



Convergence and round-off errors in a two-dimensional eigenvalue problem using spectral methods and Arnoldi–Chebyshev algorithm

Lorenzo Valdettaro^{a,*}, Michel Rieutord^b, Thierry Braconnier^{c,1}, Valérie Frayssé^{d,1}

^a*Dipartimento di Matematica, Politecnico di Milano, Piazza L. da Vinci 32, I-20133 Milano, Italy*

^b*Observatoire Midi-Pyrénées, 14 av. E. Belin, F-31400 Toulouse, France*

^c*35 Rue de Tautia, Apt 47, F-33800 Bordeaux, France*

^d*Kvasar Technology LLC., 372 Marlborough Street, # 10, Boston, MA 02115, USA*

Received 10 May 2006

Abstract

An efficient way of solving 2D stability problems in fluid mechanics is to use, after discretization of the equations that cast the problem in the form of a generalized eigenvalue problem, the incomplete Arnoldi–Chebyshev method. This method preserves the banded structure sparsity of matrices of the algebraic eigenvalue problem and thus decreases memory use and CPU time consumption.

The errors that affect computed eigenvalues and eigenvectors are due to the truncation in the discretization and to finite precision in the computation of the discretized problem. In this paper we analyze those two errors and the interplay between them. We use as a test case the 2D eigenvalue problem yielded by the computation of inertial modes in a spherical shell. This problem contains many difficulties that make it a very good test case. It turns out that single modes (especially most-damped modes i.e. with high spatial frequency) can be very sensitive to roundoff errors, even when apparently good spectral convergence is achieved. The influence of roundoff errors is analyzed using the spectral portrait technique and by comparison of double precision and extended precision computations. Through the analysis we give practical recipes to control the truncation and roundoff errors on eigenvalues and eigenvectors.

© 2006 Elsevier B.V. All rights reserved.

1. Introduction

The first step in studying the stability of the solutions of a non-linear problem, is to solve the eigenvalue problem associated with infinitesimal perturbations which are superposed to the equilibrium state. Even if the equations of these perturbations are linear, solving the eigenvalue problem may be a formidable task. The difficulties arise in general when the variables do not separate. In such a case, the eigenvalue problem cannot be reduced to a set of smaller (1D) eigenvalue problems and one is left with a 2D or 3D problem.

In most cases, after discretization of the equations, the temporal stability problem reduces to a generalized eigenvalue problem. A method to solve such problems is to use the QZ algorithm. Such an algorithm gives the full spectrum of eigenvalues/eigenvectors, but the price to pay for obtaining this very rich information is very high in terms of memory

* Corresponding author.

E-mail address: lorenzo.valdettaro@polimi.it (L. Valdettaro).

¹ The work was done while the author was at CERFACS, 42, Avenue Coriolis, F-31057 Toulouse, France.

requirement and CPU time consumption. Moreover, the QZ algorithm does not preserve the sparsity of the matrices. On the other hand it is seldom needed to know the full spectrum: typically one is interested in the few eigenvalues corresponding to the least stable or most unstable modes. For the foregoing reasons it is important to be able to solve the generalized eigenvalue problem with an iterative method which preserves the sparsity of the matrices and converges quickly and accurately to a small subset of the whole spectrum. Three types of iterative methods exist to solve such eigenproblems [8,9]. The first method is the subspace iterations method which is just a generalization of the well-known power method. A second method is the Jacobi–Davidson algorithm. The third one is to use Krylov based methods such as the Arnoldi method or the unsymmetric Lanczos method. Comparisons [11] between the subspace iterations method and the Krylov-type ones tend to show that the second ones are more efficient when applied on large sparse matrices. We have chosen to use the Arnoldi method because it is easy to implement. Its backward stability is now well understood and it does not require any heuristics whereas numerical difficulties such as serious breakdowns can be encountered using the unsymmetric Lanczos method.

In this paper we consider as a model problem the computation of the inertial modes of a rotating spherical shell. This problem contains many difficulties that make it a very good test case. A first difficulty is that the problem is essentially 2D, because variables such as the radial distance and the polar angle cannot be separated. The size of the matrices thus grows very quickly with the resolution; for parameters of physical interest matrices are of order 10^5 or larger. The second difficulty is that the partial differential equations become of hyperbolic-type and therefore yields, with boundary conditions, an ill-posed problem [14,13]. Solving this eigenvalue problem is therefore demanding numerically. The third difficulty is that matrices are highly non-normal. The eigenvalue spectrum is thus very sensitive to machine precision and special tools must be used to analyze and control the roundoff errors. Our analysis revealed some interesting aspects from the viewpoint of numerical precision. In particular we think that our results on the behaviour of roundoff and spectral errors and their interplay are useful in many fields of physics where 2D eigenvalue problems appear.

We organized the paper as follows: we first describe the physics of our test-problem and how we discretize it using spectral methods (Section 2). We also briefly recall the principle of the incomplete Arnoldi–Chebyshev algorithm (Section 3). We then discuss the role of spectral resolution (Section 4) and presents our results about the behaviour of round-off errors (Section 5); conclusions and outlooks follow.

2. The test-problem

2.1. Formulation

We consider the problem of finding the modes of oscillation of a rigidly rotating fluid contained in a spherical shell; it is investigated for its astrophysical and geophysical applications (see [14,13] for a detailed discussion). Eigenmodes of this system are called inertial modes.

The fluid is contained between two spheres of radii ηR and R ($\eta < 1$) and rotates at an angular velocity Ω around the z -axis. Choosing R as the length scale and $(2\Omega)^{-1}$ as the time scale, the non-dimensional form of the equations for the equations governing perturbations are

$$\begin{cases} E\Delta\nabla \times \mathbf{u} - \nabla \times (\mathbf{e}_z \times \mathbf{u}) = \lambda \nabla \times \mathbf{u}, \\ \nabla \cdot \mathbf{u} = 0, \end{cases} \quad (1)$$

where λ is the eigenvalue (a non-dimensional frequency) and $E=v/2\Omega R^2$ is the Ekman number. E is the non-dimensional measure of the kinematic viscosity v and is usually a small parameter ($E < 10^{-4}$).

