Offline Trajectory Package

John Methven

January 23, 1997

Contents

1 Introduction

This report describes the operation and performance of the trajectory package "Offline", which will be available for use through the British Atmospheric Data Centre (BADC). Details concerning access to Offline and ECMWF analyses can be found on the World Wide Web site run by the BADC (*http://www.badc.rl.ac.uk/*). The term **offline** indicates that the advecting wind fields are obtained from an external source. Currently Offline can cope with ECMWF upper air analyses and output from the University of Reading's spectral model. The trajectories are integrated forward by interpolating these 4-D wind records to the current particle positions, in time and space, and then using an **integrator** scheme. Additionally, the values of meteorological fields can be interpolated to the particle positions and assigned as **attributes** for the particles.

The next section outlines the basic operation of the trajectory package. The details of algorithms and accuracy tests with are given in Section 3. Several spatial interpolation methods are described in Section 3.2 and their accuracy is analysed in Section 3.3. In addition the magnitude of errors arising from the horizontal, vertical and temporal truncation of the wind record are compared. The reversibility of trajectories in an offline calculation is discussed in Section 3.4, where it is shown that trajectories are reversible to a very high degree of accuracy. This knowledge enables a demonstration that the differences between gridded PV-like tracers and fields obtained through domain filling trajectory calculations are due to non-conservation in the Eulerian tracer advection rather than errors in the trajectory calculations. Finally, in Section 4, a comparison is made between the new package and the back trajectory calculations used for the existing ECMWF datasets.

2 Outline of Operation

1. The records containing information on winds and attributes are read; upper air spectral records from the ECMWF are the default. These contain temperature, vorticity, divergence and the logarithm of surface pressure. They are transformed to grid point fields of

$$\dot{\lambda}_L, \quad \dot{\phi}_L, \quad \dot{\eta}_{L+\frac{1}{2}} \tag{1}$$

   where λ, ϕ and η describe the longitude, latitude and vertical coordinate of a particle and the dot denotes the Lagrangian derivative. The suffix L refers to the level of the data and $L+\frac{1}{2}$ indicates that vertical velocity is calculated on half-levels using the equation of continuity (see Simmons and Burridge (1981) and ECMWF Research Manual 2 (1988)).

2. Particles are positioned as desired. Clusters of trajectories can be released on model levels, pressure surfaces or isentropic surfaces.

3. The next wind record is read.

4. The position of each particle is integrated between the two wind records. This is achieved by interpolating the winds in space and time to the particle's position and using a suitable "integrator" scheme (Section 3.1). Note that forward or backward trajectories can be performed.

5. Attributes are assigned to particles by the interpolation of selected fields. Simple modifications allow the use, as an attribute, of any field which can be calculated from the data record.

6. Repeat steps 3 to 5 until the trajectories reach the desired length.

7. Particle positions and attributes are output.

3 Details of Algorithms and Accuracy Experiments

3.1 Integrator Schemes

Trajectory calculations involve solving the initial value problem for the ordinary differential equation,

$$\frac{d\mathbf{x}}{dt} = \mathbf{u}(\mathbf{x}, t). \tag{2}$$

The simplest method of numerical integration is to use the Euler method for each component. For example:–

$$x_{n+1} = x_n + u_n \delta t. \tag{3}$$

However this method is generally inaccurate compared to other methods using the same step-size, δt, and can be unstable (Press *et al.*, 1992) . For most problems a considerable improvement is made by using a 4th order Runge-Kutta method which involves evaluation of the velocity at four points for every time-step. Following Press *et al.* (1992):–

$$\begin{aligned}
k_1 &= u(t_n, x_n)\ \delta t \\
k_2 &= u(t_n + \frac{\delta t}{2}, x_n + \frac{k_1}{2})\ \delta t \\
k_3 &= u(t_n + \frac{\delta t}{2}, x_n + \frac{k_2}{2})\ \delta t \\
k_4 &= u(t_n + \delta t, x_n + k_3)\ \delta t \\
x_{n+1} &= x_n + \frac{k_1}{6} + \frac{k_2}{3} + \frac{k_3}{3} + \frac{k_4}{6}
\end{aligned} \tag{4}$$

The easiest way to implement this is to use a constant integrator time-step, δt. Note that δt is generally far smaller than the interval between wind records. The velocity at each point is determined by interpolation in space and time from the wind record. Linear interpolation in time, between the records bracketing the particles' positions (t^- and t^+), is used. All of the trajectories are integrated from t^- to t^+ before moving on to the next record so that only two time records are retained at any point in the calculation. Note that the winds are interpolated to the particle's position **at that time**:–

$$\mathbf{u}(\mathbf{x}(t), t) = (1 - s)\ \mathbf{u}(\mathbf{x}(t), t^-) + s\ \mathbf{u}(\mathbf{x}(t), t^+) \tag{5}$$

where the weight $s = \frac{(t-t^-)}{(t^+-t^-)}$. This method assumes that the Lagrangian spatial correlation scale of the velocity field is larger than the distance travelled by a particle over the interval $(t^+ - t^-)/2$. Work discussed in Section 3.3 shows that this is not too severe an approximation. The 4th order Runge-Kutta method has been used, with linear time interpolation, for trajectory calculations by many authors (e.g. Pierrehumbert and Yang (1993), O'Neill *et al.* (1994) and Sutton *et al.* (1994)).

For prolonged integrations a trajectory can keep within a specified accuracy with greater efficiency by using an adaptive step-size scheme. The scheme used in these tests is described fully in §15.2 of Press *et al.* (1992). However, the benefits of an adaptive scheme are only realised if the required accuracy is such that the ratio $\frac{(t^+-t^-)}{\langle \delta t \rangle} \gg 1$, where $\langle \delta t \rangle$ is the average step-size along a trajectory. This is generally not the case for atmospheric trajectories since errors arising from wind field truncation dominate (see below).

