LA-UR- 95-2794

Title:

A REDUCED GRID MODEL FOR SHALLOW FLOWS ON THE SPHERE

Author(s):

J. M. Reisner, L. G. Margolin, P. K. Smolarkiewicz

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Submitted to:

Parallel CFD '95
Pasadena, CA
July 1995

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED 35

MASTERForm No. 836 R5
T 2629 10/91

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

A Reduced Grid Model For Shallow Flows on the Sphere

J.M. Reisner^a and L.G. Margolin^a

^aLos Alamos National Laboratory, Los Alamos, New Mexico

P.K. Smolarkiewicz^b

^b National Center for Atmospheric Research, Boulder, Colorado

We describe a numerical model for simulating shallow water flows on a rotating sphere. The model is augmented by a reduced grid capability that increases the allowable time step based on stability requirements, and leads to significant improvements in computational efficiency. The model is based on finite difference techniques, and in particular on the nonoscillatory forward-in-time advection scheme MPDATA. We have implemented the model on the massively parallel CM-5, and have used it to simulate shallow water flows representative of global atmospheric motions. Here we present simulations of two flows, the Rossby-Haurwitz wave of period four, a nearly steady pattern with a complex balance of large and small scale motions, and also a zonal flow perturbed by an obstacle. We compare the accuracy and efficiency of using the reduced grid option with that of the original model. We also present simulations at several levels of resolution to show how the efficiency of the model scales with problem size.

1. Introduction

We describe a numerical model for simulating shallow water flows on a rotating sphere. The model is based on Eulerian spatial differencing and nonoscillatory forward-in-time (NFT) temporal differencing. Finite difference methods have advantages over spectral methods when implemented on massively parallel computers with distributed memory because the computations are localized in space. However finite difference methods have more restrictive time step constraints and so computational efficiency becomes an important issue. Our model is explicit in time, and its computational time step is limited by the largest Courant number on the mesh. Since the fastest wave speed is that of gravity waves, and is uniform over the mesh, the largest Courant number is associated with the smallest cell dimensions. In the typical latitude-longitude mesh, these smallest dimensions are found in the cells nearest the poles.

There are several strategies available to increase the time step and thus improve the computational efficiency of a finite difference model. For example, we have developed a semi-implicit version of our model (Nadiga et al. 1996 and Smolarkiewicz and Margolin 1994) in which the gravity waves are treated implicitly, but the advective velocities (which are much slower) are treated explicitly. In typical atmospheric applications, semi-implicit methods may allow a four-fold increase in time step. However the semi-implicit formulation leads to an elliptic problem, whose solution involves inverting a matrix on the mesh. Furthermore, the matrix operator is not symmetric due to the presence of Coriolis forces. This means that standard conjugate gradient methods may not converge and less optimal methods must be explored. Another alternative to allow larger time steps is filtering the velocities at the high latitudes. Filtering, however, requires global communication, making application on a massively parallel computer with distributed memory very inefficient.

Yet another alternative is the reduced grid. Here the logical structure of the regular latitude-longitude mesh is modified by removing some of the cells near the poles, effectively making the remaining cells larger. For example, if every other cell is removed from the regions within 30° of the poles, then the time step can be increased by a factor of two. The reduced grid also requires some nonlocal communication, but its impact on efficiency is much smaller than that of filtering. In addition, in our experiments we have found that the reduced grid reduces the accuracy less than either filtering or using implicit techniques.

In the following sections, we will briefly describe the model and the reduced grid. We will then show results for two problems from the suite of test problems described by Williamson et al. 1992; these are the Rossby-Haurwitz wave and the steady-state zonal flow on the sphere. The use of larger time-steps reduces the numerical diffusion associated with truncation error. In the first case this is the dominant source of error, so that the reduced grid model yields more accurate results than the nonreduced grid. In the second case the interpolation errors dominate the diffusion errors and the reduced grid model is less accurate than the nonreduced model. We will provide error estimates for the two cases, as well as timing statistics for both the reduced and nonreduced grid in section 4. We summarize our results in section 5.

