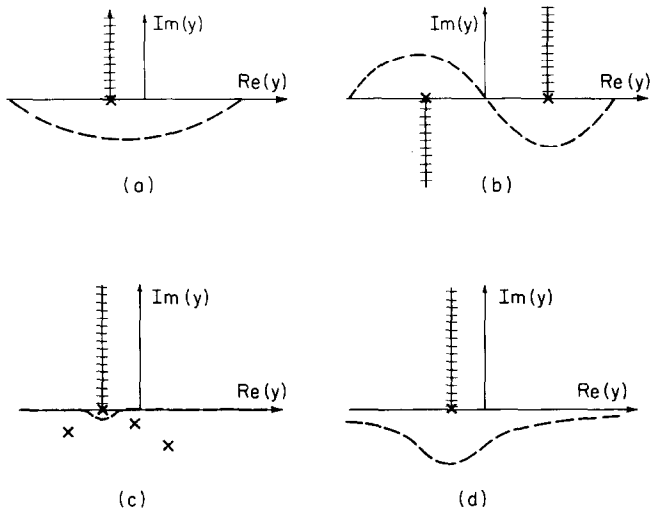
FIG. 1. Four representative complex mappings. For each case, the real and imaginary y-axes are shown as solid lines, the transformed path as a dashed line, and the branch cut away from the singularity on the real axis is depicted as a cross-hatched line. (a) The solution to the differential equation has a branch point on the real axis, and no other singularities close to the interval $[A, B]$. (b) Two singularities on $[A, B]$ with branch cuts that must be taken in opposite half-planes (typical of a symmetric jet). (c) Identical with (a)—only one branch point on $[A, B]$—except for the additional complication of complex singularities which force the mapping to hug the real axis. (d) Identical with (a) except that interval of integration is infinite. The mathematical forms of these mappings are given by Eqs. (3.5), (3.7), (3.8), and (3.10) in the text. The crosses mark the location of the singularities of the differential equation.

no critical latitudes near or on the real y-axis. Alas, the equation $c = \exp(-y^2)$ has an *infinite* number of *complex* roots lying on two hyperbolas in the complex plane. One of the complex roots (and its mirror image) lay between the new contour (3.7) and the real axis, so the student found that all his intriguing instabilities disappeared when he reduced Δ.

Another example is discussed and illustrated (Figs. 1 and 2) in Boyd and Christidis [3]. Although physically the Rossby waves have a *continuous* spectrum, the complex mapping computes discrete modes anyway. The resolution of the apparent paradox is that the computed modes have phase speeds c with small, *negative* imaginary part, which implies that they have critical latitudes between the deformed contour of integration and the real axis, and thus are numerical artifacts.

No general guide can be given except to note that sometimes simple mappings like the parabola and cubic above may need to be replaced by more complicated curves so that the new contour of integration skirts all the singularities of the differential equation, and not merely those that are critical latitudes. An interior boundary layer mapping such as

$$y = x + i\,\Delta e^{-\alpha(x-\zeta)^2} \tag{3.8}$$

where $\alpha \gg 1$ and ζ is the real part of y_c may be useful when the differential equation has other singularities not too far from the real axis as well. For large α, the contour of integration makes but a short detour around the singularity and then hugs the real axis the rest of the way.

This sort of mapping may not work with Chebyshev algorithms because a detour of narrow radius may leave the singularity close enough to the new contour of integration to still poison the convergence of the Chebyshev series. However, it should always work well for a finite difference method—to find the neutral curve of a stability problem, for example—because then one need only choose the grid spacing to be sufficiently small in comparison to $1/\alpha$.

The fourth example of a mapping shows the possibilities when the original computational domain is unbounded. Boyd [7] discusses real-valued mappings that transform an infinite or semi-infinite interval into a finite one, and also the alternative of "domain truncation," i.e., solving the problem on a large but finite interval without the use of a mapping. Either approach can be readily combined with the complex parabolic mapping introduced here. For example, to solve an equation like

$$u_{yy} + [1/y - \lambda - y^2]u = 0 \quad \text{for} \quad y \in [-\infty, \infty] \text{ with } u(\pm\infty) = 0, \quad (3.9)$$

one can use the mapping

$$y = \frac{Lx}{(1-x^2)^{1/2}} + i\Delta(x^2 - 1). \quad (3.10)$$

(Note that the square root was inadvertently omitted in (5.13) of Boyd [7], which gives the *inverse* to the real part of the mapping in (3.10).)

