Multigrid Solution of an Elliptic Boundary Value Problem from Tropical Cyclone Theory

PAUL E. CIESIELSKI, SCOTT R. FULTON AND WAYNE H. SCHUBERT

Department of Atmospheric Science, Colorado State University, Fort Collins, CO 80523

(Manuscript received 11 July 1985, in final form 22 August 1985)

ABSTRACT

We consider the multigrid solution of the transverse circulation equation for a tropical cyclone. First we develop a standard multigrid scheme (SMG) which cycles between different levels of discretization (grids) to efficiently reduce the error in the solution on all scales. Whereas relaxation is inefficient as a solution method, it is used within the multigrid approach as a smoother to reduce the high-wavenumber errors on each grid. The added cost of the coarse grids is small because they contain relatively few points. The efficiency of the SMG scheme is compared to more conventional methods: Gauss-Seidel and successive-over-relaxation (SOR). Results show that the SMG scheme solves to the level of truncation error 26 times faster than an optimal SOR method.

In the SMG scheme an unbounded domain is approximated with a wall at a finite radius which leads to significant errors in the numerical solution. To better simulate an unbounded domain, we develop a second scheme (LRMG) which naturally combines local mesh refinement with multigrid processing. In this scheme, the lateral boundary is moved far enough out that the wall boundary condition is realistic, and the grid is coarsened in the outer region so that little additional work is required. Since finer grids are introduced only where needed, the LRMG scheme maintains the usual multigrid efficiency.

1. Introduction

In a review (this issue), Fulton et al. (1986) have discussed the basic concepts and techniques of multigrid methods, concentrating on their role as fast elliptic solvers. In the present paper we apply the multigrid method to a particular problem in tropical cyclone theory. Our example originates from the balanced vortex equations of Eliassen (1952). Eliassen's equations have played an important role in attempts (e.g., Ooyama, 1969; Sundqvist, 1970) to numerically simulate the life cycle of tropical cyclones. The form of the balanced vortex equations which we shall use is that of Schubert and Hack (1983). This form is based in part on a transformation of the radial coordinate from actual radius r to potential radius R, where $\frac{1}{2} fR^2 = rv + \frac{1}{2} fr^2$, v is the tangential wind and f is the coriolis parameter. The transformed model can be used either prognostically to study the evolution of a vortex, or diagnostically, as in this study, to compute the response of a vortex to a specified forcing. In the former approach the transverse circulation equation

$$\left.\begin{aligned} &\frac{\partial}{\partial R}\left(q\frac{\partial R\psi^*}{R\partial R}\right)+\frac{\partial}{\partial Z}\left(s\frac{\partial \psi^*}{\partial Z}\right)=\frac{g}{\theta_0}\frac{\partial Q}{\partial R} \\ &R\psi^*=0 \text{ at } R=0, \quad R\psi^*\to 0 \text{ as } R\to\infty \\ &R\psi^*=0 \text{ at } Z=\hat{Z}, \quad R\psi^* \text{ specified on } Z=0 \end{aligned}\right\}, \tag{1.1}$$

is solved for the streamfunction ψ^* at every time step, where the potential vorticity q and the inertial stability s are functions of the vortex, and the diabatic heating Q is either specified or parameterized in terms of other model variables. In the diagnostic approach q, s and Q are specified functions of R and Z. The effects of frictional stresses, which we assume are confined to a thin boundary layer, are incorporated in this model through the bottom boundary condition; in the inviscid case this condition becomes $R\psi^* = 0$.

A variety of methods have been used in the past for solving finite difference analogues to elliptic equations such as (1.1). Perhaps the most straightforward solution method for such problems is Gaussian elimination. Although this direct method produces the exact solution in a finite number of computations, it becomes prohibitively expensive in terms of operations and storage as resolution increases. Ooyama's (1969) use of Gaussian elimination to solve an analogue of (1.1) at each time step in the evolution of a tropical cyclone was computationally feasible because his model contained only two vertical levels. In solving for the forced radial and vertical motions in a hurricane model, Willoughby (1979) likewise chose to use a direct method since his system was solved only once in a diagnostic sense. On the other hand, Sundqvist (1970) was forced to use iterative methods to make his model simulation of tropical cyclone development with high vertical resolution computationally efficient. Employing both the successive-over-relaxation (SOR) and the alternating-direction-implicit (ADI) procedures, Sundqvist found the latter indirect method increasingly more efficient than SOR as the resolution of the model increased.

Fulton et al. (1986) describe a still faster procedure for solving elliptic partial differential equations such as (1.1). This procedure, referred to as the multigrid method, cycles between different levels of discretization (grids) to efficiently reduce the error in the solution on all scales. Relaxation is employed as a smoother to reduce the high-wavenumber errors on each grid, and the added cost of the additional grids is small because they contain relatively few points.

The purpose of this paper is to show how the multigrid approach can be used to solve (1.1) efficiently. In section 2 we present the discretized form of (1.1) on which our subsequent numerical methods are based and formulate two standard methods for its solution. In section 3 the basic multigrid concepts are applied to construct a standard multigrid scheme for solving (1.1). In section 4 we compare the efficiencies of three numerical schemes (Gauss-Seidel, successive-over-relaxation and multigrid) for a particular test case. For this comparison, the outer boundary condition is approximated by a wall at finite radius which leads to significant differences between the numerical and analytical solutions. In section 5 we show how this problem can be overcome in an efficient and natural way using the multigrid method in conjunction with local mesh refinement. Our conclusions are summarized in section 6.

2. Standard solution methods

To compare the efficiency of a multigrid method against more conventional methods, we now formulate two standard solution methods for (1.1) on the domain $\Omega = [0, \hat{R}] \times [0, \hat{Z}]$. Introducing the grid

$$\Omega_{\Delta R,\Delta Z} = \{(R_j, Z_k) = (j\Delta R, k\Delta Z) \quad 0 \leqslant j \leqslant J, 0 \leqslant k \leqslant K\},$$

where $\Delta R = \hat{R}/J$ and $\Delta Z = \hat{Z}/K$, we let $\Psi_{j,k}$ denote the discrete approximation to $R\psi^*$ at the gridpoint (R_j, Z_k). The numerical methods in this paper are all based on the usual second-order centered difference approximation

$$a_{j-1/2,k}\Psi_{j-1,k} - (a_{j-1/2,k} + a_{j+1/2,k})\Psi_{j,k} + a_{j+1/2,k}\Psi_{j+1,k} + b_{j,k-1/2}\Psi_{j,k-1} - (b_{j,k-1/2} + b_{j,k+1/2})\Psi_{j,k} + b_{j,k+1/2}\Psi_{j,k+1} = f_{j,k} \quad (2.1a)$$

to (1.1), where

$$\left.\begin{aligned} a_{j+1/2,k} &= \frac{q_{j+1/2,k}}{R_{j+1/2}(\Delta R)^2} \\ b_{j,k+1/2} &= \frac{s_{j,k+1/2}}{R_j(\Delta Z)^2} \\ f_{j,k} &= \frac{g}{\theta_0 \Delta R}(Q_{j+1/2,k} - Q_{j-1/2,k}) \end{aligned}\right\}. \quad (2.1b)$$

Equation (2.1a) holds on the interior gridpoints of Ω, that is, $(0 < j < J, 0 < k < K)$. The corresponding boundary conditions are

$$\left.\begin{aligned} \Psi_{j,k} &= 0, && j = 0 \text{ and } j = J, 0 \leqslant k \leqslant K \\ \Psi_{j,k} &= 0, && 0 \leqslant j \leqslant J, k = K \\ \Psi_{j,k} &= \text{specified}, && 0 \leqslant j \leqslant J, k = 0 \end{aligned}\right\}. \quad (2.2)$$

The form (2.1) preserves the energy constraint of (1.1), that is, the net outward flux of energy at the boundary of the domain equals the net source of energy (integral of the right-hand side over Ω). In addition, the matrix associated with (2.1)–(2.2) is symmetric.

