An Efficient Numerical Method for Solving the Eigenvalue Problem Associated with a Quasi-Geostrophic Fluid Flow

Sushil Shetty, SID: 15578395
MATH 221 Class Project, Spring 2003
University of California, Berkeley

Abstract

The study of instabilities in a fluid flow is an important area of fluid dynamics. A linear stability analysis of a typical fluid flow leads to the definition of an eigenvalue problem. The solution of the resulting eigenvalue problem usually requires a numerical method. However, the numerical solution of these problems is often fraught with difficulties that arise because an infinite dimensional operator is approximated by a finite dimensional one and/or the inconsistent application of boundary conditions. Consequently, many of the eigenvalues that are obtained from a numerical solution are spurious with no physical relevance. These eigenvalues have to be identified effectively and efficiently.

In this paper we examine an eigenvalue problem associated with a quasi-geostrophic fluid flow. An efficient numerical method is presented to solve such a problem. This method exploits the tridiagonal structure associated with a finite difference representation of the operator involved. We show that it is efficient with respect to required computations by comparing it with a method typically used for solving stability problems of this nature.

1 Introduction: Physical Motivation

Under certain approximations, a fluid flow is governed by the 2-D Quasi-Geostrophic(QG) equations. The potential vorticity q,

$$q = \nabla^2\psi - \frac{\psi}{L_r^2} + \beta y$$

satisfies the equation

$$\frac{\partial q}{\partial t} + (\vec{v}\cdot\nabla)q = 0 \tag{1}$$

where ψ is the stream function, $\vec{v} = \hat{z}\times\nabla\psi$ is the velocity, $\hat{z}$ is the vertical unit vector, β is the Coriolis parameter, y is the north-south coordinate, and L_r is the Rossby deformation radius, a measure of the atmospheric stability(see [2]).

Fluid flows governed by these equations occur in geophysical systems. In general, a given steady-state fluid flow that satisfies these equations is not always stable. Random perturbations that exist in any real system are liable to cause a particular steady-state solution to evolve into a new steady-state solution. They are also equally likely to die away leaving the flow unaffected or simply remain in the flow without fundamentally changing it. Of particular interest to fluid dynamicists is the nature of these instabilities. It is therefore desirable to quantify these instabilities in order to study their evolution in time.

2 Formulation of an Eigenvalue Problem

An eigenvalue problem is formulated by performing a linear stability analysis about a base flow that satisfies the QG equation(1). The boundary conditions are periodic in x and periodic in y. The problem is then transformed from the continuous space of x,y to collocation space by discretizing the domain of interest.

2.1 Linear Stability Analysis

The natural first step in examining the stability of any system is to consider infinitesimal disturbances, so that all terms in the equations involving products of small perturbations may be neglected. This technique makes the analysis simpler. One should keep in mind however, that there may also be instabilities that arise only if a certain threshold disturbance amplitude is exceeded.

The deformation radius, L_r, and the Coriolis parameter, β, are functions of y alone. This and the fact that the boundary conditions are periodic in x(and y) lead us to choose an initial steady-state fluid flow that is a function of y alone and satisfies the QG equations:

$$q(y) = \bar{q}(y), \quad \vec{v}(y) = U(y)\hat{x}$$

where $\bar{q}$ and U are periodic functions of y.

Consider a perturbation $\varepsilon q'(x,y,t)$ to this base state where ε is a small number. Substituting in the QG equation (1) and retaining terms of order ε,

$$\frac{\partial q'}{\partial t} + U(y)\frac{\partial q'}{\partial x} + v'_y \frac{d\bar{q}}{dy} = 0 \tag{2}$$

Representing $q'(x,y,t)$ by a Fourier series in x,

$$q'(x,y,t) = \sum \hat{q}_k(y) exp(ik(x - c_k t)) \tag{3}$$

where k is the wavenumber, real(c_k) is the wave speed and img(c_k) is called the growth rate.

Substituting the expression for q' from (3) in the linearized QG equation (2) we obtain,

$$-c_k\hat{q}_k + U\hat{q}_k + \frac{d\bar{q}}{dy}\left(\frac{d^2}{dy^2} - k^2 - \frac{1}{L_r^2}\right)^{-1}\hat{q}_k = 0 \tag{4}$$

which is an eigenvalue problem. The eigenvalues are given by c_k and the corresponding eigenfunctions are given by $\hat{q}_k(y)$. For a particular component of the perturbation given by its wavenumber k, if

1. img(c_k) > 0, it will grow exponentially in time
2. img(c_k) < 0, it will decay exponentially in time
3. img(c_k) = 0, it will neither grow nor decay

In the linearized equation, feedback from the perturbation to the base state, as happens for the fully nonlinear equations of motion, is suppressed. This is a detriment because it restricts the model to the early stages of the instability when the perturbation is small. However, it is also a virtue. The linearized stability problem is completely specified by the specification of the base state.

