Accurate numerical methods for simulating advective and wave processes in atmospheric and oceanic dynamics

David G. Dritschel

Department of Mathematics
University of Warwick
Coventry CV4 7AL, UK

Abstract

This paper reviews recent progress in improving the simulation of both advective and wave processes in atmospheric and oceanic dynamics. This progress stems from the development of a hybrid Lagrangian-Eulerian algorithm, called the "Contour-Advective Semi-Lagrangian" (CASL) algorithm, which has proven to be both remarkably accurate and efficient compared to conventional (pseudo-spectral and semi-Lagrangian) algorithms. So far, the algorithm has been developed for relatively simple flow geometries and equation sets. The most recent developments include an unforced primitive-equation algorithm in horizontally doubly-periodic geometry, a diabatically-forced quasi-geostrophic algorithm, and an algorithm for spherical geometry. Moreover, a new method has been developed for the accurate representation of freely-propagating gravity waves. The success of this approach may make it attractive as part of a future global circulation model.

1. Introduction

The numerical simulation of the motion of the atmosphere and the oceans is a highly-developed science, driven by socio-economic pressure and generally by a practical desire to predict a future state with sufficient accuracy. Vast computational resources are expended in this venture — which is not surprising considering the full complexity of the system modelled and what it takes to properly observe it — yet from the public point of view and from a critical scientific one, numerical simulation is not sufficiently accurate for many fields, e.g. precipitation, and for not long enough, e.g. for climate change. This is not to say that advances in numerical simulation have not been forthcoming, they have indeed. And, great advances have also been made in understanding many of the salient features of this system. Simply, the problem at hand is hugely complex, and not easily grasped by the public. The success of numerical prediction has been remarkable after the inauspicious beginning of L. F. Richardson.

There is room for improvement. This paper reviews a new idea that could lead to this further improvement. It is a new idea in its infancy, and preliminary research suggests that it is promising.

The new idea is to combine a purely Lagrangian method for the advection of nearly conservative fields with a conventional Eulerian method for other operations, i.e. for computing the velocity field, for wave propagation, etc. The Lagrangian method, "contour surgery" (Dritschel 1989), represents the nearly conservative fields, the most important of which is the potential vorticity q, by *contours*. These contours are simply advected by the velocity field, i.e. one solves the set of o.d.e.'s $d\mathbf{x}/dt = \mathbf{u}(\mathbf{x},t)$ for the points $\mathbf{x}$ lying on the q contours — this is the basis for "contour advection"; see Norton (1994) and Waugh & Plumb (1994). This advection is formally equivalent to material conservation of q. Though strict material conservation of potential vorticity is not a uniformly good approximation in the atmosphere and oceans, particularly over long time scales, it can give deep insight into the fundamental dynamical behaviour of these systems (Hoskins *et al* 1985, McIntyre & Roulstone 1998, etc...). Moreover, the new method described here is not limited to materially conserved fields, as discussed in §5 below.

There are two key advantages of using contour advection for the evolution of nearly conservative

fields. The first is to avoid the small time step normally required for numerical stability. For advection, the time step is chosen for accuracy reasons alone, as in the semi-Lagrangian algorithm now commonly used in weather forecasting. A larger time step of course means a more efficient computation. But accuracy must also be taken into account, and semi-Lagrangian schemes are normally no more accurate (and can even be less accurate) than conventional algorithms using grid-based advection (e.g. pseudo-spectral or finite-element, for instance see Gravel 1996, Leslie & Dietachmayer 1997, Bartello 1996 and Dritschel *et al* 1998a). The second key advantage of contour advection is that it is much more accurate than semi-Lagrangian advection. This gain in accuracy comes from not having to interpolate q. The inaccurate step in semi-Lagrangian advection is the interpolation of q, at points $\mathbf{x}(t)$ that arrive at grid points at time $t + \Delta t$, from gridded values of q. Contour advection keeps q in contour form and can efficiently represent q features (e.g. filaments and sharp gradients) well below grid scale (that such sub-grid scale features are sensible has been argued by Waugh & Plumb (1994) and Methven & Hoskins (1998)).

