Available online at www.sciencedirect.com
ScienceDirect
Journal of Computational Physics 227 (2007) 1209–1224
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Spectral radial basis functions for full sphere computations

Philip W. Livermore [a,*], Chris A. Jones [a], Steven J. Worland [b]

[a] School of Mathematics, Leeds University, Leeds, LS2 9JT, UK
[b] Formerly at School of Engineering, Computing and Mathematics, University of Exeter, EX4 4QF, UK

Received 27 June 2007; received in revised form 24 August 2007; accepted 31 August 2007
Available online 11 September 2007

Abstract

The singularity of cylindrical or spherical coordinate systems at the origin imposes certain regularity conditions on the spectral expansion of any infinitely differentiable function. There are two efficient choices of a set of radial basis functions suitable for discretising the solution of a partial differential equation posed in either such geometry. One choice is methods based on standard Chebyshev polynomials; although these may be efficiently computed using fast transforms, differentiability to all orders of the obtained solution at the origin is not guaranteed. The second is the so-called one-sided Jacobi polynomials that explicitly satisfy the required behavioural conditions. In this paper, we compare these two approaches in their accuracy, differentiability and computational speed. We find that the most accurate and concise representation is in terms of one-sided Jacobi polynomials. However, due to the lack of a competitive fast transform, Chebyshev methods may be a better choice for some computationally intensive timestepping problems and indeed will yield sufficiently (although not infinitely) differentiable solutions provided they are adequately converged.

Keywords: Coordinate singularity; Jacobi polynomials; Chebyshev polynomial; Spectral method; Polar coordinates

1. Introduction

In a cylindrical or spherical geometry, the origin is a singular point. This is manifested not only in the convergence of grid points in many numerical schemes severely restricting tractable timesteps, but in certain regularity conditions that any solution must satisfy in order to remain differentiable to all orders at the origin. This issue may be regarded in a positive or negative light depending on the point of view. On the one hand, one may be concerned that any numerically derived function might not be sufficiently differentiable (and therefore will not be physically meaningful) at the origin; on the other, the extra constraints may be exploited to hone the numerical scheme and consequently speed up convergence. A thorough review of these issues may be found in [1].

* Corresponding author.
E-mail address: phil@ucsd.edu (P.W. Livermore).

0021-9991/$ - see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2007.08.026

In a 2D polar geometry, any smooth (i.e. infinitely differentiable) function $f(r, \phi)$, depending on the radius $r \in [0,1]$ and polar angle $\phi \in [0, 2\pi]$ has a Fourier expansion of the form

$$f(r,\phi) = \sum_n \mathrm{e}^{\mathrm{i}n\phi} f_n(r), \quad f_n(r) = r^n g_n(r),$$

where the rightmost equation expresses the regularity condition and $g_n(r)$ is necessarily itself both even and smooth [1–3]. Similarly in the spherical case, an expansion in terms of spherical harmonics $Y_l^m(\cos\theta)\mathrm{e}^{\mathrm{i}m\phi}$ where (θ, ϕ) are respectively colatitude and longitude implies that the multiplying radial function must be of the form $r^l g_{lm}(r)$. Thus in both cases (and for the remainder of this paper) we may speak of the regularity condition being that $f_l(r) = r^l g_l(r)$ (with g_l smooth and even), although the physical interpretation of the wavenumber l (referring either to polar angle or colatitude) depends on the coordinate system. An immediate consequence of the regularity condition is that each radial function has a definite parity, a property that in fact may be derived independently of differentiability by, for example, identifying the point $(-r, \phi + \pi)$ with (r, ϕ) in plane polar coordinates [4]; an analogous result holds in the spherical polar case. Note that the parity and regularity conditions are neither required nor in general satisfied by any solution on a domain that excludes the origin.

Any smooth solution of a partial differential equation automatically satisfies the regularity conditions. It follows that any solution produced by a convergent numerical scheme will also satisfy these conditions, at least in the limit of infinite truncation and infinite precision. In practice however, highly differentiable solutions will only be obtained in general with numerical schemes that converge quickly in truncation level, for the imprecision caused by the accumulation of roundoff errors will violate the regularity conditions.

Methods based on Chebyshev expansions, both in spectral [5–8] and interpolation [9–12] forms have been widely used to represent the radial structure, the major advantage being the availability of a fast transform (FFT). The potential clustering of the associated grid points in physical space close to the origin may be removed by either expanding over the double interval $[-1, 1]$ (rather than $[0, 1]$ with the concomitant halving of the angular extent) or by exploiting their parity. However, in order to satisfy the regularity conditions at the origin, the Chebyshev polynomials must effect a perfect cancellation of all monomial terms of the form r^i with $i < l$. Assuming that the unknown solution has a nonzero projection onto each Chebyshev polynomial (which might be exponentially small) and noting that each Chebyshev polynomial $T_n(r)$ has a nonzero projection onto all monomials r^i of the same parity and degree $i < n$, regularity will never be achieved exactly at any given finite truncation. This therefore raises the question of how much the differentiability of the solution at the origin is compromised using such a method. One of the aims of this paper, addressed in Section 2, is to tackle this very question – to determine whether or not such a Chebyshev scheme provides a regular solution, or at least one that is sufficiently regular. Precisely how regular a solution needs to be to remain physically meaningful will depend on the problem: for most cases, only the first few derivatives need to be everywhere smooth; for others, the solution may need to be infinitely differentiable.

In such methods, regularity may be significantly affected by numerical imprecision introduced through the accumulation of roundoff errors. An important related issue therefore is the speed of convergence of the numerical scheme. Gottlieb and Orszag [5] claimed that the coordinate singularity degraded the convergence of Chebyshev methods and that introducing additional "pole" conditions (of the form $y'(0) = 0$) speeded up convergence. In Section 2 we briefly revisit the issue of whether such extra conditions are required with our implementation.

An alternative method is to expand the unknown function in a radial basis that automatically satisfies all regularity conditions. The effect of roundoff error or otherwise lack of convergence can never degrade the differentiability which is guaranteed to all orders. In addition, the fact that the correct behaviour at the origin is built into the basis may significantly accelerate global convergence. One natural choice are Bessel (or spherical Bessel) functions being the radial part of the separable solution to the Helmholtz equation $\nabla^2 f + \lambda^2 f = 0$. However, Bessel functions are not solutions of a sufficiently singular Sturm–Liouville problem [13,5] and therefore only achieve algebraic convergence (as attested by [14]). Nevertheless, Bessel functions have been used successfully in certain applications e.g. [15]. It is worth noting that the same algebraic convergence is also obtained when using Fourier series on a non-periodic domain [16]. A further option are the Poincaré polynomials, eigenfunctions of the inertial wave equation in 3D [17] although again there is no reason to suspect that

they would perform any better than Bessel functions. A better choice is to expand in terms of the so-called one-sided Jacobi polynomials [1] of the form

$$G_n(r;\alpha,\beta,l) = r^l P_n^{(\alpha,\beta)}(2r^2-1).$$

Jacobi polynomials, $P_n^{(\alpha,\beta)}(x)$, of which Chebyshev and Legendre polynomials are familiar examples, are solutions of a singular Sturm–Liouville problem and as such exhibit spectral convergence to sufficiently smooth solutions in $x \in [-1,1]$ [5]. Various choices of (α, β) have been suggested in the literature. The so-called Robert functions, $r^l T_n(r)$, corresponding to $\alpha = \beta = -1/2$, although perhaps an obvious choice are known to be exceptionally numerically poorly behaved [1]. Both Verkley [18] and Matsushima and Marcus [19] independently proposed the parameters $\alpha = 0$, $\beta = l$ producing a family of polynomials that are orthogonal with respect to the weight function $w(r) = r$. Although this is a natural property, at least in 2D polar coordinates, Worland [14] pointed out that large oscillations close to $r = 1$ might lead to poor conditioning and slow convergence. Worland's choice of $\alpha = -1/2$, $\beta = l - 1/2$ corresponds to basis functions that oscillate within an asymptotically uniform envelope as $n \to \infty$.