The Offline package integrates trajectories on the sphere (assuming a shallow atmosphere), in pressure based coordinates. For ECMWF analyses the vertical coordinate used is η (see ECMWF Research Manual 2 (1988)) and the vertical velocity $\dot{\eta}$ is calculated from the continuity equation. The horizontal wind components in spherical coordinates ($\dot{\lambda}$, $\dot{\phi}$) are readily obtained from the wind records. However, in order to avoid inaccuracies with advection near the poles the horizontal velocity (i.e. velocity tangent to the sphere) is decomposed into 3 Cartesian components with origin centred on the sphere (the z-axis is taken to pass through the poles). The trajectory calculation then amounts to the Runge-Kutta integration of four independent coordinates (in the manner of (4)):–

$$\frac{d}{dt}\left(\frac{x}{a}, \frac{y}{a}, \frac{z}{a}, \eta\right) = \left(-\cos\phi \sin\lambda\,\dot{\lambda} - \sin\phi\cos\lambda\,\dot{\phi},\ \cos\phi\cos\lambda\,\dot{\lambda} - \sin\phi\sin\lambda\,\dot{\phi},\ \cos\phi\,\dot{\phi},\ \dot{\eta}\right) \tag{6}$$

where the RHS is obtained from the wind field record.

3.2 Interpolation Schemes and Particle Attribution

3.2.1 Horizontal Interpolation

Fields must be interpolated to the positions of particles, for the advection itself, and in order to assign particle attributes. First the grid points bracketing a particle's position must be

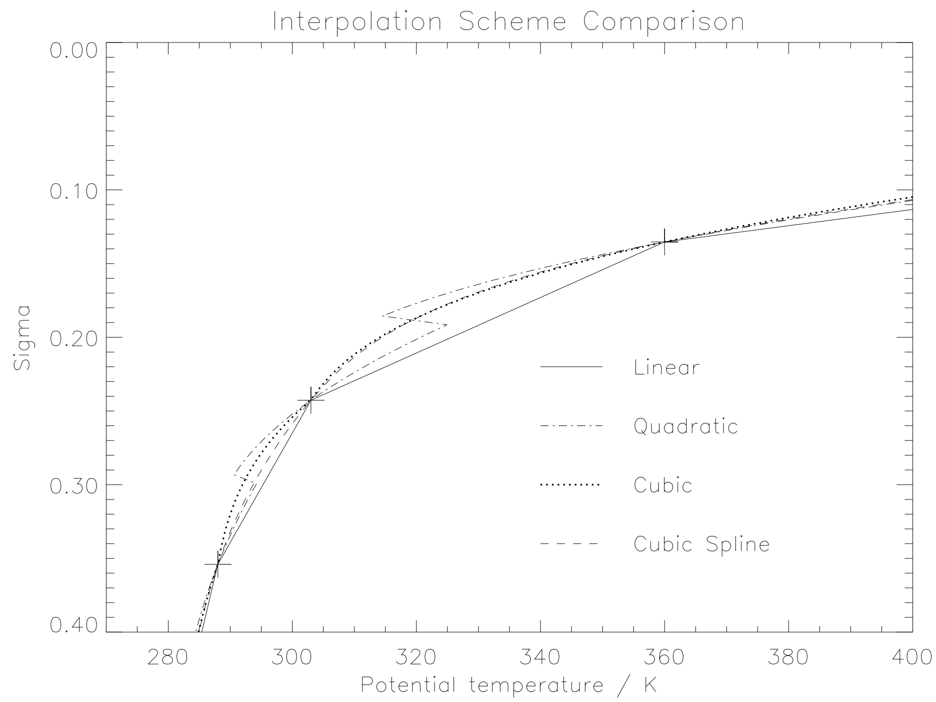

Figure 1: Part of a typical profile of potential temperature and its representation by linear, quadratic, cubic and cubic spline interpolation.

found. Model fields are regularly gridded in longitude so that the bracketing grid points can be found by dividing the particle's longitude by the grid spacing. However the Gaussian latitudes used by the model are not regularly spaced. The position of a particle in the array of latitudes is found using the 'hunt and locate' method described in §3.4 of Press *et al.* (1992). This method is also used to find the levels bracketing a particle.

Bilinear interpolation is then used to find the value of a field at the horizontal location of the particle on the model levels used for the vertical interpolation.

3.2.2 Vertical Interpolation

Linear interpolation in the vertical can be done in exactly the same way as interpolation in latitude. However this method was found to be insufficiently accurate due to the large curvature of the vertical profiles of fields near the tropopause. An improvement was sought by looking at a further 3 schemes: quadratic Lagrange, cubic Lagrange, and cubic spline interpolation. A selection of profiles of potential temperature near the top of the model in a baroclinic wave life cycle experiment (as used in Section 3.3) were used to compare the performance of each interpolation scheme (see Figure 1).

Since cubic splines enforce the first derivative of the field to be continuous at every point (§3.3 Press *et al.* (1992)) the profile obtained by this method appears to be particularly realistic, but the cubic spline calculation requires field values on all model levels for each new profile. In contrast linear, quadratic and cubic Lagrange interpolation only require the field values at the nearest 2, 3 or 4 levels respectively. For linear or cubic interpolation the points which bracket the particle change at model levels and thus the profile is continuous (for quadratic interpolation the three closest points change midway between levels). Since the cubic Lagrange and cubic spline profiles correspond far more closely to each other than to the linear profile, the most cost effective interpolation method which achieves desirable results is cubic Lagrange interpolation (calculated using Neville's algorithm, §3.1 Press *et al.* (1992)).

Near the upper and lower boundaries, the value of a field at a particle's location is found by cubic extrapolation using the closest four levels to the boundary. However for Reading spectral model data and ECMWF data the vertical velocity ($\dot{\eta}$) is calculated on half-levels (Hoskins and Simmons (1975) and Simmons and Burridge (1981)) and is identically equal to zero at the upper and lower boundaries ($\eta = 0, 1$). Therefore, when particles lie near the boundaries, the closest three half-levels and the boundary itself are used for the cubic interpolation. This makes it extremely unlikely that particles will be falsely advected out of the top or bottom of the domain (any particles which do manage this are 'switched off').

3.2.3 Isentropic Initialisation and Consistent Particle Attribution

When the potential temperature field is known on the model grid it can be found at a particle's position by interpolation. The vertical profile of potential temperature attribute is therefore defined by the interpolating polynomial. In order to initialise particles on an isentropic surface (θ_S), the vertical coordinate η is required from the knowledge that $(\theta - \theta_S) = P(\eta) = 0$. A linear interpolating polynomial can be inverted to find $\eta = P^{-1}(\theta - \theta_S)$. However this is not generally true for higher order polynomials. In this case the root of the polynomial can be found to a specified accuracy using the iterative secant method (§9.2, Press *et al.* (1992)). For the experiments using cubic vertical interpolation, a particle which is placed on surface θ_S using this method will have a θ-attribute equal to $\theta_S \pm 0.0001K$.

