

A three-dimensional, adaptive, Godunov-type model for global atmospheric flows

M.E.Hubbard¹

and

N.Nikiforakis

DAMTP, University of Cambridge

Submitted: April 2002

Revised: November 2002/January 2003

Address for correspondence

Dr. M.E.Hubbard

School of Computing

University of Leeds

Leeds, LS2 9JT, U.K.

Email: meh@comp.leeds.ac.uk

¹Currently at: School of Computing, University of Leeds

Abstract

In this paper a Godunov-type methodology is applied to three-dimensional global atmospheric modelling. Numerical issues are addressed regarding the formulation of the tracer advection problem, the application of dimensional splitting and the implementation of a Godunov-type scheme, based on the WAF approach, on spherical geometries. Particular attention is paid to addressing the problems which arise because of the convergence of the grid lines towards the poles. A three-dimensional model is then built on the sphere which is based on a uniform longitude/latitude/height grid. This provides the framework within which an Adaptive Mesh Refinement (AMR) algorithm is applied, to enhance the efficiency and accuracy with which results are obtained. These methods are not commonly used in the area of atmospheric modelling, but AMR in particular is commonly used with great success in other areas of CFD. The model is initially validated using a series of idealised case studies which have exact solutions, but is then developed into an off-line model of tracer advection, forced by data from meteorological analyses, suitable to study the evolution of trace chemical species in the atmosphere.

1. Introduction

The construction of an accurate numerical algorithm for the modelling of advection is of fundamental importance to the approximation of the equations of fluid dynamics, whether it is used as a component of the numerical model of the full dynamical equations or as a scheme in its own right for the modelling of advective transport. The form of the numerical scheme depends greatly on the nature of the most important features of the underlying problem it is being used to model. In the field of meteorology, particularly when modelling fully three-dimensional flows, it is the scheme's ability to capture flow features over a huge range of time and length scales which has driven the development. This has placed an emphasis on schemes with a high order of accuracy, to minimise the effects of numerical dissipation and dispersion.

Traditionally, spectral methods have been popular for the modelling of atmospheric dynamics, while there has been much recent effort applied to the use of semi-Lagrangian methods (Staniforth and Côté 1991) which, in idealised situations, can achieve this high degree of accuracy in time-dependent calculations using a time-step which is not restricted by stability considerations. This is of particular importance when a longitude/latitude/height grid is used on the sphere because of the way that grid nodes become clustered around the poles, creating very narrow cells which impose a prohibitively severe restriction on the time-step of most other numerical methods. Even so, the time-step which can be taken with a semi-Lagrangian scheme still has a practical limitation in terms of the scheme's accuracy.

Furthermore, the methods are not inherently either monotonic (avoid spurious oscillations) or conservative, although both these conditions may be enforced with extra work. For example, the approach of Lin and Rood (1996) recasts the semi-Lagrangian scheme within the framework of a flux-based method, thereby creating an inherently conservative method. In fact, the suggestion that fluxes can be combined with a semi-Lagrangian approach was made by van Leer (1979). The method proposed there is equivalent to the QUICKEST scheme, the foundations of which are given in Leonard (1979), which has since been made monotonic (Leonard, 1991), and recast in a semi-Lagrangian form. Much of this work is essentially one-dimensional, but more recent advances have introduced genuinely multidimensional advection schemes: the Corner Transport Upwind scheme of Colella (1990) is an early example, but the work of Leonard has since been extended in schemes such as UTOPIA (Leonard *et al.*, 1995). These methods are equally valid for modelling tracer advection, but many popular models of off-line chemical transport choose to apply schemes which are more specifically appropriate to the problem. Examples of this are the SLIMCAT and TOMCAT models (Chipperfield *et al.* 1996) which use the Prather's scheme, and the model of Plumb *et al.* (1994) which applies the contour advection method.

The desire for conservation has led to many flux-based finite volume methods being constructed and used in this area (see the work of Prather (1986) and the wide-ranging discussion paper of Leonard *et al.* (1996) for examples of high order schemes of this type). High resolution flux limited schemes, closely related to those proposed by van Leer (1979), and formalised by Sweby (1984), have so far been less

popular, but still provide a combination of second order accuracy and monotonicity. These methods are cheap to apply and, although they are formally of a lower order of accuracy than the schemes commonly used for tracer transport in meteorology, they are very flexible in how they can be generalised to model nonlinear systems of equations.

Furthermore, they can be combined straightforwardly with Adaptive Mesh Refinement (AMR) techniques to increase their efficiency and counterbalance any loss of accuracy compared to higher order schemes on grids of the same resolution. The inherent conservation and monotonicity properties of these Godunov-type schemes and the ability to reduce significantly the number of grid cells via AMR will be of immense value when chemistry is included within the model.

AMR is a widely used technique which is applied to the numerical solution of partial differential equations in order to increase the efficiency with which a numerical method can be used to obtain a solution of a given accuracy. It constructs and, in time-dependent situations, continually updates the grid on which the equations are being approximated, to achieve appropriate resolution throughout the computational domain. In some sense the grid generation seeks to equidistribute the error in the solution associated with each cell at each time-step, and in doing so it tracks interesting features of the flow. As the underlying atmospheric models become more complex, and particularly with the inclusion of chemical reactions, the value of the AMR will increase since the overhead associated with the mesh generation and communication becomes less significant.

Many variants of AMR have been developed but, given that the underlying grid

structure most commonly used in the field of meteorology is that of the standard longitude/latitude/height grid, it is a technique based on the method originally proposed by Berger and Olinger (1984), updated by Berger and Colella (1989), and contributed to by others such as Quirk (1991), which has been used as the basis of this work. Adaptive meshing of this type is not commonly used in atmospheric modelling, though Skamarock and Klemp (1993) have implemented an earlier version in smaller scale models of storm simulation. Behrens (1996,1998) and Giraldo (2000) have both carried out related research using different discretisation schemes on adaptive, unstructured triangular grids.

The combination of high resolution finite volume schemes with adaptive mesh refinement is used here to provide a basis for the accurate modelling of time-dependent off-line chemical transport (where the velocity field is supplied from some external source, analytic or observational). The next section contains a description of one such high resolution scheme, the dimensionally split WAF scheme (Toro 1989), the way in which it is applied to the scalar advection equation (as opposed to a conservation law), and the modifications required to deal with the underlying spherical geometry. Section 3. describes the AMR algorithm as applied to a general flux-based scheme for approximating hyperbolic conservation laws on the sphere. A series of test cases are then presented to illustrate the accuracy of the overall approach when applied to problems relevant to global atmospheric modelling, and the improvements achieved with the AMR over calculations carried out on single, uniform, structured grids. The paper is completed by some brief conclusions.

2. Mathematical and numerical formulation

The equation being approximated here is the three-dimensional scalar advection equation

$$\psi_t + \mathbf{u} \cdot \nabla \psi = 0, \quad (1)$$

in which ψ represents the advected scalar quantity, such as a chemical mixing ratio, and $\mathbf{u} = (u, v, w)^T$ is the advection velocity vector. There are numerous methods by which this equation can be solved directly, *e.g.* semi-Lagrangian, but we are seeking to apply standard finite volume techniques, widely used in areas outside meteorology for the numerical approximation of conservation laws. To facilitate this, note that Eq. (1) is not itself a conservation law (except in the special case when $\nabla \cdot \mathbf{u} \equiv 0$), but it may be rewritten

$$\psi_t + \nabla \cdot (\psi \mathbf{u}) = \psi \nabla \cdot \mathbf{u}. \quad (2)$$

The left hand side is in conservation form, and can be discretised using a flux-based finite volume technique, whilst the right hand side acts as a forcing term. By approximating the equation in this form, as long as the source term is discretised appropriately, the numerical method will be conservative in the special case when the discrete representation of the velocity field is divergence-free, *i.e.* precisely the situation where the advection equation and the conservation law are equivalent. Any non-zero divergence simply provides a correction to the conservation law which is consistent with modelling an advection equation. This idea is not new, and is closely related to the approaches of Easter (1993) and, further back, Petshek and Libursky (1975), which are also constructed to avoid the imbalances which can be introduced

by dimensional splitting.

a. *Finite volume representations*

The finite volume method is underpinned by the idea of discretising the integral form of a conservation law and then applying the divergence theorem to the flux integral over a series of ‘control volumes’ which cover the computational domain. An appropriate discretisation of these boundary integrals ensures that the numerical scheme retains the conservation property of the underlying equation. These ideas may be incorporated into the solution of the scalar advection equation (2) by discretising the integral form,

$$\int_{\Omega} \psi_t \, d\Omega + \oint_{\partial\Omega} \psi \mathbf{u} \cdot d\mathbf{n} = \int_{\Omega} \psi \nabla \cdot \mathbf{u} \, d\Omega, \quad (3)$$

where Ω represents an arbitrary control volume and \mathbf{n} is an outward pointing normal to its boundary $\partial\Omega$. In this work the control volumes are chosen to coincide with the mesh cells and the solution values are nominally stored at the centres of these cells. The quantity $\bar{\psi}$ is defined to be the integral average of ψ over the control volume, and it is this quantity which is considered to evolve in time, so

$$\left(\int_{\Omega} d\Omega \right) \bar{\psi}_t + \oint_{\partial\Omega} \psi \mathbf{u} \cdot d\mathbf{n} = \int_{\Omega} \psi \nabla \cdot \mathbf{u} \, d\Omega \approx \bar{\psi} \oint_{\partial\Omega} \mathbf{u} \cdot d\mathbf{n}. \quad (4)$$

The discretisation of the right hand side of (3) is a separate issue. Here the cell-average value $\bar{\psi}$ has been extracted and the divergence theorem subsequently applied to the integral of the divergence of the velocity field. This makes it easier to preserve balance between source terms and flux terms, since careful approximation of the flux

terms can now ensure that a constant scalar field is maintained by the numerical scheme whatever the velocity field.