The exact solution Ψ of the discrete equations (2.1)–(2.2) may be obtained in a finite number of operations by Gaussian elimination. However, the operation count $O(JK^3)$ for factorization and $O(JK^2)$ for solution and storage required $[O(JK^2)]$ may be prohibitively expensive. Thus, one usually turns to indirect (iterative) methods, which generate a sequence of approximations $\tilde{\Psi}$ that converge to Ψ. These methods generally require only $O(JK)$ storage locations and have the added advantage of being self-correcting, so that any round-off errors generated are reduced automatically in the iteration process.

A simple iterative method is Gauss-Seidel (GS) or successive relaxation. In this method the values $\tilde{\Psi}_{j,k}$ of an approximation $\tilde{\Psi}$ are modified point by point in lexicographic order to obtain a new approximation $\tilde{\Psi}^{\text{new}}$. The modification introduced at a point is such that the discrete equation is satisfied exactly at that point; thus $\tilde{\Psi}^{\text{new}}$ is defined by

$$a_{j-1/2,k}\tilde{\Psi}^{\text{new}}_{j-1,k} - (a_{j-1/2,k} + a_{j+1/2,k})\tilde{\Psi}^{\text{new}}_{j,k} + a_{j+1/2,k}\tilde{\Psi}_{j+1,k} + b_{j,k-1/2}\tilde{\Psi}^{\text{new}}_{j,k-1} - (b_{j,k-1/2} + b_{j,k+1/2})\tilde{\Psi}^{\text{new}}_{j,k} + b_{j,k+1/2}\tilde{\Psi}_{j,k+1} = f_{j,k}, \quad (2.3)$$

where the new approximation $\tilde{\Psi}^{\text{new}}$ is used at the points $(j-1, k)$ and $(j, k-1)$ due to the lexicographic ordering. It is convenient to write (2.3) as

$$\tilde{\Psi}^{\text{new}}_{j,k} = \tilde{\Psi}_{j,k} - \frac{r_{j,k}}{d_{j,k}} \quad (2.4)$$

where, from this point on, r denotes the residual

$$r_{j,k} = f_{j,k} - [a_{j-1/2,k}\tilde{\Psi}^{\text{new}}_{j-1,k} - (a_{j-1/2,k} + a_{j+1/2,k})\tilde{\Psi}_{j,k} + a_{j+1/2,k}\tilde{\Psi}_{j+1,k} + b_{j,k-1/2}\tilde{\Psi}^{\text{new}}_{j,k-1} - (b_{j,k-1/2} + b_{j,k+1/2})\tilde{\Psi}_{j,k} + b_{j,k+1/2}\tilde{\Psi}_{j,k+1}] \quad (2.5)$$

and $d_{j,k} = a_{j+1/2,k} + a_{j-1/2,k} + b_{j,k+1/2} + b_{j,k-1/2}$. In this form the convergence rate may be improved by scaling the correction $r_{j,k}/d_{j,k}$ in (2.4) by a relaxation parameter $(1 + \omega)$. This leads to the successive-over-relaxation (SOR) method defined by (2.5) and

$$\tilde{\Psi}_{j,k}^{\text{new}} = \tilde{\Psi}_{j,k} - (1 + \omega)\frac{r_{j,k}}{d_{j,k}}, \qquad (2.6)$$

which reduces to the Gauss-Seidel method when $\omega = 0$.

3. A standard multigrid scheme

In this section we present the details of a standard multigrid scheme, subsequently referred to as SMG, for solving (2.1)–(2.2). We define the coextensive grids

$$\Omega_l = \left\{(R_j^{(l)}, Z_k^{(l)}): \begin{matrix} R_j^{(l)} = j\Delta R_l, j = 0, \cdots, J_l \\ Z_k^{(l)} = k\Delta Z_l, k = 0, \cdots, K_l \end{matrix}\right\}, \quad l = 1, \cdots, M \qquad (3.1)$$

where $\Delta R_l = \hat{R}/J_l$ and $\Delta Z_l = \hat{Z}/K_l$. It is convenient to use the standard mesh ratio ½, i.e., $\Delta R_l = \frac{1}{2}\Delta R_{l-1}$ and $\Delta Z_l = \frac{1}{2}\Delta Z_{l-1}$; then $\Omega_1 \subset \Omega_2 \subset \ldots \Omega_M$. We will often refer to these grids as levels, with the highest level denoting the finest grid Ω_M and the lowest level denoting the coarsest grid Ω_1.

Associated with each grid is a discrete problem

$$L^l\Psi^l = f^l, \qquad (3.2)$$

where L^l is the difference operator representing the left-hand side of (2.1a) for mesh spacing ΔR_l and ΔZ_l. The difference operators on the coarse grids are computed using (2.1b) directly on each level. (In problems where the solution exhibits strong discontinuities, the coarse-grid operators should be constructed from the fine-grid operator as discussed by Alcouffe et al. (1981)). In the multigrid approach, (3.2) is relaxed on a fine grid until the error in the approximate solution $\tilde{\Psi}$ is smooth; the problem is then transferred to a coarser grid where the residual problem (Brandt, 1984) is solved to correct the fine-grid approximation.

In the remainder of this section we describe the following key elements of the multigrid method on the grid structure described above: the relaxation scheme used to smooth the error on a given level, the grid transfers used to carry information from one level to another, and the control algorithm which determines when to switch from one level to another.

a. Relaxation scheme for interior points

In a multigrid scheme the role of relaxation is to smooth the error, i.e., to reduce its high-wavenumber components. The proper choice of a smoother is critical to the success of a multigrid scheme. Fortunately, the performance of a smoother can be predicted by local mode analysis (Brandt, 1982) before it is incorporated in a multigrid context, enabling us to optimize our procedures and debug our code. In this analysis the error in the solution is expanded in a Fourier series allowing us to compute the factor μ by which the error in each Fourier component is reduced. Since low wavenumbers on a given grid are represented as higher wavenumbers on a coarser grid, μ need only be small for high wavenumbers. Thus we are interested in obtaining a suitable smoothing factor $\bar{\mu}$, where this factor is the largest μ over the high wavenumbers (i.e., the modes not visible on coarser grids).