The new algorithm discussed in this paper takes contour advection one step further by internally calculating the velocity field $\mathbf{u}$ from the potential vorticity q (and other fields for general, unbalanced, but stably-stratified flows). For this, Dritschel & Ambaum (1998) (hereafter DA) developed a fast contour-to-grid conversion scheme that supplies q on grid points for the purpose of obtaining $\mathbf{u}$. This conversion is done *only* for this purpose; advection is carried out in a fully Lagrangian way. The new algorithm, called the "Contour-Advective Semi-Lagrangian" (CASL) algorithm by DA, is semi-Lagrangian in the sense that advection is done in a Lagrangian manner whereas recovering the velocity field, wave propagation, etc, are done in an Eulerian manner. This dicotomy is arguably natural: advection is a Lagrangian notion, while inversion and wave-propagation are Eulerian (cf. McIntyre & Roulstone 1998). If one accepts this, then it is not surprising that the numerical results show a large gain in accuracy for fixed cost when using the CASL algorithm versus using the semi-Lagrangian or pseudo-spectral algorithms (Dritschel *et al* 1998a, hereafter DPM).

This short review paper presents an overview of the existing CASL algorithms, beginning with the original quasi-geostrophic algorithm of DA in §2, then describing the extensions to the shallow-water equations (DPM) in §3, to spherical geometry (Dritschel 1998) in §4, and to diabatically-forced flows (Dritschel & Ambaum 1998) in §5. Future work is discussed in §6.

2. The quasi-geostrophic algorithm

The simplest, most idealised model of atmospheric and oceanic dynamics is the quasi-geostrophic (QG) one. It is nevertheless (or as a consequence!) widely used in basic research, and often exhibits more than just a qualitative correspondence with real flows (Simmons & Hoskins 1976, Gill 1982, Holton 1982, and many others). Its validity is formally restricted to small Rossby and Froude numbers (as well as system height to width aspect ratio), but like many asymptotic systems, "small" can be $O(1)$ (see Vallis 1996).

The QG model is special in two respects. First of all, it is "balanced", that is $\mathbf{u}$ is determined by the instantaneous distribution of q. Second, the balance is linear, the simplest balance possible. The flow field is determined by inverting an elliptic Laplacian-like operator on q to find the streamfunction ψ of a layerwise-2D incompressible flow ($\mathbf{u} = \{-\partial\psi/\partial y, \partial\psi/\partial x\}$).

The CASL algorithm for the QG model is particularly straightforward as a consequence. Given the q contours, a fast contour-to-grid conversion (called the "PVCGC") is performed to get q at all grid points. The horizontally double-periodic algorithm of DA uses FFTs to convert these gridded values to spectral coefficients and furthermore projects q onto vertical modes; this enables one to calculate the transformed-projected streamfunction almost trivially (by multiplication of the inverse-squared 3D wavenumber), and likewise the transformed-projected velocity field. After inverting this projection and transformation, $\mathbf{u}$ is known on the grid points. To evolve q, the o.d.e.'s $d\mathbf{x}/dt = \mathbf{u}(\mathbf{x}, t)$ are solved for each point $\mathbf{x}$ lying on each q contour in each layer (or on each vertical level); here, $\mathbf{x}$ generally lies between grid points, so a bilinear interpolation of $\mathbf{u}$ is performed (as in previous contour advection studies — see Norton (1994) and Waugh & Plumb (1994)). The time integration is performed with a 4th-order Runge-Kutta scheme.

The inevitable cascade of PV to ever finer scales through advection is controlled in the CASL algorithm by removing filamentary PV at a *tenth* of the grid scale. Note that sharp PV gradients are not dissipated at all. This greatly reduces the dissipation compared to conventional algorithms, including semi-Lagrangian ones, which cannot avoid, in fact require, significant grid-scale dissipation. This is the primary reason why the CASL algorithm turns out to be so much more accurate than conventional ones.

Much of this fine-scale PV, however, does not contribute to the velocity field $\mathbf{u}$. It would only weakly contribute anyway due to the coarse-graining effect of inversion, but in the CASL algorithm, the PV lying below $1/m_g$ times the basic grid scale Δ does not contribute at all. When finding the gridded PV field, a grid m_g times finer in each direction is used in the contour-to-grid conversion, and then this fine-grid PV is averaged to obtain PV values on the original basic grid (m_g is a power of 2, normally equal to 2 or 4 — see DA). Thus, some sub-grid PV features do contribute to the velocity field. DA has shown that this measurably increases the accuracy of the algorithm, despite the fact that the averaging spreads out these sub-grid features. The point is that averaging does not remove these features, and their contribution to $\mathbf{u}$ is insensitive to averaging (again, due to the coarse-graining effect of inversion). See DA for a demonstration of this.