Given that the control volume is polyhedral, the integrated equations (4) may be approximated by

$$\bar{\psi}_t + \frac{1}{V_\Omega} \sum_{m=1}^{N_f} A_m (\mathbf{f}_m^* \cdot \hat{\mathbf{n}})^n = \frac{1}{V_\Omega} \bar{\psi}^n \sum_{m=1}^{N_f} A_m (\mathbf{u}_m^* \cdot \hat{\mathbf{n}})^n, \quad (5)$$

in which N_f is the number of faces of the polyhedron, V_Ω is its volume, A_m are the face areas, $\hat{\mathbf{n}}$ their unit outward normals, and the superscripts \cdot^n and \cdot^* respectively represent the current time level of the approximate solution and the evaluation of some approximate integral-average value over a face (\mathbf{f}^* is commonly called the numerical flux). Note that so far no coordinate system has been chosen for the evaluation of these integrals. The time derivative in (5) is discretised using a simple forward Euler approximation: a more accurate approximation would be used with a higher order spatial discretisation. Note that from now on ψ (with appropriate sub- and superscripts) will be used in place of $\bar{\psi}$ for clarity, and will represent a cell-centre value of the solution.

We assume from now on that the equations are to be discretised on a three-dimensional structured grid, in which case Eq. (5) can be rewritten as

$$\begin{aligned} \psi_{ijk}^{n+1} &= \psi_{ijk}^n - \frac{\Delta t}{V_{ijk}} \times \{ [A(\psi u)^*]_{i+1/2jk}^n - [A(\psi u)^*]_{i-1/2jk}^n \\ &\quad + [A(\psi v)^*]_{ij+1/2k}^n - [A(\psi v)^*]_{ij-1/2k}^n \\ &\quad + [A(\psi w)^*]_{ijk+1/2}^n - [A(\psi w)^*]_{ijk-1/2}^n \\ &\quad + \psi_{ijk}^n [A u^*]_{i+1/2jk}^n - \psi_{ijk}^n [A u^*]_{i-1/2jk}^n \\ &\quad + \psi_{ijk}^n [A v^*]_{ij+1/2k}^n - \psi_{ijk}^n [A v^*]_{ij-1/2k}^n \end{aligned}$$

$$+ \psi_{ijk}^n [A w^*]_{ijk+1/2}^n - \psi_{ijk}^n [A w^*]_{ijk-1/2}^n \}, \quad (6)$$

in which (i, j, k) indexes the grid cells in the three coordinate dimensions, half indices indicate faces between cells and $(u, v, w)^T$ is the advection velocity in the appropriate coordinate system (later assumed to be spherical polars). The grid faces have also been assumed to be perpendicular to the axes of the coordinate system being used, which has simplified Eq. (6) considerably.

b. Dimensional splitting

The equations (5) are solved here using an operator splitting approach which means that the updates corresponding to each of the three space dimensions are carried out successively. Hence, instead of using Eq. (6), three updates are made,

$$\begin{aligned} \psi_{ijk}^* &= \psi_{ijk}^n - \frac{\Delta t}{V_{ijk}} \times \{ [A \psi^n u^{n+1/2}]_{i+1/2jk} - [A \psi^n u^{n+1/2}]_{i-1/2jk} \\ &\quad + \psi_{ijk}^n [A u^{n+1/2}]_{i+1/2jk} - \psi_{ijk}^n [A u^{n+1/2}]_{i-1/2jk} \} \\ \psi_{ijk}^{**} &= \psi_{ijk}^* - \frac{\Delta t}{V_{ijk}} \times \{ [A \psi^* v^{n+1/2}]_{ij+1/2k} - [A \psi^* v^{n+1/2}]_{ij-1/2k} \\ &\quad + \psi_{ijk}^n [A v^{n+1/2}]_{ij+1/2k}^n - \psi_{ijk}^n [A v^{n+1/2}]_{ij-1/2k}^n \} \\ \psi_{ijk}^{n+1} &= \psi_{ijk}^{**} - \frac{\Delta t}{V_{ijk}} \times \{ [A \psi^{**} w^{n+1/2}]_{ijk+1/2} - [A \psi^{**} w^{n+1/2}]_{ijk-1/2} \\ &\quad + \psi_{ijk}^n [A w^{n+1/2}]_{ijk+1/2} - \psi_{ijk}^n [A w^{n+1/2}]_{ijk-1/2} \}, \quad (7) \end{aligned}$$

where \cdot^* and \cdot^{**} represent the intermediate states of the update procedure. The superscripts in Eq. (7) should be noted carefully since they do not indicate a pure dimensional splitting approach; the velocity divergence terms are all evaluated at the initial time level (the velocity itself can be evaluated either at this time or part-way

through the time-step, as indicated by the $n + 1/2$ superscript), taking no account of the intermediate states. This is done so that the overall scheme is conservative when the discrete representation of the velocity divergence is identically zero (not necessarily true for an analytically defined divergence-free velocity field). In fact, the velocity field is never evaluated at the intermediate states, even within the flux terms. This means that as long as the value of ψ at each cell interface is defined to be the same within both the flux and the divergence terms in (7), any consistent approximation to the solution will maintain $\psi = \text{constant}$ indefinitely (even when the velocity field has non-zero divergence).

It is also the use of this splitting approach that necessitates solving the advection equation (1) rather than the equivalent conservation law in the case of a divergence-free velocity field. Essentially, a problem arises because even though the discrete velocity field might satisfy

$$\sum_{m=1}^{N_f} A_m (\mathbf{u}_m^* \cdot \hat{\mathbf{n}})^n = 0, \quad (8)$$

so it makes no contribution to (5), the individual components of (8) corresponding to the three coordinate directions are not necessarily zero and may therefore contribute to the updates given by Eq. (7). As a consequence, when the dimensional splitting is applied to the flux terms, compression or rarefaction can occur in any of the coordinate directions even when they cancel out overall. Since the solution is updated after each of the dimensional sweeps, the intermediate states (\cdot^* and \cdot^{**} in (7)) would differ from the solution at time level n if the divergence terms were not included in the intermediate updates. As a result, the three flux components

which appear in (7) would each be evaluated from different solution states and the required cancellation would no longer occur, a problem which wouldn't appear if an unsplit approach was taken.

Note that only the simplest approach to splitting is used here, which involves successive updates using the full time-step in each of the three coordinate directions. Strictly speaking this approach is only first order accurate in time (but still second order accurate in space), but this seems to have little detrimental effect on the quality of the solutions at the grid resolutions which have been used. Higher order accuracy may be achieved through the likes of Strang splitting (Strang 1968) but only at the expense of extra computational complexity.

c. The WAF approach

The numerical fluxes are approximated by following the WAF approach, originally proposed by Toro (1989) for the modelling of hyperbolic conservation laws, although this could easily be replaced by any other finite volume schemes which operate on a structured, logically rectangular grid. In one dimension the Weighted Average Flux (WAF) scheme for the conservation law

$$\psi_t + f(\psi)_x = 0, \quad (9)$$

(where, in this case, $f = \psi u$ is the flux of interest) discretised by the conservative scheme

$$u_i^{n+1} = u_i^n - \frac{\Delta t}{\Delta x} (f_{i+1/2}^* - f_{i-1/2}^*), \quad (10)$$

is defined by the numerical flux

$$f_{i+1/2}^* = \frac{1}{2}(1 + \phi_{i+1/2})f_i + \frac{1}{2}(1 - \phi_{i+1/2})f_{i+1}. \quad (11)$$

ϕ incorporates a flux limiter in the form of an amplification factor applied to the local CFL number, $\nu = u\Delta t/\Delta x$, and can be thought of as an internal parameter which locally tunes the numerical diffusion of the scheme. Putting $\phi = \nu$ returns the Lax-Wendroff scheme, which is second order accurate in both space and time but doesn't preclude spurious oscillations, while $\phi = 1$ gives first order upwinding, stable but very diffusive. In order to find an optimal combination of these two schemes (accurate and oscillation-free) $\phi(r, \nu)$ is defined to be a function not only of the local CFL number, but also of a local flow parameter $r = \Delta\psi_{\text{upwind}}/\Delta\psi_{\text{local}}$, which indicates where oscillations are liable to occur in the unlimited solution, most importantly at local extrema where $r < 0$. It is simple to derive constraints which must be satisfied by the limiter functions in order for the scheme to satisfy a Total Variation Diminishing (TVD) condition and the solution to be free of spurious oscillations (see (Sweby 1984) for a more detailed discussion). Many different limiters have been devised, the most commonly used, in order of decreasing numerical diffusivity, being minmod, van Leer and superbee. In this work the superbee limiter will always be used. This is defined by

$$\phi = \text{sgn}(\nu)[1 + (|\nu| - 1)b] \quad (12)$$

in which

$$b = \max[0, \min(2r, 1), \min(r, 2)]. \quad (13)$$

The WAF approach is being applied here to the advection equation rather than a conservation law, so the numerical flux is not used in precisely the form given by (11). Note that in moving from (6) to (7), the aforementioned balance between the flux derivatives and the velocity divergence has been preserved by assuming that the numerical flux takes the form $f^* = (\psi u)^* = \psi^* u^*$, *i.e.* the velocity is decoupled from the solution as it was on the right hand side of (5), so the WAF flux becomes

$$(\psi u)_{i+1/2}^* = \frac{1}{2}(1 + \phi_{i+1/2})u_{i+1/2}\psi_i + \frac{1}{2}(1 - \phi_{i+1/2})u_{i+1/2}\psi_{i+1}, \quad (14)$$

where $u^* = u_{i+1/2} = (u_i + u_{i+1})/2$ (with appropriate temporal indices). This is equivalent to (11) when the velocity field is constant. For a constant scalar field, $\psi \equiv \psi^*$, the use of (14) ensures that the overall update to the solution is zero, whatever the local velocities are. The expressions for the numerical fluxes, (14) and those related to the other two space dimensions, are substituted directly into (7) to give the WAF-type scheme used here, using the appropriate values of ψ .