2.2 Discretization of the Eigenvalue Problem

For most $\bar{q}(y)$ and U(y) that are physically interesting, it is not possible to obtain an analytic solution for the eigenvalues and eigenfunctions. A numerical approach is required. (For comparison with the numerical solution, an analytic solution which can be obtained under certain special circumstances is described in appendix A)

The domain of interest has length L_x and width L_y. It is discretized into N_x grid points along the x direction and N_y grid points along the y direction. With a collocation space thus defined, the eigenvalue problem given by equation(4)is transformed into a matrix problem,

$$[T]\underline{\hat{\Psi}}_k = [D]\underline{\hat{\Psi}}_k \tag{5}$$

$$[D] = \begin{bmatrix} \frac{\frac{d\bar{q}}{dy}(y_1)}{c_k - U_1} & 0 & \dots & 0 \\ 0 & \frac{\frac{d\bar{q}}{dy}(y_2)}{c_k - U_2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \frac{\frac{d\bar{q}}{dy}(y_{N_y})}{c_k - U_{N_y}} \end{bmatrix} \tag{6}$$

$$[T] = [\frac{d^2}{dy^2}] - (k^2 + \frac{1}{L_r^2})[I] \tag{7}$$

$$\underline{\hat{q}}_k = [T]\underline{\hat{\Psi}}_k \tag{8}$$

where $\underline{\hat{q}}_k = [\hat{q}_k(y_1), \hat{q}_k(y_2), ..., \hat{q}_k(y_{N_y})]^T$, $\underline{\hat{\Psi}}_k = [\hat{\Psi}_k(y_1), \hat{\Psi}_k(y_2), ..., \hat{\Psi}_k(y_{N_y})]^T$, $[T]$ is a called a Helmholtz operator and $[I]$ is the $N_y * N_y$ Identity Matrix.

2.3 Required Computations

In general, a subset of perturbation components, given by their wavenumbers k, is sought. The components in this subset are the most unstable ones, i.e., they have eigenvalues with the largest growth rates. These components are then added to the base state and used as the initial condition for an initial-value solver. Formally,

1. for $k = \frac{2\pi i}{L_x}$, $\quad (i = 1, 2, ..., \frac{N_x}{2})$
2. determine the set of eigenvalues $\{c_1^i, c_2^i, ..., c_j^i, ..., c_{N_y}^i\}$
3. choose an eigenvalue c_j^i from this set for which $\text{img}(c_j^i) \geq \text{img}(c_m^i)$, $\quad (m = 1, 2, ..., N_y)$
4. let the chosen c_j^i be called σ^i
5. end for

Thus a set $\{\sigma^1, \sigma^2, ..., \sigma^{\frac{N_x}{2}}\}$ is determined corresponding to each value of k. Depending on the physical requirements of the problem one or more components may be chosen. For example, choosing two components k^a, k^b, corresponding to σ^a, σ^b where σ^a, σ^b are the two greatest σ values in the set yields the following initial condition for an initial-value solver,

$$q(x, y, t = 0) = \bar{q}(y) + \varepsilon_a \hat{q}_{k^a} \exp(ik^a x) + \varepsilon_b \hat{q}_{k^b} \exp(ik^b x)$$

2.4 A Note on Spurious Eigenvalues

An infinite-dimensional operator is characterized by a spectrum comprising a countably infinite number of discrete eigenvalues and/or a continuous spectrum of eigenvalues. The spectrum of a finite-dimensional approximation to such an operator consists of a finite number of discrete eigenvalues. As the dimension of the approximation is increased, one would expect its spectrum to tend more and more towards some portion of the discrete part of the infinite-dimensional spectrum. However, at any given resolution, we would also expect the presence of spurious eigenvalues - eigenvalues that are unresolved by the resolution, eigenvalues that are attempting to mimic the continuous part of the spectrum and eigenvalues that are a figment of incorrect boundary conditions. The issue then is to be able to effectively and efficiently filter out these eigenvalues.

Eigenvalues that are unchanged on increasing the resolution can be considered well resolved eigenvalues that belong to the discrete portion of the infinite-dimensional spectrum. It is also helpful to compute the corresponding eigenvector and ensure that it is a smooth function of the collocation variable.

