Kinematic roll dynamo computations at large magnetic Reynolds number

Henrik Latter David Ivers*

(Received 8 August 2003, revised 9 February 2004)

Abstract

The behaviour of the magnetic field produced by spherical kinematic dynamos at large magnetic Reynolds number for two stationary axisymmetric roll flows is examined. Five numerical techniques to solve the large banded eigenvalue problems, which arise for the growth rate of the magnetic field, are compared, particularly their performance when eigenvalues are closely clustered. The five eigenvalue methods are inverse iteration, orthogonal iteration, lop-sided iteration, the implicitly restarted Arnoldi method and the non-symmetric Lanczos method. A shift and invert strategy was employed to obtain the growth rates of largest real part.

*University of Sydney, School of Mathematics & Statistics, Australia.
mailto:hl278@cam.ac.uk

See http://anziamj.austms.org.au/V45/CTAC2003/Latt/home.html for this article, Published August 24, 2004. ISSN 1446-8735

Contents

1 Introduction

We consider the magnetic stability of the steady state $\mathbf{B} = \mathbf{0}$ for a sphere V of radius l filled with incompressible, homogeneous, electrically conducting fluid undergoing a prescribed axisymmetric roll motion $\mathbf{v}$, and surrounded by an insulating exterior. The magnetic field $\mathbf{B}$ satisfies the non-dimensionalized

problem:

$$\begin{aligned} &\frac{\partial \mathbf{B}}{\partial t} = \nabla^2 \mathbf{B} + R_m \nabla \times (\mathbf{v} \times \mathbf{B}) \quad \text{in } V; \\ &\nabla \cdot \mathbf{B} = \mathbf{0} \quad \text{in all space;} \\ &[\mathbf{B}]_{\partial V} = \mathbf{0}\,; \quad \mathbf{B} \to \mathbf{0} \quad \text{as } \mathbf{r} \to \infty\,; \end{aligned} \tag{1}$$

where $R_m = vl/\eta$ is the dimensionless magnetic Reynolds number, v is a typical speed and η is the magnetic diffusivity. The non-dimensionalized radius of V is 1. We seek magnetic solutions of the form $\mathbf{B}(\mathbf{r}, t) = \mathbf{B}_0(\mathbf{r})e^{\lambda t}$, in which case (1) becomes an eigenproblem for the growth rate λ of largest real part.

The standard method to solve the discretised form of the problem, as described in Section 2, is inverse iteration, since only one eigenvalue is sought. However, for the roll flows considered, the eigenvalues cluster for large R_m and, moreover, the imaginary parts are $\mathcal{O}(R_m^{1/2})$ larger than the real parts. Thus inverse iteration has great difficulty in distinguishing the various modes without an accurate estimate for the eigenvalue. Use of the QR-algorithm to find all eigenvalues is infeasible due to the size of the problem. Instead, we compare the sparse eigenvalue methods described in Section 3, which calculate several eigenvalues at once. Results are given in Section 4 and conclusions in Section 5.

2 The spectral/finite-difference approach for spherical dynamos

The solenoidal fields $\mathbf{B}$, $\mathbf{v}$ and $\nabla \times (\mathbf{v} \times \mathbf{B})$ are first decomposed into toroidal and poloidal fields,

$$\mathbf{B} = \mathbf{T}[T] + \mathbf{S}[S]\,, \quad \mathbf{v} = \mathbf{T}[t] + \mathbf{S}[s]\,, \quad \nabla \times (\mathbf{v} \times \mathbf{B}) = \mathbf{T}[\tilde{T}] + \mathbf{S}[\tilde{S}]\,.$$

This reduces the magnetic induction equation (1) to two scalar equations. Next the poloidal and toroidal scalar potentials are expanded in complex-valued spherical harmonics.

$$f = \sum_{\gamma} f_\gamma(r,t) Y_\gamma(\theta,\phi)\,, \quad \text{where } f = s,\, t,\, S,\, T,\, \tilde{S},\, \tilde{T}\,.$$