Eq. (1) are completed by boundary conditions on the velocity taken at $r=\eta$ and $r=1$. We impose stress-free boundary conditions, namely that

$$u_r = \frac{\partial}{\partial r} \left(\frac{u_\theta}{r} \right) = \frac{\partial}{\partial r} \left(\frac{u_\varphi}{r} \right) = 0$$

on the boundaries (r, θ, φ are the usual spherical coordinates).

2.2. Numerical method

We discretize the preceding partial differential equations using spectral methods because of their efficiency at convergence [10,3].

For obvious geometrical reasons, the angular part of the fields is expanded on spherical harmonics; hence, we set

$$\mathbf{u} = \sum_{l=0}^{+\infty} \sum_{m=-\ell}^{m=+\ell} u_m^\ell(r) \mathbf{R}_\ell^m + v_m^\ell(r) \mathbf{S}_\ell^m + w_m^\ell(r) \mathbf{T}_\ell^m,$$

where

$$\mathbf{R}_\ell^m = Y_\ell^m(\theta, \phi) \mathbf{e}_r, \quad \mathbf{S}_\ell^m = \nabla Y_\ell^m, \quad \mathbf{T}_\ell^m = \nabla \times \mathbf{R}_\ell^m$$

and where $Y_\ell^m(\theta, \phi)$ are normalized spherical harmonics (gradients in the definitions of \mathbf{S}_ℓ^m and \mathbf{T}_ℓ^m are taken on sphere of unit radius).

Following some simple rules, given in [12], the equation of vorticity (1a) may be projected rather easily on spherical harmonics. The radial functions $u_m^\ell(r)$ and $w_m^\ell(r)$ then obey the following system:

$$\begin{cases} \left(E \Delta_\ell + \frac{im}{\ell(\ell+1)} \right) w_m^\ell + A_m(\ell) r^{\ell-1} \frac{d}{dr} \left(\frac{u_m^{\ell-1}}{r^{\ell-2}} \right) + A_m(\ell+1) r^{-\ell-2} \frac{d}{dr} (r^{\ell+3} u_m^{\ell+1}) = \lambda w_m^\ell, \\ \left(E \Delta_\ell + \frac{im}{\ell(\ell+1)} \right) \Delta_\ell(r u_m^\ell) - B_m(\ell) r^{\ell-1} \frac{d}{dr} \left(\frac{w_m^{\ell-1}}{r^{\ell-1}} \right) - B_m(\ell+1) r^{-\ell-2} \frac{d}{dr} (r^{\ell+2} w_m^{\ell+1}) = \lambda \Delta_\ell(r u_m^\ell), \end{cases} \quad (2)$$

where we have eliminated the v_m^ℓ 's using $\nabla \cdot \mathbf{u} = 0$. The following notations have also been introduced:

$$A_m(\ell) = \frac{1}{\ell^2} \sqrt{\frac{\ell^2 - m^2}{(2\ell-1)(2\ell+1)}}, \quad B_m(\ell) = \ell^2(\ell^2 - 1)A_m(\ell), \quad \Delta_\ell = \frac{1}{r} \frac{d^2}{dr^2} r - \frac{\ell(\ell+1)}{r^2}.$$

System (2) is an infinite set of differential equations where the coupling between radial functions of indices $\ell - 1$, ℓ and $\ell + 1$ is due to the Coriolis force. Note that different m 's are not coupled.

For the discretization in radial coordinate we approximate the radial functions by truncated expansions of $N + 1$ Chebyshev polynomials. Thus, each of the functions may be represented either by its spectral components or by its values on the Gauss Lobatto collocation nodes. We use the latter representation. In such case, differential operators d^k/dr^k are represented by full matrices of order $(N + 1)$. As system (2) couples radial functions of indices $\ell - 1$, ℓ and $\ell + 1$, it yields a generalized eigenvalue problem with tridiagonal block matrices which we write symbolically

$$\begin{pmatrix} \hat{A}_{-1,\ell}^I & \hat{A}_{0,\ell}^I & \hat{A}_{1,\ell}^I & 0 \\ 0 & \hat{A}_{-1,\ell}^{II} & \hat{A}_{0,\ell}^{II} & \hat{A}_{1,\ell}^{II} \end{pmatrix} \begin{pmatrix} \hat{w}_{\ell-1} \\ \hat{w}_\ell \\ \hat{w}_{\ell+1} \\ \hat{w}_{\ell+2} \end{pmatrix} = \lambda \begin{pmatrix} \hat{B}_{-1,\ell}^I & \hat{B}_{0,\ell}^I & \hat{B}_{1,\ell}^I & 0 \\ 0 & \hat{B}_{-1,\ell}^{II} & \hat{B}_{0,\ell}^{II} & \hat{B}_{1,\ell}^{II} \end{pmatrix} \begin{pmatrix} \hat{u}_{\ell-1} \\ \hat{u}_\ell \\ \hat{u}_{\ell+1} \\ \hat{u}_{\ell+2} \end{pmatrix}. \quad (3)$$

For each value of ℓ , $\hat{A}_{-1,\ell}^I$, $\hat{A}_{0,\ell}^I$ and $\hat{A}_{1,\ell}^I$ are $(N + 3) \times (N + 3)$ matrices. They correspond to the discretization of the l.h.s. of the first equation in (2) at the $N + 1$ Gauss Lobatto nodes, and of the boundary condition $r(dw_m^\ell/dr) - w_m^\ell = 0$ imposed at the two radial boundaries. ℓ runs from m to L by steps of two when m is even and it runs from $m + 1$ to L by steps of two when m is odd. Similarly $\hat{A}_{-1,\ell}^{II}$, $\hat{A}_{0,\ell}^{II}$ and $\hat{A}_{1,\ell}^{II}$ are $(N + 5) \times (N + 5)$ matrices corresponding to the discretization of the l.h.s. of the second equation in (2) at the Gauss Lobatto nodes, plus boundary conditions $u_m^\ell = d^2 u_m^\ell/dr^2 + (2/r) du_m^\ell/dr = 0$ at the two radial boundaries.