d. *Extension to spherical geometry*

Until now, no mention has been made of the coordinate system in which Eq. (2) is to be approximated. For a spherical shell representing the earth's atmosphere, the most obvious choice is spherical polar rather than Cartesian coordinates. Within the framework defined above, this simply entails an appropriate definition of the cell volumes and face areas which are substituted into (7). For a regular longitude(λ)/latitude(θ)/height(r) structured grid the volumes are given exactly by

$$V_{ijk} = \frac{1}{3}\Delta\lambda [\sin \theta_{j+1/2} - \sin \theta_{j-1/2}] [3r_k^2\Delta r + (\Delta r)^3/4]. \quad (15)$$

The values of $\Delta\lambda$, $\Delta\theta$ and Δr are taken to be constant over the whole computational domain so they are fixed from the start of the calculation. The face areas satisfy a similar property when they are calculated in the spherical polar coordinate system, and are given by

$$\begin{aligned} A_{i+1/2jk} &= r_k \Delta\theta \Delta r \\ A_{ij+1/2k} &= \cos \theta_{j+1/2} r_k \Delta\lambda \Delta r \\ A_{ijk+1/2} &= r_{k+1/2}^2 \Delta\lambda [\sin \theta_{j+1/2} - \sin \theta_{j-1/2}]. \end{aligned} \quad (16)$$

Note that these areas are all independent of longitude λ , so only two-dimensional arrays of values have to be calculated and stored. Although the presentation given here is slightly different, the resulting scheme is precisely the same as would be derived by applying the finite volume method to the transformed differential equations written in curvilinear (spherical polar) coordinates.

The WAF scheme as presented above was originally derived in a manner designed to solve problems on regular Cartesian grids. In spherical polar coordinates when a regular longitude/latitude/height grid is used there is a convergence of the grid lines towards the poles which creates a smooth distortion of the grid. Even so, the WAF-type flux given by (14) is used without any modification, except that the local dual cell CFL number used in the calculation of ϕ is now taken to be

$$\nu_{i+1/2} = u_{i+1/2} \Delta t \frac{2A_{i+1/2}}{V_i + V_{i+1}}, \quad (17)$$

to account for the variation in the cell sizes.

The problem which occurs due to the singularities in both the grid and the underlying coordinate system at the north and south poles is dealt with naturally

by using a cell-centre flux-based approach, and it is not necessary to leave small holes in the grid at the poles which cut the offending region out. The reason for this can be seen by considering the dimensional sweep in the south-north direction: the other two sweeps remain unaffected because they do not involve crossing the singularity, although the coordinate system does vary from cell to cell so there may be some deterioration in the accuracy of the approximation. The sweep over the poles (illustrated by the double arrow in Fig. 1), leaving aside superscripts for the moment, takes the form

$$\begin{aligned} \psi_{ijk} \rightarrow \psi_{ijk} - \frac{\Delta t}{V_{ijk}} \times \{ [A(\psi v)^*]_{ij+1/2k} - [A(\psi v)^*]_{ij-1/2k} \\ + \psi_{ijk} [A v^*]_{ij+1/2k} - \psi_{ijk} [A v^*]_{ij-1/2k} \}. \end{aligned} \quad (18)$$

At first glance this appears to require communication to take place across the poles in order to calculate the fluxes necessary to update the neighbouring cells. However, the face areas $A_{ij\pm 1/2k}$ which coincide with the pole are zero (the corner of a shaded triangle in Fig. 1) so the flux through these edges is nonexistent. Thus the only cross-polar communication required is in the retrieval of the data to substitute into the dummy cells at the poles, which are considered as though they were outer boundaries of the grid. The data can simply be taken from the cells directly opposite, across the pole, remembering that the velocity field must be negated (also shown in the figure). The cell-centre representation used here ensures that no solution value is actually stored at either pole. The only requirement is that there is an even number of coarse grid cells in the zonal direction. The overall effect is precisely that of having a single layer of prismatic (triangular in two dimensions) cells surrounding

the poles, although they are actually degenerate hexahedra. The boundaries in the zonal direction are considered to be periodic, while the upper and lower surfaces are currently treated with a very simple (extrapolated) transmissive boundary condition.

3. Adaptive Mesh Refinement

The Adaptive Mesh Refinement (AMR) algorithm of Berger and Olinger (1984) and Berger and Colella (1989) provides a straightforward process by which the resolution of structured Cartesian meshes can be adapted locally to efficiently give the desired accuracy when approximating partial differential equations. The Cartesian AMR code, which solves the compressible Euler equations in a rectangular box, and on which this work has been built can be obtained via Berger (2002).

The method employs a hierarchical system of grids, the coarsest being represented by G_0 , *cf.* Fig. 2, which is shown as a two-dimensional representation for clarity. For simplicity, in this work the coarse grid consists of a single mesh, the term ‘mesh’ being used to define a three-dimensional, logically rectangular patch of cells. Each of the finer grid levels (G_1 , G_2 , etc.) may consist of any number of meshes, each of which is embedded in the meshes of the coarser level grid. The fine meshes are constructed so that no two meshes within the same grid are allowed to overlap. The relative level of refinement between two grid levels can take any positive integer value (even 1) and need not be the same in any of the three coordinate directions. In its application to atmospheric flows this is particularly useful because the refinement in the vertical direction can be made independent of that in

the horizontal directions.

The application of AMR within a spherical shell (for global atmospheric modelling) requires little modification from the standard procedure. A single coarse mesh is used to cover the whole of the sphere via a transformation into spherical polar coordinates. (In the presence of variable orography, which is not dealt with here, either a ‘cut-cell’ approach can be implemented, which will fit the grid to the boundary defined by the earth’s surface, or the calculation can be carried out using an appropriate surface following vertical coordinate.) The only changes seen by the AMR algorithm when such a transformation is applied are in the interpolation required for communication between coarse and fine grid levels and the application of the boundary conditions at the poles. Initial conditions are either defined analytically or interpolated trilinearly from externally supplied data on to the initial adapted grid. Similar interpolation (trilinear in space, linear in time) is applied to externally supplied velocity fields which may be used to force the advective flow. This is done for simplicity and robustness.

AMR uses a sequence of embedded regular grids to create the discretisation (see Fig. 2) so it is simple to apply an adaptive time-stepping strategy. This avoids unnecessarily prohibitive restrictions on the time-step used in the integration of the equations on the coarser grids, which not only wastes cpu time but can also degrade the accuracy. The resulting sequence of events is most easily illustrated by example, and the following two-dimensional situation is taken from Quirk and Karni (1996).

Given a three level grid hierarchy G_0, G_1, G_2 with refinement ratios of 2 and 4 in both coordinate directions, the AMR scheme carries out the following series of grid

operations during a single integration of the equations on the coarsest grid G_0 :

Integrate	:	G_0 (with time-step Δt)
Integrate	:	G_1 (with time-step $\Delta t/2$)
Integrate	:	G_2 (4 times with time-step $\Delta t/8$)
Project	:	G_2 on to G_1
Adapt	:	G_2
Integrate	:	G_1 (with time-step $\Delta t/2$)
Integrate	:	G_2 (4 times with time-step $\Delta t/8$)
Project	:	G_2 on to G_1
Project	:	G_1 on to G_0
Adapt	:	G_2
Adapt	:	G_1

First, the coarse grid integration is carried out (for simplicity all solution values are updated and stored even where they are overlaid by a more accurate solution on a finer grid) with the coarse grid time-step Δt and then the next finer grid G_1 is integrated once with half this time-step. Now the finest grid level G_2 has been reached and this is integrated four times with a further refined time-step $\Delta t/8$ to advance the solution to the same time level as grid G_1 . The next step is to project the solution on the finest grid on to the coarser grid (overwriting the underlying coarse solution with volume-weighted, conservative cell-averages of the fine grid values) to give a more accurate representation of the solution on grid G_1 at the current time level. The finest grid is then adapted (the meshes at this level are completely

recalculated according to the flagging and clustering process described below) before a second complete integration of the middle level grid G_1 is carried out. The fine grid solutions are projected successively on to the coarser grids at the new coarse grid time level, to ensure that each grid has the most accurate information available to it before both of the fine grids, G_2 and G_1 , are adapted. This cycle is repeated until the required time level has been reached.

Each mesh is surrounded by a double layer of dummy cells, as shown by the dashed and dotted lines in Fig. 3, (more would be needed for a numerical scheme with a larger stencil) which, prior to each grid integration, are primed with information about the solution around the boundary of the mesh. This allows the mesh boundaries to be invisible to the mesh integrator being used by ensuring that no special treatment is required beyond the overwriting of the solution in the dummy cells. Most of these mesh boundaries are likely to lie within the computational domain, in which case the solution in the dummy cells is taken from the finest mesh underlying those cells (by direct copying if the underlying mesh has the same level of refinement, or by interpolation - minmod limited MUSCL-type (van Leer 1979) in space, linear in time - if the mesh is one level coarser: no other possibility is allowed). The minmod limiter is used because it is the only one which guarantees a monotonicity preserving interpolant for *arbitrary* refinement ratios. When the mesh boundaries coincide with domain boundaries the dummy cells are supplied with information appropriate to the specified boundary type. For spherical shell calculations this involves transferring information from other parts of the grid for the periodic zonal and polar meridional boundaries (simply a matter of extra book-keeping appended

to the method used for internal boundaries) and a simple transmissive or reflective condition at the inner and outer shells.