When $a = b$ in (2.1a), local mode analysis gives the smoothing factor $\bar{\mu} = \frac{1}{2}$ for Gauss-Seidel relaxation with lexicographic ordering, which is quite satisfactory (Brandt, 1977). However in the anisotropic case (i.e., where a and b differ), the smoothing rate for Gauss-Seidel degrades significantly. For example, in the case of a strong vortex analyzed by Schubert and Hack (1983), the ratio of b/a in (2.1a) varied by a factor of ~ 100, resulting in $\bar{\mu} = 0.98$ for Gauss-Seidel relaxation. In such cases, that is where b/a is quite large or quite small, line relaxation provides efficient smoothing. For example, when $b/a \gg 1$ in (2.1a) Z-line relaxation is appropriate; in Z-line relaxation the values along a line of points in the Z-direction (constant R_j) are replaced by new values which simultaneously satisfy all the equations on that line. Thus, Z-line relaxation is defined by

$$b_{j,k-1/2}\tilde{\Psi}_{j,k-1}^{\text{new}} - d_{j,k}\tilde{\Psi}_{j,k}^{\text{new}} + b_{j,k+1/2}\tilde{\Psi}_{j,k+1}^{\text{new}} = f_{j,k} - [a_{j-1/2,k}\tilde{\Psi}_{j-1,k}^{\text{new}} + a_{j+1/2,k}\tilde{\Psi}_{j+1,k}] \qquad (3.3)$$

where we assume the lines to be taken in lexicographic order and the notation is the same as in (2.3). Each line to be solved for results in a tridiagonal and diagonally dominant system of equations, which is easy and efficient to solve. Z-line relaxation results in $\bar{\mu} = 0.45$ for any $b/a \geq 1$. However, in regions where $b/a \to 0$ the smoothing rate degrades significantly, and R-line relaxation, defined analogously but for constant Z_k, is more suitable. Following one sweep of R-line relaxation with a sweep of Z-line relaxation results in alternating direction line relaxation (ADLR), which has the smoothing rate $\bar{\mu} = 0.45$ for *any* b/a. Such a scheme is needed for our purposes due to the spatial variation of b/a in a vortex of hurricane intensity as mentioned earlier.

Although line relaxation as presented above is not directly vectorizable, by changing the order in which lines are relaxed one obtains a vectorizable algorithm known as Zebra relaxation. One sweep of Zebra relaxation consists of the following two stages: In the first stage the tridiagonal systems for all white (even) lines are solved simultaneously with new values of the unknown replacing the old ones; then using these updated values, the same procedure is repeated for all the black (odd) lines. Each stage involves solving a set of tridiagonal systems; since they are decoupled they can be solved in parallel. For problems, such as ours, where b/a is large in some regions and small in others, alternating direction Zebra (ADZ) is appropriate. In the SMG scheme, one sweep of ADZ relaxation is implemented as the following four steps:

1) R-line relaxation for all even k simultaneously,
2) R-line relaxation for all odd k simultaneously,
3) Z-line relaxation for all even j simultaneously,
4) Z-line relaxation for all odd j simultaneously.

In contrast to schemes with standard ordering, the local mode analysis for ADZ relaxation is more complicated in that several modes are coupled by this scheme (Stüben and Trottenberg, 1982). We consider here two theoretical estimates of the performance of ADZ relaxation for the discrete problem (2.1a) with constant (frozen) coefficients. These estimates were obtained as described in Brandt (1982). The first and simplest measure of multigrid efficiency is the smoothing factor $\bar{\mu}$, which is the error reduction due solely to the relaxation scheme; the second predictor, which considers the effects of grid transfers as well, is the two-level convergence factor. In Fig. 1 these predictors of efficiency show that ADZ relaxation is robust (relatively insensitive to changes in the coefficients a and b) and efficient ($\bar{\mu} < 0.23$ for *any* ratio b/a). The third curve in this figure represents asymptotic convergence factors obtained numerically by solving (2.1)–(2.2) with the SMG scheme. The details of the numerical results are discussed in section 4.

b. Grid transfers

Multigrid methods require frequent switching between grid levels to maintain the efficiency of relaxation. This switching is accomplished by the fine-to-coarse transfer of the residual $r^l = f^l - L^l\tilde{\Psi}^l$ and the coarse-to-fine interpolation of the coarse-grid correction. For the SMG scheme we use bilinear interpolation for the coarse-to-fine transfer.

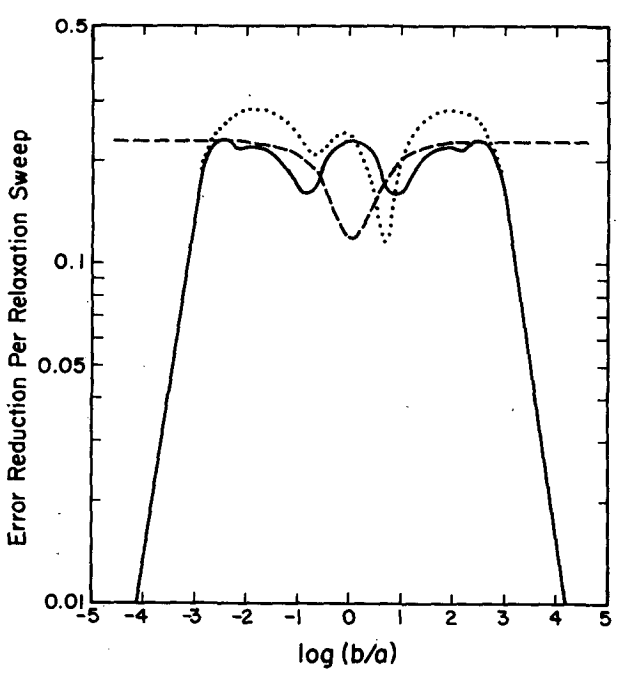

FIG. 1. Theoretical smoothing factor (dashed), and two-level convergence factor (solid) plotted as a function of the base 10 logarithm of b/a. The dotted curve is the observed asymptotic convergence factor of the SMG scheme using V-cycles.

After one sweep of ADZ relaxation the resulting residual field contains a large high-frequency component, since the residual is zero on every other line due to the order of relaxation. Therefore, to transfer a representative value of the residual to the coarse grid, one must select a transfer operator which uses all available information in the fine grid residual field. Full residual weighting (Brandt, 1984) accomplishes this task. Since we know a priori that residuals on lines for j odd are zero after an ADZ relaxation sweep, full residual weighting in this case at the point (j, k) reduces to

$$r^l_{j,k} = \frac{1}{4} r^{l+1}_{2j,2k} + \frac{1}{8}(r^{l+1}_{2j,2k-1} + r^{l+1}_{2j,2k+1}). \qquad (3.4)$$

c. Control algorithm

The final element needed to carry out multigrid processing is a cycling algorithm which controls when to switch from one level to another. For the SMG scheme we have chosen to use an accommodative algorithm referred to as CYCLE C (Brandt, 1977). In contrast to fixed algorithms, which switch levels after a preassigned number of relaxation sweeps (e.g., V-cycle or W-cycle), the accommodative CYCLE C algorithm switches grids using internal checks based on the relative magnitude of the residuals. The CYCLE C algorithm was used in a correction scheme mode, in which a coarse grid variable stores a correction to the finer grid approximation. With Dirichlet boundary conditions, such as (2.2), relaxation is performed only on the interior points; in the correction scheme this means using specified values on the boundaries of the finest grid and zeros along the boundaries of the coarser grids where no correction is needed.