Although fast transforms from spectral representation to physical space do exist for Jacobi polynomials (or indeed any polynomial related by a three-term relation [20]), they are no faster than a standard slow transform for n less than O(500) due to high overheads (c.f. the fast Legendre transform [21,22]). Thus a major drawback of such methods is the potentially large computational cost and their usefulness depends entirely on how their rate of convergence compares with a standard Chebyshev expansion. Should they converge slower as suggested by Mohseni and Colonius [9], then a Chebyshev series might be the preferred choice although infinite differentiability is not guaranteed. However, should they exhibit superior convergence (as may be expected and indeed as confirmed in Section 3) then they would be ideally suited for problems where compactness of spectral representation is more important than computational speed, such as optimisation problems that do not involve timestepping [23,24]. An additional issue is to determine which choices of (α, β), if any, generate a family G_n that is spectrally convergent. In fact, as we show in Appendix C, any family of G_n, with $\alpha > -1$, $\beta > -1$, converges exponentially fast to functions of the form $r^l g(r)$ where g is even and smooth.

The aim of this paper is to provide a quantified comparison between radial expansions in Chebyshev polynomials and various choices of one-sided Jacobi bases, in order to provide the reader with some guidance as to which is the optimum choice for any particular case. We benchmark the various methods against the least negative eigenvalue (and corresponding eigenvector) of Bessel's equation, an eigenvalue problem stemming from the 2D Helmholtz equation, rather than solving a multidimensional problem such as the Navier–Stokes equations in full. Bessel's equation contains all the crucial ingredients with which we can test the numerical schemes: the exact solutions are known analytically so it is straightforward to determine convergence, and as the solutions are everywhere differentiable, reasonable regularity is expected of any numerical approximation.

We begin in Section 2 by analysing the rate of convergence of a Chebyshev expansion, paying particular attention to the degree of regularity of the solution at the origin. Although such numerical solutions can never be infinitely differentiable (as has been already noted), we investigate what effect this lack of differentiability has on the behaviour of the solutions. In Section 3 we discuss one-sided Jacobi polynomials, detailing the motivation for the various choices of (α, β), including the choice of Worland only recently appearing in the literature, and providing a comparison of their convergence with the Chebyshev method of Section 2. We also investigate how well the different one-sided Jacobi bases resolve boundary layers, and show that the new Worland polynomials are analogous to Chebyshev polynomials in being close to the approximation attaining minimum maximum pointwise error (minimax). We conclude with a discussion in Section 4.

2. Expansions in Chebyshev polynomials

We consider solving Bessel's equation

$$[\mathcal{D} + \lambda^2]y(r) = 0, \quad \mathcal{D} = \frac{\mathrm{d}^2}{\mathrm{d}r^2} + \frac{1}{r}\frac{\mathrm{d}}{\mathrm{d}r} - \frac{l^2}{r^2}, \tag{1}$$

with the boundary conditions $y(0) = 0$, $y(1) = 0$. The analytic solutions are $J_l(\lambda_i r)$ where J_l is a Bessel function of order l and λ_i is the ith positive root of J_l. These solutions satisfy the regularity conditions, being of the form $r^l g_l(r)$ with g_l even and smooth.

In order to construct a solution based on Chebyshev polynomials, it is a poor idea to build a prefactor of r^l into the expansion due to severe numerical ill-conditioning especially when l is large [1]. One can, of course, premultiply by r^2 or r depending on whether l is even or odd, thus at least partially building in regularity, although again even this can lead to numerical problems [6]. Much more successful are expansions exploiting the parity of y, expanding in one of the two forms

$$y_N(r) = \sum_{n=1,l \text{ even}}^{n=N} b_n T_{2n-2}(r), \quad y_N(r) = \sum_{n=1,l \text{ odd}}^{n=N} b_n T_{2n-1}(r),$$

depending on whether l is an odd or even integer. We may construct a generalised eigenvalue problem by using a tau method, imposing the orthogonality condition

$$2\int_0^1 (1-r^2)^{-1/2} T_{2i-1}(r) R_N(r) = 0, \quad R_N(r) = (\mathcal{D} + \lambda^2) \sum_n b_n T_{2n-1}, \tag{2}$$

for each $i = 1, 2, \ldots, N-1$ and l is here supposed odd. Note that by virtue of construction the boundary condition $y(0) = 0$ is satisfied automatically; $y(1) = 0$ is imposed on the last row of the matrix system. We have exploited parity of both $T_{2n-1}(r)$ and R_N to reduce the orthogonality range $[-1, 1]$ to $[0, 1]$ as the integrand is even; such integrals may either be computed by fast transform or quadrature. Importantly, although we work only on $[0, 1]$ the grid points do not cluster near $r = 0$ and we avoid any concomitant numerical problems. Such a method is certainly not new and performance for $l = 7$ is provided in [5] although they report that the apparently poor convergence is improved if the additional "pole" condition $y'(0) = 0$ is implemented, replacing one of the tau-orthogonality rows of the matrix system. The reported necessity of such supplementary conditions is, at the very least, somewhat unnerving as of course for functions that are $O(r^l)$ as $r \to 0$, all of the first $l - 1$ derivatives are similarly zero; if $l - 1 > N$, we cannot possibly use all the conditions as this would entail more constraints than unknowns. In attempting to reproduce their results, contrary to their original study, we find that in fact convergence is marginally degraded rather than improved with the addition of extra boundary conditions. This is demonstrated in Fig. 1 which shows convergence to the least negative eigenvalue of Bessel's equation with $l = 7$ and $l = 101$. These (and subsequent) calculations were performed in double precision arithmetic with round-off errors of $O(10^{-14})$. For any particular truncation, the grey curves (with the additional pole condition) are of similar or greater error than those without (shown as black).

Thus the Chebyshev polynomials need no assistance from additional pole conditions in converging to the analytic solution, excellent convergence (within an error of $O(10^{-10})$) being achieved both in the eigenvalue

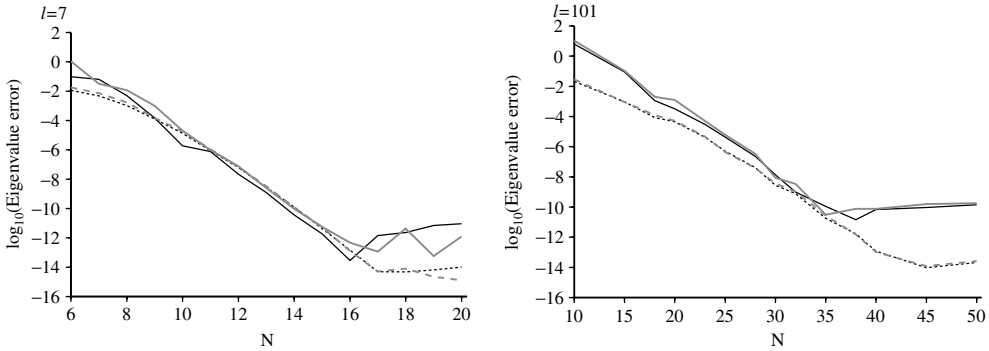

Fig. 1. Convergence to the least negative eigenvalue of Bessel's equation of order $l = 7$ and $l = 101$ using a Chebyshev-tau method with N odd polynomials. The solid lines show $\log_{10}$ of the absolute eigenvalue error, dashed show $\log_{10}$ of the rms error in the eigenvector, the functions being normalised to attain a maximum value of unity over $[0, 1]$. The grey curves have an extra pole condition imposed: $y'(0) = 0$ that replaces the last tau-orthogonality row of the matrix system; the black curves have no additional conditions.

and eigenvector by $N = 14$ and $N = 35$ for $l = 7$ and $l = 101$ respectively. The latter case is somewhat surprising, given that we require $N > 50$ to have any chance of generating a solution that is formally $O(r^{101})$; for $N = 35$ the solution is at most $O(r^{69})$. There are two conclusions we can draw from this. Firstly, low degree polynomials can approximate high degree polynomials extremely well (at least in an average sense), as shown by the convergence of both the eigenvalue and eigenvectors. Secondly, regularity does not automatically follow from convergence: it is impossible that the solution (at least with $N = 35$) can fulfill the regularity conditions and so, for some $k < l$, $y_N(r)/r^k$ is singular at $r = 0$. The solution y_N shows little sign of changing (at least in its convergence) as N increases from 35, indeed there are signs that the error starts to increase slightly with truncation (due to the effects of accumulated roundoff error), so it would appear that for no choice of N is a regular solution attainable.