2. Shallow Water Model

The equations expressing conservation of mass and momentum for a shallow fluid on a rotating sphere are as follows:

$$\frac{\partial G\Phi}{\partial t} + \nabla \cdot (\mathbf{v}\Phi) = 0, \quad (1a)$$

$$\frac{\partial GQ_x}{\partial t} + \nabla \cdot (\mathbf{v}Q_x) = GR_x, \quad (1b)$$

$$\frac{\partial GQ_y}{\partial t} + \nabla \cdot (\mathbf{v}Q_y) = GR_y, \quad (1c)$$

where $G = h_x h_y$, and h_x and h_y represent the metric coefficients of the general orthogonal coordinate system, $\Phi = H - H_o$ is the thickness of the fluid with H and H_o denoting the height of the free surface and the height of the bottom, $\mathbf{v}$ is the horizontal velocity vector, and $\mathbf{Q} = (\Phi u h_x, \Phi v h_y)$ is the momentum vector. The right-hand-side forcings are

$$R_x = -\frac{g}{h_x} \Phi \frac{\partial(\Phi + H_o)}{\partial x} + fQ_y + \frac{1}{G\Phi} \left(Q_y \frac{\partial h_y}{\partial x} - Q_x \frac{\partial h_x}{\partial y} \right) Q_y \quad (2a)$$

$$R_{xy} = -\frac{g}{h_y} \Phi \frac{\partial(\Phi + H_o)}{\partial y} - fQ_x + \frac{1}{G\Phi} \left(Q_y \frac{\partial h_y}{\partial x} - Q_x \frac{\partial h_x}{\partial y} \right) Q_x, \quad (2b)$$

where g is the acceleration of gravity and f is the Coriolis parameter.

The integration in time of the discretized approximations to (1) is described in Smolarkiewicz and Margolin (1993). The basic element of our nonoscillatory forward-in-time (NFT) algorithm is the sign-preserving advection scheme MPDATA (Smolarkiewicz 1984). The use of two-time-level integration schemes is a departure for Eulerian global atmospheric models where three-time-level or leapfrog schemes are traditionally used. However two-time-level schemes are widely preferred in most other fields of computational fluid dynamics. Some of the advantages of NFT schemes include a larger computational time step, reduced memory usage, and less numerical dispersion. In addition, the nonoscillatory property is crucial for preserving the sign of the layer thickness and of the thermodynamic scalars, and further controls the nonlinear stability of the computations. The model is implemented on a rotating sphere, and allows for arbitrary bottom topography as well as a free surface on top of the layer.

We have ported the model to the CM-5. It runs in data parallel mode, with the horizontal dimensions being spread across processors. In a typical problem, the speed

(measured in Megaflops) depends on problem size. For 32 nodes, a problem with a 64×128 mesh yields performance equivalent to 0.5 CRAY YMP processors, whereas a problem with 256×512 nodes runs at a speed equivalent to 1.5 CRAY YMP processors.

The reduced grid that we have implemented is adapted from the work of Rasch (1994). We use a sequence of nonoverlapping domains, where the number of grid points along circles of latitude decreases as one approaches the poles. The number of points from one domain to the next decreases by multiples of two, both for accuracy of interpolation as well as efficiency on the CM-5 architecture. One difference from Rasch (1994) is that the minimum number of grid points at a given latitude that is allowed for the top domains was fixed at 32 and not 4. This choice results both in increased accuracy and efficiency. Initially the latitude at which the reduction occurred was chosen as suggested by Rasch (1994); however sensitivity tests have revealed that the most accurate results occur when the reductions occur only in the vicinity of the pole—with only a slight increase in CPU time (with 2.8° resolution about 1 s for 7 days)—of a simulation being noted with this strategy. For example, in a simulation with resolution of 2.8° resolution at the equator three grids of $128 \times 58 \times 1$, $64 \times 4 \times 2$, and $32 \times 4 \times 2$ are used to cover the globe (see Fig. 1 for a visual representation of the reduced grid). The ghost points at the top and bottom of each domain are simply interpolated values from other domains to be used in the calculations of north-south derivatives within a given domain. We use our NFT advection scheme for this interpolation, thus maintaining the positivity of the height and thermodynamic fields (Smolarkiewicz and Grell, 1992) and preserving the overall second-order accuracy of the model. This interpolation does require some nonlocal communication. At present, no temporal interpolation is used between the meshes of the reduced grid. The ratio between the time step required for the regular grid versus the reduced grid ranges from 4 with 2.8° resolution to 50 with 0.35° resolution. Note that as the resolution is increased in the nonreduced grid, the velocity field becomes more finely resolved, and locally may exceed the maximum values found in the less resolved solution. To ensure stability in such cases, the time step must be reduced by more than a factor of 4 as we increase the nonreduced grid's resolution from 0.7° to 0.35° —see Rasch (1994). We will demonstrate in section 4 that the nonlocal communication associated with interpolation does not allow for a corresponding ratio of CPU time between the reduced versus nonreduced grid on the CM5.

Fig. 1 Upper left portion of reduced grid.