The reader can doubtless invent his own variations, but the point is that paths of integration can be devised to deal with almost any combination of critical latitudes and other singularities of a differential equation. In the next section, we will deal with the harder question of how to optimize the disposable parameters of the mapping.

4. Optimizing the Mapping

Choosing a good value of the map parameter Δ for a finite difference method is easy: pick the real and imaginary parts of Δ so that the contour is two to three grid points away from the singularity. For the parabolic mapping, for example, the imaginary part of y is largest at $x = 0$, and one would like this to correspond to the real part of y_c so that the contour is at its greatest distance from the real y-axis at the point where this separation will do the most good. Since

$$y(x = 0) = i\Delta_r - \Delta_{im} \quad (4.1)$$

where Δ_r and Δ_{im} are the real and imaginary parts of Δ, respectively, it follows that, letting h be the grid spacing in x and assuming that critical latitude y_c is on the real

y-axis, we can obtain a maximum separation of the old and new integration paths of $2h$, directly opposite the singularity, by taking

$$\Delta_r = 2h; \qquad \Delta_{im} = -y_c. \tag{4.2}$$

Since this article was originally motivated by the limitations of Chebyshev algorithms, it is ironic that complex mapping should in some ways work even better for finite difference methods. Because the justification of a finite difference method is a *local* Taylor series expansion whose radius of convergence need extend only to the next grid point, one can employ a complex contour like that suggested in (4.2) which bends only a small distance ($\leqslant 2h$) away from the real axis. This hugging-of-the-real-axis avoids the problem of other complex singularities of the differential equation which can be a major complication for Chebyshev methods as explained in the previous section. On the other hand, it is still proper that this paper should emphasize Chebyshev algorithms: finite difference methods need the complex mapping trick only when the critical latitude is on the real axis itself (the "neutral curve") or when it is within $2h$ of the real axis where h is the smallest affordable grid spacing.

Chebyshev methods are less robust, and the proof follows from the well-known formula [7] that the coefficients of a Chebyshev series are asymptotically of the form

$$b_n \sim (\text{constant})\, n^{-k} \delta^{-n} \tag{4.3}$$

where k and δ are constants. The algebraic factor of n is much less important than the term involving δ and will be ignored hereafter; parenthetically, it can be noted that $k = 0$ for a simple pole, $k = 1$ for a logarithm, and so on.

The constant δ is given by

$$\delta = S \pm (S^2 - 1)^{1/2} \tag{4.4}$$

where S is the (complex) location of the convergence limiting singularity and where the sign in (4.4) is that which makes $|\delta| > 1$. Ignoring the algebraic factor n^{-k} and exploiting the well-known bound on the Chebyshev polynomials [9], $|T_n(x)| \leqslant 1$ for all n and all real $x \in [-1, 1]$, it follows that (asymptotically) a Chebyshev series converges like a geometric series. (More precisely, it can be tightly bounded from above by a geometric series.)

It is instructive to consider the special case of a singularity very close to the real axis. After shifting from $[A, B]$ to the standard interval of $[-1, 1]$ via

$$x_c = (2/[B-A])\{y_c - [A+B]/2\}, \tag{4.5}$$

the location of the singularity is

$$x_c = X + i\varepsilon; \qquad -1 < x < 1, \; \varepsilon \ll 1. \tag{4.6}$$

One can then show, using the equivalent form $\delta = \exp[\text{arccosh}(S)]$, that

$$\delta \sim 1 + \frac{|\varepsilon|}{\sin[\cos^{-1}(X)]}, \tag{4.7}$$

provided that X is not within $O(\varepsilon)$ of ± 1.