2.1. Regularity of Chebyshev solutions

Having shown that it is in general difficult to obtain fully regular Chebyshev solutions, it is of interest to quantify by how much such solutions violate the required conditions; we do this using two different methods. Firstly, we follow [3] and compute the monomial expansion of the truncated Chebyshev series approximation to the solution of Bessel's equation of order $l = 21$:

$$y_N(r) = \sum_{n=1}^{N} b_n T_{2n-1}(r) = \sum_{i=0}^{2N-1} a_i r^i.$$

We measure the degree of nonregularity by the squared power of the coefficients that lead to a nonregular solution, $\sum_{i<l} a_i^2$. The intermediate value of $l = 21$ was chosen to show typical behaviour; in comparison to Fig. 1, convergence within $O(10^{-10})$ is achieved by $N = 25$ for this value of l.

The results are shown in Table 1. In the second column from the left, the squared power in the "nonregularity" coefficients reaches a minimum at $N = 32$. The coefficients a_n in general individually become very large, as can be seen by comparing with the third column, so to achieve a value as low as 6.523e + 01 a remarkable degree of cancellation has occurred. Thus the solution is "most regular" at a truncation higher than that for which the eigenvector and eigenvalue have converged (at $N = 25$). It is also apparent that the solution is never fully regular; indeed, as the truncation increases, the nonregularity increases without bound. A clue for why this occurs may be gleaned from the third and fourth columns, which list respectively the squared power in all monomial and Chebyshev coefficients as a function of N. Although the Chebyshev series

Table 1
Squared power of coefficients as a function of N, the total number of odd Chebyshev polynomials used in the Chebyshev-tau approximation to the least negative eigenvalue of Bessel's equation of order $l = 21$

N	$\sum_{n<21} a_n^2$	$\sum_{n=0}^{n=2N-1} a_n^2$	$\sum_{n=1}^{n=N} b_n^2$
10	1.238e + 08	1.238e + 08	1.713769
15	7.557e + 08	1.330e + 10	1.713999
20	2.277e + 10	6.668e + 12	1.713999
25	9.315e + 05	3.838e + 14	1.713999
30	1.322e + 04	7.517e + 14	1.713999
31	7.217e + 02	9.307e + 14	1.713999
32	6.523e + 01	1.100e + 15	1.713999
33	2.064e + 03	8.266e + 16	1.713999
35	4.174e + 04	1.790e + 20	1.713999
40	3.125e + 05	3.205e + 27	1.713999
45	2.499e + 09	2.331e + 35	1.713999
50	6.530e + 11	1.381e + 43	1.713999

The b_n are the coefficients of $T_{2n-1}(r)$, and a_n are the coefficients of r^n (which are incidentally 0 when n is even in this case). The second column shows the squared power in the "nonregularity coefficients" giving rise to the singularity of some derivative at $r = 0$; the third gives a comparison with the total squared power of the monomial coefficients and the fourth the sum of the squared Chebyshev coefficients. The absolute error in both the eigenvalue and eigenvector is $O(10^{-10})$ for $N \geqslant 25$.

converges quickly (so that the coefficients fall off exponentially), due to numerical imprecision (even in double precision) the magnitude of the b_n will fall no lower than $O(10^{-14})$. At the same time, the monomial coefficients of $T_{2n-1}(r)$, when n is large, become increasingly significant; for example, the largest coefficients of $T_{49}(r)$ are $O(10^{17})$. It follows that at large truncations there is great deal of numerical "noise" that prevents the monomial coefficients from converging and the solution from becoming regular.

However, all that we can conclude from this type of analysis is that the solution is never formally regular, there is no indication of how bad the nonregularity actually is. For example, it is impossible to say whether or not the lowest order Cartesian derivative to become singular is its second or tenth. To discriminate between these cases, Fig. 2 shows the power $|a_n|$ as a function of n, for the truncations $N = 32$ and $N = 10$. Somewhat reassuringly, the spectra are shaped as to minimise the monomial coefficients at small n, and hence to maximise the regularity of the function. Indeed, as the truncation is increased from $N = 10$ to $N = 32$, the magnitude of the coefficients for $n \leqslant 5$ drops from $O(1)$ to $O(10^{-10})$ indicating that we might expect only a "weak" singularity at the origin in the first five Cartesian derivatives in this latter case. It is clear that attaining a converged solution is essential for any kind of regular behaviour.

A more direct method is to compute numerically the error in the various contributions to the Cartesian derivatives. In plane polar coordinates (r, ϕ),

$$\frac{\partial}{\partial x} = \cos\phi \frac{\partial}{\partial r} - \frac{\sin\phi}{r}\frac{\partial}{\partial \phi}, \quad \frac{\partial}{\partial y} = \sin\phi \frac{\partial}{\partial r} + \frac{\cos\phi}{r}\frac{\partial}{\partial \phi},$$

and it is clear that for any function f, the computation of its kth order Cartesian derivatives requires the accurate calculation of terms of the form $r^{-j}\left(\frac{\partial}{\partial r}\right)^{k-j} f, j = 0, \ldots, k$. For our Bessel function benchmark, we quantify the error in two representative terms $Q_k(r) = |y_N(r) - y(r)|/r^k$ and $D_k(r) = |\frac{\mathrm{d}}{\mathrm{d}r}(y_N(r) - y(r))\ |\ /r^{k-1}$, where $y(r)$ is the analytic solution.

The integer $k < l$ is in general chosen to be of physical significance. For example, if y represents temperature, then we might expect that not only y but also its first order Cartesian derivatives should be well behaved – thus both $Q_1(r)$ and $D_1(r)$ should be small everywhere, including arbitrarily close to $r = 0$. If y_N has any nonzero constant in its Chebyshev expansion, then $Q_1(r)$ will be singular at $r = 0$. However, by using an odd Chebyshev expansion (if l is odd) then this is avoided and the solution is sufficiently regular to be physically meaningful. Perhaps a more useful example is that of solving the incompressible Navier–Stokes equations in a spherical geometry, in which one may express the unknown velocity field u in terms of poloidal and toroidal components [25]:

$$\mathbf{u} = \nabla \times \nabla \times [s(r,\theta,\phi)\hat{\mathbf{r}}] + \nabla \times [t(r,\theta,\phi)\mathbf{r}],$$

where $\hat{\mathbf{r}}$ is the unit vector in the radial direction. If the scalar field s is expanded in terms of spherical harmonics $\sum Y_l^m(\theta,\phi)s_l^m(r)$ (and similarly for t), then for example, the radial and latitudinal components of velocity are given by

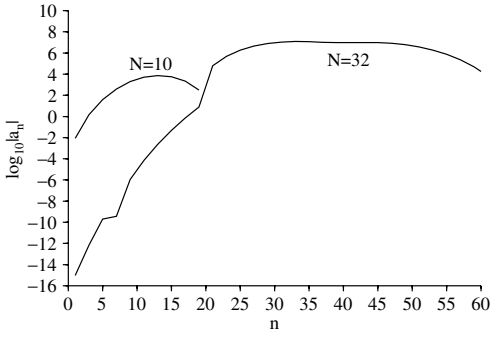

Fig. 2. Spectrum of the monomial coefficients, $|a_n|$, of the numerical solution to the leading eigenvector of Bessel's equation of order 21 for truncations $N = 32$ and $N = 10$. Low values of $|a_n|$ for small n indicate only "weak" singularities of the first few Cartesian derivatives.