3. The incomplete Arnoldi—Chebyshev method

For efficiency reasons and memory requirements, the generalized eigenvalue problem (3) should be solved using an iterative method because the matrices are large and sparse. As previously stated, a good method is the incomplete Arnoldi—Chebyshev method which we now briefly recall.

Let $K(u, A) = \{u, Au, \dots, A^{m-1}u\}$ be the Krylov subspace built from the initial vector u , $V_m = \{v_i\}_{i=1,\dots,m}$ of size $n \times m$ be an orthonormal basis of this subspace.

For applications to stability problems, one is mostly interested in the least-stable (or most unstable) eigenmodes which are associated with the generalized eigenvalues λ with the greatest real part. Since these eigenvalues obviously do not belong to the outside part of the spectrum, we have to perform a spectral transformation. Let (μ, \bar{y}) be the solutions of

$$([A] - \sigma[B])^{-1}[B]\bar{y} = \mu\bar{y}. \quad (4)$$

Then, one easily shows that $(\lambda = \sigma + 1/\mu, \bar{x} = \bar{y})$. Thanks to this spectral transformation, the eigenvalues near the shift (the guess) σ are sent to the outside part of the spectrum and the Arnoldi method can now deliver the desired eigenpair very efficiently. The derived method can be summarized as follows:

Algorithm 1. Parameter: integers r (number of desired eigenpairs), m (number of Arnoldi steps), with $r \leq m \ll n$, Arnoldi starting vector u and degree k of Chebyshev acceleration polynomial.

1. Perform m steps of the Arnoldi method starting from u to compute V_m and H_m :

$$([A] - \sigma[B])^{-1}[B]V_m = V_m H_m + v_{m+1}e_m^T.$$

2. Compute the eigenpairs $(\mu_i, y_i)_{i=1:m}$ by applying the QR algorithm to H_m :

$$H_m y_i = \mu_i y_i.$$

3. If the stopping criterion is satisfied for the r wanted eigenvalues then go to step 6.

4. Compute the parameters of the ellipse containing the $m - r$ unwanted eigenvalues of H_m and set $z_0 = \sum_{i=1}^r \alpha_i V_m y_i$ where

$$\alpha_i = \frac{\|([A] - \sigma[B])^{-1}[B]V_m y_i - \mu_i V_m y_i\|}{(\|([A] - \sigma[B])^{-1}[B]\| \|V_m y_i\|)}.$$

5. Perform k steps of the Chebyshev acceleration starting from z_0 to obtain a better starting vector u for the Arnoldi method; go to step 1.

6. Set $\{\lambda_i = \sigma + 1/\mu_i, \bar{x}_i = V_m y_i\}_{i=1:r}$.

This algorithm requires a matrix–vector product involving the matrix $[B]$ and a linear solver to compute z_2 solution of $([A] - \sigma[B])z_2 = z_1$. In our application, $[A]$ and $[B]$ are banded matrices, so that a band linear solver from LAPACK has been used. The internal dense eigensolver in step 2 has been taken from EISPACK. The interested reader is referred to [4–6,2] for more details.

4. The role of spatial resolution

We first study the convergence of eigenvalues and eigenvectors as a function of the resolution. The two relevant parameters are the degree of the largest Chebyshev polynomial (equal to the number of radial nodes N minus one) and the degree L of the last spherical harmonic. We deal only with axisymmetric $m = 0$ modes and therefore drop the index m (no additional difficulty arises when $m \neq 0$).

In the following we use the notation ω for the imaginary part of the eigenvalue (the frequency), and τ for the real part. Thus, $\lambda \equiv \tau + i\omega$. All the modes of our test-problem are stable, i.e. $\tau < 0$, and $|\tau|$ is the damping rate.

We define the Chebyshev and Legendre spectra of the field u with spectral components $u(\ell, n)$ in the following way:

$$\mathcal{C}(n) = \frac{\max_\ell |u(\ell, n)|}{\max_{\ell,n} |u(\ell, n)|}, \quad \mathcal{L}(\ell) = \frac{\max_n |u(\ell, n)|}{\max_{\ell,n} |u(\ell, n)|}.$$