Particles are 'switched off' (ignored from then on) if the isentropic surface, as described by the interpolating polynomial, lies below ground ($\eta > 1$).

3.3 Accuracy Experiments

The velocities from an adiabatic, frictionless baroclinic wave life cycle experiment (the LC1 experiment of Thorncroft *et al.* (1993)) were used to conduct a series of accuracy experiments. Since adiabatic trajectories should stay on an isentropic surface, potential temperature was assigned as an attribute at the particles' positions. Particles were initially evenly spread across $\frac{1}{6}$[1] of a hemisphere on the 360K surface. They were then advected by the winds which were interpolated to the particle positions using cubic interpolation in the vertical, and linear interpolation in the horizontal and time (see last section). The deviation of trajectories from the 360K isentropic surface is used as an absolute error measure, ϵ_a, given by:–

[1] The baroclinic wave had six-fold zonal symmetry.

Label	NN	NL	Δt (hr)	Integrator Scheme	δt (hr)	Vertical Interpolation	ϵ_a (K)	ϵ_r (K)	δ_r (km)	CPU %
A1	85	60	3	Runge-Kutta	0.6	cubic	0.36	—	—	100
A2	85	60	6	Runge-Kutta	0.6	cubic	0.51	0.42	22	74
A3	85	60	12	Runge-Kutta	0.6	cubic	1.54	1.60	93	62
A4	85	60	24	Runge-Kutta	0.6	cubic	3.67	3.72	304	59
A5	85	30	3	Runge-Kutta	0.6	cubic	1.48	1.47	200	89
A6	85	15	3	Runge-Kutta	0.6	cubic	3.96	3.98	475	83
A7	42	60	3	Runge-Kutta	0.6	cubic	0.67	0.72	41	72
A8	21	60	3	Runge-Kutta	0.6	cubic	2.47	2.51	214	67
A9	42	15	3	Runge-Kutta	0.6	cubic	3.15	3.16	455	69
B1	85	30	6	Runge-Kutta	adaptive	cubic	1.38	—	—	100
B2	85	30	6	Runge-Kutta	0.2	cubic	1.38	0.00	0.0	25
B3	85	30	6	Runge-Kutta	0.6	cubic	1.38	0.01	0.1	10
B4	85	30	6	Runge-Kutta	2.0	cubic	1.39	0.06	1.1	6
B5	85	30	6	Runge-Kutta	6.0	cubic	1.50	0.68	12	4
B6	85	30	6	Euler	0.6	cubic	1.73	1.17	82	5
B7	85	30	6	Runge-Kutta	0.6	linear	1.52	0.52	38	4

Table 1: A table showing the horizontal, vertical and temporal truncation of the LC1 wind field records, as well as changes in the integrator scheme, used in the accuracy experiments. The set labelled A investigates the effect wind field truncation, whilst set B investigates changes in integrator scheme. The RMS deviation from the initial isentropic surface (over a two day integration) is used as an absolute error measure, ϵ_a. The RMS deviation in θ, ϵ_r, and the mean deviation in distance along a great circle, δ_r, are used as relative error measures. The computer time taken by each experiment is shown as a percentage of the time taken by the control.

$$\epsilon_a = \sqrt{\frac{\sum_i (\theta_i - 360)^2}{n}} \tag{7}$$

where n denotes the total number of particles (here $n = 471$). In order to further distinguish experiments, one run is designated as a control and the RMS difference between the θ-attributes for this run and a comparison run is used as a relative error measure, ϵ_r :–

$$\epsilon_r = \sqrt{\frac{\sum_i (\theta_i - \theta_i^{control})^2}{n}}. \tag{8}$$

The final comparative tool is the mean deviation in great circle distance, δ_r :–

$$\delta_r = \frac{\sum_i a \cos^{-1}(\frac{1}{a^2}\mathbf{r}_i.\mathbf{r}_i^{control})}{n} \tag{9}$$

where $\mathbf{r}_i$ is the position vector (from the centre of the Earth) of the particle with label i, and a is the Earth's radius.

In addition to integration errors, the truncation of the advecting wind field inevitably degrades the accuracy of any trajectory calculations. The use of less frequent records results in larger inaccuracies in time interpolation. Furthermore, records on a coarse grid exacerbate the spatial interpolation errors. In the first set of experiments (labelled A in Table 1) the records from a T85, L60 resolution LC1 experiment (storing records at intervals of 3 hours from day 6 to day 8) are further truncated temporally and spatially to estimate the errors incurred by truncation. The period investigated lies within the mature phase of the nonlinear baroclinic wave when its behaviour closely resembles that of synoptic scale weather systems (Thorncroft *et al.*, 1993) . Importantly the life cycle evolution is identical in all the derived records because they are truncated versions of the control. Horizontal truncation is achieved by chopping off the higher wavenumbers to the lower truncation limit. Vertical truncation is achieved by picking out every other level from the control record[2]. Note that the model levels in the control run are not evenly spread in η but are concentrated near the tropopause.

The error measures in Table 1 clearly illustrate that truncating the wind field temporally or spatially results in a degradation of trajectory accuracy. A systematic increase in severity of one aspect of truncation (e.g. temporal truncation only) results in an increase in all error measures which is faster than linear. Vertical truncation is seen to have the greatest impact on trajectory accuracy. Examination of ϵ_r for experiments A5$\rightarrow$A8 indicates that the impact of halving horizontal truncation is roughly half that of halving vertical truncation. Furthermore, ϵ_a and δ_r demonstrate that truncation to T42 has even less impact than ϵ_r might suggest. The relationship between the horizontal truncation and the displacement errors in trajectory calculations is connected to the scale effect of potential vorticity inversion and its consequences for the great influence of large scale features in the flow on the strain field. This was investigated in detail using the contour advection technique in Waugh and Plumb (1994) and Methven (1996).