3 Numerical Method: Finite difference/Root finding

The numerical method described in this paper takes advantage of the sparse tridiagonal structure that is associated with the Helmholtz operator when a finite differencing scheme is used to represent the second derivative operator. A tridiagonal matrix has a number of properties that make it attractive from a computational perspective. Firstly, its sparse nature implies less storage requirement. This means that computations run at certain resolutions will benefit greatly from cache related speed up. For example, the Helmholtz operator for a 512*512 computation involving doubly precise 8 byte floating point numbers would fit entirely into the cache of today's computer. Secondly, the determinant of a tridiagonal operator can be cheaply computed with a three term recursion relationship. We will show in section 3.2 how this is an important feature. Finally, its sparse banded structure enables cheap computation of an eigenvector corresponding to a particular eigenvalue.

3.1 Representing the Helmholtz Operator

The second derivative operator in equation(7), $[\frac{d^2}{dy^2}]$, is represented using a second order accurate central differencing scheme. This results in a tridiagonal structure for the Helmholtz operator $[T]$,

$$[T] = \frac{1}{\delta y^2}\begin{bmatrix} -2-(k^2+\frac{1}{L_r^2})\delta y^2 & 1 & 0 & \dots & 0 & 1 \\ 1 & -2-(k^2+\frac{1}{L_r^2})\delta y^2 & 1 & 0 & \dots & 0 \\ 0 & \ddots & \ddots & \ddots & \dots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 & -2-(k^2+\frac{1}{L_r^2})\delta y^2 & 1 \\ 1 & 0 & \dots & 0 & 1 & -2-(k^2+\frac{1}{L_r^2})\delta y^2 \end{bmatrix}$$

where $\delta y = \frac{L_y}{N_y}$

3.2 Root Finding to Determine Eigenvalues

Once the Helmholtz operator $[T]$ has been represented by a tridiagonal matrix, from equation(5), for non-trivial vectors $\underline{\hat{\Psi}}_k$, we require that,

$$f(c_k) = det([T] - [D]) = 0$$

where $f(c_k)$ is a rational function of c_k with poles along the real axis given by $\{U_1, U_2, ..., U_{N_y}\}$

A root finding technique such as the *Secant Method*(see [5]) can be used to determine the complex roots of this rational function.

$$c_k^{n+1} = c_k^n - \frac{(c_k^n - c_k^{n-1})}{f(c_k^n) - f(c_k^{n-1})} f(c_k^n) \tag{9}$$

The Secant method requires two initial guesses per eigenvalue. Note that in order for the iterations to converge to a complex root, the initial guesses have to be complex. For the results presented in this paper, the initial guesses are chosen such that their real parts lie between two consecutive poles on the real axis and their imaginary parts are perturbed by a tiny number.

The value of the function $f(c_k^n)$ in (9) at each step of the iteration is computed with a three term recursion relationship. For a general $N*N$ tridiagonal matrix $[a_{ij}]$ without periodic boundary conditions, the determinant can be computed with a three term recursion relationship

$$\begin{aligned} det_0 &= 1 \\ det_1 &= a_{NN} \\ det_j &= a_{N-j+1} * det_{j-1} - a_{N-j+1,N-j+2} * a_{N-j+2,N-j+1} * det_{j-2} \quad (j = 2, ..., N) \end{aligned}$$

For a $N*N$ tridiagonal matrix with extra elements bc_1 in the top right hand corner and bc_2 in the bottom left hand corner, corresponding to periodic boundary conditions, the determinant is computed as follows,

$$\begin{aligned} f(c_k) &= g(c_k) + bc_1 * bc_2 * h(c_k) - bc_1 * L - bc_2 * U \\ U &= \prod a_{j,j+1} \quad (j = 1, ..., N-1) \\ L &= \prod a_{j+1,j} \quad (j = 1, ..., N-1) \end{aligned}$$

where $g(c_k)$ is the determinant of the tridiagonal matrix ignoring the extra elements, $h(c_k)$ is the determinant of the tridiagonal submatrix obtained by omitting the first, last row and the first, last column. L is the product of the elements on the subdiagonal and U is the product of the elements on the superdiagonal(See Appendix B for a derivation of this recursion relationship).

Every $f(c_k^n)$ computation using the recursive relations requires $o(N)$ operations. For convergence to a particular eigenvalue within a small tolerance of say, 10^{-16}, the Secant Method requires order 10 iterations. Since N is usually order 100 or greater, the total cost of computing an eigenvalue is $o(N)$ operations.

Every time a complex root is found, $f(c_k)$ can be deflated by the root and its complex conjugate. This forces the root finder to look for roots other than the one just computed, preventing different initial guesses from converging to the same root. This is likely to happen given the arbitrary nature of the initial guesses. Future work would involve finding an analytic foundation for the initial guesses.

If one is solely interested in complex(and hence unstable) eigenvalues, which is usually the case, a further optimization could be to discard eigenvalues which appear to be converging to eigenvalues with "small" imaginary parts after a certain number of iterations("small" being decided by the user).