Table 2
Behaviour of the numerical solution to the leading eigenfunction y_N of Bessel's equation of order $l = 7$ and $l = 101$ close to the origin as a function of N, the total number of odd Chebyshev polynomials used

N	$l = 21$		$l = 101$		
	$\log_{10} Q_3(10^{-2})$	$\log_{10} Q_3(10^{-3})$	$\log_{10} Q_3(10^{-2})$	$\log_{10} Q_3(10^{-3})$	$\log_{10} D_3(10^{-3})$
10	1.61	3.62	2.45	4.46	4.46
15	−0.97	1.04	0.18	2.19	2.19
20	−3.83	−1.80	−0.04	1.99	1.99
25	−7.69	−5.64	−2.20	−0.15	−0.15
30	−11.1	−8.73	−4.47	−2.38	−2.39
35	−9.69	−7.66	−6.45	−4.34	−4.34
40	−10.6	−8.41	−9.02	−6.91	−6.91
45	−9.97	−8.00	−10.4	−8.62	−8.46
50	−9.61	−7.46	−10.1	−8.15	−8.06

$Q_k(r) = |y_N(r) - y(r)|/r^k$ where $y(r)$ is the analytic solution and $D_k(r) = |\frac{\mathrm{d}}{\mathrm{d}r}(y_N(r) - y(r))|/r^{k-1}$. In all cases $y(r)$ has a zero of order l at $r = 0$; $Q_k(r)$ and $D_k(r)$ therefore provide measures of how singular the numerical solution is as $r \to 0$. To achieve a satisfactorily nonsingular solution at $r = 10^{-3}$ with $k = 3$, we require $N = 20$ for $l = 21$ and $N = 35$ for $l = 101$.

$$u_r = \sum_{l,m} \frac{l(l+1)s_l^m Y_l^m}{r^2}, \quad u_\theta = \sum_{l,m} \frac{1}{r}\frac{\mathrm{d}s_l^m}{\mathrm{d}r}\frac{\partial Y_l^m}{\partial \theta} + \frac{t_l^m}{r\sin\theta}\frac{\partial Y_l^m}{\partial \phi}.$$

A reasonable demand of the numerical solution might be that $\mathbf{u}$ is everywhere well defined; it follows that the analogous condition in our test case is that both $Q_2(r)$ and $D_2(r)$ should be small as $r \to 0$; indeed, if the vorticity is required to be physically meaningful then $Q_3(r)$ and $D_3(r)$ should also be small. In the preceding discussion, we have already made the point that that evaluation of $Q_3(r)$ is in general impossible at the origin, for the Chebyshev series will have nonzero (though possibly tiny) coefficients r^i with $i < 3$, rendering $Q_3(r)$ formally singular. At best, we might hope to replace $Q_3(0)$ by $Q_3(r)$ for r sufficiently close to zero to remain a consistent representation.

Table 2 shows the error in evaluating y_N/r^3 (denoted by $Q_3(r)$) and a related derivative (denoted by $D_3(r)$) close to the origin for the case of our Chebyshev approximation to the leading eigenvector of Bessel's equation of orders $l = 21$ and $l = 101$. We tabulate $\log_{10} Q_3(r)$ as a function of the truncation N used for $r = 10^{-2}$ and $r = 10^{-3}$, these radial values being close enough to $r = 0$ to represent the behaviour at the origin in a reasonable way. In order to get a solution that is "regular enough", that is, for which Q_3 is (say) $O(10^{-3})$, one requires $N = 30$–40 for $l = 101$ for the two choices of r. Thus if one is prepared to approximate the solution at $r = 0$ by the solution at $r = 10^{-3}$, then sufficiently regular solutions can be obtained by moderate truncation levels. Notice also that for the values indicated in the table, $D_3(r)$ closely approximates $Q_3(r)$; thus the error from the related contributions to the Cartesian derivatives are comparable. Note also that the accuracy of y_N/r^k does not increase monotonically with N: for large truncation levels, the accuracy becomes degraded due to the effects of accumulated roundoff error. Additionally, as k gets larger the value of r for which $Q_k(r)$ is smaller than any prescribed threshold increases. Thus as the demands of regularity on the solution increase, the region $r \leq r^c$ inside which we are unable to compute a solution accurately will grow large enough to prevent the approximation of $y_N(0)$ by any $y_N(r)$. Thus although a Chebyshev method may be able to meet mild regularity conditions, more stringent demands, for instance, $Q_k(10^{-3}) \ll 1$ with $k > 3$, will be out of reach. It is also noteworthy that even a slightly under resolved solution may exhibit strongly nonregular behaviour. For example, Fig. 1 shows that when $l = 101$ an error in the eigenvalue of magnitude $O(10^{-3})$ is obtained when $N = 20$ (indicating a marginally converged solution), but Table 2 indicates than an error of $O(10^2)$ can occur when evaluating derivatives at $r = 10^{-3}$.

3. One-sided Jacobi polynomials

A parameterised family of polynomials that span functions of the form $r^l g(r)$ where $g(r)$ is smooth and even may be written

$$G_n(r;\alpha,\beta,l) = r^l P_n^{(\alpha,\beta)}(2r^2-1),$$

where $P_n^{(\alpha,\beta)}(x)$ is a Jacobi polynomial [26]. Jacobi polynomials are solutions of a singular Sturm–Liouville problem, and as such not only are they orthonormal (when appropriately normalised):

$$\int_{-1}^{1}(1-x)^{\alpha}(1+x)^{\beta}P_n^{(\alpha,\beta)}(x)P_m^{(\alpha,\beta)}(x)\,\mathrm{d}x = \delta_{nm}, \tag{3}$$

but are also exponentially convergent to any sufficiently smooth function [5,13]. Under the change of coordinate $x = 2r^2-1$ that maps the domain $[-1,1]$ to $[0,1]$ and renders $P_n^{(\alpha,\beta)}(2r^2-1)$ an even polynomial, under an appropriate normalisation the G_n are orthonormal with weight function $w(r)$:

$$\int_0^1 G_n G_m w(r)\,\mathrm{d}r = \delta_{nm}, \quad w(r;\alpha,\beta,l) = (1-r^2)^{\alpha}r^{2(\beta-l)+1}. \tag{4}$$

Various choices of the parameters have been made in the existing literature. We have already described in Section 2 the so-called Robert functions, $r^l T_n(r)$, which correspond to the parameters $\alpha=\beta=-1/2$ and are orthogonal with weight function $w(r)=(1-r^2)^{-1/2}r^{-2l}$. What makes this basis so ill conditioned is the fact that the polynomials are zero over almost all the interval $[0,1]$; since $r^l \approx \exp(-l(1-r))$ for $r\approx 1$, each polynomial decays exponentially away from $r=1$. Close to $r=1$ where the functions are large, the low degree Chebyshev polynomials are numerically almost linearly dependent [1]. Another way of expressing the failure of the formal orthogonality to create a well conditioned scheme is to note that when $l \gg 1$, the weight function $w(r)$ weights most heavily the part of the interval closest to $r=0$ where the functions themselves are tiny. Thus the orthogonality provides little constraint for the part of the functions that dominates on $[0,1]$. In a typical calculation, such numerical ill-conditioning is manifested in extremely large basis coefficients when representing a function that is O(1) for most of the interval. The coefficients necessarily must be large since the basis functions themselves are small away from $r=1$, but must effect a near perfect cancellation for the solution to remain O(1) at $r=1$. This discussion motivates us to suggest a weight function that is at most only mildly singular at $r=0$; thus β must increase with l.

The independent studies of both Verkley [18] and Matsushima and Marcus [19] used $\alpha=0$, $\beta=l$ (although note the different definition of the parameters in [19]). This leads to the weight function $w(r)=r$, a natural choice in plane polar coordinates. For brevity, we will refer to these as Verkley polynomials in the remainder of the paper. Note that the equivalent weight function in a spherical geometry, that of $w(r)=r^2$, would be obtained by $\alpha=0$, $\beta=l+1/2$.

Fig. 3 shows plots of Verkley polynomials as solid lines for various choices of n and l. It is clear that there is a large jump in the amplitude of the oscillations close to $r=1$ for large n. It was suggested by Worland [14] that such polynomials might be mildly ill-conditioned, particularly when representing boundary layers, and he suggested the different choice of $\alpha=-1/2$, $\beta=l-1/2$. The associated weight function is $w(r)=(1-r^2)^{-1/2}$, the same as that for Chebyshev polynomials (although only over $[0,1]$). These parameters generate a family of polynomials (shown by the dashed curves in Fig. 3) for which oscillations are asymptotically (as $n\to\infty$) contained within a uniform envelope. The motivation for this choice of (α,β) is not immediately obvious and a derivation is given in Appendix A. Some examples of Worland and Verkley polynomials for low l and n are given in Table 3, normalised to unity at $r=1$. Adopting this normalisation, the definition of the Worland polynomials is then

$$W_n^l(r) = \frac{2^{2n}(n!)^2}{(2n)!}r^l P_n^{(-1/2,l-1/2)}(2r^2-1), \tag{5}$$

with orthogonality relation

$$\int_0^1 (1-r^2)^{-1/2}W_n^l W_m^l\,\mathrm{d}r = h_n\delta_{nm}, \quad h_n = \frac{\pi(n+l)}{2^{2l+1}(2n+l)}\frac{(2n+2l)![n!]^2}{[(n+l)!]^2(2n)!}, \tag{6}$$

where $P_n^{(\alpha,\beta)}$ follows the standard normalisation as in [26]. The Verkley polynomials are