The double index γ represents the duple (n_γ, m_γ). Substituting these into the toroidal and poloidal magnetic induction equations, and collecting the coefficients of Y_γ yields the Bullard and Gellman [1] equations,

$$(\partial_t - D_\gamma) S_\gamma = R_m \tilde{S}_\gamma\,, \quad (\partial_t - D_\gamma) T_\gamma = R_m \tilde{T}_\gamma\,, \tag{2}$$

with $D_\gamma = r^{-2}\{\partial_r(r^2\partial_r) - n_\gamma(n_\gamma+1)\}$ the spherical Bessel operator.

The interaction terms on the right side must be expressed in terms of the spectral components of **B** and **v**. The ensuing relations are complicated and involve the evaluation of seven coupling integrals composed of Wigner coefficients and products of the magnetic and velocity potentials, and their radial derivatives:

$$\tilde{S}_\gamma = \sum_{\alpha,\beta} \{(s_\alpha S_\beta S_\gamma) + (s_\alpha T_\beta S_\gamma) + (t_\alpha S_\beta S_\gamma)\}\,;$$
$$\tilde{T}_\gamma = \sum_{\alpha,\beta} \{(s_\alpha S_\beta T_\gamma) + (s_\alpha T_\beta T_\gamma) + (t_\alpha S_\beta T_\gamma) + (t_\alpha T_\beta T_\gamma)\}\,.$$

Each term is a coupling integral, which represents the contribution to the γ mode of S or T from the interaction of the fluid velocity's α toroidal/poloidal mode with the magnetic field's β toroidal/poloidal mode. The sum then includes the contribution from every possible interaction between the toroidal and poloidal modes of **v** and **B**. Note that a toroidal velocity field cannot generate poloidal magnetic field from toroidal magnetic field. We omit the explicit expressions for the coupling integrals. For full details see [1, 6].

We actually employed a slightly alternative formulation of the spectral equations in which we make use of vector spherical harmonic expansions of

both $\mathbf{B}$ and $\mathbf{v}$. A hybrid of the vector spherical harmonic approach and the toroidal-poloidal approach yields spectral equations which are algorithmically simpler than the latter. They involve only two relatively simple coupling integrals making the computation of the linear equation matrix easier to program and less prone to error. We omit the full details, see [4, 5].

After discretising the radial interval into J segments we introduce into D_γ and the coupling integrals centred finite differences for the interior points and a one-sided right boundary scheme for $r = 1$. Truncate the number of harmonics, allocating N each to both S and T. The choice of which of these modes to include can be an intricate business. We determine which ones to include via observing the 'selection rules' of the coupling integrals which are closely related to those of the Wigner coefficients. For axisymmetric velocity flow we find that only magnetic modes of the same azimuthal wave number m contribute to each other. In other words $m_\beta = m_\gamma = m$. The azimuthal dimension separates out and we install m as a constant parameter. We may not choose $m = 0$. Cowling's theorem ensures decaying modes in this case because $\mathbf{v}$ is also axisymmetric.

This done (2) reduces to a large numerical eigenproblem,

$$(\mathbf{A} + R_m\mathbf{B})\mathbf{x} = \lambda\mathbf{x}\,, \tag{3}$$

in which $\mathbf{A}, \mathbf{B}$ are complex, non-symmetric, and nearly block tridiagonal for second-order finite-differences. The size of the blocks is $2N$ except one N sized block arising from the boundary condition. The total matrix sizes are $N(2J-1)$. The eigenvalues $\{\lambda_i\}$ determined from this system depend on J, N and R_m. We regard these as valid approximations if increasing J and N induces negligible change in them.