Equation (4.6) shows that a singularity close to the interval of integration will invariably wreck a Chebyshev method unless the complex mapping or some other trick is used. If $\varepsilon = 1/100$, for example, then $N = 100$ polynomials will suffice to make the last retained coefficient, b_{100}, smaller than b_0 by only about a single factor of $e = 2.718$. And yet, for the QR algorithm, which has an operation count of $O(N^3)$, taking N even this large is extremely expensive. Put another way, to obtain good results with an unmodified Chebyshev method, one must choose N sufficiently large so that $N/\varepsilon \gg 1$, which is usually too expensive unless $\varepsilon \sim O(1)$, i.e., the singularity is so far from the real axis that it does not matter.

TABLE I

Optimum Values of Δ^a

y_c	δ	Re(Δ)	Im(Δ)
0.0	2.41	0.500	0.0
0.1	2.40	0.500	0.050
0.2	2.39	0.490	0.100
0.3	2.34	0.480	0.150
0.4	2.31	0.460	0.200
0.5	2.26	0.450	0.260
0.6	2.23	0.400	0.300
0.7	2.11	0.370	0.360
0.8	2.04	0.300	0.400
0.85	1.94	0.265	0.430
0.90	1.84	0.220	0.455
0.95	1.69	0.150	0.480
0.98	1.54	0.100	0.492
0.99	1.43	0.0712	0.499
0.998	1.27	0.0312	0.500

[a] Optimum values of the map parameter $\Delta = \text{Re}(\Delta) + i\,\text{Im}(\Delta)$ for various locations of the critical latitude y_c along the real axis where the new coordinate x is related to the original coordinate y via $y = x + i\Delta(x^2 - 1)$. Note that it is assumed that the computational domain has been normalized to $y\varepsilon[-1, 1]$. Note also that $\Delta_{\text{opt}}(y_c) = \Delta_{\text{opt}}(-y_c)$, so only $y_c \geqslant 0$ need be listed. The error E_N in the Chebyshev series in x decreases geometrically with N where N is the number of Chebyshev polynomials retained in the truncation; $\delta = E_N/E_{N+1}$ when the corresponding optimum Δ is used in the mapping.

Computing that value of the map parameter Δ which makes δ as large as possible is analytically difficult but numerically easy: one merely evaluates $\delta(\Delta)$ from (4.4) for various Δ and a particular location of the singularity to see which Δ works best. One can simplify the task by noting that (4.4) is symmetric about $S = 0$, which means that it is sufficient to consider $\text{Re}(y_c) \geqslant 0$, and also that one can always rescale the interval to $[-1, 1]$ as in (4.5). If we make the additional restriction of real y_c, since singularities on the real axis are the worst, it is then enough to examine singularities located on the real interval $[0, 1]$ alone.

Table I shows the results for the parabolic mapping (3.5). What is remarkable is that for $|x_c| < 0.8$, it is possible to achieve $\delta > 2.0$—in other words, with the optimized mapping, each successive Chebyshev coefficient is (asymptotically) less than half its predecessor so that one is roughly halving the error with each additional term kept in the series. Table II shows that it is possible to achieve these good results in practice: the lowest eigenvalue for (2.1), for example, is computed to an accuracy of better than 1 part in 100,000 with just ten Chebyshev basis functions.

Even when the singularity is rather near the boundaries, Table I shows that one can still obtain fairly good results. Although the "wall modes" of the Orr–Sommerfeld equation do have critical latitudes close to the endpoints, it is possible that complex mapping would improve numerical efficiency in spite of the warning given in Section 2. The reason for this success near $x = 1$ in Table I is that the Chebyshev polynomials oscillate more and more rapidly near ± 1; they have a "built-in quadratic variable stretching" near the endpoints as discussed in Boyd [10]. Crude curve-fitting suggests that

$$\delta \sim 1 + 1.3(1 - x_c)^{1/4}; \qquad \Delta \sim (1/2)i + 0.7(1 - x_c)^{1/2} \tag{4.8}$$

for $|1 - x_c| \ll 1$, but no attempt has been made to justify (4.8) analytically because a refined approximation to the limits as $x_c \to 1$ would be useful only where the complex mapping would fail, and one would be forced to use viscous damping instead. Nonetheless, Table I shows that one can go very close to the boundaries without sacrificing the benefits of the complex change of coordinates.