$$V_n^l(r) = r^l P_n^{(0,l)}(2r^2-1), \tag{7}$$

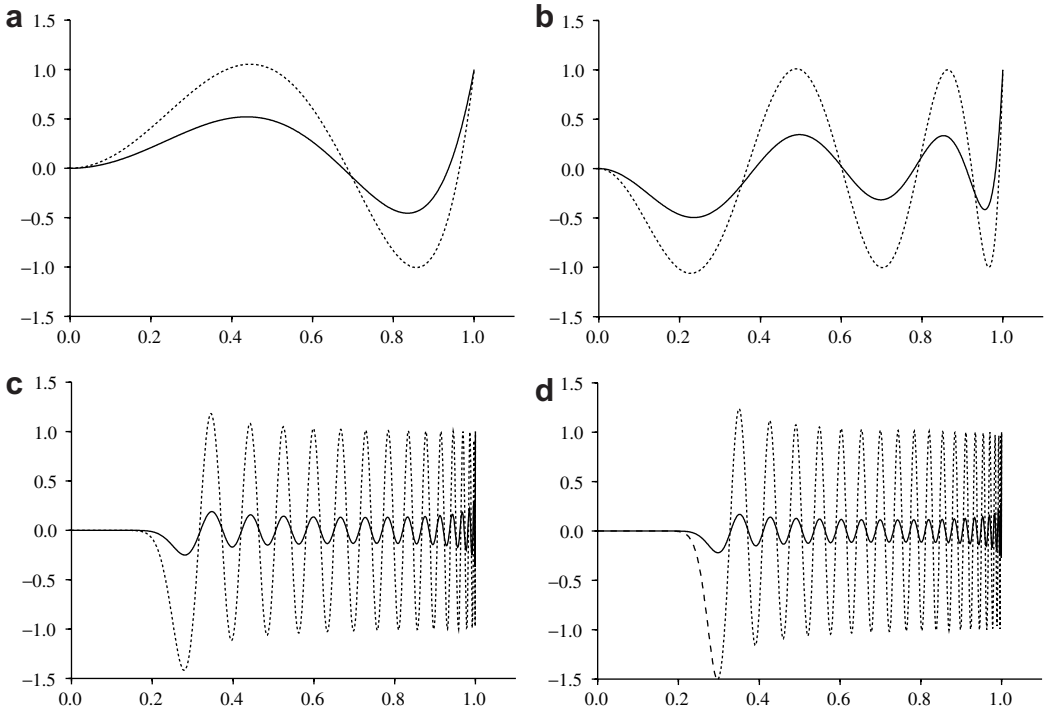

Fig. 3. Comparison plots of Verkley polynomials (shown by the solid lines) and Worland polynomials (shown by the dashed lines); each curve is normalised to attain l at $r = 1$. (a) $l = 2$; $n = 3$; (b) $l = 2$, $n = 6$; (c) $l = 20$, $n = 30$; and (d) $l = 30$, $n = 40$. The Worland polynomials show asymptotically uniformly modulated oscillations as $n \to \infty$; the Verkley polynomials show a jump in modulation near $r = 1$.

Table 3
Verkley and Worland polynomials for $l = 2,3$ and $n = 0,1,2$, normalised equal to unity at $r = 1$

n	Verkley polynomials		Worland polynomials	
	$l = 2$	$l = 3$	$l = 2$	$l = 3$
0	r^2	r^3	r^2	r^3
1	$r^2(4r^2 - 3)$	$r^3(5r^2 - 4)$	$r^2(6r^2 - 5)$	$r^3(8r^2 - 7)$
2	$r^2(15r^4 - 20r^2 + 6)$	$r^3(21r^4 - 30r^2 + 10)$	$\frac{1}{3}r^2(80r^4 - 112r^2 + 35)$	$r^3(40r^4 - 60r^2 + 21)$

with orthogonality relation

$$\int_0^1 V_n^l V_m^l r \,\mathrm{d}r = h_n \delta_{nm}, \quad h_n = \frac{1}{4n + 2l + 2}. \tag{8}$$

The efficient computation of general one-sided Jacobi polynomials is discussed in Appendix B.

Lastly, it is perhaps not obvious that the one-sided Jacobi basis inherits the spectrally convergent property of the Jacobi polynomials on which they are based. We provide a proof in Appendix C.

3.1. Convergence

We now investigate the convergence of the Verkley and Worland polynomials in solving Bessel's equation (Eq. (1)) of order $l = 7$ and $l = 101$. After expanding the unknown solution in N polynomials,

$$y_N = \sum_{n=0}^{N-1} a_n G_n(r; \alpha, \beta, l),$$

we may determine a system of equations for the coefficients in two ways: tau or collocation. To implement the tau method, we impose orthogonality of the residual R_N in the analogue of Eq. (2) to each appropriate G_n (with suitable weight function). Numerically, this may be achieved by a quadrature method (Gauss–Legendre in the Verkley case and Gauss–Chebyshev in the Worland case, reflecting the form of the weight function) using some minimum number M of abscissae. In either case, $G_n\mathcal{D}G_m$ is a polynomial (if $l \geqslant 1$) of degree at most $2(l+2(N-1))-2 = 2l+4N-6$ (if $n,m \leqslant N-1$) and including the weight function may be integrated exactly using M abscissae where $M \geqslant l+2N-2$. In order to project a nonlinear term back onto the basis, the standard use of a transform to convert between physical and spectral space would require the same radial grid for all values of l. We would then be restricted by the highest degree integrand, requiring $M = \mathrm{O}(3L_{\max}/2+3N)$ abscissae, where $L_{\max}$ is the largest value of l [18]. In a typical application with $L_{\max} = 200$ and $N = 40$, this slow transform involves multiplication by a matrix of approximate side 420, potentially crippling for a timestepping problem.

An alternative is to adopt a collocation scheme, whereby we impose $R_N(r_i) = 0$ on a discrete grid $\{r_i\}$. The grid points, r_i, are usually chosen to be identical to the associated quadrature points of the related tau scheme (thus rendering both methods equivalent) [1]. In this case however there are more quadrature points than unknown polynomial coefficients (since $M \geqslant l+2N-2 > N$ in general) and no collocation method will be equivalent to a tau scheme. Motivated by the quadrature method, we may choose r_i to be the first $N-1$ zeros of G_N. As mentioned above, to project an unknown (for example nonlinear) function back onto the basis will in general require the same grid points for all l; thus these l-dependent collocation points would not be applicable in such a case. However, in view of the above discussion, the choice of collocation points is somewhat arbitrary and we may consider instead the first $N-1$ zeros of $G_{N-1}(r;\alpha,\beta,L_{\max})$ rendering the collocation points independent of l, a choice that Worland [14] actually showed to be preferable. We additionally note that in this case the slow transform between physical and spectral space would be considerably faster than the quadrature methods described above, although still of complexity $\mathrm{O}(N^2)$.

Fig. 4 shows analogous plots to Fig. 1, depicting the convergence to the least negative eigenvalue of Bessel's equation for both tau (solid line with circles) and collocation (dashed line with squares) methods. The tau methods converge faster than the collocation schemes (as perhaps may be expected), and there is some evidence that the Verkley polynomials converge faster than the Worland polynomials although there is little to choose between them. The apparent floor on the error of $\mathrm{O}(10^{-10})$ in the $l = 101$ case, for example, is due to numerical imprecision (in double precision); the solid lines with no symbols show an identical tau-method calculation but performed in Maple at much higher precision and are not bounded from below. The rms error in the eigenvectors shows almost no deviation between the various schemes (not shown).