3 Linear eigenvalue methods

Throughout we assume that our generic matrix $\mathbf{M}$ is non-defective and of fixed order n. Its eigenvalues $\lambda_1, \ldots, \lambda_n$ are indexed so that

$$|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_n| .$$

For a fixed integer $h < n$ with $|\lambda_h| > |\lambda_{h+1}|$ there is a unique invariant subspace $\mathcal{U}_h$, for which the first h dominant eigenvectors of $\mathbf{M}$ are a basis.

Inverse iteration is the standard method if only one eigenvalue is sought. Under reasonable restrictions it produces good approximations to the dominant eigenvector of $(\mathbf{M} - \mu)^{-1}$ and hence the eigenvalue closest to the given shift μ. If we begin with a random vector $\mathbf{v}_0$ the approximation generated by the kth step is $\mathbf{v}_{k+1} = \alpha(\mathbf{M} - \mu)^{-1}\mathbf{v}_k$ where α is some normalization. Of course the inverse is not explicitly formed, $\mathbf{v}_{k+1}$ is determined by a direct solving routine.

It only takes one eigenvalue with positive real part to reverse the magnetic stability. One may then think that inverse iteration is the only method needed in this context. For low magnetic Reynolds numbers this is usually true. But for large R_m we find that the leading modes begin to cluster. Not only does this slow the convergence of inverse iteration, but it also makes it difficult to isolate the eigenvalue of largest real part. Due to the size of our problems the complete eigensolution (via the QR-algorithm) is out of the question and we resort to partial sparse eigenvalue methods.

The methods discussed here approximate the dominant (extremal) eigenvalues first. To obtain approximations to eigenvalues of largest real part we employ a shift and invert strategy (as in inverse iteration). Consequently our algorithms require direct solving steps and hence an LU-decomposition. The matrices with which we deal are banded and large, but not so large as to make this computation prohibitively expensive. The banded LU-decomposition costs about $2npq$ complex flops where $p = 2N$ is the upper band width and $q = 4N$ is the lower bandwidth.

In the following section we outline the standard algorithms only. For the shift-inverse versions replace $\mathbf{M}$ with $(\mathbf{M}-\mu)^{-1}$ and λ_i by $(\mu-\lambda_i^*)^{-1}$ where λ_i^* is the ith closest eigenvalue to the shift μ.

3.1 Simultaneous iteration methods

A simultaneous iteration method is a power method for subspaces. Let h starting vectors make up the columns of the initial matrix $\mathbf{Q}_0$ upon which we iterate.

3.1.1 Orthogonal iteration (OI)

This variant consists primarily of two steps, a power step $\mathbf{MQ}_k = \mathbf{Z}_k$ and a QR factorization $\mathbf{Z}_k = \mathbf{Q}_{k+1}\mathbf{R}_k$. The second step maintains the independence of the columns of $\mathbf{Q}_k$. Without it we would end up with h copies of the dominant eigenvector. It can be proved that the columns of $\mathbf{Q}_k$ converge to the first h dominant Schur vectors of $\mathbf{M}$. The rate of this convergence is variable. For the ith vector it goes like $\mathcal{O}(|\lambda_{i+1}/\lambda_i|^k)$ just like the power method. For clustered eigenvalues this rate is generally unacceptably slow. Stewart [10] proposed an additional 'SRR' step, which is performed intermittently and has a convergence rate of $\mathcal{O}(|\lambda_{r+1}/\lambda_i|^k)$. It projects $\mathbf{M}$ onto the subspace $\mathbf{Q}_k$ and computes its Schur form. The new subspace $\mathbf{Q}_{k+1}$ is determined by right multiplying $\mathbf{Q}_k$ by the matrix of calculated Schur vectors.