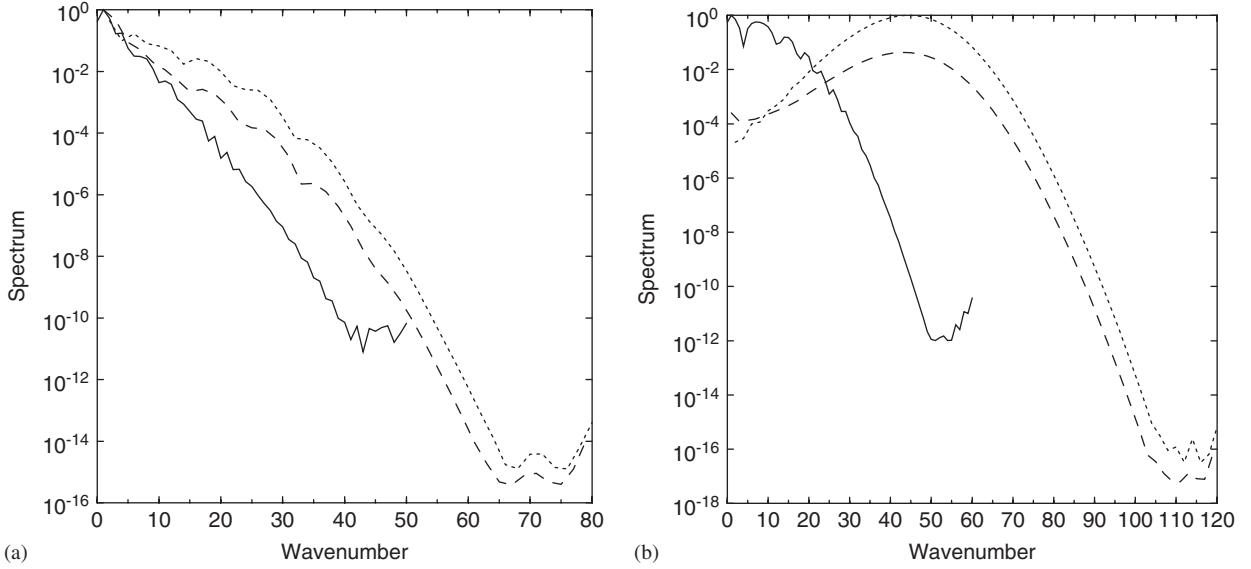


Fig. 1. Chebyshev (solid line) and Legendre (dashed line for ru and dotted line for w) spectra. (a): mode at $E = 10^{-4}$ with $\omega = 0.657976$ and $\tau = -0.00875$. (b): mode at $E = 10^{-4}$ with $\omega = 0.654580$ and $\tau = -0.51$.

Both spectra are filled because inertial modes display very fine structures (see [14] for typical spectra and eigenfunctions occurring at Ekman numbers as low as $E = 10^{-8}$).

Here, we take a moderately small value of the Ekman number: $E = 10^{-4}$ so that the full eigenvalue spectrum can be explored with an affordable resolution. As it may be expected, eigenvalues with smaller $|\tau|$ require less resolution to converge than those with large damping rate (see Figs. 1(a) and (b)). This is easily understood since eigenvectors with small $|\tau|$ tend to have a smoother pattern, which is well approximated by a small number of spectral modes.

The convergence of eigenvalues as a function of spatial resolution goes together with that of eigenvectors: unless all the scales present in the eigenvector are resolved, both the eigenvector and the eigenvalue are not well approximated. This observation allows us to give a simple rule to check the convergence of the eigenvalue. Let us define the ratio f_L between the lowest spherical harmonics coefficient and the largest one, and define g_N as the same ratio but for the Chebyshev expansion.

$$f_L = \frac{\min \mathcal{L}(l)}{\max \mathcal{L}(l)}, \quad g_N = \frac{\min \mathcal{C}(n)}{\max \mathcal{C}(n)}. \quad (5)$$

These two ratios measure the truncation error in the spherical harmonic expansion (f_L) and in the Chebyshev expansion (g_N). We next define ε as the absolute value of the difference between the computed eigenvalue and the converged one (i.e. obtained with a large resolution).

In Fig. 2 we plot ε as a function of f_L . The number of Chebyshev polynomials was chosen large enough to resolve completely the radial dependence. We clearly see that ε follows the law $\varepsilon \propto f_L^2$ until a plateau is reached. The plateau appears at the largest resolutions and indicates that no better approximation to the eigenvalue can be obtained by increasing the resolution. It gives thus a measure of the round-off error of the computation. From the curves obtained for different eigenmodes, we see that *the round-off error is a rapidly increasing function of the damping rate*.

In Fig. 3 we plot ε as a function of the parameter g_N . Here the number of spherical harmonics was set large enough to fully resolve the angular dependence. The Chebyshev convergence appears to be governed by the law $\varepsilon \propto g_N$. Here too, good convergence is obtained only for least-damped modes. We note also that the plateau values are very close to those of the preceding figure.

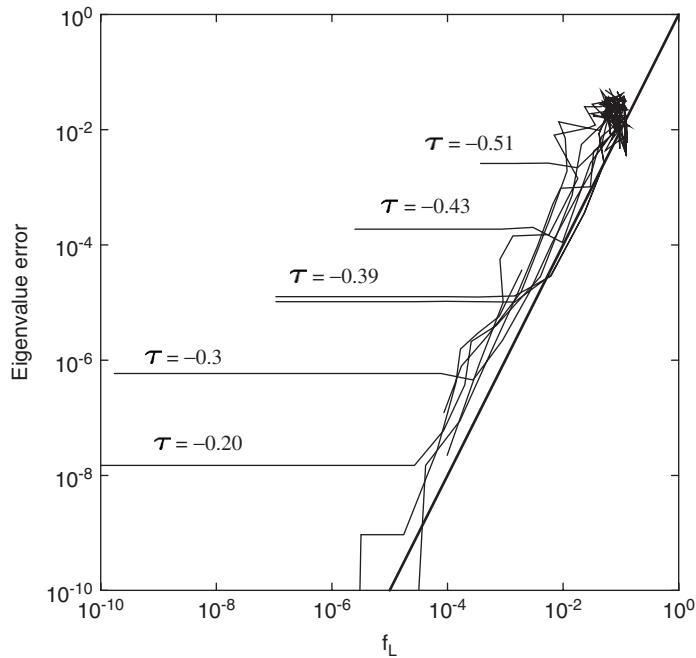


Fig. 2. Error ε of the computed eigenvalue plotted as a function of the Legendre truncation error of the eigenvector f_L (Eq. (5)). Different curves correspond to different eigenmodes. Thick line corresponds to the law $\varepsilon = f_L^2$. Note the horizontal plateau at large resolutions, due to round-off errors.

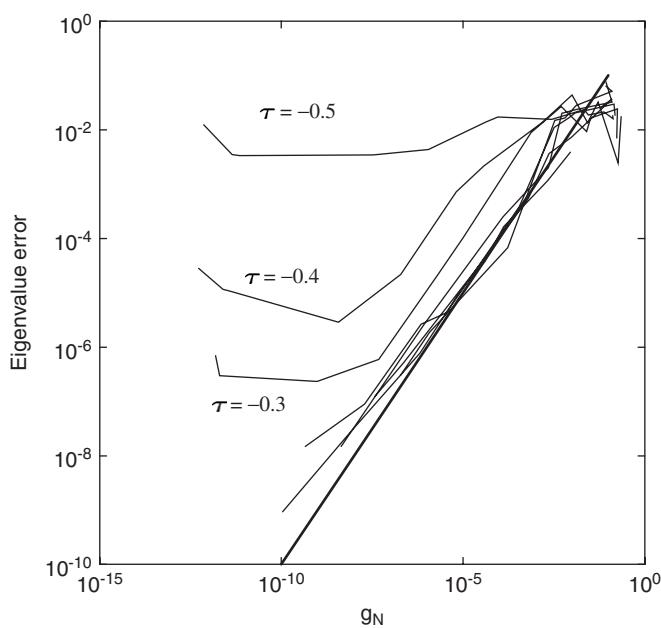


Fig. 3. Error ε of the computed eigenvalue plotted as a function of the Chebyshev truncation error of the eigenvector g_N (Eq. (5)). Thick line corresponds to the law $\varepsilon = g_N$.