Measure ϵ_a shows that halving the number of levels (from 60 to 30) is also roughly equivalent to using an interval between records (Δt) of 12 hours. Likewise, coarse-graining to 15 levels has a similar effect to using $\Delta t = 24\,\text{hr}$. Moreover one can see that the use of $\Delta t = 6\,\text{hr}$ is considerably better than halving horizontal truncation to T42, especially if the distance error is examined, whilst $\Delta t = 12\,\text{hr}$ is significantly worse than T42. Interestingly, coarsening the control in more than one aspect of truncation at once can result in an apparent improvement in accuracy. For instance, experiment A9 (T42 truncation) is more accurate that experiment A6 (T85 truncation). Furthermore, experiment B1 ($\Delta t = 6\,\text{hr}$) is more accurate than experiment A5 ($\Delta t = 3\,\text{hr}$). This is likely to be related to a mismatch between truncations. For instance, a reduction in the number of levels whilst retaining high horizontal resolution will result in jumps in θ near tropopause troughs, which then impact upon interpolation and trajectory accuracy. When the horizontal resolution is also decreased the steps are blurred out somewhat and trajectory accuracy may increase to some extent. Typically vertical and horizontal scales in the flow are related by Prandtl's ratio (f/N); spatial and temporal truncations are related by tight regions of high velocity (i.e. jets).

In the second set of experiments (labelled B in Table 1), three integrator schemes (Euler and 4th order Runge-Kutta (RK) with and without adaptive step-size) were tested along side

[2] The even numbered levels were picked because the large increase in θ near the top of the model ($L = 1$) resulted in a non-monotonic profile in θ-attribute if the odd numbered levels were taken.

each other using the wind field at T85, L30 resolution, storing records at intervals of 6 hours. The adaptive step-size scheme was used as the control[3]. Their performance, as gauged by the error measures is also shown. It is immediately apparent that increases in the step-size of the Runge-Kutta scheme (δt) incur errors which are insignificant when compared to wind field truncation errors (until $\delta t \approx \Delta t$), although they have a large impact on cost. It was also found (not shown here) that the use of the 3D Cartesian components for the integration of longitude and latitude (6) had an insignificant impact in this experiment. However there is virtually no flow near the poles; the scheme was found to make a difference when using ECMWF analyses where trajectories can cross the poles. Experiment B6 illustrates that an Euler stepping scheme performs very poorly compared to a Runge-Kutta scheme, incurring relative distance errors which were almost as severe as using $\Delta t = 12\,\mathrm{hr}$. Linear vertical interpolation was also found to perform poorly, compared to cubic interpolation, resulting in relative distance errors greater than those incurred by increasing Δt from 3 hr to 6 hr. At lower vertical resolutions the linear scheme was found to be even worse (with T42, L15 winds ϵ_a was over twice as large when using linear vertical interpolation rather than cubic).

All the above results apply to the upper levels of a synoptic scale weather system; other flows are not considered here. However Waugh and Plumb (1994) found, by qualitative comparison with accurate contour dynamics calculations of flow around the stratospheric vortex, that $\Delta t = 12\,\mathrm{hr}$ was sufficient to accurately simulate small scale tracer filaments using the contour advection technique. Furthermore O'Neill *et al.* (1994) qualitatively compared the use of $\Delta t = 6\,\mathrm{hr}$ and $\Delta t = 24\,\mathrm{hr}$ in the advection of PV in the mid-stratosphere and found that $\Delta t = 24\,\mathrm{hr}$ was sufficiently accurate for their needs. Tracer advection in the stratosphere is dominated by large scale strain to an even greater extent than for synoptic scale systems. Also the polar vortex evolves more slowly than tropospheric systems so it comes as no surprise that tracer studies in the stratosphere can make do with coarser temporal and spatial resolution of the wind field than required for the baroclinic life cycles.

This contrasts with the results of Doty and Perkey (1993) who performed 3-D trajectory calculations using winds from a hydrostatic, mesoscale model simulation of the vigorous extratropical cyclone ERICA IOP4. As a control trajectory calculation they used $\Delta t = 0.25\,\mathrm{hr}$ for a 3 day integration. They found that the median horizontal distance error (comparing with the control) increased from 21 km for $\Delta t = 1$ hr to 174 km for $\Delta t = 3$ hr and thus concluded that the largest interval that should be used for such trajectory calculations is 1 hr. However their trajectories require a wind record which is updated more frequently because their model resolves strong, vertical motions of up to $20 cm\, s^{-1}$ (associated with latent heat release during cyclogenesis) which have short spatial and temporal scales. Also their trajectory calculations only use linear vertical interpolation (with $\Delta z = 1\,km$) and an integrator scheme which is less accurate than the RK one and therefore their errors from sources other than temporal truncation will be larger.

[3] The adaptive step-size scheme used an initial time-step of 0.48 hr and the fractional accuracy in step distance, $u\,\delta t$, was set at 10^{-7} (see p.557, Press *et al.* (1992)), which was found to keep the step-size, averaging over trajectories, near 0.5 hr.

3.4 Reversibility of Trajectories

The reversibility of trajectories was investigated using wind field records from the LC1 experiment (T42, L15). A cluster of 25 particles (with a 5° spread) was started at day 5, just behind the surface warm front. This cluster rapidly dispersed and the trajectories were calculated to day 12, by which time the baroclinic wave had occluded. Back trajectories were then calculated by reversing the order of the wind field record and changing the sign of velocity. The particles were advected from their final day 12 positions back to day 5. The RMS deviations of the back trajectories from the forward trajectories (at day 5) for longitude, latitude and σ-coordinate are only 0.017°, 0.008° and 3.7×10^{-5} respectively.

The trajectories are therefore reversible in the sense that the distance errors accumulated over the trajectory for 7 days forwards and 7 days backwards are very small compared to the wind field grid spacing. Sutton *et al.* (1994) have also found trajectory calculations to be reversible in the sense implied here. By advecting particles backwards for 5 days and then forwards 5 days, in the mid-stratosphere, they found that 78% of trajectories had net displacements of less than $100\,km$ and 95% had net displacements of less than $400\,km$. Reversibility does not directly imply much about the accuracy of wind interpolation, or the method of integration, since the errors in position at day 12 cannot be deduced by this method. However, in the next section, the reversibility of trajectories is used, in conjunction with other results concerning θ and PV particle attributes, to conclude that errors in the domain filling trajectory picture of PV resulting from wind field truncation are significantly smaller than diabatic effects on the PV **field**[4] arising from hyper-diffusion.