The particularly novel feature of the CASL algorithm is the PV contour-to-grid conversion, the "PVCGC". DA originally developed a fast method for this in horizontally doubly-periodic geometry that requires only $O(n)$ operations, where n is the number of points on the PV contours. Since then, this method has been extended to channel geometry, a circular domain, and the surface of a sphere. The latter is described in §4 below. The basic idea of the method is to determine where the contour segments between adjacent points on a contour cross grid lines y = constant. Knowing the PV jump across the contour, and the direction of the contour (e.g. y increasing or y decreasing), one can work out in very few operations that the PV must jump by a certain amount from one grid point to the next along this grid line y = constant. Repeating this procedure for every pair of adjacent nodes, one finally obtains a "PV jump array", from which a sweep from left to right (in x)

allows one to find the PV at any grid point given the PV at the left hand edge of the domain. That PV is obtained in the same manner as above with x and y interchanged; i.e. one finds the contour crossings with the left hand edge of the domain and therefore the PV jump from one y grid point to the next. Sweeping in y gives the PV along this edge to within a constant, and therefore the PV everywhere to within a constant. The constant is then fixed by requiring the horizontally-averaged PV to have a prescribed, fixed value in each layer (or on each level).

In simulations of QG turbulence, DA have estimated that the CASL algorithm is more than 1000 times faster than either a pseudo-spectral or a contour dynamics algorithm having comparable accuracy. This has clarified the inherent three-dimensionality of QG turbulence at horizontal scales below the Rossby radius, and has lent strong support to Charney's (1971) theory of QG turbulence, in particular his prediction of isotropy (after scaling height on Prandtl's ratio) — see Dritschel *et al* (1998b).

3. The shallow-water algorithm and the "wave equation"

The basic idea behind the CASL algorithm is to represent advected fields characterised by a shallow spectral decay (and much fine-scale structure) in a Lagrangian way, and to represent all other fields characterised by a steep (or steeper) spectral decay in an Eulerian way. Such fields, in QG flow, are the "inverted" fields ψ and $\mathbf{u}$; they are clearly of broader scale than the PV. In shallow-water flow, and generally in the primitive equations, there are additional wave phenomena, gravity waves, that do not owe their existence to PV; these freely-propagating gravity waves are part of the "unbalanced" flow (in some suitably defined sense), and contribute to the advecting velocity field in a way that cannot be deduced from the instantaneous distribution of PV. These waves are most clearly visible in the horizontal divergence field $\delta = \nabla \cdot \mathbf{u}$, though much of δ may consist of balanced motions (derivable from the PV). Now, if these waves either have a steep spectral decay (i.e. δ has a steep spectral decay), or if these waves are negligible (as in the case of a nearly balanced flow), then it makes sense to model these wave motions in an Eulerian way, for instance using the pseudo-spectral method.

This behaviour has been verified in the works of Dritschel *et al* (1998a) (DPM) who implemented and tested the shallow-water CASL algorithm, and Mohebalhojeh & Dritschel (1998), who developed a more accurate means to represent gravity waves. DPM developed the CASL algorithm using the primary variables q, δ and h, where h is the height field (or rather the perturbation from the mean) in doubly-periodic geometry. In this algorithm, PV is represented by contours and evolved by contour advection, while δ and h are evolved in spectral space using a standard second-order, semi-implicit time-stepping scheme with the usual Robert-Asselin time filter. Notably, no dissipation was added to the right hand side of the divergence equation. This was found to be unnecessary, since in all the simulations conducted, divergence did not tend to pile up at the smallest scales.

The main purpose of DPM was to show that the CASL algorithm leads to a significant improvement with respect to PV conservation than its leading contenders, the pseudo-spectral (PS) and semi-Lagrangian (SL) algorithms. In an example flow beginning with an unstable jet and characterised by smooth PV, the CASL algorithm was found to produce results as accurate as those produced by the PS and SL algorithms using 8–10 times finer resolution in each direction, representing a speed-up factor of about 1000 and 100, respectively. Accuracy was measured by the ability to conserve mass (the area integral of h) between PV isolevels. This is a stringent measure, but arguably an important one for assessing the accuracy of tracer transport.