It is noteworthy when comparing Figs. 1 and 4, for any particular error tolerance, that both the one-sided Jacobi bases shown require a lower truncation (with the tau scheme) than the Chebyshev scheme. In particular, for $l = 101$, both the Verkley and Worland schemes achieve an absolute error of less than 10^{-6} with just

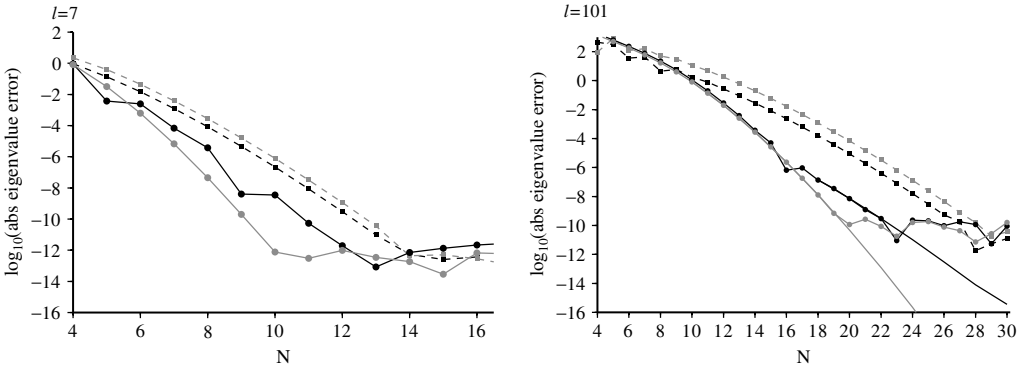

Fig. 4. Convergence to the least negative eigenvalue of Bessel's equation of order $l = 7$ and $l = 101$ using N polynomials of either Worland (black) or Verkley (grey) type and either a collocation with l-dependent grid points (dashed line, squares) or tau (solid line, circles) method. The lines without symbols show tau-method calculations performed in Maple to high precision and show the numerically induced error floor when computing in double precision.

17 polynomials. The same problem solved using Chebyshev polynomials required $N = 28$, roughly 1.5 times as many. Thus the one-sided Jacobi polynomials converge faster and offer a much more concise representation of the solution than Chebyshev polynomials. When using collocation, the Jacobi method achieves a similar accuracy to the Chebyshev scheme using roughly the same number of polynomials [9].

3.2. Minimax representation

One common method of function approximation is that of least squares; it is well known that the truncated Legendre and Chebyshev expansions of a function $f(r)$ up to degree N correspond to the polynomial of degree N, K_N say, that minimises the integrated weighted squared error over $[-1,1]$ with weight function $w(x) = 1$ and $w(x) = (1 - x^2)^{-1/2}$ respectively [27]. For instance, we may determine

$$\min \int_{-1}^{1} (f(x) - K_N(x))^2 \, \mathrm{d}x, \quad K_N(x) = \sum_{0}^{N} a_n P_n(x),$$

where the minimisation is taken over all polynomials of degree N that, without loss of generality, may be written in terms of the spanning set of Legendre polynomials, P_n. Taking partial derivatives with respect to each of the coefficients a_i in turn leads to the result that

$$a_n = \int_{-1}^{1} f(x) P_n(x) \, \mathrm{d}x.$$

Analogously, it follows that the truncated expansion of a function $f(r) = r^l g(r)$ (with g even and smooth) in terms of any family $G_n(r; \alpha, \beta, l)$ up to degree N minimises the integrated squared error between f and any polynomial $B_N(r) = r^l h(r)$ where h is an even polynomial

$$\int_{0}^{1} w(r)(f(r) - B_N(r))^2 \, \mathrm{d}r,$$

with respect to the appropriate weight function (from Eq. (4)). Thus in this precise sense, any choice of (α, β) produces a family G_n that results in an optimal expansion of f.

A second method of approximation is to find the so-called minimax polynomial, $M_N(x)$, the polynomial of degree N that minimises the maximum pointwise error over the domain [27]. However, this is in general a substantially more difficult formulation than weighted least squares and there is no known algorithm for finding M_N. However, if M_N can be found it can be shown that the error $e(r) = f(r) - M_N(r)$ must oscillate uniformly over the required interval [27]. Thus if we attempt an expansion of an arbitrary function f in terms of Chebyshev polynomials,

$$f(r) = \sum_{n=0}^{N} b_n T_n(r), \tag{9}$$

the error would be largely given by the first excluded term, $T_{N+1}(r)$, which does oscillate uniformly. A Chebyshev expansion therefore is known to be close to the minimax polynomial approximation of a function.

Given that Worland polynomials oscillate within an asymptotically uniform envelope, it is tempting to draw some analogy to the minimax property of the Chebyshev polynomials. Although we provide no proof, by applying Eq. (9) it seems likely that the Worland polynomial projection is a good approximation to the minimax polynomial. We demonstrate this by considering the error in the projection of the boundary layer function $f(r) = r^{20}\exp(20(1 - r^2))$ on both the Worland (black) and Verkley (grey) polynomials, as shown in Fig. 5(a). It is clear that although the error in the Verkley projection is smaller in the interior, the Worland polynomials do a far better job close to the boundary. Indeed, over the domain the Worland projection achieves a smaller maximum error than the Verkley case.

The large jump in the oscillation amplitude of the Verkley polynomials close to $r = 1$, a property shared by the Legendre polynomials, merits further discussion (see also Appendix A). In the Legendre case, the jump is associated with degraded (although still exponential) convergence; compared to Chebyshev polynomials, the coefficients are $O(N^{1/2})$ larger in magnitude [1]. However, as Fig. 5(b) shows, the Verkley polynomials con-

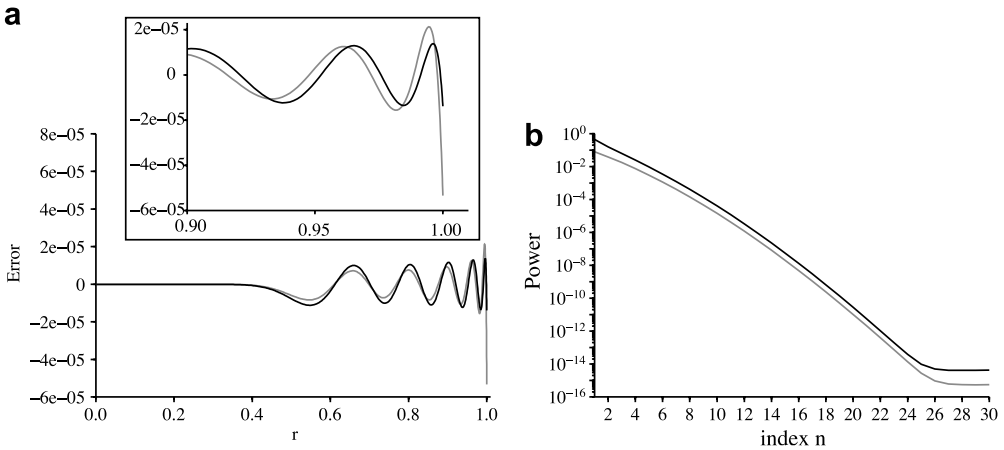

Fig. 5. Projection of the boundary layer function $r^{20}\exp(20(r^2-1))$ onto orthonormal Verkley (grey) and Worland (black) polynomials $G_n(r;\alpha,\beta,20)$. (a) The pointwise error for expansions truncated with 10 polynomials, the inset shows a magnification of the interval [0.9, 1]; (b) the spectral power as a function of degree. Although the Verkley coefficients are smaller in magnitude than those for Worland, the two rates of convergence are essentially the same. The maximum error is larger for the Verkley projection; the Worland projection looks close to the minimax case as the error is almost uniformly oscillating.

verge at the same rate as the Worland polynomials (indeed, their coefficients are actually smaller by an approximately constant factor) and so the behaviour close to $r=1$ has no analogous degrading effect on the global convergence.

4. Discussion

In this paper, we considered two different spectral methods for representing functions of radius in a planar or spherical polar geometry that includes the origin: Chebyshev series of definite parity and one-sided Jacobi polynomials. Crucially, the regularity conditions yielding solutions that are everywhere infinitely differentiable are not in general satisfied by Chebyshev methods (in finite precision), but are satisfied explicitly (by construction) when using one-sided Jacobi polynomials. However, the availability of a fast transform may swing the balance in favour of Chebyshev methods and there is clearly a tradeoff between regularity and computational speed between the two approaches.

We discussed in some detail the problems associated with evaluating a numerically derived Chebyshev expansion close to the origin. For instance, it is not possible to evaluate the limit y_N/r as $r\to 0$ (and therefore compute its first order Cartesian derivatives at $r=0$) even if the numerical solution y_N approximates a function that is $O(r^2)$, due to the fact y_N will have a constant term that is not precisely zero (although it may be at the limit of numerical roundoff). Thus we are resigned to approximating the solution at the origin by the solution at some $r=r_0$ close by. The question is whether such a point, being sufficiently close to the origin to remain consistent, yet sufficiently far away to avoid the singularity, can be found. In general, we seek such a point for which $y_N(r_0)/r_0^k$ for $k\geqslant 1$ is accurate, this being necessary for the first k Cartesian derivatives to be well defined close to the origin; the positive integer k is defined by the physics of the problem. In fact, we showed in Section 2 that this is in general possible; for example, it is reasonably straightforward to generate solutions $y_N(r_0)$ such that $y_N(r_0)/r_0^3$ is accurate for values $r_0\geqslant 10^{-3}$ provided it is adequately converged. Thus if one is willing to approximate the solution at the origin by the solution at $r=10^{-3}$ then sufficient regularity is in general obtainable. However, it is worthwhile adding a comment on real-life problems where codes may be run at relatively extreme parameter values. If for whatever reason the solution is not fully converged then it is unlikely that any kind of regularity is assured. Table 2 shows that even if the solution is marginally under resolved, that it may be impossible to evaluate the solution and required derivatives accurately close to the origin.