3.3 Determining an Eigenvector

For a particular eigenvalue, the corresponding eigenvector can be determined by Gaussian elimination. Partial pivoting is required since the tridiagonal operator is not guaranteed to be diagonally dominant. Gaussian elimination with partial pivoting is $o(N)$ operations when coded so as to exploit the tridiagonal structure of the matrix operator. While the input matrix consists of a lower diagonal, main diagonal, upper diagonal and an extra block for periodic boundary conditions, the upper triangular matrix obtained from Gaussian elimination consists only of a main diagonal, two upper diagonals and some non-zero elements in the last column because of the periodic boundary conditions. The eigenvector $\underline{\hat{\psi}}_k$is obtained by backward solving. $\underline{\hat{q}}_k$ is obtained from $\underline{\hat{\psi}}_k$ using equation(8).

3.4 Algorithm for Required Computations

We present an algorithm that uses this numerical method for the required computations described earlier. The algorithm also filters undesirable spurious eigenvalues.

1. Choose a reference resolution, say N_1, usually the resolution used by the initial-value solver
2. For a particular wavenumber k, compute all the eigenvalues using suitable initial guesses. Discard all eigenvalues with very small or negative imaginary parts
3. Change the resolution to N_2, where $N_2 > N_1$
4. Use the previously retained eigenvalues as initial guesses. Keep converged eigenvalues that agree with the initial guesses within a user defined tolerance
5. Repeat step 4 for successively higher resolutions $N_i > ..... > N_3 > N_2 > N_1$
6. After several resolution increases, the set of chosen eigenvalues are most likely to be "real" eigenvalues that belong to the spectrum of the infinite dimensional operator
7. Compute eigenvectors corresponding to these "real" eigenvalues at the reference resolution
8. Repeat the above steps for different wavenumbers, k, to determine the most unstable components for use in an initial-value based time evolution

3.5 An Example Computation using the Algorithm

The base states $\bar{q}(y)$ and $U(y)$ to be used for this computation are shown in figure(1) and figure(2) respectively. Let $N_1 = 256$ be the reference resolution. As described in section 2.3, we are interested in eigenvalues with the largest imaginary parts that are well resolved at this resolution. In order to determine if the candidate eigenvalues are well resolved, i.e. not spurious, it is necessary to "track" them over successively increasing resolutions.