5. NUMERICAL ILLUSTRATIONS

Table II presents the lowest ten eigenvalues for (2.1) as computed using the parabolic mapping with optimum Δ. The Chebyshev–QR algorithm employed is identical to that described in Boyd [10, 11] except that the differential equation (2.1) was first transformed via (3.2) and (3.5). The complex mapping is a striking success. In the absence of the transformation, the Chebyshev series of $u(y)$ does not converge because $u(y)$ has a logarithmic singularity at $y = 0$, so the pseudospectral method fails. With the mapping, however, it is possible to compute the lowest eigenvalue to an accuracy of 1.4 % with just six basis functions.

TABLE II

The First 10 Eigenvalues of $u'' + [1/|y - \lambda]u = 0$ Subject to $u(-6) = u(6) = 0$ as Computed Using the Complex Parabolic Mapping $y = x + i\Delta(x^2 - 1)$ with $\Delta = \frac{1}{2}i^\alpha$

n	N = 6	%	N = 10	%	N = 20	%	N = 30	%	N = 40
1	0.1268 +i 0.2852	1.4 .09	0.125053 +i 0.284986	.0008 0.	0.125054 +i 0.284986	0. 0.	0.125054 +i 0.284986	0. 0.	0.125054 +i 0.284986
2			-0.2980 +i 0.5169	0.43 0.65	-0.296730 +i 0.520320	0. 0.	-0.296730 +i 0.520320	0. 0.	-0.296730 +i 0.520320
3			-0.655 +i 0.034	2.5	-0.638802 +i 0.000262	0. 0.	-0.638802 +i 0.000262	0. 0.	-0.638802 +i 0.000262
4					-1.21749 +i 0.563350	.0008 .0004	-1.21750 +i 0.563349	0. 0.	-1.21750 +i 0.563349
5					-1.60228 +i 0.007844	.024 .24	-1.60189 +i 0.007820	.0006 .064	-1.60190 +i 0.007825
6							-2.72604 +i 0.489975	.0022 .0076	-2.72598 +i 0.490012
7							-3.10131 +i 0.070353	.0055 .25	-3.10114 +i 0.070176
8							-4.693 +i 0.286	.47 13.0	-4.71486 +i 0.329311
9									-5.23748 +i 0.228039
10									-6.94888 +i 0.168790

[a] N is the number of Chebyshev basis functions retained in the truncation. Errors shown for $N = 30$ and smaller truncations were computed by comparison with the eigenvalues for $N = 40$. Numerical values which violate the theorem that the imaginary part of the eigenvalue is always positive, or which are wildly in error, have been omitted.

A rough rule-of-thumb in nonsingular Sturm–Liouville problems of the first kind is that when N Chebyshev basis functions are used, $O(N/2)$ of the eigenvalues of the resulting $N \times N$ matrix will be good approximations (to within a few percent) of eigenvalues of the original differential equation. The larger half of the numerical eigenvalues will be wildly inaccurate. Table II shows that the price that is paid for the singularity is that the rule-of-thumb must be amended to read that for a Sturm–Liouville problem of the fourth kind, only $O(N/4)$ of the matrix eigenvalues will be reasonable approximations to those of the differential equation.

This is still a vast improvement over none, however. As explained in Boyd [4], the WKB method gives excellent analytic approximations to the large eigenvalues, so one only needs to compute the lowest few via Chebyshev polynomials anyway. Thus, the complex mapping has turned a numerical disaster into a rather easy and inexpensive means of computing all the important eigenvalues.

Table II is also a dramatic, if personally embarrassing, justification of the jeremiad against local, iterative methods given in the introduction. A local, finite difference-cum-secant iteration *missed* the third and the seventh modes listed in Table II. Table IV and Fig. 10 of an earlier work by the author [4] are both incorrect in that these two eigenvalues were omitted. The suggestion that the mode number could be different from the integer m that appears in the WKB formula for a given mode is cast in doubt since the two extra modes destroy Table IV as an example of such behavior. Happily, all the other conclusions of Boyd [4] are unaffected by this mistake.