4. Test case and numerical results

In this section we apply three numerical methods (Gauss-Seidel, successive-over-relaxation, and the multigrid scheme SMG) to obtain the solution of (2.1)–(2.2) for a particular test case, and examine the results and efficiencies in detail.

a. Test case

We consider a test case on the domain $\Omega = [0, \hat{R}] \times [0, \hat{Z}]$, where $\hat{R} = 960$ km and $\hat{Z} = 16$ km. In order to have an analytic solution to (1.1) against which we can compare our numerical results, we consider a physical situation in which an inviscid, resting atmosphere is heated. Under such conditions the potential radius (R) equals the actual radius and q and s in (1.1) are constants given by 2.3×10^{-9} m^3 s^{-2} kg^{-1} and 1.3

$\times 10^{-4}$ m^3 s^{-2} kg^{-1}, respectively. For this case the analytical solution to (1.1) is

$$\psi^*(R, Z) = \psi_0^* \alpha R e^{-\alpha R} \sin\left(\frac{\pi Z}{\hat{Z}}\right), \qquad (4.1)$$

where $\psi_0^* = \dfrac{gQ_0}{\alpha q \theta_0}$, $\theta_0 = 300$°K and $g = 9.81$ m s^{-2}. Here Q_0 (1.5×10^{-10} K s^{-1}) and α (1.6×10^{-6} m^{-1}) are parameters which specify the magnitude and e-folding distance of the solution (ψ^*), respectively. From the field $R\psi^*$ (Fig. 2) one can define the components of the transverse circulation as follows

$$(u^*, w^*) = \left\{-\frac{1}{\rho}\frac{\partial \psi^*}{\partial Z}, \frac{1}{\rho}\frac{\partial R\psi^*}{R\partial R}\right\}, \qquad (4.2)$$

where for simplicity we assume the pseudo-density is a constant given by $\rho = 1.1$ kg m^{-3}. The analytical u^* and w^* fields associated with (4.1) are shown in Figs. 3 and 4, respectively. Using (4.1) in (1.1) we obtain the forcing

$$Q(R, Z) = Q_0 e^{-\alpha R}[\alpha^{-1}(\lambda^2 - \alpha^2)(\alpha^{-1} + R) + 3] \sin\left(\frac{\pi Z}{\hat{Z}}\right), \qquad (4.3)$$

where

$$\lambda = \left(\frac{s}{q}\right)^{1/2} \frac{\pi}{\hat{Z}}$$

is the inverse of the generalized Rossby radius. Knowing Q and w^*, we can compute the local warming ($\partial\theta/\partial T$) from the thermodynamic equation

$$\frac{\partial \theta}{\partial T} = Q - \frac{\theta_0}{g}\rho q w^* . \qquad (4.4)$$

This equation shows that the local warming is the difference between the specified heating and the adiabatic cooling associated with the forced secondary circulation; the local warming field is of primary interest to us since its spatial distribution is critical to intensity changes in the vortex. The warming field derived analytically by using (4.2) and (4.3) in (4.4) is shown in Fig. 5. From thermal wind balance the distribution in this warming field results in cyclonic wind tendencies at lower levels and anticyclonic tendencies above the level of maximum warming. Further details on how the warming fields and corresponding wind tendencies change as the vortex evolves are given in Schubert and Hack (1982).

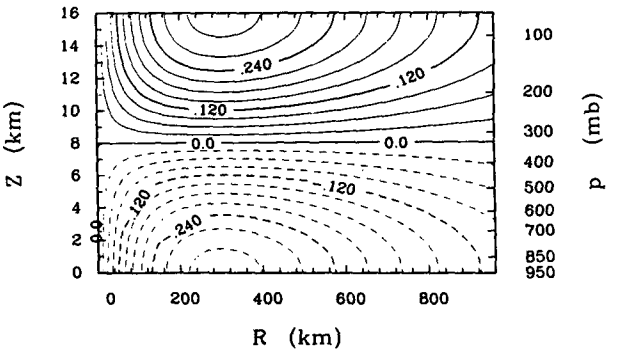

FIG. 3. Analytical u^* (m s^{-1}) with contour interval 0.03 m s^{-1} and dashed lines for negative values.

b. Results and efficiencies of numerical schemes

We consider the numerical solution of (2.1)–(2.2) over the discretized version of Ω, where $\Delta R = 10$ km and $\Delta Z = 250$ m resulting in a grid of $J = 97$ by $K = 65$ points. For the SOR results presented in this section we used an optimal relaxation parameter $\omega = 0.93$. This value was determined by trial and error; in a prognostic model optimizing ω repeatedly as the coefficients a and b change would be expensive.

Figure 6 shows the scaled discrete l_2 norm

$$\|r\| = [\sum_{j,k} (r_{j,k})^2 \Delta R \Delta Z]^{1/2} \qquad (4.5)$$

of the residual r as a function of execution time for the three numerical schemes. Included in this execution time is the computation of the residuals which were

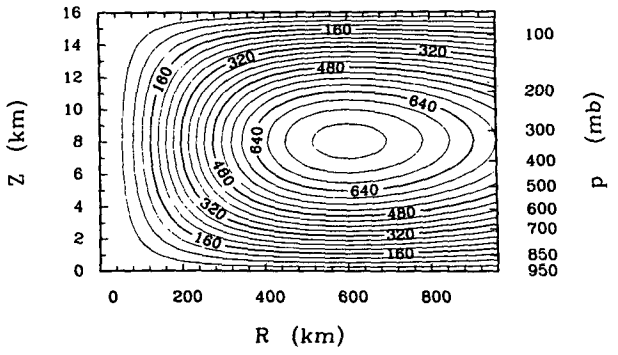

FIG. 2. Analytical streamfunction field $R\psi^*$ (10^6 kg s^{-1}) with contour interval 40×10^6 kg s^{-1}.

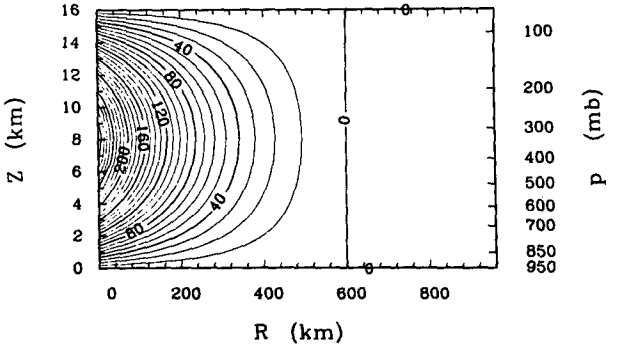

FIG. 4. Analytical w^* (10^{-4} m s^{-1}) with contour interval 10^{-4} m s^{-1}.

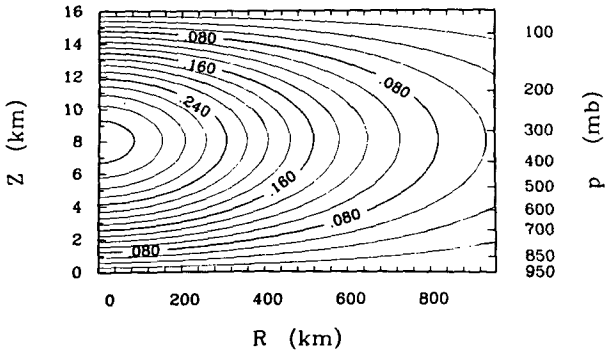

FIG. 5. Analytical warming field $\partial\theta/\partial T$ (K d^{-1}) with contour interval 0.02 K d^{-1}.

obtained dynamically (i.e., during the course of the relaxation sweep). In this figure the values plotted along the SOR and GS curves represent the number of relaxation sweeps; the values along the SMG curve are work units where a work unit is defined as the computational work of one ADZ sweep on the finest grid. The truncation error τ, defined as the residual which results from inserting the true continuous solution $R\psi^*$ in the discrete equations (2.1)–(2.2), is $\sim 10^{-8}$ m s^{-3} for the test case. To solve the discrete problem to the level of truncation error, SMG required ~4 work units or 0.033 seconds of execution time on a CRAY1 computer. In contrast the SOR scheme took 100 relaxation sweeps (0.85 seconds) to achieve this same level of accuracy, and the GS scheme required 2200 relaxation sweeps (19 seconds). Thus, in solving the test case to the level of truncation error the SMG scheme was 26 times faster than SOR and 575 times faster than GS. The curve associated with SMG becomes horizontal after 19.9 work units indicating that the solution has reached machine accuracy.