Despite this significant improvement for PV conservation, no measurable improvement was seen in the accuracy of the gravity-wave motions. In fact, gravity waves are poorly modelled in most numerical models, particularly at small Froude number (when fluid speeds are slow compared to gravity wave speeds). This is not a serious problem generally because such gravity waves in

numerical models are widely regarded as noise, and forecast initialisations have tended to eliminate them as much as possible to avoid erroneous flow developments later in the forecast. Should one want accurate gravity wave fields in a future GCM, then current model formulations (using mass and divergence or the horizontal velocity components as fundamental model variables) have to change — at least this is the conclusion of our latest work (Mohebalhojeh & Dritschel 1998, MD). The work of MD was originally motivated by trying to see how far the concept of "balance" could be taken, i.e. in Froude number, using the most accurate algorithm available at the time, the CASL algorithm of DPM. However, we soon discovered that the divergence field was not sufficiently accurate; in fact, for small Froude numbers ($\lesssim 1/4$), numerical errors were found to dominate the unbalanced part of the divergence field, and the situation got worse for smaller Froude numbers. The source of this numerical error was found to be the mass tendency term, $\nabla\cdot(h\mathbf{u})$. Since h tends to the QG streamfunction as the Froude number tends to zero (as well as the Rossby number), h is essentially balanced at small Froude number. Indeed, this is clear from the work of DPM. Now, any error in computing $\nabla\cdot(h\mathbf{u})$, such as that arising from the spatial truncation of both the gradient operator and the fields it operates on, may project onto the spectrum of freely-propagating gravity waves, and not just add to the balanced flow. In fact, careful flow analysis has verified that this truncation error does project onto gravity waves at all scales, particularly at small scales, and the only way to avoid it is to not use the mass tendency equation!

MD then reformulated the shallow water equations in a way that *diagnoses* the mass, formally just as one diagnoses mass in the QG model by a linear inversion relation. In the simplest scheme described by MD, the mass equation was eliminated by taking a time derivative of the divergence equation and replacing occurences of $\partial h/\partial t$ by $-\bar{h}\delta - \nabla\cdot(h\mathbf{u})$, where $\bar{h}$ is the mean depth. This results in the second-order in time "wave equation"

$$\frac{\partial\delta}{\partial t} - \mathcal{H}\delta = S$$

where $\mathcal{H}$ is a linear operator that generates the dispersion relation for gravity waves, namely

$$\mathcal{H} = c^2\nabla^2 - f^2$$

where $c = \sqrt{g\bar{h}}$ is the phase speed of non-rotating gravity waves and f is the Coriolis parameter (a constant in this geometry). Note, the ratio c/f is the mean Rossby radius L_R, and one can recognise $\mathcal{H}/c^2 = \nabla^2 - 1/L_R^2$ as the elliptic operator occuring in QG theory. The source term S is a complicated function of q, h, δ and $\partial\delta/\partial t$. But given all of these quantities at one time, one can integrate this wave-equation for δ and $\partial\delta/\partial t$ at the next time. As for the height field h, this is deduced from the original divergence equation re-arranged, which formally looks like

$$\mathcal{H}h = B$$

where B is a function of q, h, δ and $\partial\delta/\partial t$. In the QG limit, B tends a constant times the PV anomaly, $\bar{h}q - f$. In this limit, then, h is recovered directly from the PV (this, by the way, shows the inconsistency of trying to evolve h directly in this limit). In general, however, B depends on h in a nonlinear manner, so h must be found by iteration. The number of iterations required for convergence is sensitive to the Froude number, as is to be expected, but often only one or two iterations are sufficient for Froude numbers $\lesssim 1/4$.

The results of MD show a very large reduction in divergence error using this method relative to that produced by the original algorithm of DPM. This development opens the way for a careful assessment of the interaction between gravity waves and mean, balanced motions.

Both the algorithms of DPM and MD have been extended to the primitive equations, so far with the constraint that the vertical boundaries are isentropes or isopycnals. Details will be given in a forthcoming paper. See also §6 for related extensions to the nonhydrostatic equations.

4. Spherical geometry

This section describes how the CASL algorithm is adapted to spherical geometry (see Dritschel (1998) for further details and numerical tests). The only novel feature of the spherical CASL algorithm is the PV contour-to-grid conversion, and this will be illustrated in the context of the single-layer QG equations. Extensions to more realistic sets of equations are as straightforward to implement in the CASL algorithm as they are in any conventional algorithm.