It should be noted that for the scheme to remain conservative across interfaces between fine and coarse meshes, and to maintain consistency with the projection of the solution from fine to coarse meshes, it is necessary to overwrite the coarse grid fluxes with the corresponding fine grid fluxes accumulated while advancing the solution to the next coarse grid time level. This is equivalent to applying the difference between the two sets of fluxes over the given time-step to recover conservation. It is applied at all cell faces where there is a change in mesh resolution, and ensures that the numerical flux entering/leaving a coarse cell through a fine/coarse interface over a coarse time-step is the same as that leaving/entering the adjacent fine cells over the same time interval.

The adaptation step is designed to ensure that the fine meshes (which give the high grid resolution) follow the movement of the flow features they are intended to resolve. It is simple to construct monitor functions which detect phenomena of interest. A number of options have been used for the definition of the monitor, the simplest being

$$\xi_{ijk} = \max(|\psi_{i+1jk} - \psi_{i-1jk}|, |\psi_{ij+1j} - \psi_{ij-1k}|, |\psi_{ijk+1} - \psi_{ijk-1}|), \quad (19)$$

which considers only the local cell-to-cell difference of the tracer values, so the mesh is refined where the solution is changing rapidly. It is simple to implement alternative monitors, which can be based on quantities such as local vorticity, a predefined range of the tracer values, the geometry of the computational domain, or

more sophisticated error estimates of the solution. For the purposes of robustness, the flagged cells are augmented by flagging additional layers of buffer cells to move the fine/coarse interface further away from the region of interest.

During adaptation, mesh cells are flagged for refinement when the monitor function ξ exceeds a specified value. In order to ensure that the adapted meshes remain properly nested, *i.e.* every cell is underlain by a mesh at every coarser grid level, a cell is also flagged if it is about to be overlaid by a newly created finer mesh of any level. The flagged cells are ‘clustered’ into new rectangular meshes (Bell *et al.* 1994) which are constructed in a manner which completely covers the flagged region, subject to a specified tolerance η which dictates the proportion of unflagged cells allowed in each new mesh. Setting this parameter to 1 ensures that the new meshes contain only flagged cells, while a value of 0 will create a single mesh which constitutes the smallest rectangle covering the flagged cells. The parameter η is chosen to provide a balance between creating many small meshes (with the consequent communication overhead between them) and few large meshes (which contain a large proportion of redundantly refined cells) and typically takes a value of between 0.6 and 0.9.

The clustering is carried out according to the method described by Bell *et al.* (1994), in which the smallest rectangular block covering all of the flagged cells is subdivided and tested recursively. This could be done by simple bisection but here their more sophisticated algorithm is used to reduce the tendency of the grid generation to create large numbers of very small meshes, which would be very inefficient.

The technique involves the calculation of the signatures of a cluster,

$$\sigma_i = \left(\sum_{jk} \text{flags} \right)_i, \quad \sigma_j = \left(\sum_{ik} \text{flags} \right)_j, \quad \sigma_k = \left(\sum_{ij} \text{flags} \right)_k, \quad (20)$$

which simply count the number of flagged cells in each plane of the three-dimensional rectangular block. If any signature contains a zero then this indicates a possible partition. Otherwise, a plane (of cell faces) is sought where the second difference of the signature passes through zero and the plane with the largest such crossing determines the partition (or the closest to the centre of the block if two are of equal strength). The new clusters are taken to be the smallest enclosing rectangular blocks covering the flagged cells in each partition, and they are accepted as new meshes if they satisfy the clustering criterion defined by η . This is repeated recursively until every cluster satisfies the criterion (bisection in the longest direction being used if all else above fails) and the result gives the new mesh structure. This minimises the size of the fine grid, subject to this clustering criterion.

Having adapted the mesh, the solution is transferred from the old meshes to the new ones, by copying it across directly in the regions where the new mesh occupies the same space as an old one of the same level, or by once more using minmod limited MUSCL-type interpolation from the underlying coarser grid (for monotonicity). The old meshes can now be discarded and the corresponding computer memory made available for further calculations.

Note that the remeshing process, as presented, is global in nature and hence may not be particularly suited to parallelisation. This is not necessary: modifications can be made which apply the remeshing on a mesh-by-mesh basis and, although

some monitoring would still be needed to balance the loads between processors as the calculation proceeds, this is likely to be far outweighed by the savings achieved through the use of the adaptation, especially when more species are included in the model. Further details are given in Quirk (1996) and Rendleman *et al.* (2000).

a. *Issues on the Sphere*

The application of the adaptive mesh refinement algorithm remains essentially unchanged when applied to advection over a sphere. The only modifications required are

- The transfer of solution information from fine meshes to coarse meshes must be carried out through a volume weighted averaging procedure to ensure conservation, since the cell volumes now vary over the grid. This is simplified considerably when exact expressions are available for the cell volumes, and since these do not vary in the zonal direction they can be stored efficiently in two-dimensional arrays.
- The transfer of solution information from coarse meshes to fine meshes (used to fill both the dummy cells surrounding the mesh boundaries and the newly created mesh structure) is conservative when piecewise constant interpolation is used for the solution, but a MUSCL-type (or one of higher order) reconstruction must allow for the possibility that the integral-average of the linear reconstruction is not necessarily equal to the reconstructed solution value at the cell's centre.

Here, a standard MUSCL reconstruction (passing through the cell-average value at the cell centre) is found, leading to cell-centre values for the overlying fine cells. These are then scaled by the ratio of the total ‘mass’ in the fine cells to the ‘mass’ in the coarse cell. The scaled reconstruction is no longer guaranteed to be monotonic, but the overall procedure will not change the sign of the solution. In the future a form of interpolation which is inherently both conservative and monotonic over spherical geometries will be sought.

- The communication between meshes through layers of dummy cells which span either of the poles must take into account the reversal of the coordinate system across the singularity. As a result the velocity field swaps sign in the dummy cells and any filling of strips of cells is carried out in reverse. Otherwise, the scheme is simply that which is applied on a single mesh. Zonal boundaries require periodic conditions to fill the dummy cells.

4. Numerical Experiments

a. *Idealised case studies*

i) SOLID BODY ROTATION

The first problem used here to test the code is essentially the two-dimensional solid body rotation test case of Williamson *et al.* (1992), in which a cosine bell shaped profile is advected, without distortion, around the sphere at various angles to the

equator. The velocity components of the advecting wind field are given by

$$\begin{aligned} u &= u_0(\cos \alpha \cos \theta + \sin \alpha \cos \lambda \sin \theta) \\ v &= -u_0 \sin \alpha \sin \lambda \end{aligned} \tag{21}$$

where α is the angle between the axis of solid body rotation and the polar axis of the spherical coordinate system. The period of the rotation is taken to be 12 days, and the radius of the earth to be 6371.22km, giving a value of $u_0 \approx 38.61\text{ms}^{-1}$. The initial scalar distribution is of the form

$$\psi = \begin{cases} (1 + \cos(\pi r/R))/2 & \text{if } r < R \\ 0 & \text{if } r \geq R \end{cases} \tag{22}$$

where

$$r = \cos^{-1}[\sin \theta_c \sin \theta + \cos \theta_c \cos \theta \cos(\lambda - \lambda_c)] \tag{23}$$

is the great circle distance between (λ, θ) and the bell centre, initially taken as $(\lambda_c, \theta_c) = (\pi/2, 0.0)$. The bell radius R is set to $7\pi/64$, as proposed by Williamson and Rasch (1989), and the error is measured using the standard l_1 , l_2 and l_∞ norms, as defined in Williamson *et al.* (1992). All numerical experiments were run with a maximum Courant number of 0.9.

The AMR requires specification of certain control parameters which dictate the outcome of the remeshing procedure. In this case the condition $\xi \geq 0.1$ (using (19) as the monitor ξ) was used for the cell flagging, including one additional layer of buffer cells, and the remeshing was carried out every time-step. The cell clustering parameter η is taken to be 0.9 so a high proportion of cells in the fine grids have been flagged.

Table 1 shows the errors obtained for each of the four angles of rotation suggested in Williamson *et al.* (1992) and indicates the significant savings can be made in terms of cpu time, particularly in a problem such as this where a relatively small region of the domain is covered by fine meshes. The corresponding results obtained after a single revolution for $\alpha = \pi/2$ are also shown in Figure 4. It should also be noted that the results on the 128×64 (2.8125°) uniform mesh compare favourably with others presented in the literature (see, for example, Williamson and Rasch (1989), Rasch (1994), Li and Chang (1996), Lin and Rood (1996) and Nair *et al.* (1999)), despite the fact that the scheme presented here has a lower order of accuracy than most of these other schemes. Of particular interest is the fact that there is no discernible distortion of the profile as it passes over the poles (see Fig. 4).

The speed is increased, in the case $\alpha = \pi/2$, by a factor of nearly 5 when a single level adaptive grid (factor of 2 refinement) is used in place of a fixed 128×64 grid, which has the same effective resolution. This improves to a factor of 22 when a doubly refined (factors of 2 and 3) 64×32 mesh is used in place of a fixed 384×192 mesh. In three dimensions these factors can be even larger. When the adaptive mesh is tracking a relatively small-scale feature it is commonplace to be able to decrease the cpu time necessary for a calculation by a factor of 50 or more simply through the use of AMR. There are two main reasons for these huge improvements: the reduction in the number of computational cells on which the integration is carried out, and the relative increase in the stable time-step for the period of the adaptive calculation when the fine grid does not cover either of the poles (increased by the fact that the coarse grid cell centres at the poles are further away from the singularity

and therefore provide a less stringent bound on the time-step).