5. The importance of round-off errors

The foregoing results indicate that round-off errors play a major role in the accuracy of the numerical solution, especially for strongly damped modes. We shall now investigate this point more thoroughly.

First, we stress the fact that good spectral convergence is not at all a guarantee against round-off errors. This point can be made very clear using the mode displayed on Fig. 1(b) for example. No doubts that for such a mode the spectral expansion has converged: there are 12 decades in the Chebyshev spectrum and 16 decades in the Legendre spectrum; however, the whole spectrum is subject to large round-off error at all wavenumbers. To illustrate this point we consider two different computations where we only change the value of the shift σ (see Eq. (4)) of the Arnoldi–Chebyshev algorithm: $\sigma = -0.51 + i0.65458$ in the first case and $\sigma = -0.51 + i0.65558$ in the second case. In both computations the Ekman number is $E = 10^{-4}$, $L = 120$ and $N = 64$; the Arnoldi–Chebyshev algorithm converges to the same eigenmode. The Chebyshev and Legendre spectra for the first case are those represented in Fig. 1(b); the two spectra for the second case are similar. We plot in Fig. 4 the relative difference of the spectral coefficients, defined as

$$\delta\mathcal{C}(n) = \frac{|\mathcal{C}_2(n) - \mathcal{C}_1(n)|}{0.5(\mathcal{C}_2(n) + \mathcal{C}_1(n))}, \quad \delta\mathcal{L}(n) = \frac{|\mathcal{L}_2(n) - \mathcal{L}_1(n)|}{0.5(\mathcal{L}_2(n) + \mathcal{L}_1(n))},$$

where subscript 1 (resp. 2) corresponds to first (resp. second) eigenvector. We see that the relative difference is spread almost uniformly throughout the wavenumbers, until round-off error in the spectrum is reached, where necessarily the relative error grows to $\mathcal{O}(1)$. This uniform spreading is not surprising: in [1] it is shown that for Chebyshev expansions the spectral round-off error of differential operators is distributed uniformly among wavenumbers.

The round-off error may be investigated quite systematically by computing the spectral portrait of this eigenvalue problem.

Spectral portraits and pseudospectra have recently attracted the attention as a tool of choice for investigating spectral properties of non-normal matrices (see [7,15,17]). It consists in the representation of the map

$$z \rightarrow \text{spp}(z) = \log_{10}[\|(A - zB)^{-1}\|_2(\|A\|_2 + |z|\|B\|_2)]$$

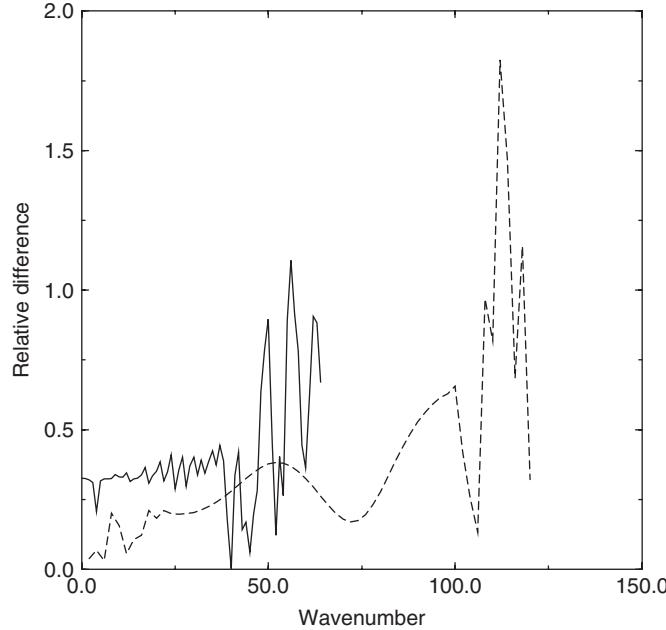
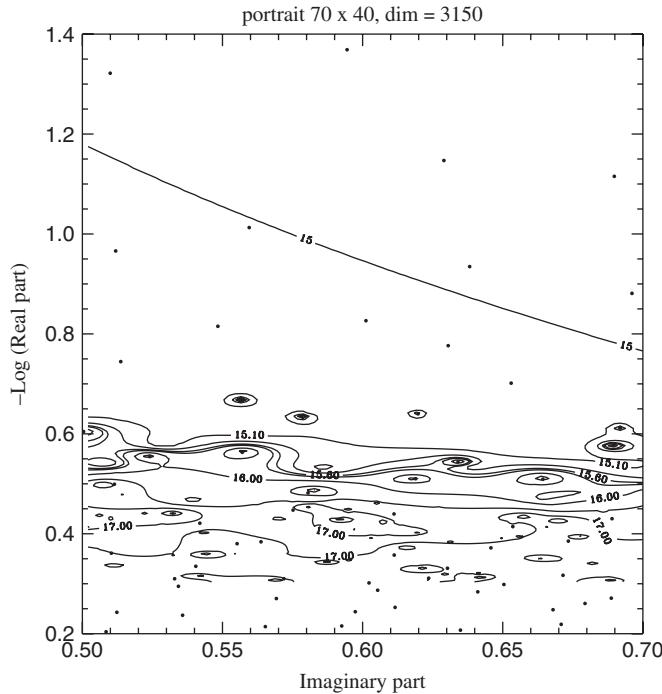


Fig. 4. Relative difference of the spectral coefficients obtained with two computations where the only difference is a change of the Arnoldi–Chebyshev shift σ of the Arnoldi–Chebyshev algorithm. The mode is that of Fig. 1(b). Dashed line: Legendre coefficients $\delta\mathcal{L}(n)$. Solid line: Chebyshev coefficients $\delta\mathcal{C}(n)$.