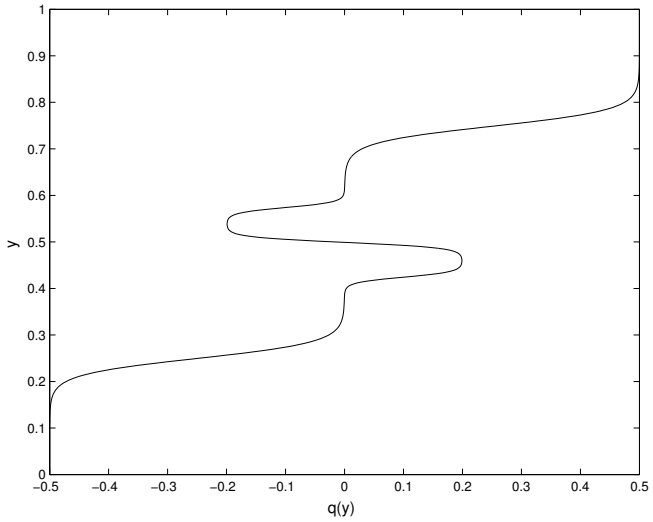

Figure 1: base state potential vorticity $\bar{q}(y)$

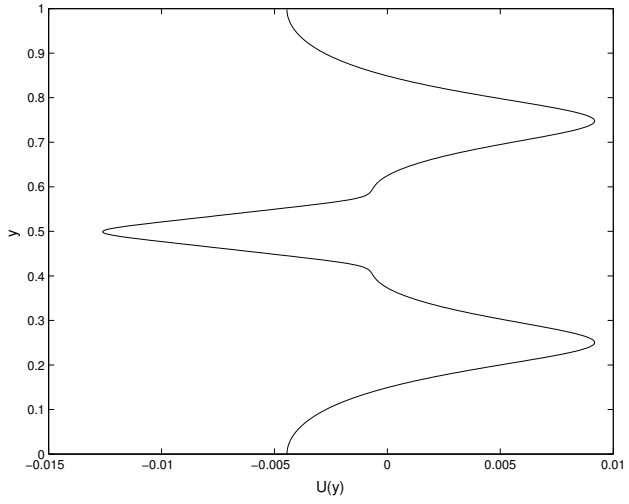

Figure 2: base state velocity $U(y)$

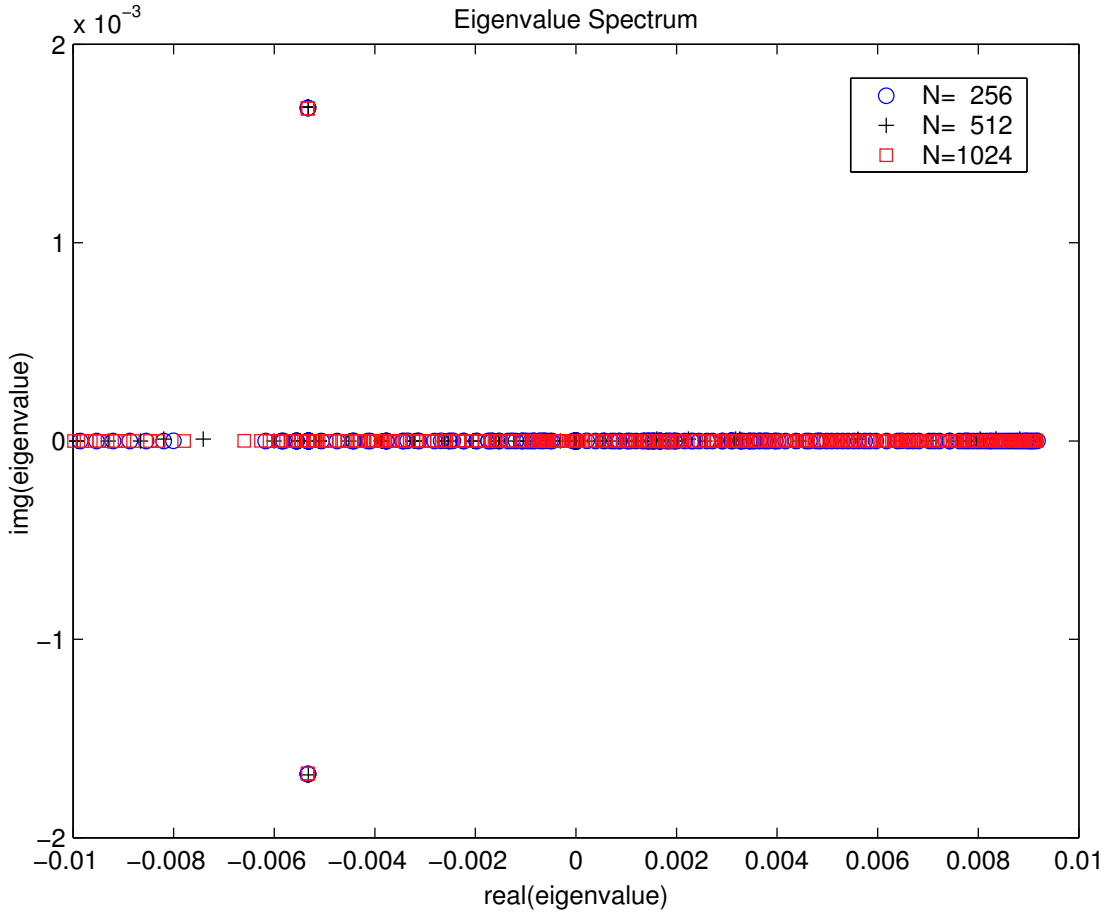

Figure 3: Eigenvalue spectrum at each resolution

To this end, let our "filtering" resolutions be $N_2 = 512$ and $N_3 = 1024$. The wavenumber $k = 6 * \frac{2\pi}{L_x}$. The spectrum for each resolution is plotted on the same figure(3). This gives us an an eye-ball estimate for the variation in eigenvalues as the resolution is increased. For this particular computation, the spectrum is characterized by a stretch of eigenvalues on the real axis and a single discrete complex eigenvalue. Moreover, this discrete eigenvalue changes by less than 0.001% with increasing resolution(see table below). The eigenvector corresponding to this eigenvalue at each resolution is shown in figures (4), (5) and (6). The eigenvectors are similar and smooth(because the Fourier coefficients fall off rapidly with increasing wavenumber). We therefore conclude that the eigenvalue is "real" and that it belongs to the spectrum of the infinite-dimensional operator.