The cpu times in the table are obtained on a SUN ULTRA 10 workstation, but it should be emphasised that the method used makes no use of any additional acceleration techniques to counteract the prohibitive restriction on the time-step imposed on the scheme by the narrow cells close to the poles. The efficiency of the method can be improved dramatically by the simple technique of applying multiple sweeps in the zonal direction during the operator splitting procedure, increasing the number of sweeps as the poles are approached.

The effect is illustrated in Table 2, which shows a series of results obtained on the 384×192 grid. The level of reduction referred to in the first column, is related to the idea of a reduced grid, commonly used in meteorology (as in Rasch (1994)), and indicates the number of times the grid would have been ‘reduced’ by a factor of 2 in the zonal direction to gain an increase in the global time-step equivalent to that gained by using multiple sweeps as part of the dimensional splitting. A level of n indicates that cells near to the poles are updated 2^n times by zonal sweeps during each full time-step in which a single sweep is made at the equator. The number of sweeps decreases from the poles to the equator by a factor of 2 at equispaced intervals of $\cos \theta$ (θ being latitude), except for the case denoted ‘optimal’ where the number of sweeps can take any positive integer value, not just a power of 2, and is calculated separately for every layer of cells. It is actually taken to be the maximum of 1 and integer part of $\cos^{-1} \theta$. Initially the accuracy of the solution actually increases with the number of sweeps, probably because when the cosine bell is away from the poles it is being advected with a higher CFL number, but

eventually the error introduced by the multiple sweeping starts to dominate. Also the extra sweeps cease to be beneficial when the time-step is no longer restricted by the cells near to the pole and the calculations start to take longer again. Even so, the loss of accuracy suffered with the optimal sweeping strategy is extremely small in comparison with the huge increase in the time-step and the consequent factor of nearly 50 by which the cpu time of the calculation is reduced. The simple multiple sweeping used here produces less speed-up on adapted grids because the adaptation (along with adaptive time-stepping) is often doing the same job, *i.e.* giving a lower effective grid resolution at the poles.

A slight alteration to this test case can be used to show that the solution is not distorted as it traverses the poles. The non-zero region of the initial cosine bell profile is simply replaced by a solution value of $\psi = 1$ (giving a piecewise constant solution since it is zero elsewhere), and this is advected over the poles ($\alpha = \pi/2$). Given a 64×32 coarse mesh and two levels of refinement, by a factor of 4 each time, the results in Figure 5 show meshes and solution contours corresponding to $\frac{1}{8}$, $\frac{3}{16}$ and $\frac{1}{4}$ of a revolution, together with some solutions at other times. The polar singularity has no visible effect on the quality of the solution which is unperturbed from its initial value of 1 across the circle (to machine accuracy).

Note that although the grid appears to contain triangular cells, this is simply an artefact created by converting the hierarchy of structured rectangular meshes into a single unstructured grid whose vertices correspond to the original cell centres. This is done to simplify the application of the graphical tools to the data. The resulting grid is uniquely determined but some of the rectangular elements degenerate into

triangles. A similar effect is seen with this conversion in three dimensions, where degenerate hexahedra can appear.

ii) DEFORMATIONAL FLOW – IDEALISED CYCLOGENESIS

The second test case presented here is a generalisation of the idealised cyclogenesis problem of Doswell (1984) to spherical geometry, as presented by Nair *et al.* (1999). Given a rotated coordinate system (λ', θ') with its north pole at (λ_0, θ_0) with respect to the regular spherical coordinate system (λ, θ) , a steady circular vortex is defined by zero normal velocity $v' \equiv d\theta'/dt$ and tangential velocity

$$u'(\theta') \equiv \cos \theta' \frac{d\lambda'}{dt} = \frac{3\sqrt{3}}{2} a \operatorname{sech}^2(\gamma\rho) \tanh(\gamma\rho), \quad (24)$$

where a is the radius of the sphere,

$$\rho(\theta') = \frac{2 \cos \theta'}{1 + \sin \theta'}, \quad (25)$$

and $\gamma = 3/2$ is a stretching parameter that controls the length scale of the vortex. ρ can be interpreted as the distance from the north pole of the polar stereographic projection of the point (λ', θ') . The amplitude of the vortex has been normalised to have a maximum tangential velocity of unity, which occurs when

$$\theta' = 2 \tan^{-1} \left(\frac{\gamma - c}{\gamma + c} \right), \quad (26)$$

where

$$c = \frac{1}{4} \ln \left(\frac{\sqrt{3} + 1}{\sqrt{3} - 1} \right) \approx 0.3292395. \quad (27)$$

The initial conditions are taken to be

$$\psi(\lambda', \theta', 0) = -\tanh \left[\frac{\rho}{\delta} \sin \lambda' \right], \quad (28)$$

in which $\delta = 0.01$ is the characteristic width of the frontal zone. This problem has an analytic solution, which is given by

$$\psi(\lambda', \theta', t) = -\tanh\left[\frac{\rho}{\delta}\sin(\lambda' - \omega t)\right], \quad (29)$$

where

$$\omega(\theta') = \frac{u'(\theta')}{a \cos \theta'} \quad (30)$$

is the angular velocity. Further details of the stereographic projection in the rotated coordinate system can be found in Nair *et al.* (1999), which also presents results with which these can be compared.

The experiment was run for 2.5 time units and the vortex was centred on a latitude of 65.24N, which places the north pole at a position where the wind speed is at its greatest, and at a longitude of 90.0E. Note that the region of high solution gradient is too narrow (in terms of grid cells) for the scheme to exhibit the appropriate order of accuracy. In spite of this, the error estimates in Table 3 show an encouraging rate of decrease in the l_1 error as the grid is refined. The improvements are less noticeable in the l_∞ error, which is effectively a local, pointwise, error measure and so is the most sensitive to small perturbations in the solution at very steep gradients. The limiter is also active in the critical region, where the method perceives a discontinuity at this grid scale, so the scheme is, at best, first order here. The increase in speed expected of the AMR (used here with $\xi \geq 0.2$ and $\eta = 0.9$) is also less visible here (see Table 3), because a larger proportion of the mesh has been refined and the region of fine mesh invariably covers the north pole, where the velocity field has its highest magnitude. The mesh generation has a significant overhead in this

simple case of a single scalar advection equation, which is only partly mitigated by the improved efficiency of the resulting grid, but the adaptive meshing would show its value if more physical and (particularly) chemical processes are included in the model. Also, it should be noted that refining the mesh by a factor of 3 rather than 2 gives a much greater improvement in speed, even in this simple case, with no loss of accuracy.

Figure 6 shows the solution obtained for $t = 2.5$ on the 64×32 ($\times 2 \times 3$) grid, as projected stereographically on to the plane tangent to the vortex centre. The same figure also shows the corresponding grid and illustrates how the refinement follows the regions of high solution gradient picked out by the monitor function (19). The triangular cells appear for the same reason as before. Close examination of the solution reveals no spurious overshoots or undershoots ($\psi_{\max} = 1$ and $\psi_{\min} = -1$) and the steep interface is captured very sharply, in the correct position.

5. A three-dimensional off-line model

A small amount of additional work is necessary for modelling the evolution of data obtained from meteorological analyses. The main requirement is the interpolation of the data supplied for such problems on to the computational mesh. This is used to set up both the initial conditions and the velocity field with which the off-line calculation is being forced. Trilinear interpolation is used in space along with linear interpolation in time (where necessary).

The method has been used to model both two- and three-dimensional meteoro-

logical datasets, of which one is presented below. Detailed analysis of the results is left to future publications. Note that where the radial (vertical) coordinate is not altitude, potential temperature or pressure for example, an assumption has been made that the atmosphere is very shallow compared to the radius of the earth, so cell geometries have been calculated without varying their size in the radial direction.

a. *A case-study of stratosphere/troposphere exchange*

The test case studied here covers a stratosphere-troposphere exchange event which took place over the North Atlantic during June 1996. Atmospheric data from the ERA dataset was used to provide three-dimensional velocity fields and values of potential vorticity at intervals of 6 hours for the period between 0600 17th June and 1800 23rd June. The data was given on a T107 (320×160) Gaussian longitude/latitude grid in the horizontal with 17 isentropic surfaces in the vertical (so diabatic heating is used as the vertical velocity component), at 10K intervals between 240K and 400K (although the numerical calculation only used the data between 300K and 360K). Note that in some regions these isentropic surfaces intersect the earth's surface, so no data is available there. For simplicity these regions have been assigned a nominal (constant) value of ψ and zero velocity, but are still included in the calculation. This is bound to contaminate the later solution to some extent but over the time scale considered here it does not interfere with the strong features higher up in the atmosphere which have been used to validate the code. This problem will be dealt with in the future either by considering data on a terrain following coordinate system or applying a cut-cell approach to the lower boundary

of the domain.

This test case is simply being used as part of the initial validation process for the numerical model, to provide evidence that the scheme can approximate advection processes in practical meteorological calculations. For this reason potential vorticity (PV) has been chosen as the advected variable, it being less sensitive to other physical and chemical processes than most chemical species. Even so, only the strongest features would be expected to be modelled accurately.