If we were solely interested in the solution Ψ, solving (2.1)–(2.2) to its truncation error τ would be adequate. However the warming field is often a small difference between two large terms in (4.4): Q and a term involving w^* which is a derivative of the Ψ field. Errors in either of these terms are magnified in the warming field. Thus to maintain a reasonable hope for computing the warming field to the level of accuracy of its truncation error, the Ψ field, in practice, is solved to several orders of magnitude below τ and w^* is computed with fourth-order finite differences.

A further comparison of the efficiencies of the SMG and SOR schemes is presented in Table 1, where the execution time per grid point is listed for several cases with different grid resolutions. These results were computed by solving (2.1)–(2.2) for the test case to a tolerance of 10^{-8} in the l_2 norm of the residual. The ex-

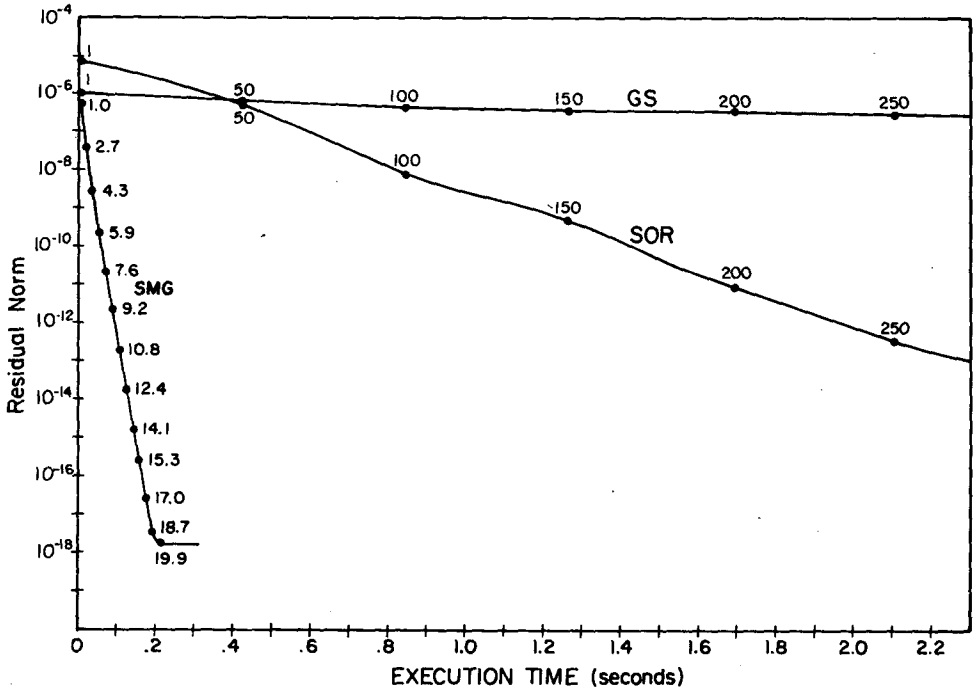

FIG. 6. Scaled discrete l_2 norm of the residual r (m s^{-3}) as a function of CRAY1 execution time in seconds for three different numerical schemes: Gauss-Seidel (GS), optimal successive-over-relaxation (SOR), and the standard multigrid scheme (SMG). Values labeled along the GS and SOR curves represent the number of relaxation sweeps; values along the SMG curve are work units, where a work unit is defined as the computational work of one ADZ sweep on the finest grid.

TABLE 1. Efficiencies for the SOR and SMG schemes in terms of CRAY1 execution time (in sec) per grid point for five cases with different resolutions. The last column gives the ratio of the execution times for the two schemes.

Number of grid points			Grid resolution (km)		Number of grid levels in SMG	Execution time per grid point		
J_M	K_M	$J_M \times K_M$	ΔR_M	ΔZ_M	(M)	SOR	SMG	SOR / SMG
17	17	289	60.0	1.0	4	8.59 (−5)	1.31 (−5)	6.6
33	33	1089	30.0	0.5	5	1.02 (−4)	8.98 (−6)	10.2
65	65	4225	15.0	0.25	6	1.31 (−4)	7.23 (−6)	18.1
97	65	6305	10.0	0.25	6	1.67 (−4)	6.49 (−6)	25.8
129	65	8450	7.5	0.25	6	3.22 (−4)	5.88 (−6)	54.8

ecution time per grid point for SOR increases as the grid resolution increases, consistent with the theoretical asymptotic convergence rate, $1 - O(h)$, for SOR where h represents the mesh spacing. On the other hand, as the resolution increases in the SMG runs, the execution time per grid point decreases. Since multigrid smoothing rates are independent of grid resolution (Brandt, 1977), this behavior reflects the efficiency of vectorization as a function of vector length. The overhead, that is, the time required for setting up the problem, was 2% more in the case of the SMG runs reported in Table 1 in comparison to the SOR runs; this additional expense of the SMG scheme resulted from computing the coefficients (a and b) of (2.1a) and performing an LU decomposition of the tridiagonal systems on each grid level.

The asymptotic convergence rate computed (over a V-cycle) for the SMG solution of the test case was 0.28. This rate is considerably larger than predicted by the theoretical convergence factor in Fig. 1 at the corresponding ratio of b/a for the test case, in which $\log(b/a) = 4.75$. This discrepancy can be understood by realizing that the coefficients a and b in the test case are functions of R, whereas the theoretical results were computed by assuming constant coefficients. In cases with slowly varying coefficients (such as the test case) local mode analysis is still valid on the finer grids where coefficients are locally constant; however, on coarser grids the changes in the coefficients are more abrupt, invalidating a basic assumption in the local mode analysis. On the other hand, the numerical results in Fig. 1, which were computed with a and b constant, closely match the theoretical convergence rates. At ratios of b/a where the numerical results differ from the theoretical estimates, the two curves can be brought into agreement by using W-cycles, which solve the problem more accurately on the coarser grids; the theoretical estimates assume the problem was solved exactly on a coarse grid. The asymmetry in the numerical results in Fig. 1 reflects the order of the ADZ relaxation; first Zebra relaxation is done along Z-lines, then along R-lines. Switching this order of relaxation causes the asymmetry to reverse itself about $\log(b/a) = 0$.

Figures 7–10 present the numerically computed solution fields for comparison with their analytical counterparts in Figs. 2–5. Inspection of these fields reveals a poor numerical simulation which one would expect from using a wall (i.e., $\Psi = 0$) at a finite radius as an outer boundary condition. This condition acts to suppress the forced secondary circulation and to concentrate the warming within the finite radius specified by the wall; this effect is seen in the numerical warming field (Fig. 10) which shows significantly more warming than in the analytical case (Fig. 5).