3.1.2 Lop-sided iteration (LSI)

We outline the basic structure of the algorithm given by Lehoucq et al. [7]. Step 1 forms the projection of $\mathbf{M}$ onto $\mathbf{Q}_k$, $\mathbf{H} = \mathbf{Q}_k^{-H}\mathbf{M}\mathbf{Q}_k$. Because the $\mathbf{Q}_k$ are non-unitary in general, this requires a linear solving substep. Step 2

eigensolves for the projected matrix $\mathbf{VH} = \mathbf{DV}$, where the columns of $\mathbf{V}$ are $\mathbf{H}$'s right eigenvectors. Step 3 constructs the new subspace $\mathbf{Q}_{k+1} = \mathbf{MQ}_k\mathbf{V}$. The algorithm forces the $\mathbf{Q}_k$ to converge to the matrix of dominant eigenvectors, the ith column converging with rate $\mathcal{O}(|\lambda_{h+1}/\lambda_i|^k)$.

3.2 Krylov subspace methods

These methods exploit the attractive properties of Krylov subspaces,

$$\mathcal{K}_m(\mathbf{u}_0, \mathbf{M}) \equiv \text{span}\{\mathbf{u}_0, \mathbf{Mu}_0, \ldots, \mathbf{M}^m\mathbf{u}_0\}.$$

Each method computes an orthogonal basis for $\mathcal{K}_m(\mathbf{u}_0, \mathbf{M})$, adding an extra basis element per iteration and stopping when the space becomes $\mathbf{M}$-invariant. At this point the eigenvalues extracted from $\mathbf{M}$'s projection in $\mathcal{K}_m(\mathbf{u}_0, \mathbf{M})$ are excellent approximations. However, round-off error usually ensures the iteration never terminates. We truncate the process when a specified number of Ritz values (projected eigenvalues) converge sufficiently.

3.2.1 Non-Hermitian Lanczos method (NHLM)

This method computes two bi-orthogonal bases for $\mathcal{K}_m(\mathbf{u}_0, \mathbf{M})$ so that $\mathbf{M}$'s oblique projection is tridiagonal. Because of the ensuing simple three term recurrence relation a practical and speedy algorithm simply determines the projected matrix entry by entry negating the need for the storage of both basis sets. This allows us to build very large projections with access to a sizable portion of the spectrum. However, leaving the bases unstored renders their orthogonality extremely vulnerable to round-off error. This manifests itself in the intermittent recomputation of those basis elements spanning the most dominant sections of the Krylov subspace. Thus the projected matrix will be 'encoded' with multiple copies (or 'ghosts') of the most dominant eigenvalues. The existence of spurious eigenvalues is another problem, these

arise because of the truncation of the Lanczos iteration. They are in fact unconverged approximates. For more details see [3].

3.2.2 Implicitly restarted Arnoldi method (IRAM)

The basic Arnoldi iteration computes a single orthogonal basis so that $\mathbf{M}$'s projection is upper-Hessenberg. This means the entire Arnoldi basis is required at every step. Its storage requirements consequently prohibit very large subspaces being generated. Also the information about $\mathbf{M}$'s extremal eigenvalues need not emerge within the generated subspace particularly quickly. Convergence largely depends on the choice of the initial vector $\mathbf{u}_0$. This coupled with restrictions on the number of iterations may compromise the convergence of the Ritz values.

One response is the restart idea by which the Arnoldi factorization is begun anew with a fresh initial vector chosen from the span of the previously computed Arnoldi vectors. This restart vector is chosen so that better approximations to desirable eigenvectors can be found in the new Krylov space. There are a number of ways this can be done. We employed the method of Sorenson which determines the restart vector implicitly using the shifted QR iteration (for details see [8, 9, 3]).