3.5 Relative Magnitudes of Errors in Representations of Closely Conserved Tracers

In this section two representations of the potential vorticity (PV) on the 360K isentropic surface are compared. In the first case a tracer field was set equal to PV at day 0 of the LC1 experiment (at T42, L15 resolution). The advection-diffusion equation was then integrated (to day 8) using the same diffusion coefficient as used for vorticity, divergence and temperature in the primitive equation model (equivalent to a 4 hour diffusion timescale on the highest retained wavenumber). This tracer will be referred to as a PV-like tracer. At day 5 its distribution is wave-like (Figure 2a) and by day 8 its distribution is similar to Figure 2d. In the second case particles were positioned on the 360K surface on day 5 of the LC1 experiment. Their arrangement was such that they were homogeneously spread on a polar stereographic projection with an inter-particle spacing of about 1.33° along a great circle. The values of the PV-like tracer field (and θ) at the particle positions, were assigned as particle attributes (see Figure 2a). The offline trajectory code (with $\Delta t = 6\,\mathrm{hr}$) was then used to advect the particles and the second PV picture obtained by using the value of PV-attribute from day 5 at the positions of the particles at day 8 (Figure 2c). This representation will be referred to as the fixed attribute picture or domain filling trajectory picture.

[4] i.e. diagnosed from the vorticity, divergence and temperature.

The fixed attribute picture was seen to exhibit far tighter gradients than the PV-like tracer field. In addition the attribute picture showed no sign of a cut-off blob of low PV at the poleward end of the ridge and showed a clear signature of cyclonic wrapping of PV to the poleward side of the jet. Unfortunately this experiment alone is insufficient to assign the differences in the pictures to truncation and diffusion of the gridded fields; the particle advection itself may be too inaccurate. However by calculating some back trajectories and invoking the time reversibility of the trajectories some more positive conclusions can be drawn.

Assume that hypothetical surfaces exist which describe the "true" surfaces of PV and θ in the absence of diabatic processes. The surfaces of PV and θ, as described by the model's fields, are displaced by the effects of diffusion and truncation. The θ-surface, as described by the domain filling forward trajectories, is influenced by two factors: the position of the model's θ-surface at the time of particle initialisation and errors in particle advection over the course of the integration. Figure 3 illustrates the vertical location of each of these surfaces, following the horizontal location of a 'true' parcel of fluid over the 3 day time interval. For instance, by day 8 the forward trajectory (curve OA) has diverged from the model isentropic surface (now at point B), and both of these endpoints have been influenced by non-conservative processes. Moreover the height of the PV surface has also been influenced by non-conservative processes and ends up at point D.

The adiabatic and frictionless flow in the life cycles should be reversible in a kinematic sense; the only processes which are not strictly reversible are the numerical diffusion and truncation. Back trajectories were started at day 8, using the same initial horizontal distribution of particles as for the forward trajectories but placing them on the model's 360K surface **at this time** (point B in Figure 3). The integration was run backwards to day 5 (arriving at point C). The value of the PV-like tracer field at each particles' position was recorded every 6 hours (when the winds were updated), just as for the forward trajectory. Now the following notation will be used to describe particle attributes:–

$F_{t_a}^{f/b}$ where F is an attribute field (PV (P) or θ),
t_a denotes the day on which the attribute was assigned,
and f/b indicates a forward trajectory (OA) or a backward trajectory (BC).

The relative magnitudes of the distances OC, AB and BD can be deduced from three salient observations:–

1. An examination of the deviations of θ_5^b from 360K shows them to be the same as the deviations of θ_8^f from 360K, but with the opposite sign (i.e. $(\theta_5^b - 360) \approx -(\theta_8^f - 360)$). This observation can only be explained if the vertical deviation of the model's isentrope from the position of the forward trajectory (distance AB) is small compared to the height scale of the vertical wind shear. In this situation the horizontal winds at points A and B will be very similar. Since trajectories are reversible to a high degree of accuracy