5. Multigrid with local refinement

To improve the numerical solution shown in Figs. 7–10, let us consider the following changes in how the outer boundary condition is imposed. First, we could apply a more realistic boundary condition such as that discussed by Schubert and Hack (1983), who simulate an infinite domain by applying (at $R \sim 1000$ km) a condition derived from the analytical far-field solution of (1.1). This method appears quite acceptable, although it is complicated in the sense of being nonlocal, i.e., it must be applied mode-by-mode in the vertical. Thus, it requires the solution of an associated (discretized) vertical structure problem and repeated projections onto the resulting basis functions. A second al-

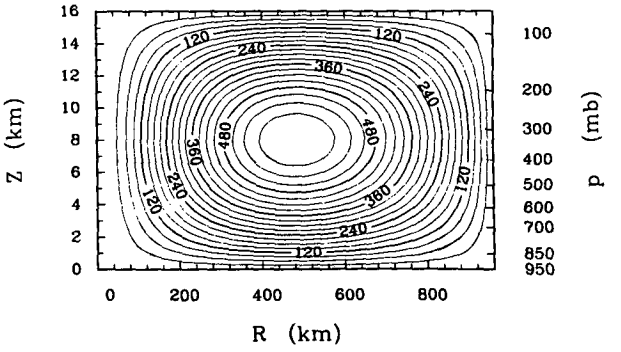

FIG. 7. Streamfunction $R\tilde{\Psi}$ (10^6 kg s^{-1}) computed numerically with the SMG scheme with contour interval 40×10^6 kg s^{-1}.

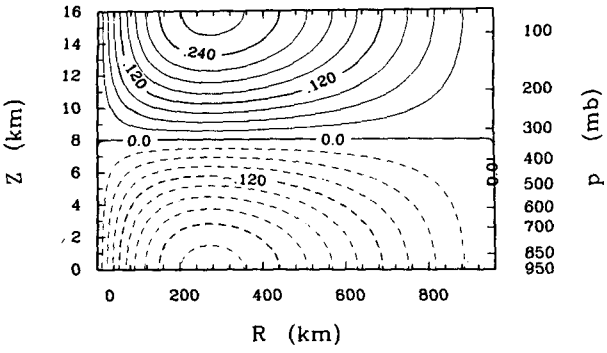

FIG. 8. Plot of u^* (m s^{-1}) computed numerically from streamfunction field in Fig. 7 with contour interval 0.03 m s^{-1} and dashed lines for negative values.

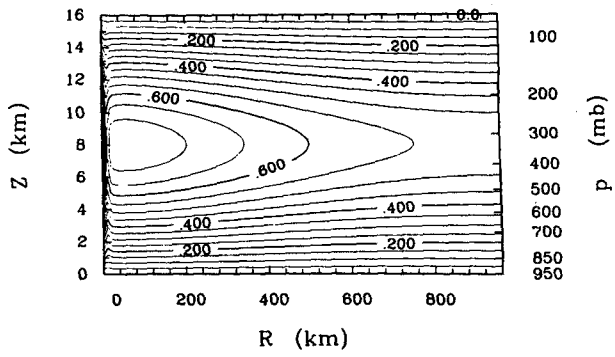

FIG. 10. Warming field $\partial\theta/\partial T$ (K d^{-1}), with contour interval 0.05 K d^{-1}, computed numerically using (4.3) and the streamfunction field in Fig. 7.

ternative is based on the fact that solutions of (1.1) become smoother as R increases. Thus, the lateral boundary can be moved far enough out such that $\Psi = 0$ is a realistic assumption, while the grid is coarsened in the outer regions so that not much additional computational work is required. In this approach we still make use of the analytical far-field solution, but only as a rough guide to the rate at which the grid is coarsened as R increases.

a. Far field solution and radial extent of various grids

To examine the far-field solution of (1.1), let us consider R large enough ($R > \bar{R}$, say) such that

$$\rho q \to \frac{g}{\theta_0}\frac{\partial\theta}{\partial Z} \equiv N^2,\ \rho s \to f^2,$$

and $\partial Q/\partial R \to 0$. In addition, if we make the Boussinesq approximation and assume N is a constant, the far-field equation becomes

$$\frac{\partial}{\partial R}\left(\frac{\partial R\psi^*}{R\partial R}\right) + \frac{f^2}{N^2}\frac{\partial^2\psi^*}{\partial Z^2} = 0 \quad \text{for} \quad R > \bar{R}. \tag{5.1}$$

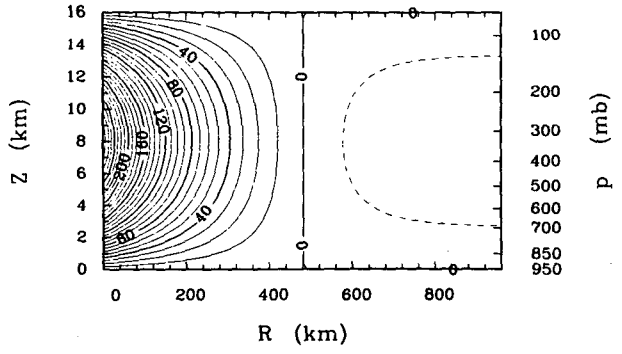

FIG. 9. Plot of w^* (10^{-4} m s^{-1}) computed numerically from streamfunction field in Fig. 7 with contour interval 10^{-4} m s^{-1} and dashed lines for negative values.

For the frictionless case the solution of (5.1) is

$$\psi^*(R, Z) = \sum_{n=1}^{\infty} A_n K_1(\lambda_n R)\sin\left(\frac{n\pi Z}{\hat{Z}}\right) \tag{5.2}$$

where $\lambda_n = f/c_n$, $c_n = N\hat{Z}/n\pi$ and K_1 is a modified Bessel function of the first order. The coefficients A_n are determined by the details of the solution in the inner region ($R < \bar{R}$). For the present purpose the crucial feature of (5.2) is that the various modes decay with radius at different rates. This is illustrated by the characteristic decay scales, λ_n^{-1}, given in the third column of Table 2, where the higher vertical modes decay faster with R.

Suppose that a single mode n is excited and that the lateral boundary $\hat{R}$ of our finite difference grid is placed far enough away so that the solution decays to about 10% or less of its value at $R = \bar{R}$. Then $\hat{R}$ must satisfy

$$\frac{K_1(\lambda_n\hat{R})}{K_1(\lambda_n\bar{R})} \leqslant 0.1. \tag{5.3}$$

Convenient values of $\hat{R}$ satisfying (5.3) for different n with $\bar{R} = 800$ km are given in the last column of Table 2. If, in addition to the internal modes ($n > 0$), an

TABLE 2. Characteristics of the vertical modes in (5.2): c_n is the phase speed, λ_n^{-1} the characteristic decay scale, and $\hat{R}$ a convenient radius at which the solution has decayed to less than 10% of its value at $R = \bar{R}$.

n	$c_n = \frac{N\hat{Z}}{n\pi}$ (m s^{-1})	$\lambda_n^{-1} = \frac{c_n}{f}$ (km)	$\hat{R}$ (km)
1	61.12	1220	2880
2	30.56	611	1920
4	15.28	306	1440
8	7.64	153	1160
16	3.82	77	960

TABLE 3. Grid domains and resolutions required to adequately resolve vertical modes in Eq. (5.2).

Grid level l	Vertical mode resolved (n)	Number of grid points			Resolution (km)		Extent of domain (km)
		J_l	K_l	$J_l \times K_l$	ΔR_l	ΔZ_l	$\hat{R}_l$
1	0	16	3	48	320.0	8.0	4800
2	1	19	5	95	160.0	4.0	2880
3	2	25	9	225	80.0	2.0	1920
4	4	37	17	629	40.0	1.0	1440
5	8	51	33	1683	20.0	0.50	1160
6	16	97	65	6305	10.0	0.25	960

external type response is possible due to boundary-layer frictional effects, an even larger $\hat{R}$ of about 4800 km is needed. Assuming that eight grid points per vertical wave length are necessary to resolve a given mode, Table 3 summarizes the grid domains and resolutions required for the various modes discussed above. These grids are depicted in Fig. 11.