N	Eigenvalue
256	-0.00532276+i0.00167384
512	-0.00533049+i0.00167885
1024	-0.00533358+i0.00167656

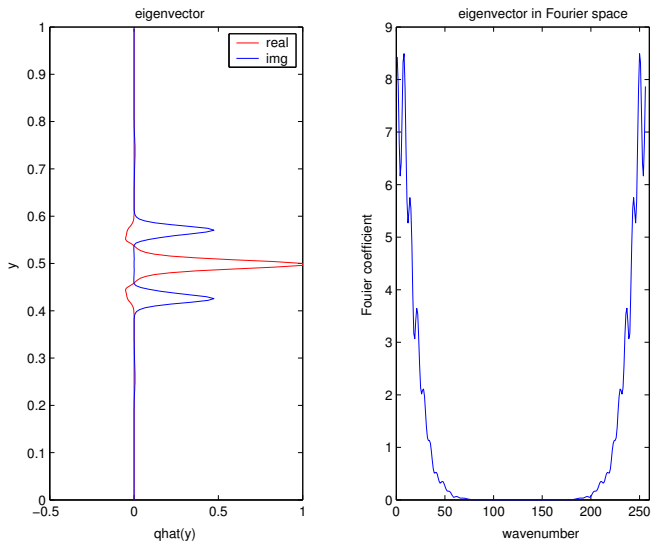

Figure 4: Eigenvector $\hat{\underline{q}}_k(y)$ at $N_1 = 256$

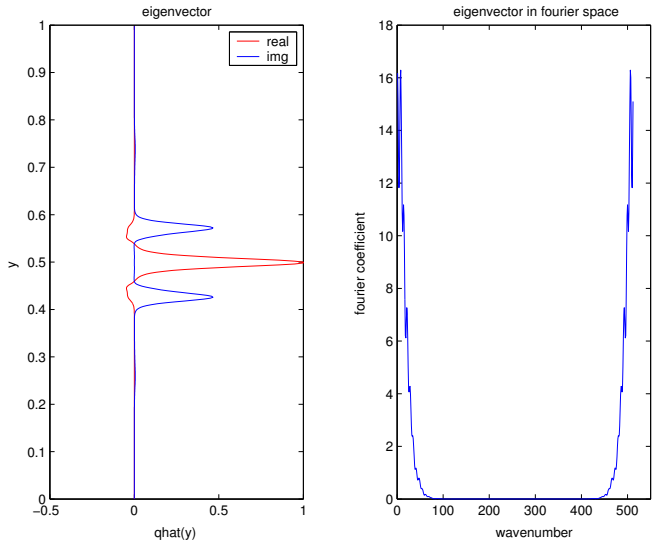

Figure 5: Eigenvector $\hat{\underline{q}}_k(y)$ at $N_2 = 512$

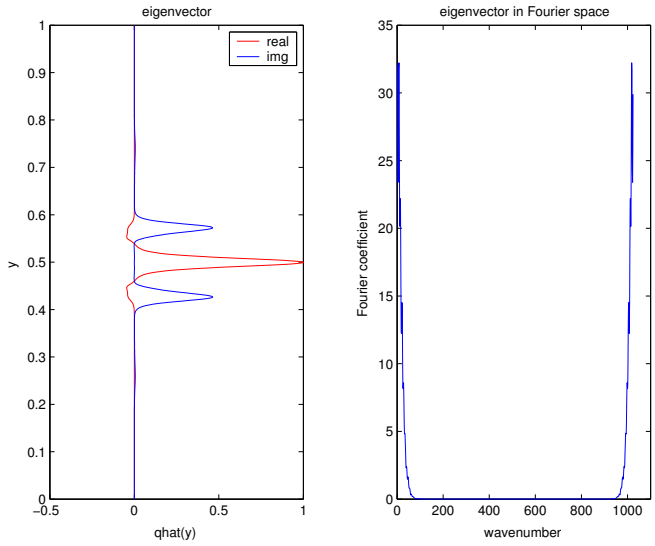

Figure 6: Eigenvector $\underline{\hat{q}}_k(y)$ at $N_3 = 1024$

4 Alternative Numerical Method: Spectral method/QR iteration

A method typically used for solving stability problems of this nature uses a spectral representation for the Helmholtz operator $[T]$(see [3]). The resulting operator is a full matrix and all the eigenvalues and the corresponding eigenvectors are determined by the QR iteration. A spectral representation would accurately resolve a greater number of modes than a finite difference representation, for the same number of grid points. However, the QR iteration is relatively expensive requiring $o(N^3)$ operations per iteration(see [1], [4]). This coupled with the fact that it would have to be run at increasing resolutions in order to identify spurious eigenvalues implies that it would be well worth ones while to investigate the possibility of using an alternative approach such as the one described in this paper.

4.1 Comparison of the Finite difference/Root finding method with the Spectral/QR method

For the particular base state used to obtain the results in this paper, the infinite-dimensional spectrum appears to be characterized by a continuous region along the real axis and a single, discrete, complex eigenvalue. The finite-dimensional approximation did just as well as the spectral representation in resolving this isolated eigenvalue for the same number of grid points. It should be noted that such spectra are by no means exclusive to this problem. There are a number of fluid flow instabilities that are characterized by a similar spectrum.