The results of the earlier paper were computed by careful use of the so-called "continuation method," varying the parameters in small steps away from a set of eigenvalues calculated using a linear friction/interior boundary layer mapping (and 60 basis functions!) as in Boyd [7]. However, eigenvalues with small imaginary parts—true of both missed modes—are difficult to compute by iteration because the eigenfunction jumps to a different branch when the sign of the imaginary part of λ reverses, so the radius of convergence of any iterative procedure goes to 0 with $\mathrm{Im}(\lambda)$. Thus, the singularity of the eigenfunction which wrecks the Chebyshev–QR algorithm also greatly *weakens* the competing local *iterative* methods. Although safe enough, if used wisely, for self-adjoint, nonsingular Sturm–Liouville problems of the first kind, local finite difference iterative algorithms are very risky and liable to miss modes when applied to the much harder case of singular Sturm–Liouville problems of the fourth kind or to hydrodynamic instability equations—as the author learned to his sorrow.

Boyd [7] explored the usefulness of purely real mappings for solving problems via Chebyshev polynomials on an *unbounded* interval on the basis of Chebyshev convergence theory without providing any explicit examples. To remedy this lack of examples and simultaneously illustrate the principle that *complex* mappings are

problem of barotropic stability. (Haltiner and Williams [12] give a good discussion of the physics of this example.) Because the differential equation is symmetric about

TABLE III

The Eigenvalue c^a

L	Symmetric mode		Antisymmetric mode	
	c	rel. error (%)	c	rel. error (%)
Exact	-0.700764	—	-0.787731	—
	$+i \quad 0.190689$	—	$+i \quad 0.042011$	—
10	No unstable modes			
5	-0.6692	-4.5	-0.735	-6.7
	$+i \quad 0.1902$	0.23	$+i \quad 0.0815$	-94.0
2	-0.699084	-0.24	-0.764	-3.0
	$+i \quad 0.19197$	-0.13	$+i \quad 0.0518$	-23.3
1	-0.700571	-0.027	-0.7938	0.77
	$+i \quad 0.190585$	0.055	$+i \quad 0.0474$	-12.9
0.75	-0.700398	-0.052	-0.78517	-0.33
	$+i \quad 0.190350$	0.18	$+i \quad 0.0478$	-13.8
0.50	-0.705	0.71	-0.7958	1.0
	$+i \quad 0.1888$	0.99	$+i \quad 0.0454$	-8.1
0.25	-0.687	-1.9	-0.7857	-0.25
	$+i \quad 0.229$	-20.0	$+i \quad 0.0304$	27.6

[a] The eigenvalue c for the two unstable modes of the Bickley jet [12] with $U(y) = -\text{sech}^2(y)$, $\beta = 0.4$, $k = 1.0$ (all nondimensional) as computed with 20 Chebyshev basis functions (10 symmetric and 10 antisymmetric) and the mapping $y = Lx/(1 - x^2)^{1/2}$ where x is the argument of the Chebyshev polynomials. These values of the zonal wavenumber k and parameter β give the largest growth rates for any k, β [12], so the *complex* coordinate transformation is unnecessary.

$y = 0$, one can halve the number of basis functions by using only symmetric Chebyshev functions to compute the unstable mode which is symmetric about $y = 0$, and similarly only antisymmetric functions to find the antisymmetric eigenvalue. Note that although only ten basis functions of the proper symmetry are used, excellent results can be obtained with the proper choice of L because the disposable parameters β and k of the differential equation were chosen to correspond to the region of *largest* instability as shown in [12].

As predicted by the theory of Boyd [7], both algebraic mapping and domain truncation work, but the mapping is more successful here because, again as shown by the theory, it is much less sensitive than "domain truncation" to the criticial latitude singularities for complex y, and also much more indifferent to the choice of the parameter L. Even though the singularities for this choice of parameters are sufficiently far from the real axis so that we can solve the problem without a complex transformation, the critical latitudes still control the choice of numerical method—in the absence of singularities, Boyd [7] shows that domain truncation