In practice, of course, all vertical modes are excited and the lateral boundary of the computational domain should be placed at about 4800 km. A grid with uniform resolution to this distance with $\Delta R = 10$ km and $\Delta Z = 0.25$ km would require 31 850 points, which is not an attractive prospect. A more practical grid construction would include nonuniform resolution such as would result from overlapping the grids in Fig. 11.

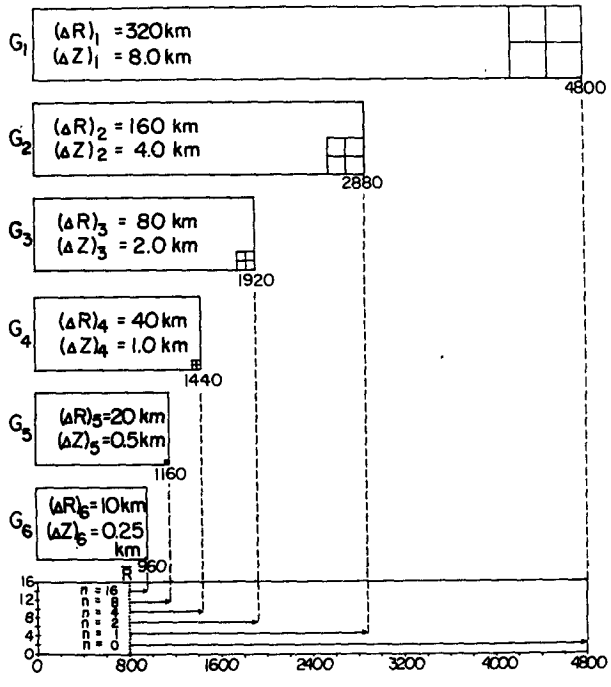

FIG. 11. Grid domains and resolutions required to resolve the vertical modes n which result from the solution of the far-field equation (5.1). The outer boundary of a given grid G_l is placed at the radius $\hat{R}_l$ given in Table 3 such that the solution for the mode n decays away to 10% or less of its value at $R = \hat{R}$.

A grid constructed in this manner would provide adequate resolution where it is needed without an excess of additional grid points. In fact, for the case considered here the number of points exterior to G_6 (i.e., the region of finest resolution) is 9% of the number interior to G_6; this compares to a 405% increase if a grid with the uniform resolution of G_6 was used out to 4800 km.

Computing with standard finite differences on a grid with nonuniform mesh spacing is an awkward procedure. An advantage of the multigrid method is that it allows nonuniformity to be organized through the use of uniform grids on various levels of discretization, but with each of these uniform grids extending over different domains. In the discussion that follows we demonstrate how the concepts of multigrid and local mesh refinement can be merged in an efficient and natural way to obtain a numerical scheme which accommodates nonuniform resolution.

b. *Local refinement*

We now present an alternative to the SMG scheme, which will subsequently be referred to as LRMG (i.e., local refinement with multigrid). The multigrid approach with local refinement involves relaxing (2.1)–(2.2) on a certain grid G_l which has uniform resolution. If convergence is slow on that grid, the problem is switched to a coarser but noncoextensive grid G_{l-1} having different but uniform resolution. On the other hand, if the solution has converged on G_l, the problem is transferred to a finer grid G_{l+1} where the outer boundary values are obtained naturally by interpolating from the coarser grid. This process is repeated until the solution has converged to a specified tolerance on the finest grid level. Such a scheme is computationally convenient in that each grid has uniform resolution. Since finer grids are introduced only where needed, this approach solves the problem with the usual multigrid efficiency.

The grid structure chosen for the LRMG scheme is summarized in Table 3. Due to the noncoextensive structure of these grids, the parts of the domain on a coarse grid not covered by a finer grid must certainly carry the full solution. This is accomplished by using the Full Approximation Scheme (FAS), in which the

full current approximation is stored on each grid level (as opposed to the Correction Scheme, where a coarse grid stores a correction to the finer grid approximation). Details of the FAS mode of multigrid processing are given in Fulton et al. (1986).

For the LRMG scheme we have chosen to use Full Multigrid (FMG) as the control algorithm. The FMG algorithm, unlike cycling algorithms, works itself from the coarsest level ($l = 1$) to the finest level ($l = M$). In this manner a good first approximation is achieved on each grid in turn with minimal expense. Figure 12 shows a flowchart of the accommodative FAS-FMG algorithm used in LRMG. In this flowchart m denotes the current finest level, that is, the finest level for which an approximate solution $\tilde{\Psi}$ has already been computed; l refers to the current working level. The difference operator L^l as well as the initial right hand side f^l are computed using (2.1b) directly on each level. As in the SMG scheme, we again choose ADZ as the relaxation scheme for our current application. However, in contrast to the SMG scheme, the values inserted along the outer boundaries of the finer grids ($l > 1$) represent approximate solutions interpolated from coarser grids. When transferring the problem to a coarser grid, the FMG algorithm generates a new right hand side, F^l, at coarse grid points which are interior to the finer grid. The transfer operators are denoted in Fig. 12 as follows: I^l_{l+1} (full weighting of residuals), $\tilde{I}^l_{l+1}$ (injection), I^l_{l-1} (bilinear interpolation) and II^m_{m-1} (bicubic interpolation). A detailed explanation of the steps in this FMG algorithm is given in Brandt (1979).

An advantage of the FMG algorithm in Fig. 12 is that on each grid level $m > 1$ it generates an appropriate convergence tolerance ϵ_m internally. Setting the parameter $\gamma = 1$ makes ϵ_m an approximation to the truncation error; setting $\gamma < 1$ allows us to solve below truncation error, as discussed in section 4b. For problems where the truncation error is not known in advance this feature is very useful. When (2.1)–(2.2) was solved for the test case described in section 4 with the LRMG scheme, the computed approximate truncation error on the finest grid (ϵ_M) was 7.4×10^{-9} m s^{-3}, whereas the actual truncation error (τ) was 2.1×10^{-8} m s^{-3}.

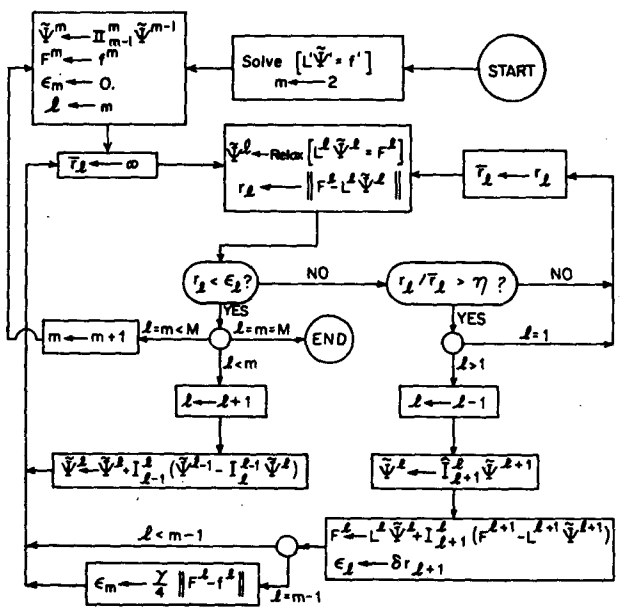

FIG. 12. Accommodative FAS-FMG algorithm (after Brandt, 1979).

c. *Results from the LRMG scheme*

We now apply the LRMG scheme to solve (2.1)–(2.2) for the test case described in section 4. For the results to be presented the parameters in the FMG algorithm were set as follows: $\delta = 0.3$, $\gamma = 1.0$, and $\eta = -1.0$. By specifying η to be negative we ignore the test for fast convergence which avoids a second (and in our case, unnecessary) relaxation sweep; one ADZ sweep substantially reduces the error in high-wavenumber modes, and further relaxation on the same level will reduce the remaining low-wavenumber modes very slowly.