If the required computations are performed at resolutions $N_1 < N_2 < N_3$, the number of computations required by the algorithm described in section 3.4 for a finite difference/root finding method, is approximately $o(N_1^2) + o(N_2) + o(N_3)$ per wavenumber k. An equivalent algorithm for a spectral method /QR iteration based numerical method would require approximately $o(N_1^3) + o(N_2^3) + o(N_3^3)$ operations per wavenumber k. This is without taking into account the cache related speed increase that would characterize a finite difference/root finding method for most resolutions of interest.

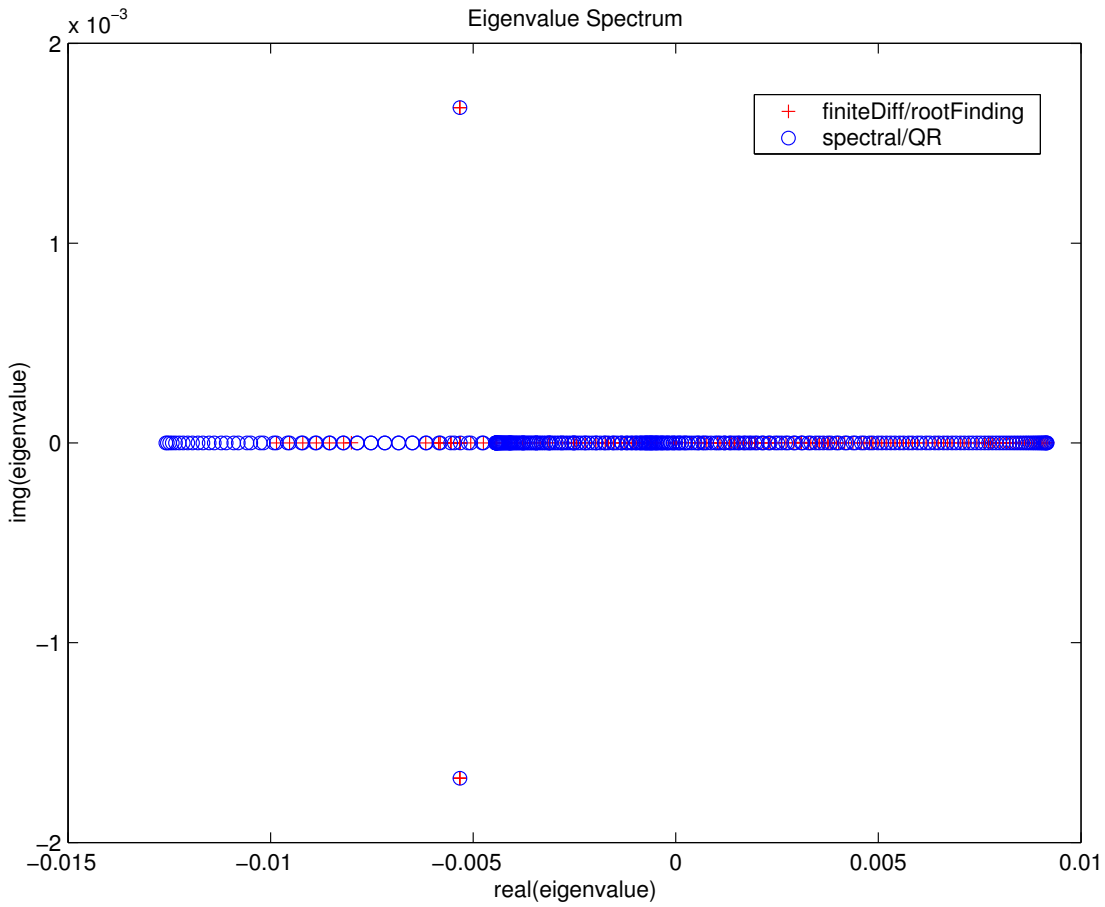

Figure 7: The eigenvalue spectrum of each method, $N_1 = 512$

A computation was carried out using each numerical method for the base state described in section 3.5, with reference resolution $N_1 = 512$ and filtering resolutions $N_2 = 1024$ and $N_3 = 2048$. The table below indicates the time[2] taken by each method for a single value of $k = 6 * \frac{2\pi}{L_x}$ as per the algorithm described in section 3.4.

Res	Spectral Time (secs)	Finite Diff Time (secs)
$N_1 = 512$	45	30
$N_2 = 1024$	300	0.1
$N_3 = 2048$	2000	0.2
Total time	2345	30.3

The spectrum obtained by the spectral/QR method and the spectrum obtained by the finite difference/root finding method, for the reference resolution, are shown in figure(7). The corresponding eigenvectors for each method are shown in figure(8) and figure(9).