The parameterised family of one-sided Jacobi polynomials give an extremely concise spectral representation and any solution using such an expansion, even if not fully converged, is automatically regular at the origin

and differentiable to all orders. Of the infinite set of specifying parameters (α, β), we compared two such choices generating the Verkley and Worland polynomials, and found that they behave similarly and converge much faster than an expansion in standard Chebyshev polynomials. We argued that a truncated expansion in Worland polynomials is close to the minimax polynomial (being analogous to the Chebyshev polynomials), that of minimising the maximum pointwise error over the interval. The major drawback of such methods however is the lack of a competitive fast transform, degrading their efficiency when performing nonlinear transforms.

We now return to the main point of the paper, to offer a recommendation to the reader as to which method to choose for any particular application. It is clear that some tradeoff must be made between the speed of computing Chebyshev transforms and the accuracy of the one-sided Jacobi polynomials. If the problem involves repeated nonlinear transforms between spectral and physical space, for example when timestepping the Navier–Stokes equations, and provided that the solution is sufficiently converged (which may well limit the parameter values available), then a Chebyshev method is probably preferable. However, the Cheyshev method has a serious potential pitfall for the unwary. A frequently adopted practice is to choose the truncation level adopted on the basis of the convergence of global quantities such as an eigenvalue in a linear problem or the energy in a nonlinear problem. This is an inadequate method of proceeding if quantities such as vorticity are needed, as these involve derivatives and division by powers of r that may produce large errors near the origin. While these difficulties can be overcome by using more polynomials, a conservative approach to deciding on truncation levels is strongly recommended. For other problems where conciseness of representation and regularity are more pressing, then the one-sided Jacobi polynomials are the obvious choice. Of the two families studied in this paper, both the Worland and Verkley polynomials are excellent such choices and there is little in practice to choose between them.

Appendix A. Worland polynomials: WKBJ motivation

In this section we show that the one-sided Jacobi polynomials $G_n(r;\alpha,\beta,l) = r^l P_n^{(\alpha,\beta)}(2r^2-1)$ oscillate within an asymptotically uniform envelope (as $n \to \infty$) if we choose the parameters $\alpha = -1/2$, $\beta = l - 1/2$.

Jacobi polynomials themselves $y(x) = P_n^{(\alpha,\beta)}(x)$ satisfy the following differential equation

$$(1-x^2)y''(x) + [\beta - \alpha - (\alpha+\beta+2)x]y'(x) + n(n+\alpha+\beta+1)y(x) = 0. \tag{A.1}$$

We shall look for a solution of the form [28]

$$y(x) = A(x)\exp\left(\frac{\mathrm{i}\int k(x)\,\mathrm{d}x}{\epsilon}\right),$$

for large $n = \hat{n}/\epsilon \gg 1$ where $\epsilon \ll 1$, and seek constraints on (α, β) such that $r^l y(2r^2-1)$ oscillates with uniform amplitude, that is, $r^l A(2r^2-1) = 1$. It follows that

$$\begin{aligned} y'(x) &= \left(A' + \frac{\mathrm{i}kA}{\epsilon}\right)\exp\left(\frac{\mathrm{i}\int k(x)\,\mathrm{d}x}{\epsilon}\right), \\ y''(x) &= \left(A'' + \frac{2\mathrm{i}kA'}{\epsilon} + \frac{\mathrm{i}k'A}{\epsilon} - \frac{k^2A}{\epsilon^2}\right)\exp\left(\frac{\mathrm{i}\int k(x)\,\mathrm{d}x}{\epsilon}\right). \end{aligned} \tag{A.2}$$

On substituting these into Eq. (A.1) we obtain at $\mathrm{O}(\epsilon^{-2})$

$$k(x) = \hat{n}(1-x^2)^{-1/2},$$

(up to sign), yielding the expected ever more frequent oscillations as $x \to \pm 1$. At $\mathrm{O}(\epsilon^{-1})$ we find

$$(1-x^2)(2\mathrm{i}kA' + \mathrm{i}k'A) + [\beta - \alpha - (\alpha+\beta+2)x]\mathrm{i}kA + \hat{n}(\alpha+\beta+1)A = 0.$$

After some algebra, the solution may be written

$$A(x) = C(1-x)^{-(2\alpha+1)/4}(1+x)^{-(2\beta+1)/4}\exp\left(\mathrm{i}\frac{(\alpha+\beta+1)}{2}\cos^{-1}x\right),$$

and so

$$|A(2r^2-1)| = D(1-r^2)^{-(2\alpha+1)/4} r^{-(2\beta+1)/2}, \tag{A.3}$$

with C and D unknown constants. It follows that in order for $|r^l A(2r^2-1)| = 1$ that we must choose $\alpha = -1/2$ and $\beta = l - 1/2$. Thus as $n \to \infty$ the Worland polynomials oscillate with an increasingly uniform amplitude; in particular, the bulk of their oscillations close to $r = 1$ become bounded by a uniform envelope. Note that the Verkley choice, $\alpha = 0$, leads to a mild singularity in the envelope as $r \to 1$, giving rise to the behaviour there seen in Figs. 3 and 5.

It is noteworthy that although the Worland polynomials oscillate asymptotically uniformly as $n \to \infty$, they can never be obtained by a scaled stretched Chebyshev polynomial (as would be required to make use of the fast cosine transform). That is, it is not possible to represent

$$G_n(r, -1/2, l-1/2, l) = S(r) T_{2n}(f(r)),$$

where the stretching function $f(r)$ is a bijective map: $[0,1] \to [-1,1]$ and $S(r)$ provides the (non uniform) modulation of the uniform Chebyshev oscillations. This is because f would have to map the zeros of the Chebyshev polynomials (of multiplicity one) to the zeros of the G_n (of which $r = 0$ is a zero of multiplicity l). Thus f cannot be bijective.

Appendix B. Computation of one-sided Jacobi polynomials and their derivatives

In this section we introduce a general method of stably computing the one-sided Jacobi polynomials $G_n(r;\alpha,\beta,l) = r^l P_n^{(\alpha,\beta)}(2r^2-1)$ with their derivatives, supplementing those given in [18,19]. The Jacobi polynomials themselves are simply computed by the three term recurrence:

$$a_1 P_{n+1}^{(\alpha,\beta)}(x) = (a_2 + x a_3) P_n^{(\alpha,\beta)}(x) - a_4 P_{n-1}^{(\alpha,\beta)}(x), \tag{B.1}$$

where a_1, a_2, a_3, a_4 are functions of α, β and n [26], along the starting values

$$P_0^{(\alpha,\beta)}(x) = 1, \quad P_1^{(\alpha,\beta)}(x) = \frac{1}{2}[(\alpha-\beta) + (\alpha+\beta+2)x]. \tag{B.2}$$

However, the standard recurrence relations for their derivatives are singular at $r = 1$, and although limiting forms can be used on the boundary, we show that this difficulty can be sidestepped by exploiting the relationship with the Hypergeometric function.

Before doing so however, we note in passing that it is possible, at least in theory, to compute the polynomials by using an iterative Gram–Schmidt procedure. By writing $G_n = \sum_{i=0}^{n} \alpha_i r^{l+2i}$ and assuming that $\{G_{n+1}\} \cup \{G_0, G_1, \ldots, G_n\}$ form an orthonormal basis, the coefficients a_i can be found. However, this process is numerically poorly conditioned, as the coefficients quickly become unmanageably large. For example, the largest monomial coefficient in $G_{16}(r; -1/2, 19/2, 10)$ is $O(10^{13})$ even though the polynomial itself is O(1). Thus almost exact cancellation must take place between the monomial terms, a recipe for numerical disaster.