Fig. 5. Spectral portrait. $E = 10^{-4}$, $L = 70$, $N = 40$.

in a prescribed region of the complex plane. The contour lines of level ε of the spectral portrait are the borders of the ε -pseudospectrum of the matrix pair (A, B) : they enclose all the eigenvalues of the matrix pairs $(A + \Delta A, B + \Delta B)$ and with $\|\Delta A\|_2 \leq \varepsilon \|A\|_2$ and $\|\Delta B\|_2 \leq \varepsilon \|B\|_2$.

If ε is chosen as the backward error for a computed eigenvalue $\tilde{\lambda}$, then the contour line of level ε encloses all the complex numbers with the same backward error ε for the pair (A, B) . The larger the enclosed area, the worse-conditioned the eigenvalue. The diameter of the enclosed area gives an idea of the largest possible relative error on $\tilde{\lambda}$. For a semisimple eigenvalue, it is always possible to bound the error on $\tilde{\lambda}$ by the product of the condition number and the backward error. This is not possible for multiple defective eigenvalues and the spectral portrait is a useful alternative.

On the example studied here, the computation is backward stable: we can then look at the contour line of level machine precision that is 10^{-16} (only “ $-\log_{10}\varepsilon$ ” appears on the figures). We see that this level curve encloses a large region of the spectrum, which tends to indicate a significant spectral instability in the matrix pairs under study.

We display in Fig. 5 the spectral portrait for our eigenvalue problem using a resolution of $L = 70$ and $N = 40$ which corresponds to matrices of order 3150. We superpose the eigenvalues obtained using the QZ algorithm (black points) and the isolines of spectral portrait. For values of the spectral portrait larger than approximately 16 (lower part of the figure) the computed eigenvalues are completely undetermined in double precision. This corresponds to damping rates larger than 0.25 approximately.

However, computation of pseudospectra is an expensive task and it is therefore not feasible on production runs. There have been recent developments in the algorithms whereby one can obtain an approximation to the pseudospectra in a region near the interesting eigenvalues at reasonable cost [16,18]. However, those techniques must be used with special care as they are not totally reliable in the case where the matrix is highly non-normal. A cheap technique that is used routinely to evaluate the sensitivity of eigenvalues to round-off error is to compute the eigenvalues of randomly perturbed matrices. This technique can be used without further coding on any eigenvalue solver, yet one must code the perturbations to the matrix elements. In the following we explore the impact of round-off errors by means of matrix perturbation, and present a new technique which gives the same results but does not require any coding at all.

In Fig. 6 we plot the eigenvalues obtained by making several calculations on perturbed matrices. Each point in the figure is the eigenvalue obtained by perturbing the two matrices A and B of Eq. (4) by random values uniformly

Table 1

Statistics for the computed eigenvalues of Figs. 6 and 7

	Matrix perturbation	Shift perturbation	Matrix perturbation in quadruple precision
$\bar{\tau}$	-0.38520966	-0.38521089	-0.385210900523
$\bar{\omega}$	0.65359249	0.65359283	0.653592789414
$\bar{\tau} - \tau_{QP}$	1.24×10^{-6}	8.67×10^{-9}	1.04×10^{-11}
$\bar{\omega} - \omega_{QP}$	-2.95×10^{-7}	3.67×10^{-8}	-2.65×10^{-11}
σ_τ	7.69×10^{-6}	7.80×10^{-6}	4.99×10^{-9}
σ_ω	6.83×10^{-6}	6.91×10^{-6}	4.68×10^{-9}
$\rho_{\tau\omega}$	0.173	0.142	-0.476
Skewness(τ)	-0.016	-0.004	0.004
Skewness(ω)	-0.003	-0.011	0.0004
Kurtosis(τ)	2.96	2.98	2.93
Kurtosis(ω)	2.98	2.98	2.95

We give the values of the averages $\bar{\tau}$ and $\bar{\omega}$, standard deviations σ_τ and σ_ω , cross correlation $\rho_{\tau\omega}$, skewness and kurtosis for the perturbed matrix case (column 2), the perturbed shift case (column 3) and the perturbed matrix case using extended precision. ($\tau_{QP} = -0.385210900533277$, $\omega_{QP} = 0.653592789440845$) stands for the “exact” eigenvalue computed with quadruple precision.

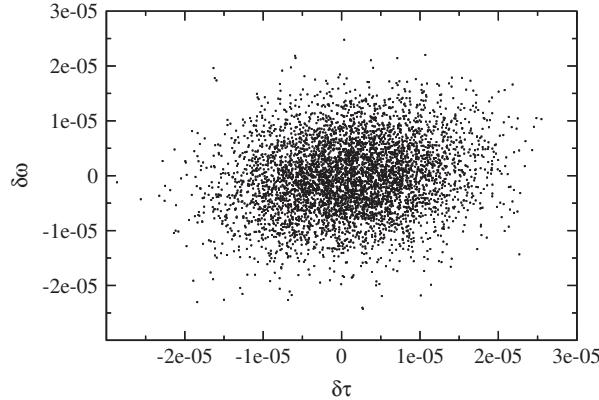


Fig. 6. Plot of several eigenvalues obtained by perturbing randomly the two matrices A and B of Eq. (4). The magnitude of the perturbation is the machine precision 2.22×10^{-16} . The shift is a fixed value near the exact eigenvalue. The Ekman number is $E = 10^{-4}$ and the resolution $L = 94$, $N = 50$. Each black dot in the plot is the difference between the computed eigenvalue and the exact one $\lambda = -0.385210900533277 + i0.653592789440845$.