4 Results

We analyse the two axisymmetric roll flows [2] with toroidal/poloidal spherical harmonic expansions:

$$\mathbf{v}_1 = \sigma\mathbf{s}_1^0 + \mathbf{t}_1^0\,, \quad \mathbf{v}_2 = \sigma\mathbf{s}_1^0 + \mathbf{t}_1^0 + \mathbf{t}_3^0\,,$$

where $s_1^0 = \sin\pi r$ for both flows and

$$t_1^0 = -\tfrac{4}{5\sqrt{3}} r^2 \sin\pi r\,, \quad t_3^0 = \tfrac{2}{15\sqrt{7}} r^2 \sin\pi r\,,$$

for $\mathbf{v}_1$ and $t_1^0 = s_1^0$ for $\mathbf{v}_2$. The parameter σ measures the poloidal flow speed compared to the toroidal speed. We set $m = 1$ in all calculations and obtain growing modes for certain values of σ. For $\mathbf{v}_1$ we let $\sigma = 0.2050$ and for $\mathbf{v}_2$ $\sigma = 0.1373$. R_m varies from 0 to 500,000. At the highest R_m fourth-order finite-differences are used (which doubled the bandwidth), at lower R_m second-order finite-differences suffice.

The eigenvalue methods were programmed in Fortran using routines from the LAPACK library. For the IRAM we used the routines from ARPACK. Our machine was a Compaq/Digital Alpha 2 CPU, 3 GB memory, ES40.

4.1 Behaviour of eigenvalues

Converged results were obtained for the first few leading eigenvalues but our truncation levels had to be increased for larger R_m. This suggests that the length scales of the eigenvectors decrease and more modes couple as R_m increases.

To obtain converged results for two leading modes at $R_m = 500,000$ we required $N = 30$ and $J = 1000$, yielding matrices of size 60,000. For smaller R_m more leading eigenvalues had converged sufficiently. At the extremal end of the spectrum (eigenvalues of very large negative real part) convergence could not be obtained at our maximum truncation level for any R_m. This suggests possibly that the faster a mode decays the smaller its length scale. This picture is complicated though at lower R_m when the growth rates follow particularly tangled trajectories in the complex plane on increasing R_m. However, we find that for $R_m > 5,000$ and for both flows the trajectories of the leading modes converge to the branches of closely spaced parabolas centred at the origin. Moreover, their imaginary parts approach one another. As the imaginary parts are $\mathcal{O}(R_m^{1/2})$ larger than the real parts, inverse iteration has great difficulty in distinguishing the various modes without an accurate estimate for μ. For example, the leading four $\mathbf{v}_2$-eigenvalues

at $R_m = 50,000$ are $201.5+26147.8\,i$, $-245.8+25672.4\,i$, $-325.4+24678.9\,i$, and $-250.8+23573.1\,i$.

4.2 Performance of eigensolvers

We compare the robustness, maximum number of eigenvalues computable and computational speed of each method as their accuracy levels and R_m are varied.

The IRAM uncovered a maximum of 40 leading eigenvalues and LSI a maximum of 35 for any R_m at relative accuracy 10^{-6}. Method OI could manage 14 when $R_m < 10,000$, but convergence slowed to prohibitive levels when searching for more than 8 when R_m was larger. OI's speed of convergence was very sensitive to choice of shift and the distribution of eigenvalues. Despite the SSR step if a shift lay equally between two eigenvalues convergence was extremely slow or non existent. Storage requirements (other than for $\mathbf{M}$) were of $\mathcal{O}(nh)$ for these three methods, this limited the maximum number of eigenvalues obtainable.

The NHLM produced large numbers (≥ 500) of 'eigenvalues' but (depending on accuracy levels and the shift) at best nearly a half were true, the rest being copies or spurious. These had to be sorted. Convergence was ascertained from a calculation of their residual norms. However, this provides only a guide as $\mathbf{M}$ is non-Hermitian. The maximum residual norm tolerated was 10^{-6} which ensured most accepted eigenvalues had converged to at least 10^{-6} in relative error. There were a few exceptions as expected. The ambiguity in convergence is a significant detraction. We found the larger the projection the more converged the dominant eigenvalues. The number of copies depended greatly on the choice of shift. A shift too close to a true eigenvalue yielded excessive copies of that eigenvalue. The shift and inversion in this case would make the eigenvalue extremely dominant and thus any round off error would be swiftly amplified in the dominant eigenvector's