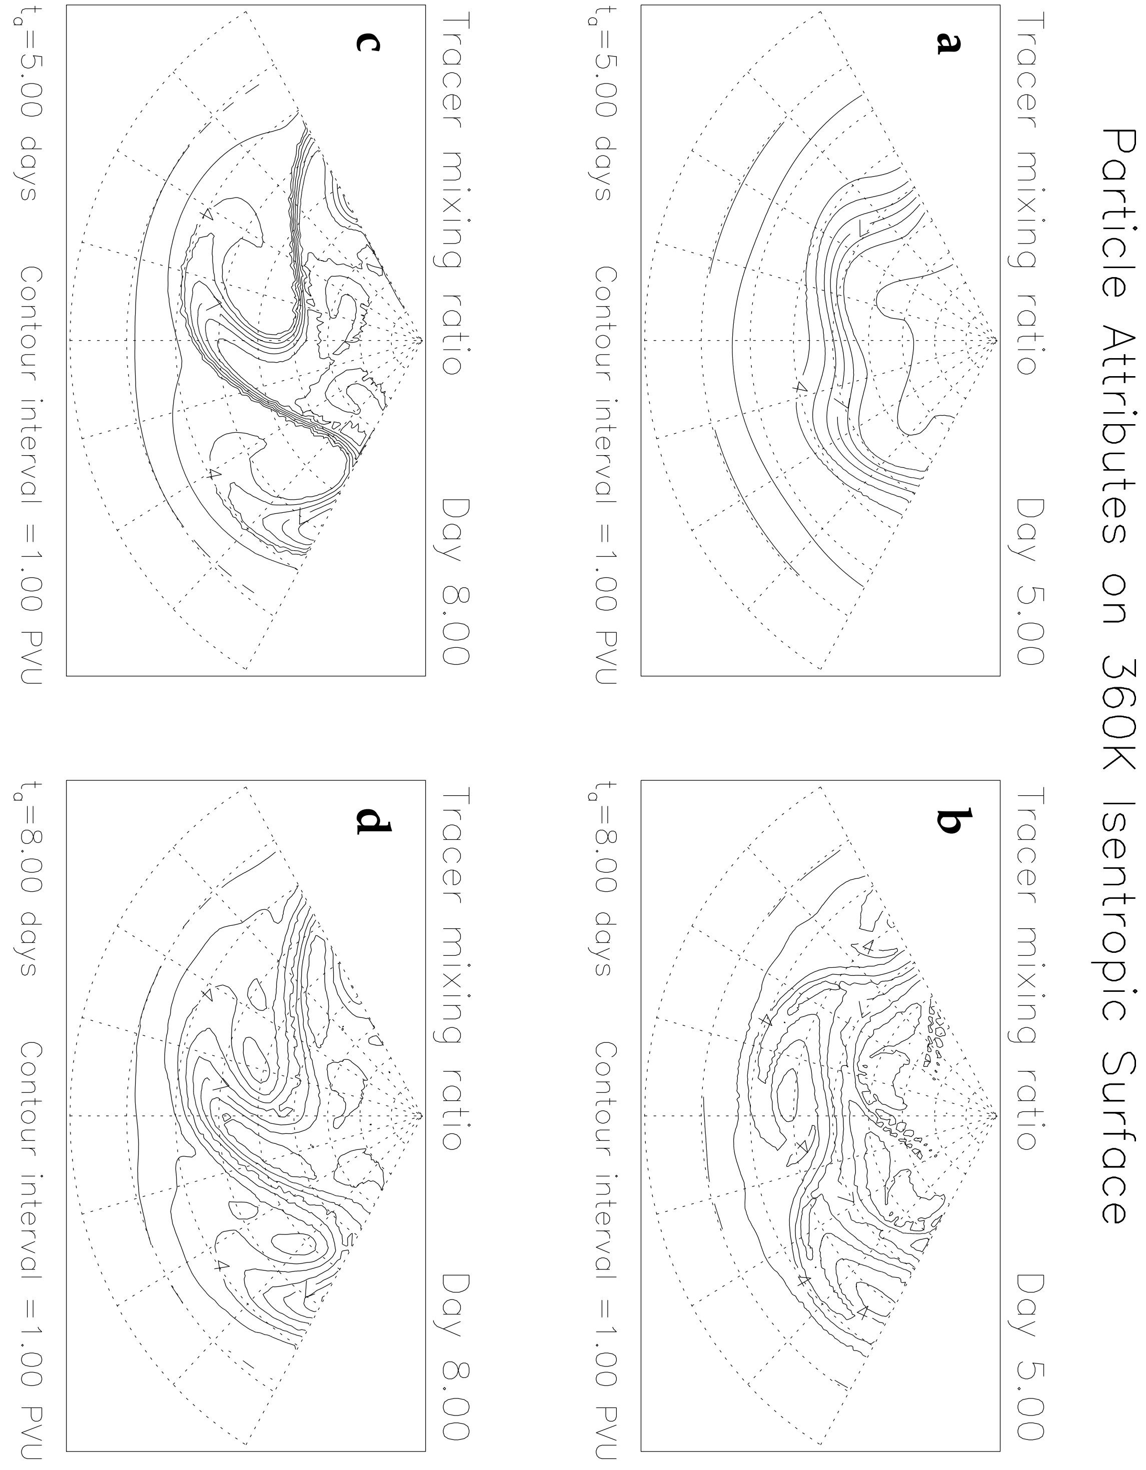

Figure 2: (**a**) and (**b**) show the PV attributes P_5^f and P_8^f (notation in text) at the particles' positions for day 5 of the forward trajectories. (**c**) and (**d**) show the same two attributes but at the particles' positions for day 8 of the forward trajectories. The PV-contours are determined by interpolating the particle attributes onto a regular lat-lon grid before contouring.

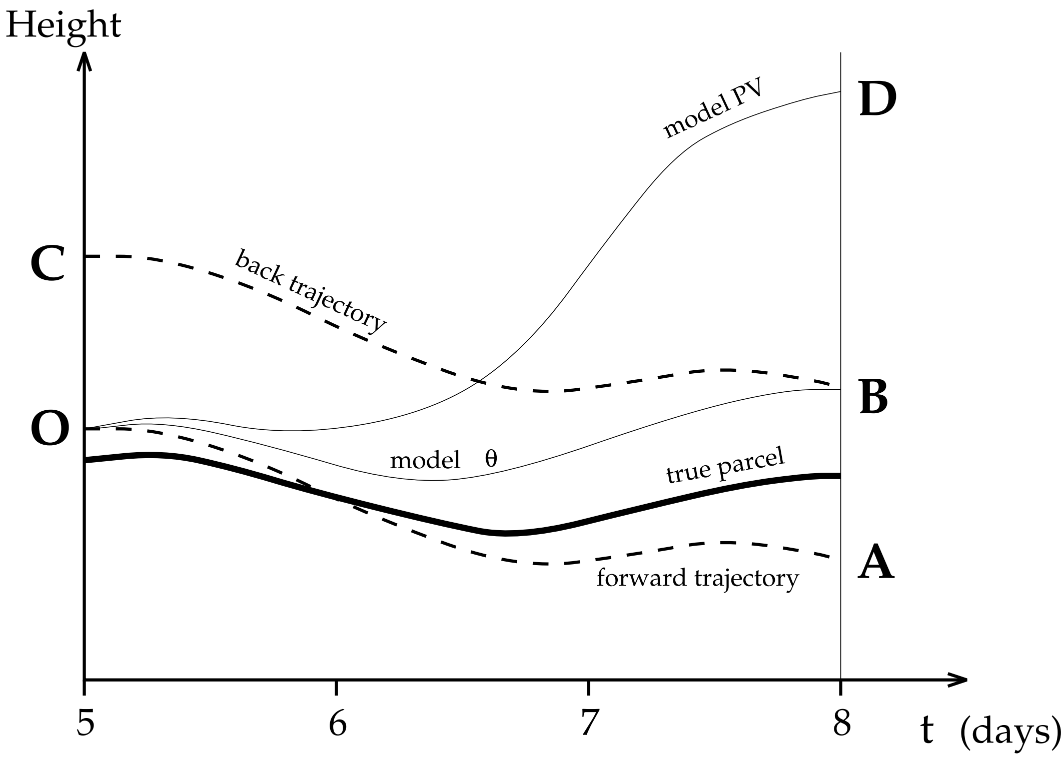

Figure 3: The vertical location of PV and θ surfaces, as described by the model variables and trajectories, following the horizontal location of a fluid parcel.

the backward trajectory (BC) should lie almost parallel to the forward trajectory (OA), and thus distance OC is then approximately equal to AB, assuming a roughly uniform static stability following the parcel.

2. The PV-attributes assigned at day t_a are approximately equal for the forward and backward trajectories (i.e. $P^f_{t_a} \approx P^b_{t_a}$). This observation can be explained if OC and AB and both are small compared to a scale height in the background PV variation.