Figures 7 and 8 can be used to compare the evolution of (PV) predicted by the numerical experiment with that of the supplied data. Equispaced contours of PV are shown and the mesh is shaded according to the local value of PV. The computational results have been obtained on a $480 \times 240 \times 24$ mesh, initialised by interpolation of the potential vorticity field of the given ERA dataset at 0600 17th June. The figures focus on a particular vortex structure which has formed and is shown at 1200 21st June. One quadrant of the northern hemisphere is shown, from which a chunk of the grid has been removed to reveal a two-dimensional surface which slices through the feature at the 325K level. It is readily seen that the shape and size of the feature shows encouraging agreement between the two figures. It is more sharply defined in the numerical results, partly because they are calculated on a finer mesh than the meteorological data is supplied on. Exact correspondence will not happen anyway, because the numerical model is simply of advection, and ignores all the other physical and chemical processes. The comparison is also affected by the unreliable nature of the vertical velocity field, which is particularly difficult to measure, and possibly also by the simple modelling of the upper and lower boundaries.

Figure 9 shows the results obtained for the same case using the full adaptive algorithm. The computational results have been obtained on an underlying coarse $120 \times 60 \times 8$ mesh with two levels of refinement, the first by a factor of 2 in each dimension and the second by a factor of 3. A value of 0.9 was again used for η but a slightly different monitor was used, given by

$$\xi_{ijk} = \frac{\max(|\psi_{i+1jk} - \psi_{i-1jk}|, |\psi_{ij+1k} - \psi_{ij-1k}|, |\psi_{ijk+1} - \psi_{ijk-1}|)}{\max(1, \sqrt{\psi_{ijk}})}, \quad (31)$$

where ψ represents PV here. The grid cells flagged when this expression satisfies $\xi \geq 1.0$, along with the condition $1 \leq \psi \leq 10$, which ensures that the grid is only adapted in the northern hemisphere. The shape and size of the feature agrees closely with both the observed and experimental results shown in Figures 7 and 8. The denominator includes $\sqrt{\psi}$ to ensure that the monitor can pick out steep gradients of the solution as PV decreases in magnitude towards the equator, without introducing too many extra mesh cells.

The important aspect of these results is that the AMR has placed a fine mesh over the regions where PV is varying rapidly (so the AMR can be thought of as tracking the tropopause) but in doing so does not distort the solution, allowing instead a finer structure to be seen than on the coarser meshes. The solution is not notably different to that obtained on a uniform mesh of equivalent fine resolution ($720 \times 360 \times 48$), however the adaptive solution is obtained using one sixth of the cpu time and, at any given time, involves fewer than one quarter of the number of cells of the uniform mesh calculation (the size of the adapted grid varies between 2.7 and 3.1 million cells during the experiment).

6. Conclusions and Future Work

In this paper a dimensionally split WAF scheme has been modified to approximate the three-dimensional scalar advection equation and combined with an adaptive mesh refinement algorithm for use on a spherical geometry. An appropriate discrete form has been presented for the divergence of the velocity field which is used to balance the flux terms and ensure that a constant scalar field is maintained indefinitely by the computational scheme, whatever the velocities.

The basic (unadapted) numerical scheme has been applied on a uniform longitude/ latitude/height grid but, since a cell-centred flux-based numerical scheme is used, no special treatment is required at the pole beyond extra care being necessary in the application of the boundary conditions. The tight restriction imposed on the time-step caused by the convergence of the grid towards the poles is alleviated by a simple technique which applies multiple sweeps within each time-step in the zonal direction.

The WAF scheme, combined here with a superbee limiter, is shown to accurately model scalar advection without creating spurious oscillations in the solution for a variety of test cases for which exact solutions are available. This is the case for both solid body rotations and deformational flows. In some cases the accuracy of the solution was comparable to that of other schemes presented elsewhere in the literature with a higher formal order of accuracy. The method has also been used to model the evolution (through advection) of data supplied through meteorological analyses, as would be used in off-line chemical transport models. It successfully

predicts the advection of potential vorticity, retaining the main features detected in the supplied data.

The addition of adaptive mesh refinement has no detrimental effect on the accuracy of the solution on the finest grid level compared with a single grid of equivalent resolution. However, the AMR allows such solutions to be obtained with much greater efficiency. The improvement is particularly dramatic in three dimensions, and when the flow features of interest are clearly defined. However, it is also simple to reduce cpu times by a factor of 5 or more when modelling the evolution of typical meteorological data.

The benefits of using AMR will increase with the introduction of additional physical and chemical processes into the model. There is only a single overhead associated with constructing the grid at each time-step, and this will decrease in significance when large numbers of tracers are used, or when chemistry is included. Any chemical process which involves solving an ordinary differential equation locally within each grid cell is simple to implement and, since this is usually the stage which dominates the speed of the calculation, this is where AMR will be most valuable. The work will also be extended to encompass the approximation of dynamical models of the atmosphere, initially using a two-dimensional shallow water model, followed by more sophisticated three-dimensional situations. These extensions should allow the Godunov-type methodology and the AMR should show their true value.

Acknowledgements

The authors would like to thank Prof. Marsha Berger of the Courant Institute, New York, for supplying the Cartesian ‘box’ AMR code on which this work has been based, UGAMP for supplying funding for this work, and Dr. Paul Berrisford of the University of Reading for supplying the ERA data.

References

Behrens, J., 1996: An adaptive semi-Lagrangian advection scheme and its parallelization. *Mon. Weather Rev.*, **124**, 1386–2395.

Behrens, J., 1998: Atmospheric and ocean modeling with an adaptive finite element solver for the shallow-water equations. *Appl. Numer. Math.*, **26**, 217–226.

Bell, J., M. Berger, J. Saltzman, and M. Welcome, 1994: Three-dimensional adaptive mesh refinement for hyperbolic conservation laws. *SIAM J. Sci. Comput.*, **15**, 127–138.

Berger, M. J., cited 2002: Adaptive mesh refinement software for hyperbolic conservation laws. [Available online at <http://cs.nyu.edu/cs/faculty/berger/amrsoftware.html>]

Berger, M. J., and P. Colella, 1989: Local adaptive mesh refinement for shock

hydrodynamics. *J. Comput. Phys.*, **82**, 67–84.

Berger, M. J., and J. Oliger, 1984: Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.*, **53**, 482–512.

Chipperfield, M. P., M. L. Santee, L. Froidevaux, G. L. Manney, W. G. Read, J. W. Waters, A. E. Roche, and J. M. Russell, 1996: Isentropic 3D chemical transport model and comparison with UARS data in southern polar vortex, September 1992. *J. Geophys. Res.*, **101**(D13), 18861–18881.

Colella, P, 1990: Multidimensional upwind methods for hyperbolic conservation laws. *J. Comput. Phys.*, **87**, 171–200.

Doswell, C. A., 1984: A kinematic analysis of frontogenesis associated with a non-divergent vortex. *J. Atmos. Sci.*, **41**, 1242–1248.

Easter, R. C., 1993: Two modified versions of Bott’s positive-definite numerical advection scheme. *Mon. Weather Rev.*, **121**, 297–304.

Giraldo, F. X., 2000: The Lagrange-Galerkin method for the two-dimensional shallow water equations on adaptive grids. *Int. J. Numer. Meth. Fluids*, **33**(6), 789–832.

Leonard, B. P., 1979: A stable and accurate convection modeling procedure based on quadratic upstream interpolation. *Comput. Methods Appl. Mech. Eng.*, **19**, 59–98.

Leonard, B. P., A. P. Lock, and M. K. MacVean, 1996: Conservative explicit unrestricted-time-step multidimensional constancy-preserving advection schemes. *Mon. Weather Rev.*, **124**, 2588–2606.

Leonard, B. P., M. K. MacVean, and A. P. Lock, 1995: The flux integral method for multidimensional convection and diffusion. *Appl. Math. Model.*, **19**, 333–342.

Leonard, B. P., and H. S. Niknafs, 1991: Sharp monotonic resolution of discontinuities without clipping of narrow extrema. *Comput. Fluids*, **19**, 141–154.

Li, Y., and J. S. Chang, 1996: A mass-conservative, positive-definite, and efficient Eulerian advection scheme in spherical geometry and on a nonuniform grid system. *J. Appl. Meteorol.*, **35**, 1897–1913.

Lin, S. -J., and R. B. Rood, 1996: Multidimensional flux-form semi-Lagrangian transport schemes. *Mon. Weather Rev.*, **124**, 2046–2070.

Nair, R., J. Côté, and A. Staniforth, 1999: Cascade interpolation for semi-Lagrangian advection over the sphere. *Q. J. R. Meteorol. Soc.*, **125**, 1445–1468.

- Petshek, A. G., and L. D. Libersky, 1975: Stability, accuracy and improvement of Crowley advection schemes. *Mon. Weather Rev.*, **103**, 1104–1109.
- Plumb, R. A., D. W. Waugh, R. J. Atkinson, P. A. Newman, L. R. Lait, M. R. Schoeberl, E. B. Browell, A. J. Simmons, and M. Lowenstein, 1994: Intrusions into the lower stratospheric Arctic vortex during the winter 1991–1992. *J. Geophys. Res.*, **99**(D1), 1089–1105.
- Prather, M. J., 1986: Numerical advection by conservation of second-order moments. *J. Geophys. Res.*, **91**, 6671–6681.
- Quirk, J. J., 1991: *An adaptive grid algorithm for computational shock hydrodynamics*, PhD thesis, College of Aeronautics, Cranfield Institute of Technology.
- Quirk, J. J., 1996: A parallel adaptive grid algorithm for computational shock hydrodynamics. *Appl. Numer. Math.*, **20**(4): 427–453.
- Quirk, J. J., and S. Karni, 1996: On the dynamics of shock-bubble interaction. *J. Fluid Mech.*, **318**, 129–163.
- Rasch, P. J., 1994: Conservative shape-preserving two-dimensional transport on a spherical reduced grid. *Mon. Weather Rev.*, **122**, 1337–1350.