TABLE 1: CPU times and number of iterations for $\mathbf{v}_1$ and $h = 6$ at $R_m = 1,000$ with shift $(40, 0)$ at varying accuracies.

tol	LSI.	OI	IRAM	NSL
10^{-3}	11.3 (22)(0)	28.7 (72)	4.5 (19)	10.7 (16)
10^{-4}	12.1 (24)(0)	48.7 (138)	5.2 (23)	10.7 (16)
10^{-5}	14.1 (30)(0)	55.3 (160)	5.3 (23)	11.1 (18)
10^{-6}	17.9 (42)(0)	69.6 (207)	5.6 (27)	12.1 (20)
10^{-8}	22.9 (55)(0)	97.8 (296)	6.2 (31)	13.2 (22)

TABLE 2: CPU times and number of iterations for $\mathbf{v}_2$ and $h = 6$ at $R_m = 2,500$ with shift $(0, 1100)$ at varying accuracies.

tol	LSI.	OI	IRAM	NSL
10^{-3}	8.6 (15)(0)	24.3 (56)	4.5 (16)	9.9 (15)
10^{-4}	12.3 (18)(0)	33.0 (79)	4.9 (19)	11.0 (16)
10^{-5}	15.4 (18)(3)	47.6 (117)	5.3 (22)	11.0 (16)
10^{-6}	16.5 (19)(3)	57.3 (143)	5.6 (23)	11.0 (16)
10^{-8}	21.4 (26)(3)	76.1 (190)	6.1 (27)	12.5 (20)

direction when premultiplied by $(\mathbf{M} - \mu)^{-1}$. In general a Lanczos projection of size $3h$ yielded h eigenvalues which had converged with relative accuracy of 10^{-6}. Increasing R_m had little effect on the method's performance. Storage requirements were small, of $\mathcal{O}(h)$.

We used elapsed CPU time as the measure of computational speed. These times could vary slightly, the ones given are representative. For the first four tables we used truncation levels of $N = 20$, $J = 400$. A prediction of the imaginary part of the leading modes' growth rate was usually taken as the shift in each case.

Tables 1 and 2 exhibit the calculated CPU times versus eigenvalue accu-

TABLE 3: CPU times for $\mathbf{v}_1$ and $h = 8$ at varying R_m for tol $= 10^{-6}$.

R_m	Shift	LSI.	OI	IRAM	NSL
0	(0,0)	30.9 (0)	27.8	6.7	22.4
100	(0,0)	34.9 (3)	148.3	8.1	15.6
1000	(0,400)	39.3 (4)	xx	7.2	18.4
10,000	(0,4500)	30.3 (6)	119.8	6.0	18.1
50,000	(0,26000)	30.7 (6)	197.4	6.9	19.2
100,000	(0,52000)	27.1 (6)	xx	7.0	19.5

TABLE 4: CPU times for $\mathbf{v}_1$ and various h at $R_m = 5,000$ with shift $(0, 600)$ with tol $= 10^{-6}$.

h	LSI.	OI	IRAM	NSL
4	8.3 (0)	33.7	4.2	14.9
8	19.3 (2)	85.0	5.4	16.3
12	30.7 (2)	240.1	7.5	25.6
16	52.6 (3)	xx	10.7	26.7
20	96.2 (4)	xx	10.8	32.7

TABLE 5: CPU times for $\mathbf{v}_2$ and $h = 4$ at $R_m = 2,500$ with shift $(0, 1100)$ with tol $= 10^{-6}$ for various truncations.