distributed in the interval $(-\varepsilon_m, \varepsilon_m)$, where $\varepsilon_m = 2.22 \times 10^{-16}$ is the machine precision. The eigenvalues form a cloud of points concentrated in the neighbourhood of the exact eigenvalue $\lambda = -0.385210900533277 + i0.653592789440845$ (this eigenvalue has been obtained using extended precision; it is the “exact” eigenvalue of the truncated problem (3) and not the one of the differential problem (2)). We did a statistical analysis on a large number of eigenvalues (50 000). The results are summarized in Table 1 and Figs. 6–9. Both the real and imaginary parts follow quite well a gaussian law: this can be seen in Fig. 9 where the probability density functions of computed eigenvalues are plotted together with the gaussian curve that fits at best the data. More quantitatively, the skewness and kurtosis (Table 1) are very close to those of the normal distribution (resp. 0 and 3). The order of magnitude of the round-off error is given by the standard deviation of the data which is $\sigma_\tau \simeq 7.69 \times 10^{-6}$ for the real part and $\sigma_\omega \simeq 6.83 \times 10^{-6}$ for the imaginary part. Note that the covariance $\sigma_{\tau\omega} = \sum_{i=1}^n (\tau_i - \bar{\tau})(\omega_i - \bar{\omega})$ (like the correlation coefficient $\rho_{\tau\omega} = \sigma_{\tau\omega}/\sigma_\tau\sigma_\omega$) is small but non-zero, which means that the error distributions for the real and imaginary parts are slightly correlated. We remark that the standard deviations σ_τ and σ_ω have similar values, i.e. the round-off error on τ is of the same order of magnitude as that on ω , even though $|\tau| \ll \omega$.

The standard deviations σ_τ and σ_ω turn out to be essentially independent of the number of Chebyshev polynomials and spherical harmonics, provided that both spectra are well resolved. The values increase when the damping rate of the mode is increased, in perfect accordance with the plateaux observed in Figs. 2 and 3.

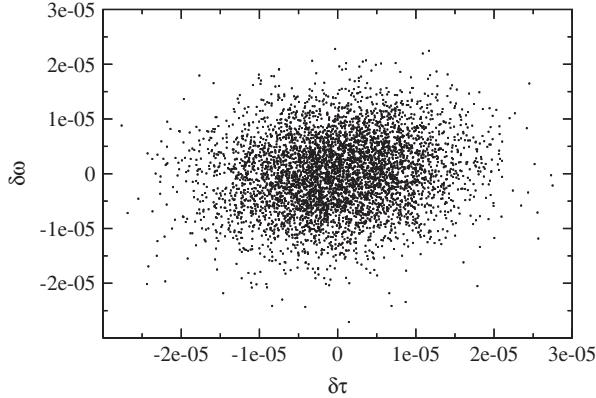


Fig. 7. Plot of several eigenvalues obtained by making different calculations where the only change is the shift parameter σ of Eq. (4). σ is changed by a random perturbation of magnitude 10^{-5} near the exact eigenvalue. The parameters are the same as in Fig. 6. We remark that the two plots are hardly distinguishable.

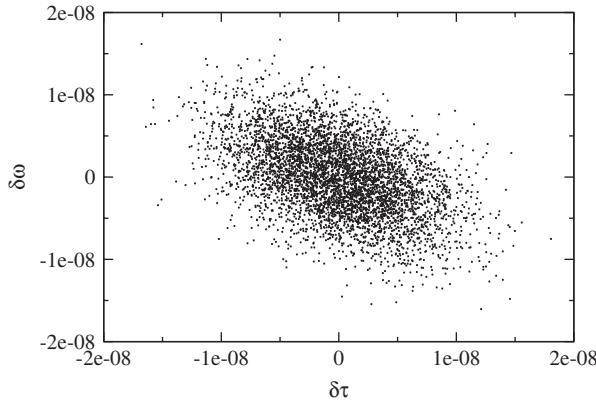


Fig. 8. Same as Fig. 6, only the calculations are done using extended precision.

In a second series of 50 000 computations we did not perturb the matrices A and B but instead we perturbed the value of the Arnoldi–Chebyshev shift σ by a small random quantity near the exact eigenvalue. With this method there is no need to modify the source code for the eigenvalue solver and/or for the construction of the matrices: we only need to change the shift parameter σ on input to the eigenvalue solver. We obtain a cloud of eigenvalues (Fig. 7) which looks almost identical to that obtained in Fig. 6. Each point in the figure is the eigenvalue obtained by changing the real and imaginary part of the shift around the exact eigenvalue by random values uniformly distributed in the interval $(-10^{-5}, 10^{-5})$. Actually we have verified that the statistics does not depend on the amplitude of the shift perturbation. So there is no need to know a priori the exact value of the eigenvalue: any value of the shift which delivers the wanted eigenmode is good. The statistical values in Table 1 confirm that the statistics of the eigenvalues obtained in the two approaches are almost the same.

In a third series of 50 000 computations the matrices A and B are perturbed as in the first series by random values uniformly distributed in the interval $(-\varepsilon_m, \varepsilon_m)$, where $\varepsilon_m = 2.22 \times 10^{-16}$. However in this series the computation is performed using extended precision. Thus, we measure directly the sensitivity of the eigenvalues to perturbation of the matrices; in other words we compute the ε_m -pseudoeigenvalue. From Table 1 and Fig. 8 we see that there are about three digits of difference between the standard deviations of the first two series and the present one: this means that the Arnoldi–Chebyshev algorithm adds an extra factor of order of magnitude 10^3 to the round-off error.