Jacobi polynomials are related to the Hypergeometric function F in the following way:

$$\frac{n!}{(\alpha+1)_n} P_n^{(\alpha,\beta)}(1-2z) = F(-n, \alpha+\beta+1+n, \alpha+1, z), \tag{B.3}$$

where $(x)_n = (x)(x+1)(x+2)\cdots(x+n-1)$ is the Pochhammer symbol [26].

Direct computation of F itself is rather inaccurate (e.g. by using the routines in [29]), but by making use of

$$\frac{\mathrm{d}}{\mathrm{d}z} F(a,b,c,z) = \frac{ab}{c} F(a+1, b+1, c+1, z), \tag{B.4}$$

$$\frac{\mathrm{d}^2}{\mathrm{d}z^2} F(a,b,c,z) = \frac{a(a+1)b(b+1)}{c(c+1)} F(a+2, b+2, c+2, z), \tag{B.5}$$

it is straightforward to show that

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}z}P_n^{(\alpha,\beta)}(1-2z) &= -(n+\alpha+\beta+1)P_{n-1}^{(\alpha+1,\beta+1)}(1-2z),\\ \frac{\mathrm{d}^2}{\mathrm{d}z^2}P_n^{(\alpha,\beta)}(1-2z) &= (n+\alpha+\beta+1)(n+\alpha+\beta+2)P_{n-2}^{(\alpha+2,\beta+2)}(1-2z).\end{aligned} \tag{B.6}$$

It follows by substituting $1-2z = 2r^2-1$ that

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}r}P_n^{(\alpha,\beta)}(2r^2-1) &= 2r(n+\alpha+\beta+1)P_{n-1}^{(\alpha+1,\beta+1)}(2r^2-1),\\ \frac{\mathrm{d}^2}{\mathrm{d}r^2}P_n^{(\alpha,\beta)}(2r^2-1) &= 4r^2(n+\alpha+\beta+1)(n+\alpha+\beta+2)P_{n-2}^{(\alpha+2,\beta+2)}(2r^2-1)\\ &\quad + 2(n+\alpha+\beta+1)P_{n-1}^{(\alpha+1,\beta+1)}(2r^2-1),\end{aligned} \tag{B.7}$$

and the derivatives of $r^l P_n^{(\alpha,\beta)}(2r^2-1)$ follow easily.

Appendix C. Spectral convergence of the one-sided Jacobi polynomials

Although the one-sided Jacobi polynomials

$$G_n(r;\alpha,\beta,l) = r^l P_n^{(\alpha,\beta)}(2r^2-1),$$

where $0 \leqslant r \leqslant 1$ are based on the standard Jacobi polynomials, $P_n^{(\alpha,\beta)}(x)$, it is not immediately apparent that the property of spectral convergence is inherited from them. By spectral convergence, we mean that in the expansion of any sufficiently smooth function $f(x)$ defined on $[-1,1]$,

$$f(x) = \sum_{n=0}^{n=\infty} a_n P_n^{(\alpha,\beta)}(x),$$

that the coefficients a_n decay asymptotically faster than any algebraic power of n [1,13]. That is, $a_n = \mathrm{O}(n^{-k})$ for any integer k as $n \to \infty$. The choices of α and β are restricted to the interval $(-1,\infty)$, a property inherent in the definition of the Jacobi polynomials themselves [26].

The key aspect here for the one-sided Jacobi basis is that we are not trying to approximate an arbitrary function, but one of the form $r^l g(r)$ with g both even and smooth. If we make the change of variable $x = 2r^2-1$, $-1 \leqslant x \leqslant 1$, then from the above relation we may expand the function $g(r(x))$ in terms of Jacobi polynomials, or equivalently

$$g(r) = \sum_{n=0}^{n=\infty} a_n P_n^{(\alpha,\beta)}(2r^2-1),$$

where spectral convergence is guaranteed. By multiplying both sides by r^l, it is immediate that

$$r^l g(r) = \sum_{n=0}^{n=\infty} a_n G_n(r;\alpha,\beta,l),$$

which exhibits the required spectral convergence.

References

[1] J.P. Boyd, Chebyshev and Fourier Spectral Methods, Dover, 2001.
[2] G. Backus, R. Parker, C. Constable, Foundations of Geomagnetism, CUP, 1996.
[3] H. Eisen, W. Heinrichs, K. Witsch, Spectral collocation methods and polar coordinate singularities, J. Comput. Phys. 96 (1991) 241–257.
[4] R.R. Kerswell, A. Davey, On the linear instability of elliptic pipe flow, J. Fluid Mech. 316 (1996) 307–324.
[5] D. Gottlieb, S. Orszag, Numerical Analysis of Spectral Methods; Theory and Applications, twenty-sixth ed., CBMS-NSF Regional Conference Series in Applied Mathematics, SIAM, Philadephia, PA, 1977.
[6] D.J. Torres, E.A. Coutsias, Pseudospectral solution of the two-dimensional Navier–Stokes equations in a disk, SIAM J. Sci. Comput. 21 (1) (1999) 378–402.

[7] J. Shen, Efficient spectral-Galerkin methods III: polar and cylindrical geometries, SIAM J. Sci. Comput. 18 (6) (1997) 1583–1604.
[8] P.W. Livermore, A. Jackson, On magnetic energy instability in spherical stationary flows, Proc. Roy. Soc. A 460 (2004) 1453–1476.
[9] K. Mohseni, T. Colonius, Numerical treatment of polar coordinate singularities, J. Comput. Phys. 157 (2000) 787–795.
[10] W. Huang, D.M. Sloan, Pole condition for singular problems: the pseudospectral approximation, J. Comput. Phys. 107 (1993) 254–261.
[11] W. Heinrichs, Spectral collocation schemes on the unit disk, J. Comput. Phys. 199 (2004) 66–86.
[12] B. Fornberg, A Practical Guide to Pseudospectral Methods, CUP, 1998.
[13] C. Canuto, M.Y. Hussaini, A. Quarteroni, T.A. Zang, Spectral Methods for Fluid Dynamics, Springer-Verlag, 1988.
[14] S. Worland, Magnetoconvection in Rapidly Rotating Spheres, PhD Thesis, University of Exeter, 2004.
[15] S. Li, D. Montgomery, Decaying two-dimensional turbulence with rigid walls, Phys. Lett. A 218 (1996) 281–291.
[16] P.W. Livermore, A. Jackson, A comparison of numerical schemes to solve the magnetic induction eigenvalue problem in a spherical geometry, Geophys. Astrophys. Fluid Dyn. 99 (6) (2005) 467–480.
[17] K. Zhang, X. Liao, P. Earnshaw, The Poincaré equation: A new polynomial and its unusual properties, J. Math. Phys. 45 (12) (2004) 4777–4790.
[18] W.T.M. Verkley, A spectral model for two-dimensional incompressible fluid flow in a circular basin, J. Comput. Phys. 136 (1997) 100–114.
[19] T. Matsushima, P.S. Marcus, A spectral method for polar coordinates, J. Comput. Phys. 120 (1995) 365–374.
[20] J.R. Driscoll, D.M.J. Healy, D.N. Rockmore, Fast discrete polynomial transforms with applications to data analysis for distance transitive graphs, SIAM J. Comput. 26 (4) (1997) 1066–1099.
[21] R. Suda, M. Takami, A fast spherical harmonics transform algorithm, Math. Comput. 71 (2002) 703–715.
[22] D.M.J. Healy, D.N. Rockmore, P.J. Kostelec, S. Moore, FFTs for the 2-sphere – improvements and variations, J. Fourier Anal. Appl. 9 (4) (2003) 341–385.
[23] P.W. Livermore, A. Jackson, Transient magnetic energy growth in spherical stationary flows, Proc. Roy. Soc. A 462 (2006) 2457–2479.
[24] P.J. Schmid, D.S. Henningson, Stability and Transition in Shear Flows, Springer, 2001.
[25] R. Hollerbach, A spectral solution of the magneto-convection equations in spherical geometry, Int. J. Numer. Meth. Fluids 32 (2000) 773–797.
[26] M. Abramowitz, I.A. Stegun, Handbook of Mathematical Functions, Dover, 1972.
[27] L. Fox, I.B. Parker, Chebyshev Polynomials in Numerical Analysis, Oxford University Press, London, 1968.
[28] E.J. Hinch, Perturbation Methods, CUP, 1991.
[29] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipies in Fortran 77, second ed., CUP, 1992.