Rendleman, C. A., V. E. Beckner, M. Lijewski, W. Y. Crutchfield, and J. B. Bell, 2000: Parallelization of structured, hierarchical adaptive mesh refinement algorithms. *Comput. Visualization in Science*, **3**.

Skamarock, W. C., and J. B. Klemp, 1993: Adaptive grid refinement for 2-dimensional and 3-dimensional nonhydrostatic atmospheric flow. *Mon. Weather Rev.*, **121**(3), 788–804.

Staniforth, A., and J. Côté, 1991: Semi-Lagrangian integration schemes for atmospheric models – A review. *Mon. Weather Rev.*, **119**, 2206–2223.

Strang, G., 1968: On the construction and comparison of difference schemes. *SIAM J. Num. Anal.*, **5**, 506–517.

Sweby, P. K., 1984: High resolution schemes using flux limiters for hyperbolic conservation laws. *SIAM J. Num. Anal.*, **21**, 995–1011.

Toro, E. F., 1989: A weighted average flux method for hyperbolic conservation laws. *Proc. R. Soc. Lond.*, **423**, 401–418.

van Leer, B., 1979: Towards the ultimate conservative difference scheme V. A second order sequel to Godunov’s method. *J. Comput. Phys.*, **32**, 101–136.

Williamson, D. L., J. B. Drake, J. J. Hack, R. Jakob, and P. N. Swartztrauber, 1992: A standard test set for numerical approximations to the shallow water equations in spherical geometry. *J. Comput. Phys.*, **102**, 211–224.

Williamson, D. L., and P. J. Rasch, 1989: Two-dimensional semi-Lagrangian transport with shape-preserving interpolation. *Mon. Weather Rev.*, **117**, 102–129.

List of Tables

1	Continued on next page.	44
1	Error measures for the solid body rotation of a cosine bell after one revolution using a superbee limited WAF scheme with varying degrees of adaptation. Results are shown for all 4 of the angles of rotation α recommended by Williamson <i>et al</i> (1992). In the first column the first two numbers give the number of cells in the coarse mesh while any in brackets indicate successive refinement factors.	45
2	Error measures for the solid body rotation of a cosine bell ($\alpha = \pi/2$) after one revolution using a superbee limited WAF scheme with varying numbers of zonal sweeps on a single uniform 384×192 mesh. . .	46
3	Error measures for the two-dimensional idealised cyclogenesis at $t = 2.5$ using a superbee limited WAF scheme with varying degrees of adaptive meshing. In the first column the first two numbers give the number of cells in the coarse mesh while any in brackets indicate successive refinement factors.	47

Grid	Equivalent resolution	l_1	l_2	l_∞	Time (s)
$\alpha = \pi/2$					
64×32	64×32	0.5828	0.4567	0.4590	25
$64 \times 32 (\times 2)$	128×64	0.1325	0.1355	0.1870	78
128×64	128×64	0.1296	0.1320	0.1794	370
$64 \times 32 (\times 2 \times 3)$	384×192	0.0661	0.0616	0.0795	1526
$128 \times 64 (\times 3)$	384×192	0.0611	0.0559	0.0677	2074
384×192	384×192	0.0466	0.0438	0.0607	32807
$\alpha = \pi/2 - 0.05$					
64×32	64×32	0.5868	0.4603	0.4730	25
$64 \times 32 (\times 2)$	128×64	0.1327	0.1342	0.1832	82
128×64	128×64	0.1298	0.1308	0.1792	378
$64 \times 32 (\times 2 \times 3)$	384×192	0.0648	0.0602	0.0744	1569
$128 \times 64 (\times 3)$	384×192	0.0600	0.0549	0.0642	2139
384×192	384×192	0.0466	0.0438	0.0610	33509

Table 1: Continued on next page.

Grid	Equivalent resolution	l_1	l_2	l_∞	Time (s)
$\alpha = 0.05$					
64×32	64×32	0.5353	0.4299	0.4372	3
$64 \times 32 (\times 2)$	128×64	0.1318	0.1255	0.1458	5
128×64	128×64	0.1446	0.1352	0.1565	27
$64 \times 32 (\times 2 \times 3)$	384×192	0.0296	0.0282	0.0465	25
$128 \times 64 (\times 3)$	384×192	0.0299	0.0288	0.0467	57
384×192	384×192	0.0302	0.0294	0.0477	1777
$\alpha = 0.0$					
64×32	64×32	0.2228	0.1833	0.1762	2
$64 \times 32 (\times 2)$	128×64	0.0576	0.0521	0.0568	3
128×64	128×64	0.0572	0.0523	0.0572	9
$64 \times 32 (\times 2 \times 3)$	384×192	0.0264	0.0211	0.0219	12
$128 \times 64 (\times 3)$	384×192	0.0220	0.0211	0.0277	19
384×192	384×192	0.0212	0.0206	0.0291	237

Table 1: Error measures for the solid body rotation of a cosine bell after one revolution using a superbee limited WAF scheme with varying degrees of adaptation. Results are shown for all 4 of the angles of rotation α recommended by Williamson *et al* (1992). In the first column the first two numbers give the number of cells in the coarse mesh while any in brackets indicate successive refinement factors.

Equivalent level of ‘reduction’	l_1	l_2	l_∞	Time (s)	Time-step (s)
0	0.0466	0.0438	0.0607	32807	20.0
1	0.0464	0.0437	0.0605	17485	39.8
2	0.0461	0.0435	0.0604	9106	79.5
3	0.0453	0.0430	0.0593	4988	159.1
4	0.0442	0.0421	0.0587	2872	318.1
5	0.0505	0.0461	0.0683	1701	636.2
6	0.0822	0.0723	0.0749	1192	1215.1
7	0.0828	0.0729	0.0758	1720	1215.1
‘Optimal’	0.0772	0.0689	0.0758	684	1215.1

Table 2: Error measures for the solid body rotation of a cosine bell ($\alpha = \pi/2$) after one revolution using a superbee limited WAF scheme with varying numbers of zonal sweeps on a single uniform 384×192 mesh.

Grid	Equivalent resolution	l_1	l_2	l_∞	Time (s)
64×32	64×32	0.0779	0.2135	1.0508	24
$64 \times 32 (\times 2)$	128×64	0.0405	0.1486	0.9894	354
128×64	128×64	0.0405	0.1486	0.9894	351
$64 \times 32 (\times 2 \times 3)$	384×192	0.0139	0.0810	0.8818	11515
$128 \times 64 (\times 3)$	384×192	0.0139	0.0810	0.8817	11362
384×192	384×192	0.0139	0.0811	0.8833	29328

Table 3: Error measures for the two-dimensional idealised cyclogenesis at $t = 2.5$ using a superbee limited WAF scheme with varying degrees of adaptive meshing. In the first column the first two numbers give the number of cells in the coarse mesh while any in brackets indicate successive refinement factors.

List of Figures

1	The grid and orientation of coordinate axes close to a polar singularity.	50
2	An example of an adapted grid.	51
3	An adapted grid with the dummy cells superimposed.	52
4	Solution contours for the solid body rotation of a cosine bell over the poles ($\alpha = \pi/2$) after one revolution. The solid lines represent the numerical solutions and the dashed lines give the exact solution. The cases shown correspond to the different grids indicated in Table 1: (a) 64×32 , (b) $64 \times 32 (\times 2)$, (c) 128×64 , (d) $64 \times 32 (\times 2 \times 3)$, (e) $128 \times 64 (\times 3)$, and (f) 384×192 (the first two numbers give the number of cells in the coarse mesh while any in brackets indicate successive refinement factors).	53
5	Adapted meshes using a $64 \times 32 (\times 4 \times 4)$ grid for the solid body rotation of a constant region over the poles ($\alpha = \pi/2$) after $\frac{1}{8}$ (top left), $\frac{3}{16}$ (top right) and $\frac{1}{4}$ (bottom left) of a revolution, together with some solutions at these and other times (bottom right).	54
6	Solution surface (top) and adapted grid (bottom) for the two-dimensional idealised cyclogenesis problem at $t = 2.5$ for a 64×32 coarse grid with two levels of refinement ($\times 2 \times 3$) using a superbee limited WAF scheme. The solution has been projected stereographically on to the plane tangent to the centre of rotation, and only the region north of 22.5°S is shown.	55

7	Contours of observed potential vorticity (obtained from ERA data) over one quadrant of the northern hemisphere for 1200 21st June 1996 (shown on a global $480 \times 240 \times 24$ mesh between the isentropic surfaces of 300K and 360K). The mesh is shown and is shaded according to potential vorticity.	56
8	Contours of potential vorticity over one quadrant of the northern hemisphere for 1200 21st June 1996 after 4 days of a computational run on a $480 \times 240 \times 24$ mesh covering the whole globe between the isentropic surfaces of 300K and 360K. The mesh is shown and is shaded according to potential vorticity.	57
9	Contours of potential vorticity over one quadrant of the northern hemisphere for 1200 21st June 1996 after 4 days of a computational run on a $120 \times 60 \times 8$ ($\times 2 \times 3$) mesh covering the whole globe between the heights of 290K and 370K (given as potential temperature). The mesh is shown and is shaded according to potential vorticity.	58




Figure 1: The grid and orientation of coordinate axes close to a polar singularity.




Figure 2: An example of an adapted grid.




Figure 3: An adapted grid with the dummy cells superimposed.