The single-layer QG equations may be stated in the following conservative form:

$$\frac{\partial q}{\partial t} + \mathbf{u}\cdot\nabla q \equiv \frac{Dq}{Dt} = 0 \tag{1a}$$

$$\nabla^2\psi - \gamma^2\psi = q - f \tag{1b}$$

$$\mathbf{u} = \hat{\mathbf{n}}\times\nabla\psi \tag{1c}$$

where $\hat{\mathbf{n}}$ is unit vector normal to the surface of the flow (e.g. $\hat{\mathbf{e}}_z$ for planar flow and $\hat{\mathbf{e}}_r$ for flow on a sphere), $\gamma = 1/L_R$ is the inverse Rossby radius, and $f = 2\Omega\sin\phi$, where Ω is the Earth's rotation rate and ϕ is latitude. (For a shallow-water flow, $L_R = c/f$, where c is the phase speed of gravity waves. See Ambaum (1997) for why constant L_R is justifiable even though f varies with latitude.)

In a CASL simulation, one begins by the PVCGC (detailed below), which gives the PV at grid points. This gridded field is then used to solve Eq.(1b) for the streamfunction ψ, and Eq.(1c) for the velocity $\mathbf{u}$ on the grid. The velocity of the points on the PV contours is then found by bilinear interpolation, and a 4th-order Runge-Kutta scheme is used to integrate the trajectory equations, $d\mathbf{x}/dt = \mathbf{u}(\mathbf{x},t)$. Periodically (generally not every time step), the points representing the contours are redistributed to ensure adequate resolution in regions of high curvature, and "surgery" is applied to remove PV filaments (or merge appropriate contours) at a tenth of the grid scale. For details of these contour operations, see Dritschel (1989).

To get the gridded velocity field $\mathbf{u}$ from the gridded q field, FFTs are used in the doubly-periodic algorithm to solve Eqs.(1b–c) by simple operations in spectral space, then $\mathbf{u}$ is recovered by inverse FFTs. In the spherical algorithm, a semi-spectral approach is adopted. FFTs are used in longitude λ while equally-spaced grid points are used in latitude ϕ, along with 2nd-order centred differences. Equally-spaced latitudes are convenient for the PVCGC, and moreover a semi-spectral approach is more efficient than a full spectral (spherical harmonic) one, though in principal the accuracy is lower (but perhaps not significantly so when dealing with spectrally-shallow fields like the PV). For each longitudinal wavenumber m, the transformed streamfunction $\hat{\psi}_m$ is found by solving the discretised version of Eq.(1b) multiplied by $\cos\phi$,

$$\frac{d}{d\phi}\left(\cos\phi\frac{d\hat{\psi}_m}{d\phi}\right) - \left(\frac{m^2}{\cos\phi} + \gamma^2\cos\phi\right)\hat{\psi}_m = \hat{\sigma}_m \tag{2}$$

where $\sigma \equiv (q-f)\cos\phi$. (Here, we take the radius of the sphere to be unity — this is just a choice of scaling that makes the notation simpler.) In Eq.(2), the second-order derivative term is approximated by standard centred differences to give

$$(\Delta\phi)^{-2}[(\hat{\psi}_{m,k+1} - \hat{\psi}_{m,k})\cos\phi_{k+1/2} - (\hat{\psi}_{m,k} - \hat{\psi}_{m,k-1})\cos\phi_{k-1/2}]$$

where the integer k denotes a discrete latitude. At the poles, $\hat{\psi}_m = 0$ and the source term $\hat{\sigma}_m = 0$ as well. This results in a tridiagonal problem, for the interior latitudes $\phi_k = k\Delta\phi - \frac{1}{2}\pi$, $k = 1,\ 2,\ \ldots,\ n_\phi - 1$ (where $\Delta\phi = \pi/n_\phi$), whose solution requires only $O(n_\phi)$ basic operations. (The number of longitudes used for the discrete Fourier transform is twice n_ϕ, because the angular

variation of the longitude is twice that of the latitude.) An exception occurs when $\gamma = 0$ (barotropic flow) and when $m = 0$. Then, Eq.(2) is simply integrated once with respect to ϕ, using the trapezoidal rule, to find the zonal component of the angular momentum $\hat{\Gamma}_0 = \cos\phi\, d\hat{\psi}_0/d\phi$. In general, the angular momentum $(\cos^2\phi\, d\lambda/dt)$ for each longitudinal wavenumber $\hat{\Gamma}_m = \cos\phi\, d\hat{\psi}_m/d\phi$ requires differentiation with respect to ϕ, and here this is done by centred differences (recall the polar values of $\hat{\psi}_m$ are zero; likewise, the polar values of $\hat{\Gamma}_m$ are zero). The vector velocity $\mathbf{u}$, being tangent to the spherical surface, can be determined by two scalar fields, Γ and the axial velocity $w = dz/dt = -\partial\psi/\partial\lambda$ (where $z = \sin\phi$). The latter is evaluated in spectral space, where differentiation with respect to λ becomes multiplication by im. With Γ and w known at all grid points (i.e. after inverse FFTs), the x and y components of $\mathbf{u}$ are given by