3. Shallow Water Test Problems and Model Setup

Shallow water is a useful testbed for evaluating and comparing numerical methods that can be extended to fully three-dimensional global circulation models (GCMs). Also, shallow water layers can be stacked, and with the addition of a hydrostatic equation and slight recoding of (1)-(2), can be extended to model fully three-dimensional flows (cf. Bleck 1984). Such models are termed isentropic in meteorology. As part of DOE's CHAMMP program, Williamson et al. (1992) have published a suite of seven test problems for shallow water. One of these, the Rossby Haurwitz (RH) wave is particularly challenging, since it represents an almost steady state flow that is a delicate balance of large and small spatial scales. (When the free surface is replaced by a rigid lid, the solution is an exact steady state.) The RH wave is also interesting because it represents the advection of a wave which may be significantly damped if low-order forward-in-time methods are used (Smolarkiewicz and Margolin, 1993). Thus, the use of a larger time step reduces the numerical dissipation so that the reduced grid may perform as well or better than the regular grid.

Another test of the robustness of the reduced grid is a zonal flow over the pole. We have modified Williamson's standard test by the inclusion of a 6 km tall mountain in an approximately 8 km depth fluid. The mountain is located on the north pole. Although the mountain introduces an unsteady component into the problem, the flow is still relatively balanced and little change of position of the height contours is noted during a simulation. Unlike the RH wave in which dissipation errors dominate interpolation errors, interpolation errors dominate this flow situation. Thus the total error in the reduced grid formulation is found to be larger than in the nonreduced grid. Another source of error in the reduced grid formulation is that the topography on the coarser meshes found near the pole is not as well-resolved as in the nonreduced grid, potentially leading to additional differences between the two solutions. Because analytic solutions are not known for either the RH wave or the perturbed zonal-flow case, we have run high resolution simulations as a standard for comparing the results from the nonreduced and reduced grid. Except for timing statistics, all results are from simulations with 2.8° resolution at the equator. The simulations were run for a period of 7 days with a time step of 160/40 s being used for the reduced/nonreduced grid.

4. Results

4.1. Rossby-Haurwitz wave

Since visual differences between the reduced grid and nonreduced grid are not discernible, and the solutions have been published elsewhere (Smolarkiewicz and Margolin 1993, Fig. 2), we show only the L_2 and L_∞ error measures (see Fig. 2a) with respect to the height of the shallow fluid as a function of time for the reduced and nonreduced grids. These are the measures of error recommended by Williamson et al. (1992). Even in the error comparisons, very little difference is apparent—the errors associated with the reduced grid are only slightly smaller than those of the nonreduced grid. A further test was conducted in which the reduced grid's time step was halved (160s to 80s) to determine the sensitivity of the solution to numerical dissipation. As expected, the errors fall in between those of the nonreduced grid and of the reduced grid that used a twice-larger time step.

Our original motivation for introducing the reduced grid was to improve computational efficiency. Table 1 demonstrates that with 2.8° resolution the total CPU time (total time is for one hour of simulated time) for the nonreduced grid (denoted by 2.8

in Table 1) and for the reduced grid (denoted by 2.8r in Table 1) are nearly equal. The ratio increases to about 35 with 0.35° resolution. Hence, at least on the CM5, the bigger gains in efficiency occur for the higher resolution grids. The main cause for the increasing ratio with grid resolution is directly related to the ratio of time steps required for each approach, which increases with decreasing resolution (see discussion at the end of section 2). Breaking down the total CPU time into four components, node CPU time, NEWS or parallel communication (e.g., cshifts), Send/Get or nonparallel communication (e.g., interpolation), and other (e.g., communication between nodes and program manager, global sums, ect...) we observe that for the nonreduced grid the primary component is NEWS; whereas for the reduced grid the primary component is Send/Get for the smaller domains and on node calculations for the larger domains. In addition, the reduced grid contains fewer grid points than the nonreduced grid, so that the amount of memory used in the larger grids of the reduced mesh is somewhat less than that of the nonreduced mesh (about 100 mb for a grid with 0.35° resolution).