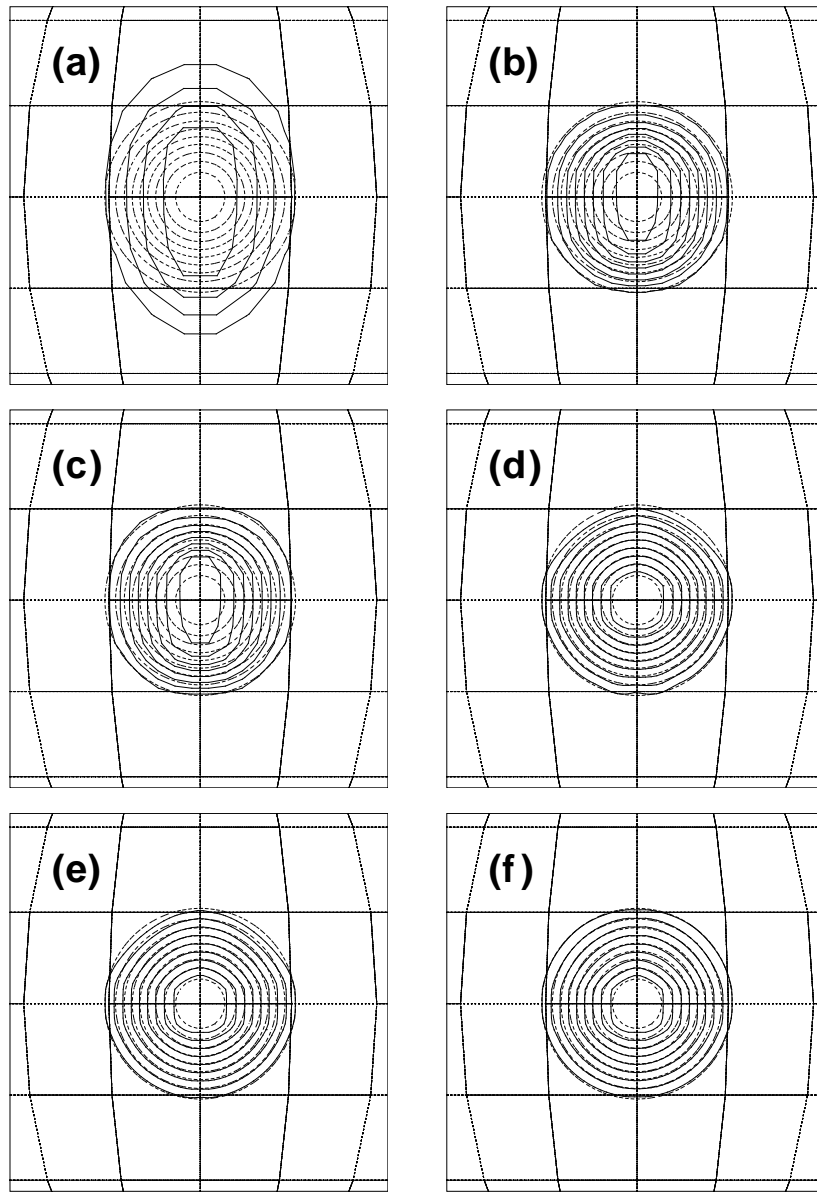


Figure 4: Solution contours for the solid body rotation of a cosine bell over the poles ($\alpha = \pi/2$) after one revolution. The solid lines represent the numerical solutions and the dashed lines give the exact solution. The cases shown correspond to the different grids indicated in Table 1: (a) 64×32 , (b) $64 \times 32 (\times 2)$, (c) 128×64 , (d) $64 \times 32 (\times 2 \times 3)$, (e) $128 \times 64 (\times 3)$, and (f) 384×192 (the first two numbers give the number of cells in the coarse mesh while any in brackets indicate successive refinement factors).




Figure 5: Adapted meshes using a 64×32 ($\times 4 \times 4$) grid for the solid body rotation of a constant region over the poles ($\alpha = \pi/2$) after $\frac{1}{8}$ (top left), $\frac{3}{16}$ (top right) and $\frac{1}{4}$ (bottom left) of a revolution, together with some solutions at these and other times (bottom right).




Figure 6: Solution surface (top) and adapted grid (bottom) for the two-dimensional idealised cyclogenesis problem at $t = 2.5$ for a 64×32 coarse grid with two levels of refinement ($\times 2 \times 3$) using a superbee limited WAF scheme. The solution has been projected stereographically on to the plane tangent to the centre of rotation, and only the region north of 22.5°S is shown.




Figure 7: Contours of observed potential vorticity (obtained from ERA data) over one quadrant of the northern hemisphere for 1200 21st June 1996 (shown on a global $480 \times 240 \times 24$ mesh between the isentropic surfaces of 300K and 360K). The mesh is shown and is shaded according to potential vorticity.

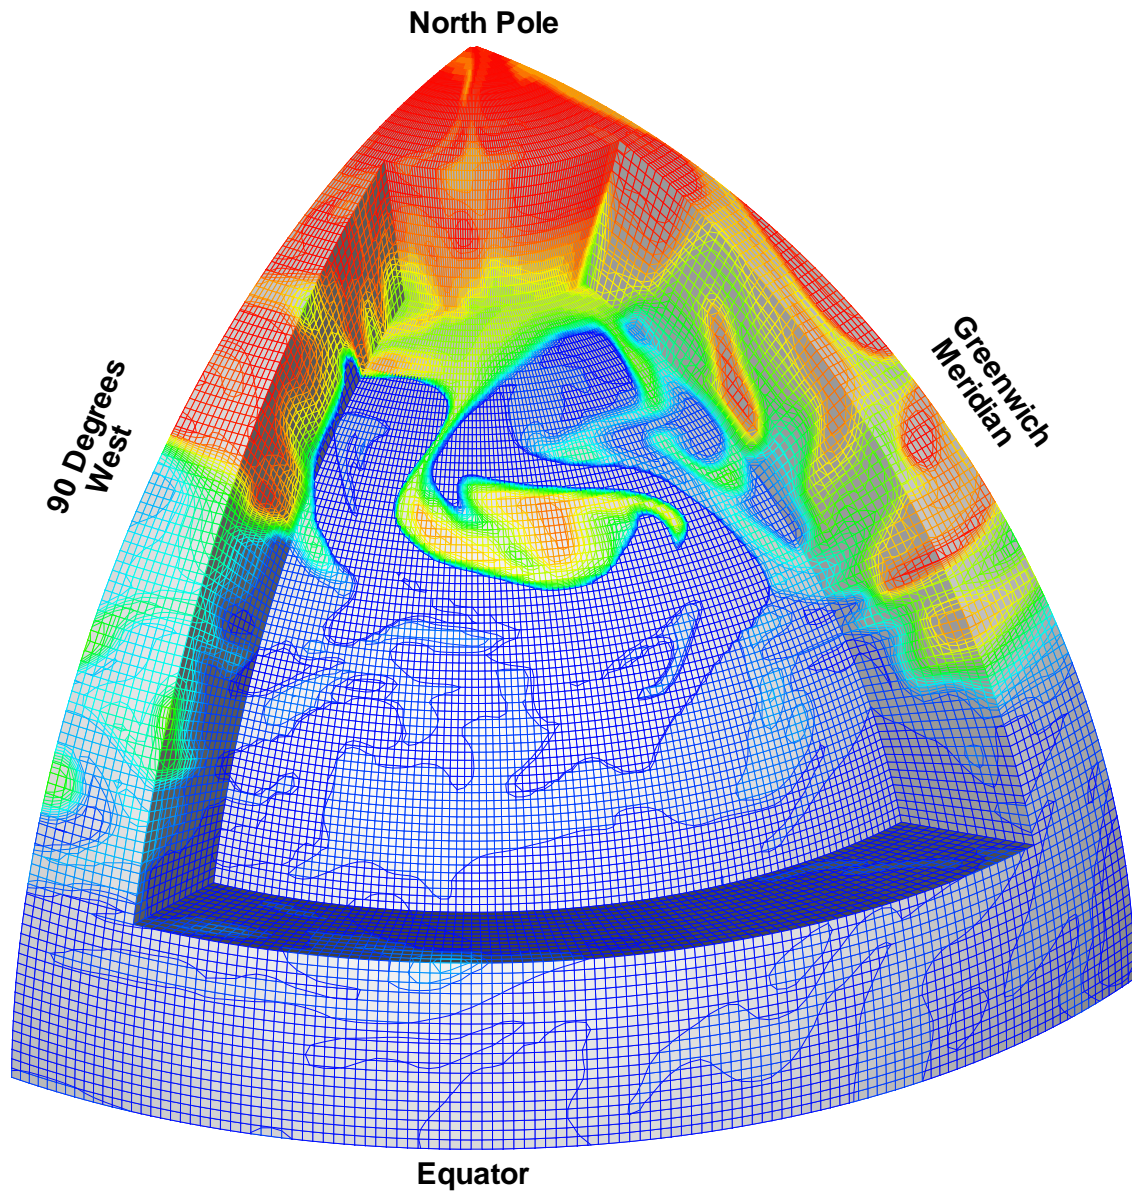


Figure 8: Contours of potential vorticity over one quadrant of the northern hemisphere for 1200 21st June 1996 after 4 days of a computational run on a $480 \times 240 \times 24$ mesh covering the whole globe between the isentropic surfaces of 300K and 360K. The mesh is shown and is shaded according to potential vorticity.




Figure 9: Contours of potential vorticity over one quadrant of the northern hemisphere for 1200 21st June 1996 after 4 days of a computational run on a $120 \times 60 \times 8$ ($\times 2 \times 3$) mesh covering the whole globe between the heights of 290K and 370K (given as potential temperature). The mesh is shown and is shaded according to potential vorticity.