$$u = \frac{-\Gamma y - wxz}{x^2+y^2} \quad ; \quad v = \frac{\Gamma x - wyz}{x^2+y^2} \tag{3}$$

— note that $\mathbf{x}\cdot\mathbf{u} = 0$, as required ($\hat{\mathbf{n}} = \mathbf{x}$ in Eq.(1c)). In Eq.(3), Γ and w are interpolated values at a given node $\mathbf{x}$. They are found by a weighted sum of the gridded Γ and w values at the corners of the latitude-longitude box containing $\mathbf{x}$. The weights used are the fractional areas of the four parts of the box formed by longitude and latitude cuts through $\mathbf{x}$.

The three *Cartesian* velocity components, though redundant, are used to integrate $d\mathbf{x}/dt$, as in previous works (Dritschel 1988 & 1989). This preserves the isotropy of advection and is more accurate than using spherical coordinates. Numerical integration of these three equations does not normally ensure that $\mathbf{x}$ will remain on the spherical surface, so after each update of $\mathbf{x}$ (this occurs four times in the 4th-order Runge-Kutta scheme), the magnitude of $\mathbf{x}$ is reset to unity. (See the above references for details.)

The last remaining feature of the algorithm to describe is the conversion of PV contours to gridded values, the PVCGC. The use of equally spaced latitudes and longitudes allows one to use much of the routine developed for the doubly-periodic algorithm. The sphere however is only periodic in longitude, and the poles need special attention. Moreover, the contours are inherently curved because of the surface they lie upon.

The first task in the PVCGC is to determine which longitudes, if any, are crossed by the contour segment connecting two adjacent nodes on a contour. For each such crossing, one calculates the latitudes between which the crossing occurs. Given the direction of the crossing and the PV jump across the contour, one knows that the PV changes across a certain longitude and between certain latitudes. After working out the crossings of every contour segment, one can sweep from the south pole to the north pole and add up these PV jumps to recover the PV field at all grid points to within an overall constant. The constant is determined by requiring the global average PV to be zero; this is not a constraint since the global average PV does not induce any motion.

Normally, this task is carried out on a grid twice finer in each direction ($m_g = 2$) to improve the accuracy of final velocity field (see DA and §2 above). On this fine grid, $\sigma = (q - f)\cos\phi$ is computed, and then a nine-point averaging is done to recover σ on the original grid, in preparation for the solution of Eq.(2). (The $\cos\phi$ factor provides the natural area weight needed for this averaging.)

The mechanics of the above is adequately described in DA. One detail differs, however, because of the curvature of the spherical surface. Between each pair of adjacent nodes, an arc of a great circle is used for interpolation to ensure that the contour lies on the sphere — the advantage of this interpolation, say compared to using straight lines between latitudes and longitudes (treated as if Cartesian coordinates), is that it is isotropic, i.e. it is invariant with respect to the location of the poles. This allows for a proper treatment of the poles, where a contour segment may cross up to half the discrete longitudes. The use of great circles is straightforward, and the PVCGC routine is no less efficient than its analogue in doubly-periodic geometry.

5. Diabatic processes

A vital extension of the CASL algorithm is the incorporation of diabatic forcing. This extension is essential to studying fundamental aspects of the development and maintainence of the eddy-driven large-scale circulation (cf. Ambaum 1997 & refs). But this extension has long been regarded as impossible for Lagrangian methods. The basic problem is that the Lagrangian representation of the PV is generally incompatable with diabatic forcing; such forcing changes the PV at a point, and contours have to move to compensate — but this is not always possible, for instance at extrema of the PV field. The attraction of incorporating diabatic forcing in the CASL algorithm, particularly diabatic forcing that tries to relax the mass field back to (a generally unstable) equilibrium, is that such forcing operates most rapidly at large scales and very weakly at small scales, i.e. in the exact opposite manner to molecular diffusion. For instance, in the idealised one-layer models used by Polvani *et al* (1994) and Ambaum (1997), the relaxation time scale of a spectral component of the PV field is proportional to $1+(kL_R)^2$, where k is the magnitude of the horizontal wavenumber. For molecular diffusion, on the other hand, the relaxation time scale is proportional to $(\nu k^2)^{-1}$, where ν is the viscosity coefficient. Hence, the fine-scale structure naturally created by PV advection is not rapidly removed by such diabatic forcing, and thus the ability to retain it could be advantageous.