3. On any one day $P^f_8 \neq P^f_5$ as mentioned earlier. Whilst at day 8 a plot of P^f_5 (Figure 2c) shows an obvious tongue of PV extending from the subtropics, P^f_8 (Figure 2d) shows cut-off, low PV which cannot arise purely from advection and therefore must be a feature of diffusion on the PV-like tracer field. Moreover, at day 5, the plot of P^f_8 (Figure 2b) is very unrealistic and does not resemble the slightly nonlinear baroclinic wave expected in PV at this time (Figure 2a).

   Since attributes assigned at day 5 give far more realistic distributions of PV than attributes assigned at day 8, $BD \gg OC \approx AB$.

In summary the non-conservation of PV arising from truncation and diffusion on PV-like tracers far exceeds the effective non-conservation due to errors in the trajectory calculations[5]. Also the errors coming from the initialisation of particles on the model's

[5] It must be pointed out that these conclusions only hold when using at least 15 levels in the vertical and

isentropic surface are much smaller than the errors from PV diffusion. Therefore domain filling trajectories are far better at reproducing the features of an approximately conserved tracer than gridded fields when the grid spacing (for winds and PV-like tracer) and the initial particle spacing are approximately equal.

3.6 Summary

In this section it was shown that, of the options tested, the best components for an accurate trajectory calculation were: a 4th order Runge-Kutta trajectory integrator (with constant step-size), cubic Lagrange vertical interpolation, and linear horizontal and temporal interpolation of the wind field. A RK step-size (δt) of 0.6 hr was sufficient to keep the integrator errors at an insignificant level. The wind fields must always be truncated at some limit. At resolutions generally considered high for analyses (T85, L30 in these experiments or T106, L31 for ECMWF analyses) the main sources of error for trajectory calculations will generally lie with the vertical truncation. Updating the winds every 6 hours was found to reduce temporal truncation errors to a magnitude less than those likely to arise from horizontal truncation when studying synoptic scale weather systems. However, the influence of latent heat release can result in a contraction of the wind field's space and time scales; a trajectory calculation would then require higher resolution analyses to achieve the same accuracy (e.g. Doty and Perkey (1993)). Finally, it was demonstrated that domain filling trajectory representations of closely conserved tracers are quantitatively more accurate than representations obtained by integrating the advection-diffusion equation at the same resolution as the wind field.

4 Comparison with Existing ECMWF Trajectory Datasets

A database of routine back trajectories has been established at the British Atmospheric Data Centre (BADC). These were calculated as part of a special project at the ECMWF, initiated by the University of Reading Meteorology Department. During the whole of 1995, 252 back trajectories were computed on a daily basis. The trajectories were released from 3 clusters centred over western Europe, the mid-Atlantic storm track region and the eastern USA; all were located at 900hPa. This project then continued until the end of September 1996, computing the trajectories once every 6 hours.

During 1995, the spectral (T213, L31), uninitialised, operational analyses from the ECMWF, were transformed directly onto a $1.5° \times 1.5°$ grid, and used for advection. These winds were updated every 6 hours and linearly interpolated in time and space to the particle positions. The vertical interpolation and integration was performed in pressure coordinates using ω (on full model levels) for the vertical velocity. A simple mid-point integrator was used with a time-step of 0.25 hr.

cubic interpolation of fields to the particles' positions. Coarser vertical resolution or less accurate vertical interpolation may result in a major degradation of the trajectory calculation accuracy and in this situation $BD \sim AB$.

Label	Description	T (days)	δ_r (km)	Δr_{max} (km)	Δp_{med} (hPa)	Δp_{max} (hPa)
96	Using 1996 scheme	—	—	—	—	—
95	Existing 1995 trajectories	2	~ 20	100–500	~ 2	40–60
		5	150–200	2000–6000	~ 8	200–500
C1	T106 $\rightarrow \dot{\eta}$, quadratic Gaussian grid	—	—	—	—	—
C2	T106 $\rightarrow \omega$, quadratic Gaussian grid	2	30	300	8	70
C3	T106 $\rightarrow \omega$, linear Gaussian grid	2	120	1100	11	230
C4	T106 $\rightarrow \omega$, $1.5° \times 1.5°$ grid	2	110	1100	12	220
C5	T213 $\rightarrow \omega$, $1.5° \times 1.5°$ grid	2	270	2800	22	250
C6	Existing 1996 trajectories	2	300	2600	20	240
		5	1200	4500	80	500
D1	Difference between C5 and C6	2	110	800	11	110

Table 2: In the top section of the table the 1995 and 1996 ECMWF trajectory schemes are compared using four error measures: the mean distance deviation, δ_r, the maximum distance deviation, Δr_{max}, the median pressure deviation, Δp_{med}, and the maximum pressure deviation, Δp_{max}. In the lower section the Offline package is compared to trajectories from the 1996 data set. In addition, several modifications are made to Offline in order to ascertain the main sources of the differences.

In 1996 a number of changes were made. Firstly, the initialised analyses were used (at the same resolution as before). The vertical interpolation scheme was changed to a cubic one (as in Section 3.2), and a fourth order Runge-Kutta integrator was used with a step-size of 0.5 hr. Tests using the LC1 experiment established that although the midpoint scheme was much better than a simple Euler step, a Runge-Kutta scheme was preferable. However, the main difference between the 1995 and 1996 schemes lies with the vertical interpolation. Table 2 shows a comparison between the two schemes using the uninitialised analyses from December 1995. Four error measures are shown: the mean distance deviation, δ_r, the maximum distance deviation, Δr_{max}, the median pressure deviation, Δp_{med}, and the maximum pressure deviation, Δp_{max}. The figures indicate the range of values seen using different release dates throughout the month. After 2 days, δ_r indicates that the differences between the schemes are smaller than those likely to arise from wind field truncaton (cf. Table 1).

The Offline trajectory code was compared with the 1996 trajectory data set by computing 5 day back trajectories from 00Z on 6/6/96. Offline differs from the 1996 trajectory scheme in several respects: the interpolation is done in η-coordinates rather than pressure, $\dot{\eta}$ (on half-levels) is used for vertical velocity, and a 3-D Cartesian coordinate basis is used for the advection on the sphere. A set of experiments (labelled C in Table 2) were performed to find the most significant differences between Offline and the 1996 scheme. C1 denotes the control experiment, using the standard features of Offline described above. Most users only have access to T106 resolution winds; these were therefore used for the control, transforming them onto a quadratic Gaussian grid[6]. Experiment C2 demonstrates that the use of ω on full levels[7] (rather than $\dot{\eta}$ on half-levels) gives rise to differences which are slightly larger than

[6] This grid is usually used for spectral transform models and at T106 resolution is approximately a $1.125° \times 1.125°$ grid.