Using the LRMG scheme, the test case was solved to the tolerance ϵ_M in 3.25 work units (0.053 seconds). In spite of the added complexity of the LRMG scheme, its computational efficiency is comparable to that of the simpler SMG scheme. Moreover, the solution fields of $\tilde{\Psi}$, u^* and w^* computed from the LRMG scheme are virtually indistinguishable from the corresponding analytical solutions (Figs. 2–4) and thus are not shown here. A comparison of the computed warming field (Fig. 13) and its analytical counterpart (Fig. 5) reveals a slight amplitude error in the numerical field inside 200 km, and a notable disparity near the boundary $R = 0$. The former error can be reduced by solving the problem to a smaller tolerance (i.e., by setting $\gamma < 1$). The latter error is the result of less accurate finite differences for w^* at the $j = 0$ and 1 points; these inaccuracies in w^*, although slight, are then magnified in the warming field.

In the test case solved above, the coefficients q and s were constant. We have also used the LRMG scheme to solve other problems which are more realistic. For example, in the case mentioned in section 3a [where s/q in (1.1) varied by a factor of 100], LRMG solved the problem to the tolerance ϵ_M (1.2×10^{-9} m s^{-3}) in 5 work units (0.067 seconds). From our experience, LRMG in a prognostic mode will solve (1.1) to the approximate truncation error ϵ_M in 6 work units or less for any ratio of s/q that occurs as the vortex evolves.

6. Concluding remarks

The multigrid schemes described in this paper were designed to handle the specific features of (1.1), namely its anisotropy and the unbounded domain. These schemes proved considerably more efficient than the

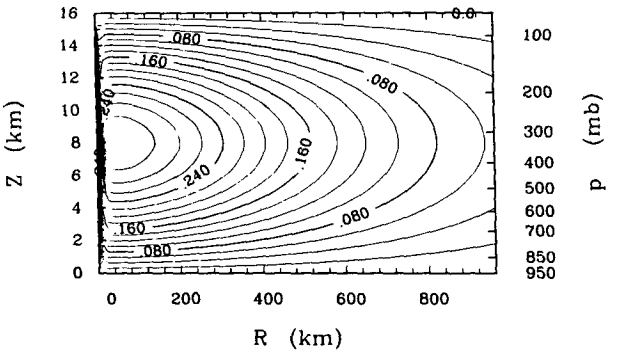

FIG. 13. Warming field $\partial\theta/\partial T$ (K d^{-1}) with contour interval 0.02 K d^{-1} obtained using (4.3) and the streamfunction field computed via the LRMG scheme.

SOR method. In solving (1.1) diagnostically, efficiency is not a practical concern; however in a prognostic mode this equation must be solved on the order of 1000 times during a typical numerical simulation of a tropical cyclone. In such cases efficiency is more crucial (especially with high resolution), and we recommend that the multigrid approach be given serious consideration.

In designing the multigrid scheme with local refinement, the gross features of the far-field solution were used as a rough guide to construct the grids. Alternately the choice of local refinements can be made totally adaptively, i.e., as required by the numerical solution as it evolves. This method is beyond the scope of the present paper, and the reader is referred to Bai and Brandt (1984) for further discussion.

In this paper the multigrid approach was used to design an efficient solver for an elliptic boundary value problem from tropical cyclone theory. Similar problems with variable coefficients occur in other areas of meteorology such as quasi-geostrophic, semi-geostrophic and nonlinear balance theory. We believe the multigrid approach has great potential for solving these problems efficiently, and we encourage others to try multigrid methods in such cases.

Acknowledgments. The authors gratefully acknowledge the help and guidance provided by Achi Brandt, especially during his sabbatical year at the Institute for Computational Studies at Colorado State University. We are also indebted to Gerald Taylor for his theoretical analysis of the relaxation scheme and to Odilia Panella for her help in preparing the manuscript.

This work was funded by the Naval Environmental Prediction Research Facility, Monterey, California, under program element 61153N, "Multigrid Elliptic Solver Development."

REFERENCES

Alcouffe, R. E., A. Brandt, J. E. Dendy, Jr. and J. W. Painter, 1981: The multi-grid method for the diffusion equation with strongly discontinuous coefficients. *SIAM J. Sci. Stat. Comput.*, **2**, 430–454.

Bai, D., and A. Brandt, 1984: Local mesh refinement multilevel techniques. [Available from Department of Applied Mathematics, Weizmann Institute of Science, Rehovot 76100, Israel.] 45 pp.

Brandt, A., 1977: Multi-level adaptive solutions to boundary-value problems. *Math. Comp.*, **31**, 333–390.

——, 1979: Multi-level adaptive techniques (MLAT) for singular-perturbation problems. *Numerical Analysis of Singular-Perturbation Problems.* P. W. Hemker and J. J. H. Miller, Eds., Academic Press, 53–142.

——, 1982: Guide to Multigrid Development. *Multigrid Methods.* W. Hackbusch and U. Trottenberg, Eds., Lecture Notes in Mathematics, Vol. 960, Springer-Verlag, 220–312.

——, 1984: Multigrid Techniques: 1984 guide with applications to fluid dynamics. [Available from GMD, Postfach 1240, D-5205 St. Augustin, 1, F.R. Germany.] 176 pp.

Eliassen, A., 1952: Slow thermally or frictionally controlled meridional circulation in a circular vortex. *Astrophys. Norv.*, **5**, 60 pp.

Fulton, S. R., P. E. Ciesielski and W. H. Schubert, 1986: Multigrid methods for elliptic problems: A review. *Mon. Wea. Rev.*, **114**, 943–959.

Ooyama, K., 1969: Numerical simulation of the life cycle of tropical cyclones. *J. Atmos. Sci.*, **26**, 3–40.

Schubert, W. H., and J. J. Hack, 1982: Inertial stability and tropical cyclone development. *J. Atmos. Sci.*, **39**, 1687–1697.

——, and ——, 1983: Transformed Eliassen balanced vortex model. *J. Atmos. Sci.*, **40**, 1571–1583.

Stüben, K., and U. Trottenberg, 1982: Multigrid methods: Fundamental algorithms, model problem analysis and applications. *Multigrid Methods.* W. Hackbusch and U. Trottenberg, Eds., Lecture Notes in Mathematics, Vol. 960, Springer-Verlag, 1–176.

Sundqvist, H., 1970: Numerical simulation of the development of tropical cyclones with a ten-level model. Part I. *Tellus*, **22**, 359–390.

Willoughby, H. E., 1979: Forced secondary circulations in hurricanes. *J. Geophys. Res.*, **84**, 3173–3183.