This appears to be the case in the one-layer diabatic CASL algorithm which we have recently developed and tested (Dritschel & Ambaum 1998). We used the simplest dynamical model, the QG equations in doubly-periodic geometry, to verify that a diabatic extension was possible. Those equations are given by Eqs.(1a–c) in §4, except that thermal forcing is added to Eq.(1a), i.e.

$$\frac{Dq}{Dt} = S$$

where

$$S = S_0(x,y,t) + \frac{\psi(x,y,t) - \psi_e(x,y)}{\tau L_R^2}.$$

Here, S_0 and ψ_e are prescribed, and ψ is the streamfunction obtained by Eq.(1b). The parameter τ is the radiative relaxation time scale (see Ambaum (1997) for details of this system).

The procedure is based on splitting the PV into a Lagrangian advective part q_a and an Eulerian diabatic part q_d which are evolved using contour advection and a conventional, semi-Lagrangian method, respectively. The goal is to uptake as much of the diabatic forcing as possible in q_a while preserving smooth contour evolution (and minimise computational cost). This is accomplished by advecting the q_a contours by an additional ficticious velocity $\mathbf{u}_d$ and allowing the q_a field to have time-varying PV jumps, i.e.

$$q_a = q_0(t) + q_1(t)\varpi(x,y,t)$$

where q_0 and q_1 are functions only of time and ϖ is the unnormalised PV field obtained from the contour-to-grid conversion with unit PV jumps on all contours. Whatever part of S that cannot be taken up by the contour motions and level changes is stored in q_d. This reservoir is transferred to q_a on a characteristic time scale (i.e. several days) through $\mathbf{u}_d$ and the level changes (dq_0/dt and dq_1/dt). The principal difficulty rests with the choice of $\mathbf{u}_d$. The choice that puts all of S into the evolution of q_a leads to singular $\mathbf{u}_d$ at the extrema of q_a. This is of course unacceptable, and it compels one to remove this singularity and therefore introduce a residual forcing for q_d. The choice of $\mathbf{u}_d$ (as well as dq_0/dt and dq_1/dt) involves a tension between trying to minimise the forcing of q_d and trying to preserve smooth contour evolution (i.e. strain rates in $\mathbf{u}_d$ no greater than those in $\mathbf{u}$). In practice, our tests show that q_d takes up only a small percentage of the total PV even when the radiative relaxation τ used to model the diabatic forcing is fast (several days) for typical atmospheric conditions. Further details, including numerical tests, may be found in Dritschel & Ambaum (1998).

We are presently working on a three-dimensional, multi-layer version of the algorithm to incorporate the diabatic forcing in a more realistic manner, and we plan to carry out a comparison with conventional numerical algorithms to determine if the retention of fine-scale PV in the CASL algorithm leads to a significant gain in accuracy for the computer cost required.

6. The future

The development of the CASL algorithm has created an opportunity for significantly improving atmospheric and oceanic modelling. It has now become possible to explicitly model a range of scales heretofore neglected, hence explicitly represent important eddy effects. More than that, the treatment of unresolved scales is far less artificial. Its use as a research tool may help increase our understanding of real atmospheric and ocean eddies and their contribution to large-scale circulations, chemical transport and mixing, ozone depletion, and climate change. Its further development may lead to its practical use as part of a global circulation model. Already, significant progress has been made on a nonhydrostatic algorithm, and the work of MD indicates how to improve the modelling of internal gravity waves in this instance. Efforts are turning to the outcropping problem, i.e. the intersection of isentropes or isopycnals with the boundaries (or other distinct fluid layers, such as the planetary boundary layer or the mixed layer). A number of other problems must be tackled as well, such as overturning isentropes or isopycnals, and generally how and where to interface aspects of the CASL algorithm with an operational global circulation model. It is clear, for example, that one would not want to use the CASL algorithm in the planetary boundary layer or elsewhere in the atmosphere or oceans where advection is secondary to other processes. Though it may be a challenging task, the development of a hybrid algorithm which deals with each process or part of the flow in the most efficient way as possible would appear to be worthwhile pursuing, considering the potential practical gains.