(N,J)	matrix size	s	LSI.	OI	IRAM	NSL
(20,400)	15980	200	14.3 (20)	28.8 (125)	5.6 (24)	8.9 (12)
(20,200)	7980	200	7.1 (21)	14.3 (124)	2.4 (21)	4.4 (12)
(10,400)	7980	100	3.4 (16)	9.3 (124)	1.3 (22)	2.2 (12)
(10,200)	3990	100	1.7 (17)	4.4 (125)	0.7 (24)	1.1 (12)

racy levels, tol is the relative accuracy used. The number in the brackets is the number of required iterations for the simultaneous iterations, the number of restarts for the IRAM and the size of the smallest Lanczos projection yielding eigenvalues of the requisite accuracy. In Table 2 the second bracketed number under Lop-sided iteration shows the optimal number of guard vectors. In Table 3 we can compare the speed of the methods as R_m was increased for a certain set of parameters. This can only give a rough indication due especially to the sensitivity of the simultaneous iterations to the shift and specific eigenvalue distribution at a particular R_m. We have presented the optimal number of guard vectors for Lop-Sided iteration next to its CPU times. Table 4 shows the dependence of speed on number of eigenvalues sought (h) for a certain set of parameters. Again we have displayed the optimal number of guard vectors for Lop-sided iteration.

Table 5 confirms that for each method and each iteration (or restart) the flop cost goes roughly like $\mathcal{O}(ns)$ for small enough h, and where n is the order of the matrix and s is the total bandwidth.

5 Conclusions

It is clear for the problems we consider and number of eigenvalues sought (< 30) that the IRAM surpasses all other methods in terms of speed. Unlike

the simultaneous iteration methods it is relatively insensitive to the choice of shift and eigenvalue distribution. Unlike the NSL it can produce unambiguous results with a specified accuracy. The NSL method should only be used if hundreds of eigenvalues are required and the simultaneous iterations should not be employed.

For extremely large problems ($n > 100,000$) direct solving may be too expensive on account of the LU-decomposition. The value of iterative solvers has yet to be explored in this context and should be an interesting area of research. A method such as the Jacobi–Davidson method for subspaces [11], which does not employ the inverse-shift technique, should also be examined. We envisage that a key problem here will be the preconditioning because for large R_m our matrices are not diagonally dominant.

References

[1] Bullard, E. and Gellman, C. (1954), "Homogeneous dynamos and terrestrial magnetism", *Phil. Trans. R. Soc. Lond. A.* **8**, 43–56. C908

[2] Dudley, M. L. and James, R. W. (1989), "Time-dependent kinematic dynamos with stationary flows", *Proc. R. Soc. Lond. A.* **425**, 407–420. C913

[3] Golub, G. H. and Van Loan, C. F. (1996), *Matrix Computations*, 3rd ed., The Johns Hopkins University Press, Baltimore. C913

[4] Ivers, D. J. and Phillips, C. G. (2003), "A vector spherical harmonic spectral code for linearised magnetohydrodynamics", *ANZIAM J.* **44(E)**, C423–C442. http://anziamj.austms.org.au/V44/CTAC2001/Ive2 C909

[5] Ivers, D. J. and Phillips, C. G. (2003), "Scalar and vector spherical harmonic spectral equations of rotating non-linear and linearised magnetohydrodynamics", in preparation. C909

[6] James, R. W. (1974), "The spectral form of the magnetic induction equation", *Proc. R. Soc. Lond. A.* **340**, 287–299. C908

[7] Jennings, A. & McKeown, J. J. (1992), *Matrix Computation* 2nd ed., John Wiley and Sons, Chichester. C911

[8] Lehoucq, R. B., Sorensen, D. C., Yang (1997), *ARPACK User's Guide.* C913

[9] Sorensen, D. C. (1992), "Implicit application of polynomial filters in a k-step Arnoldi method", *SIAM J. Matrix Anal. Appl.* **13**, 357–385. C913

[10] Stewart, G. W. (1976), "Simultaneous iteration for computing invariant subspaces of non-hermitian matrices", *Numerical Mathematics* **25**, 123–136. C911

[11] van der Vorst, H. (1996), "Subspace iteration for eigenproblems" *CWI Quarterly* **9**, 151–159. C919