TABLE IV

The Eigenvalue c^a

L	Symmetric mode		Antisymmetric mode	
	c	Rel. error (%)	c	Rel. error (%)
10	No unstable modes			
5	−0.6709	−4.3	−0.736	6.6
	+i 0.1919	−0.63	−0.081	93.0
2	−0.6478	−7.6	Stable	
	+i 0.1684	13.6		
1	−0.672	−4.1	Stable	
	0.041	78.5		
Exact	−0.7008	—	−0.788	—
	+i 0.1907		+i 0.042	—

^{a} Same as Table III (two unstable modes of the Bickley jet) except that domain truncation—solving the problem on the domain $y \in [-L, L]$ such that $u(\pm L) = 0$—was used instead of mapping originals of tables.

gives more accurate results. When the parameters are changed to bring us close to the neutral curve where the eigenvalue c is real, both algebraic mapping and domain truncation are ruined. By combining them with the complex mapping as suggested in Section 3, however, we can follow a mode right up to the neutral curve without trouble—and using only a small number ($O(10)$) of basis functions.

6. Computing the Eigenfunctions

As noted earlier, the use of a *complex* contour of integration makes it impossible to directly compute the eigenfunctions on the *real* y-axis. This is obvious for a finite difference method since it calculates $u(y)$ only on the grid points, and these are now complex. It is less obvious for a spectral method since the Chebyshev series can be evaluated with a complex argument in the map variable x to give $u(y)$ for real y. Unfortunately, this does not work because the Chebyshev series must diverge at all singularities. If the eigenfunction has a branch point for real y, then the Chebyshev series will still fail utterly in the neighborhood of the branch point even though it converges rapidly on the complex integration path where its argument x is real.

One is left with two options: to be satisfied with only the eigenvalues or to compute the eigenfunctions in a second, separate step of solving an ordinary boundary value problem.

The first alternative, computing the eigenvalue only, is often painless because the eigenvalue is the only quantity of physical interest anyway. In the "Bickley jet"

barotropic stability problem [12], for example, Haltiner and Williams give a contour plot of the eigenvalue for all values of the parameters and graph the neutral curve. The region of small growth rates and the neutral curve cannot be calculated via a Chebyshev–QR algorithm without the complex change of coordinates. However, their graphs of sample eigenfunctions are only for the *fastest growing* modes—and we have already seen in Tables III and IV that the strongly unstable modes can be computed *without* the complex transformation.

Alternatively, once the eigenvalues are known, the eigenfunctions can be found one at a time by solving an ordinary boundary value problem. The latter is much easier than an eigenvalue problem, so the complex mapping is valuable in cutting the problem down to size even when the eigenfunctions are needed, too.

In particular, the QR algorithm requires roughly $O(10N^3)$ operations to find the eigenvalues of an $N \times N$ matrix. In contrast, a boundary value problem generates an ordinary matrix equation which can be solved by Gaussian elimination in $O(N^3/3)$. This means that if the singularity is not right on the real axis, but only close to it, one can compute the eigenvalues via N Chebyshev polynomials with complex mapping and then compute the most unstable eigenfunction by solving a boundary value problem on the real axis with triple the number of polynomials, and still keep the cost of the second step under that of the QR step. Using $3N$ polynomials in a brute force attempt to compute the eigenvalues and eigenfunctions together via the Chebyshev–QR algorithm without complex mapping would be roughly 15 times more expensive.

When the offending singularity is actually *on* the real axis, as it is for hydrodynamic instability problems on the neutral curve and for Sturm–Liouville problems of the fourth kind, stronger measures are needed. One possibility is to combine a finite difference method with the local, Gaussian complex mapping described by (3.8) above. The Gaussian map hugs the real axis except within one or two mesh points of the singularity. If the mesh is fine enough—and it can be variable, if desired—a smooth graph can be made with extrapolation from one side of the singularity to the other, using only grid points on the real axis. (Parenthetically, it should be noted that in the problems of Boyd [6], the singularities are of the form $y \log y$ or $\log y$, depending on the quantity, so that one can make smooth, useful graphs in spite of the real-axis singularities.) Because of (4.7) above, this device of making a semi-circular detour of very small radius will work only with a grid point method.