[2]calculated using MATLAB's *cputime* function

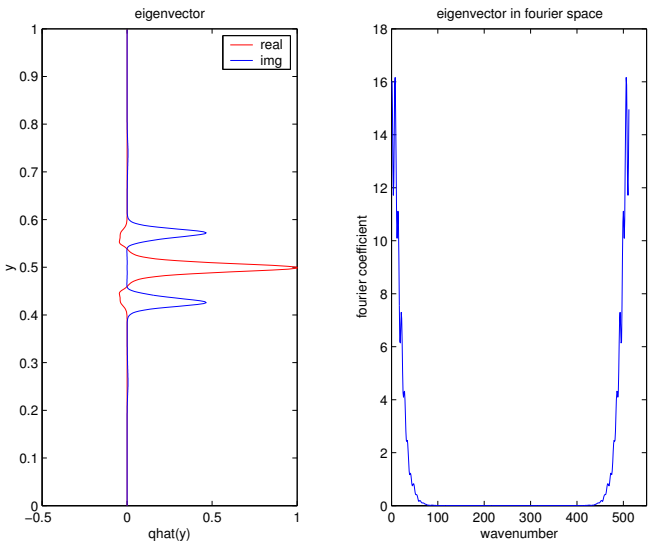

Figure 8: spectral/QR eigenvector $\underline{\hat{q}}_k(y)$ at $N_1 = 512$

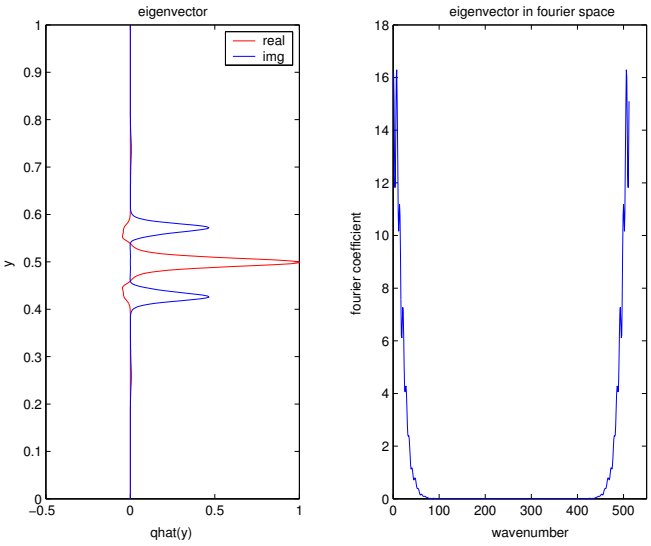

Figure 9: finite difference/root finding eigenvector $\underline{\hat{q}}_k(y)$ at $N_1 = 512$

5 Future Work

There are several unresolved problems with the proposed numerical method. The dependence of the root finder on initial guesses implies that it is possible for some eigenvalues that are resolved by the finite difference representation to go undetected. For the example considered however, the initial guesses proved sufficient to recover all the discrete eigenvalues that a spectral representation recovered at the same resolution. Future work would involve an analytic backing for the initial guesses. The convergence properties of the root finding iteration can be improved by using a higher order Hyperbolic method.

Note that the algorithm described in section 3.4 does not consider "real" eigenvalues that may appear at resolution N_2 but were absent in the reference resolution. The algorithm can easily be modified to take this into account. For the time being, we are only interested in eigenvalues that are well resolved at the reference resolution, so that a corresponding eigenvector can be used in an inital value solver running at this resolution.

Higher order accuracy can be obtained by representing the Helmholtz operator using a 4th order accurate *Numerov* scheme. This scheme would preserve the tridiagonal structure of the operator. Convergence properties of the finite-dimensional spectrum to the infinite-dimensional spectrum can be studied for this class of stability problems. An error analysis can be performed to determine the effect of round-off for this numerical method.

References

[1] W. Kahan, *Advanced Matrix Computation*, MATH 221 Class Notes, University of California-Berkeley, Spring 2003.

[2] P. S. Marcus, *Geophysical Fluid Dynamics*, ME 267 Class Notes, University of California-Berkeley, Fall 2002.

[3] P. S. Marcus, *Spectral Methods for Fluid Dynamics*, E 266B Class Notes, University of California-Berkeley, Fall 2002.

[4] G. Golub and C. F. Van Loan, *Matrix Computation*, 1st ed., John Hopkins University Press, pp. 228-238, 1983.

[5] W. H. Press, B. P. Flannery, S. A. Teukolsky and W. T. Vetterling, *Secant Method, False Position Method, and Ridders' Method. 9.2, Numerical Recipes in FORTRAN: The Art of Scientific Computing*, 2nd ed. Cambridge, England: Cambridge University Press, pp. 347-352, 1992