[7] Here ω was taken from the upper air spectral record and transformed onto the quadratic Gaussian grid.

arose from the change in interpolation scheme between 1995 and 1996. In experiments C3 and C4, T106 winds were transformed onto a linear Gaussian grid and a regular $1.5° \times 1.5°$ grid respectively; these grids are almost identical and the differences between the experiments are correspondingly small. However, when both of these experiments are compared with the control, δ_r is several times larger than that seen in the horizontal truncation experiments for the life cycle (set A in Table 1). This suggests that small scale features in the flow play a more major role in the real atmosphere than in the baroclinic wave experiment.

Unfortunately, the differences between the existing trajectory data set (C6) and Offline (C1) are still several times greater than between experiments C4 and C1. The major cause of the discrepancy has now been narrowed down to the use of T213, rather than T106, wind fields transformed directly onto the regular $1.5° \times 1.5°$ grid. In experiment C5 exactly the same winds were used as in 1996 and indeed the differences, relative to the control (C1), were found to be as large as for the existing set (C6). Moreover, the differences between C5 and C6 were two or three times smaller (see D1).

These experiments highlight the danger in inferring the origin of a trace constituent using a single back trajectory from a release point. Trajectories are sensitive not only to their initial conditions but also to errors in the trajectory calculation, particularly those due to slight changes in wind field data. The tables shown above illustrate the magnitude of distance errors which should be expected. However when considering the clusters of trajectories as a whole, their behaviour is relatively insensitive to trajectory calculation changes. For instance, Figure 4 shows the probability density of the origin of all 1176, 5 day trajectories which arrive daily at 12UT, between 1/12/95 and 7/12/95, over western Europe; the 1995 and 1996 trajectory schemes give remarkably similar results, even though individual trajectories may be several thousand kilometres in error. These kind of diagnostics would be useful for comparison with observations of chemical constituents.

References

Doty, K. and Perkey, D. (1993). Sensitivity of trajectory calculations to the temporal frequency of wind data. *Mon.Wea.Rev.*, **121**, 387–401.

ECMWF Research Manual 2 (1988). *ECMWF Forecast Model: Adiabatic Part.*

Hoskins, B. and Simmons, A. (1975). A multi-layer spectral model and the semi-implicit method. *Quart.J.Roy.Met.Soc.*, **101**, 637–655.

Methven, J. (1996). *Tracer behaviour in baroclinic waves.* Ph.D. thesis, Reading University.

O'Neill, A., Grose, W., Pope, V., MacLean, H., and Swinbank, R. (1994). Evolution of the stratosphere during Northern winter 1991/92 as diagnosed from U.K.M.O. analyses. *J.Atmos.Sci.*, **51**, 2800–2817.

Pierrehumbert, R. and Yang, H. (1993). Global chaotic mixing on isentropic surfaces. *J.Atmos.Sci.*, **50**, 2462–2480.

Press, W., Flannery, B., Teukolsky, S., and Vetterling, W. (1992). *Numerical Recipes: The Art of Scientific Computing.* Cambridge.

The interpolation to particle positions was then performed in η-coordinates.

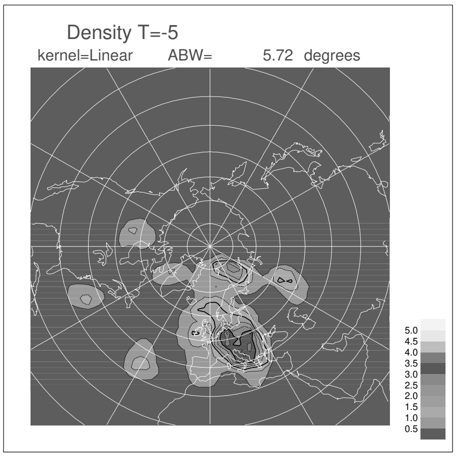

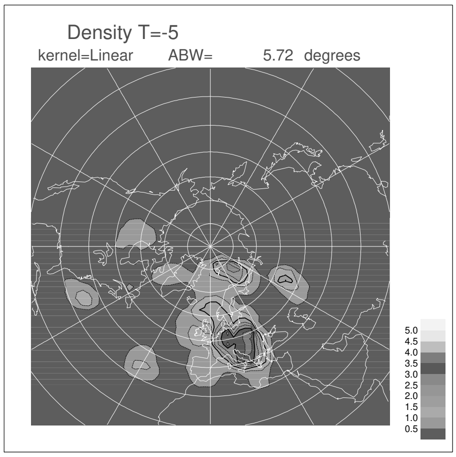

Figure 4: Both plots show the probability density of the origin of all 1176, 5 day, routine back trajectories which arrive daily at 12UT, between 1/12/95 and 7/12/95, over western Europe. The left distribution is compiled from trajectories using the 1995 scheme and the right one from trajectories using the 1996 scheme.

Simmons, A. and Burridge, D. (1981). An energy and angular-momentum conserving vertical finite-difference scheme in hybrid vertical coordinates. *Mon.Wea.Rev.*, **109**, 758–766.

Sutton, R., MacLean, H., Swinbank, R., and O'Neill, A. (1994). High-resolution stratospheric tracer fields estimated from satellite observations using Lagrangian trajectory calculations. *J.Atmos.Sci.*, **51**, 2995–3005.

Thorncroft, C., Hoskins, B., and McIntyre, M. (1993). Two paradigms of baroclinic-wave life-cycle behaviour. *Quart.J.Roy.Met.Soc.*, **119**, 17–55.

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

Acknowledgments

Thankyou to Brian Hoskins and Paul Berrisford who worked on the ECMWF trajectory project (Special Project SPGBRBT). Paul retrieved the wind data as well as calculating the routine back trajectories. Much of the code for reading ECMWF data was taken from Roger Brugge's diagnostic tool, Package X, which in turn uses decoding routines written by Jeff Cole and the ECMWF. Thanks also to Lesley Gray, Mike Bithell, Andy Smith, Chunkey Lepine, Sam Pepler and Peter Chiu at BADC for providing helpful suggestions and aiding the data retrieval. Kevin Hodges created the probability density algorithm and provided the programme for it, for which I am grateful. This work was started during my PhD, funded by a NERC studentship, and has continued as part of a project supported by the Hadley Centre.