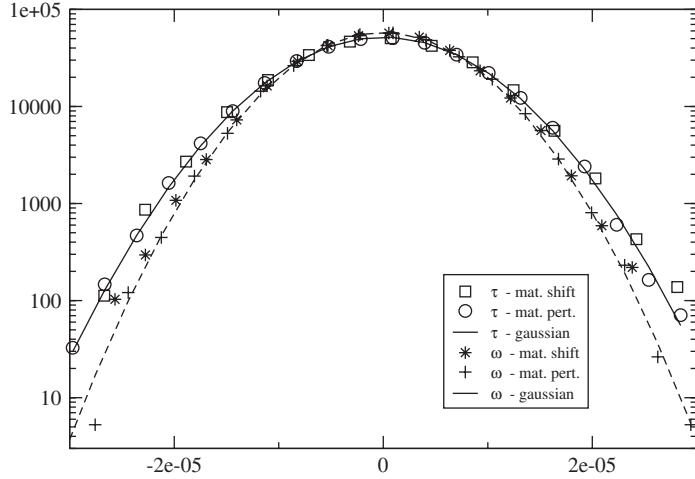


Fig. 9. Probability density functions for the computed eigenvalues of Figs. 6 and 7. In abscissas are the differences between the real part of the eigenvalue and the average value $\bar{\tau}$ for the perturbed shift case (squares) and the perturbed matrix case (circles), and the differences between the imaginary part of the eigenvalue and the average value $\bar{\omega}$ for the perturbed shift case (stars) and the perturbed matrix case (plus). The continuous and broken lines corresponds to the gaussian curves which fit at best the data. We see that gaussian fit is almost perfect.

6. Conclusions

We have analyzed in this paper the errors that arise from the discretization and numerical computation of partial differential eigenvalue problems yielding large matrices. We have chosen as a model problem the 2D eigenvalue problem yielded by the computation of inertial modes in a spherical shell.

We have solved this problem using spectral methods for discretization and the incomplete Arnoldi–Chebyshev algorithm for solving the eigenvalue problem. The combination of these methods provides an efficient solver for these large (2D) eigenvalue problems.

We have shown that the convergence of the eigenvalue and the eigenvector, with respect to spatial truncation, are tightly related: the absolute error of the eigenvalue decreases linearly with the Chebyshev truncation error, and quadratically with the spherical harmonics truncation error, until round-off error becomes dominant.

We found that most-damped modes are the most ill-conditioned and are therefore more sensitive to round-off error. This is made clear by the spectral portrait of the linear operator. Its computation is however very expensive and can be done only on small test-problems. We have shown that a good estimation of the round-off error can be done in practice by performing different computations changing only the value of the Arnoldi–Chebyshev shift parameter on input; there is no need to do extra coding and/or to use external tools. It turns out that the round-off error on eigenvalues has an almost normal distribution, a result which can be used to reduce this kind of error. If the computation of a single eigenmode is not too expensive one could take advantage of this distribution of errors and perform N computations with random shifts; one can thus reduce the round-off error of the estimated eigenvalue by a factor \sqrt{N} .

References

- [1] M. Arioli, L. Valdettaro, Round-off error analysis of the fast cosine transform and its application to the Chebyshev pseudospectral method, *East-West J. Numer. Math.* 3 (1995) 43–58.
- [2] M. Bennani, T. Braconnier, J.-C. Dunyach, Solving large-scale nonnormal eigenproblems in the aeronautical industry using parallel BLAS, in: W. Gentzsch, U. Harms (Eds.), *High-Performance Computing and Networking*, vol. 796, Springer, Berlin, 1994, pp. 72–77.
- [3] S. Bonazzola, E. Gourgoulhon, J.-A. Marck, Spectral methods in general relativistic astrophysics, *J. Comput. Appl. Math.* 109 (1999) 433.
- [4] T. Braconnier, The Arnoldi–Chebyshev algorithm for solving large complex non Hermitian generalized eigenproblems, Technical Report TR/PA/94/08, CERFACS, 1994.
- [5] T. Braconnier, Influence of orthogonality on the backward error and the stopping criterion for Krylov methods, Numerical Analysis Report 281, Department of Mathematics, University of Manchester, 1995.
- [6] T. Braconnier, F. Chatelin, J.-C. Dunyach, Highly nonnormal eigenvalue problems in the aeronautical industry, *Japan J. Ind. Appl. Math.* 12 (1995) 123–136.

- [7] F. Chaitin-Chatelin, V. Frayssé, Lectures on Finite Precision Computations, SIAM, Philadelphia, 1996.
- [8] F. Chatelin, Eigenvalues of Matrices, Wiley, Chichester, 1993 Enlarged Translation of the French Edition with Masson.
- [9] D.R. Fokkema, G.L.G. Sleijpen, H.A. van der Vorst, Jacobi–Davidson style QR and QZ for the partial reduction of matrix pencils, SIAM J. Sci. Comput. 20 (1) (1998) 94–125.
- [10] B. Fornberg, A Practical Guide to Pseudospectral Methods, Cambridge University Press, Cambridge, 1998.
- [11] R.B. Lehoucq, J.A. Scott, An evaluation of software for computing eigenvalues of sparse nonsymmetric matrices, Technical Report, MCS-P547-1195, Argonne National Laboratory, Argonne, IL, 1995.
- [12] M. Rieutord, Linear theory of rotating fluids using spherical harmonics. I, Steady flows, Geophys. Astrophys. Fluid Dyn. 39 (1987) 163.
- [13] M. Rieutord, B. Georgeot, L. Valdettaro, Inertial waves in a rotating spherical shell: attractors and asymptotic spectrum, J. Fluid Mech. 435 (2001) 103–144.
- [14] M. Rieutord, L. Valdettaro, Inertial waves in a rotating spherical shell, J. Fluid Mech. 341 (1997) 77–99.
- [15] V. Toumazou, Portraits spectraux de matrices: un outil d’analyse de la stabilité. Ph.D. Dissertation, Université H. Poincaré, Nancy, 1996.
- [16] L.N. Trefethen, Computation of pseudospectra, Acta Numer. 8 (1999) 247.
- [17] L.N. Trefethen, A.E. Trefethen, S.C. Reddy, T.A. Driscoll, Hydrodynamics stability without eigenvalues, Science 261 (1993) 578–584.
- [18] T.G. Wright, L.N. Trefethen, Large-scale computation of pseudospectra using Arpack and eigs, SIAM J. Sci. Comput. 23 (2001) 591–605.