A final remedy, which can be applied with either finite difference or Chebyshev methods, is to use a local power series approximation in the neighborhood of the singularity and match it to finite difference or Chebyshev series that handle the rest of the interval. This procedure, in conjunction with a Runge–Kutta shooting method, was used to calculate the "exact" solutions of Boyd [4], for example. This matching procedure is messy; in combination with Chebyshev polynomials, one ends up with three separate series (two Chebyshev flanking one power series with logarithms), each representing the same function on a different interval. Since the power series coefficients normally depend nonlinearly on the eigenvalue, it is

usually not possible to convert such a hodge-podge of different series into a *linear* eigenvalue problem, so the QR method cannot be used and one is again forced back to local, iterative methods to solve the resulting algebraic equations. These, as stressed above, require first guesses for the solution and may easily miss modes. However, if the eigenvalues are already known from a calculation using complex mapping, then these difficulties disappear and calculating the eigenfunctions this way is straightforward and safe.

Thus, the complex mapping method does not make it impossible to calculate the eigenfunctions; rather, it makes it easy since one has a very wide range of options once the eigenvalues are known. The price that must be paid for the complex change of coordinates is that the eigenvalues and eigenfunctions, if both are needed, must be calculated in two separate steps.

7. SUMMARY AND CONCLUSIONS

The value of a change of coordinates for resolving boundary layers [13] or an unbounded interval [7] is well known. What we have tried to show here is that a complex change of coordinates such that the differential equation is integrated along a curve in the complex plane rather than along the real axis is an equally powerful technique for computing the eigenvalues of inviscid hydrodynamic stability problems and others that have singularities ("critical latitudes" or "critical levels") on or near the original real interval.

The choice of the mapping is very flexible; some suggestions are given in Section 3 above. The complex transformation is particularly useful with Chebyshev methods since these, as a price for their high accuracy and efficiency, are very sensitive to singularities. The idea is quite general, however, and can be usefully applied with finite difference methods as well. Although it seems difficult or impossible to analytically optimize different species of mapping, Table I shows that it is possible to obtain some useful guidelines for the map parameters with very little effort.

The great strength of the complex mapping is that it veers the path of the integration away from the singularity so that it is as if the singularity never was. Its major disadvantages are (i) the lack of a general theory for optimizing the mappings (but see Table I), (ii) the need to avoid deforming the new path of integration across other complex singularities of the differential equation, and (iii) the need to compute the eigenfunctions as a second, separate step.

Still, differential equations with singularities right on the interval of integration are very difficult. The complex change of coordinates is a very practical and conceptually simple means for coping with these singularities.

ACKNOWLEDGMENTS

REFERENCES

1. S. A. ORSZAG, *J. Fluid Mech.* **50** (1971), 689.
2. J. P. BOYD AND Z. D. CHRISTIDIS, *Geophys. Res. Lett.* **50** (1982), 769.
3. J. P. BOYD AND Z. D. CHRISTIDIS, *in* "Hydrodynamics of the Equatorial Ocean" (J. C. J. Nihoul, Ed.), Elsevier, Amsterdam, 1983.
4. J. P. BOYD, *J. Math. Phys.* **22** (1981), 1575.
5. A. J. SIMMONS, *Quart. J. Roy. Meteorol. Soc.* **102** (1978), 595.
6. J. P. BOYD, *J. Atmos. Sci.* **39** (1982), 770.
7. J. P. BOYD, *J. Comput. Phys.* **45** (1982), 43.
8. D. GOTTLIEB AND S. A. ORSZAG, "Numerical Analysis of Spectral Methods: Theory and Application," Soc. Ind. Appl. Math., Philadelphia, 1977.
9. D. FOX AND L. PARKER, "Chebyshev Polynomials in Numerical Analysis," Oxford Univ. Press. London, 1968.
10. J. P. BOYD, *Mon. Wea. Rev.* **106** (1978), 1192.
11. J. P. BOYD, *J. Math. Phys.* **19** (1978), 1445.
12. G. HALTINER AND R. T. WILLIAMS, "Numerical Weather Prediction and Dynamic Meteorology," 2nd. ed., Wiley, New York, 1980.
13. E. KALNAY DE RIVAS, *J. Comput. Phys.* **10** (1972), 202.