Table 1

Resolution°	Node CPU	NEWS	Send/Get	Other	Total
2.8	0.864	2.955	0.000	0.009	3.828
2.8r	0.984	1.807	2.552	0.216	5.559
1.4	8.684	26.504	0.000	0.052	35.240
1.4r	4.152	5.076	7.800	1.932	18.960
0.7	119.504	325.216	0.000	0.336	445.056
0.7r	18.224	15.248	20.880	7.056	61.408
0.35	2860.100	6914.200	0.000	4.300	9778.600
0.35r	116.944	54.764	51.326	50.060	273.094

4.2. Zonal Flow

Unlike the RH wave, visual differences are apparent between the solutions produced by the reduced and nonreduced meshes for the perturbed zonal flow. Figs. 3a, 3b, and 3c, show the numerical solutions for the reduced mesh, the nonreduced mesh, and a high resolution simulation. Again we quantify the differences of these three solutions in terms of the L_2 and L_∞ error measures (see Fig. 2b) of the thickness of the fluid. In comparison with the errors of the RH wave, the absolute error levels for this case with respect to the highly resolved simulation are smaller—the L_2 error is about an order of magnitude smaller. Hence, as was noted in Smolarkiewicz and Margolin (1993) the smaller error measures for this case suggest a less dissipative flow regime (due to numerical discretization). The fact that the error measures for the reduced grid are greater than the nonreduced grid also suggest that the main cause of these differences is due to errors associated with the interpolation in the reduced grid.

As noted in section 2 a shallow-water model can be easily extended to include the effects of baroclinicity in the atmosphere. To investigate whether adding additional layers, and hence requiring additional interpolation will degrade the quality of the solution we ran the zonal flow problem with 8 layers for both the reduced grid and the nonreduced grid. Our analysis of the results of these runs indicate that adding additional layers to the reduced grid model does not further degrade the solution with respect to the nonreduced grid. In addition, the timing statistics suggest that as the number of levels in the vertical increases, the efficiency of the reduced grid also increases, so that with 2.8° resolution and 64 layers the reduced grid is about 1.25 times faster than the nonreduced grid.

Fig. 2 Time evolution of the error measures for the model configuration with 2.8° resolution for the (a) Rossby-Haurwitz wave and for (b) Zonal Flow.

Fig. 3 Geopotential perturbation and the flow vectors for the global nonlinear zonal geostrophic flow test problem with a tall mountain at the pole: (a) the reduced mesh; (b) the nonreduced mesh; and (c) the resolved solution.

5. Conclusions

We have described a numerical model for shallow flows on the sphere. We have developed a reduced grid model to circumvent the time step constraints, and consequent lack of efficiency associated with the convergence of meridians at the poles. At low resolution we have shown that the use of the reduced grid is as accurate as the nonreduced grid with about the same computational efficiency. The principal computational advantage of the reduced grid is realized at higher resolutions, with the reduced grid being up to 35 times faster than the nonreduced grid at a resolution of 0.35° . Due to faster communication on machines like the Cray T90 or T3D we believe that the reduced grid will lead to greater efficiency on these machines. Thus we conclude that the reduced grid framework is a useful modification for general circulation models of the atmosphere and ocean based on finite difference approximations.

6. Acknowledgements

The authors appreciate continuing discussions with Phil Rasch, and the help of the system programmers of the Los Alamos Advanced Computing Laboratory. This work was supported by in part by the CHAMMP program of the U.S. Department of Energy.

REFERENCES

- Bleck, R., 1984: An isentropic model suitable for lee cyclogenesis simulation. *Riv. Meteor. Aeronaut.*, **43**, 189-194.
- Nagida, B.T., L.G. Margolin, and P.K. Smolarkiewicz, 1995: Semi-Lagrangian Shallow Water Modeling on the CM5. submitted to special issue of Elsevier Science Publishers for Proceedings of Parallel CFD'95.
- Rasch, P.J., 1994: Conservative shape-preserving two-dimensional transport on a spherical reduced grid. *Mon. Wea. Rev.*, **122**, 1337-1350.
- Smolarkiewicz, P.K., 1984: A fully multidimensional positive definite advection transport algorithm with small implicit diffusion. *J. Comp. Phys.*, **54**, 323-362.
- Smolarkiewicz, P.K. and G.A. Grell, 1992: A class of monotone interpolation schemes. *J. Comp. Phys.*, **101**, 431-440.
- Smolarkiewicz, P.K. and L.G. Margolin, 1993: On forward-in-time differencing for fluids: extension to a curvilinear framework. *Mon. Wea. Rev.*, **121**, 1847-1859.
- Smolarkiewicz, P. K., and L. G. Margolin, 1994: Forward-in-time differencing for global dynamics. In *Preprint Vol. Fifth Symposium on Global Change Studies*, 23-28 January, Nashville, Tennessee, American Meteorological Society, 99.
- Williamson, D.L., J.B. Drake, J.J. Hack, R. Jakob, and P.N. Swarztrauber, 1992: A standard test set for numerical approximations to the shallow water equations on the sphere. *J. Comp. Phys.*, **102**, 211-224.