Acknowledgements

DGD is supported by the UK Natural Environment Research Council.

References

Ambaum, M. H. P. 1997 Isentropic formation of the tropopause. *J. Atmos. Sci.* **54**, 555–568.

Bartello, P. 1996 The cost-effectiveness of semi-Lagrangian advection. *Mon. Wea. Rev.* **124**, 2883–2897.

Charney, J. G. 1971 Geostrophic turbulence. *J. Atmos. Sci.* **28**, 1087–1095.

Dritschel, D. G. 1988 Contour dynamics/surgery on the sphere. *J. Comput. Phys.* **78**, 477–483.

Dritschel, D. G. 1989 Contour dynamics and contour surgery: numerical algorithms for extended, high-resolution modelling of vortex dynamics in two-dimensional, inviscid, incompressible flows. *Comp. Phys. Rep.* **10**, 77–146.

Dritschel, D. G. & Ambaum, M. H. P. 1997 A contour-advective semi-Lagrangian algorithm for the simulation of fine-scale conservative fields. *Q. J. Roy. Meteorol. Soc.* **123**, 1097–1130.

Dritschel, D. G. & Ambaum, M. H. P. 1998 The inclusion of non-conservative forcing into a conservative, contour advection algorithm. *In:* Numerical Methods for Fluid Dynamics VI. Ed. M.J. Baines, *ICFD*, Oxford.

Dritschel, D. G., Polvani, L. M. & Mohebalhojeh, A. R. 1998a The contour-advective semi-Lagrangian algorithm for the shallow-water equations. *Mon. Wea. Rev.* (in press).

Dritschel, D. G., de la Torre Juárez, M. & Ambaum, M. H. P. 1998b The three-dimensional vortical nature of atmospheric and oceanic turbulent flows. *Phys. Fluids* (submitted).

Dritschel, D. G. 1998 A fast hybrid algorithm for the simulation of fine-scale conservative fields on a sphere. *Mon. Wea. Rev.* (submitted).

Gill, A. E. 1982 Atmosphere-Ocean Dynamics. Academic Press, 662pp.

Gravel, S. 1996 The semi-Lagrangian method. *Centre de Recherches Mathematiques Proceedings*, Nov. 07–22 1995, Montréal.

Holton, J. R. 1982 An introduction to dynamic meteorology (3rd edition). San Diego, Academic Press, inc., 511pp.

Hoskins, B. J., McIntyre, M. E. & Robertson, A. W. 1985 On the use and significance of isentropic potential-vorticity maps. *Q. J. Roy. Meteorol. Soc.* **111**, 877–946.

Leslie, L. M. & Dietachmayer, G. S. 1997 Comparing schemes for integrating the Euler equations. *Mon. Wea. Rev.* **125**, 1687–1691.

McIntyre, M. E. & Roulstone, I. 1998 Hamiltonian balanced models: constraints, slow manifolds and velocity splitting. *J. Fluid Mech.*, (under revision).

Methven, J. & Hoskins, B. J. 1998 Spirals in potential vorticity. Part I: Measures of structure. *J. Atmos. Sci.* **55**, 2053–2066.

Mohebalhojeh, A. R. & Dritschel, D. G. 1998 On the representation of gravity waves in numerical models of the shallow water equations. *Q. J. Roy. Meteorol. Soc.* (submitted).

Norton, W. A. 1994 Breaking Rossby waves in a model stratosphere diagnosed by a vortex-following coordinate system and a technique for advecting material contours. *J. Atmos. Sci.* **51**, 654–673.

Polvani, L. M., Waugh, D. W. & Plumb, R. A. 1994 On the subtropical edge of the stratospheric surf zone. *J. Atmos. Sci.* **52**, 1288–1309.

Simmons, A. J. & Hoskins, B. J. 1976 Baroclinic instability on the sphere: Normal modes of the primitive and quasi-geostrophic equations. *J. Atmos. Sci.* **33**, 1454–1477.

Vallis, G. K. 1996 Potential vorticity inversion and balanced equations of motion for rotating and stratified flows. *Q. J. Roy. Meteorol. Soc.* **122**, 291–322.

Waugh, D. W. & Plumb, R. A. 1994 Contour advection with surgery: a technique for investigating finescale structure in tracer transport. *J. Atmos. Sci.* **51**, 530–540.