

A Lagrangian trajectory filter for constituent data assimilation

By PETER M. LYSTER^{1*}, STEPHEN E. COHN², BANGLIN ZHANG³, LANG-PING CHANG³,
RICHARD MÉNARD⁴, KEVIN OLSON⁵ and ROBERT RENKA⁶

¹*Earth Systems Science Interdisciplinary Center, University of Maryland at College Park, USA*

²*GMAO, NASA Goddard Space Flight Center, USA*

³*Science Applications International Corp., Calverton, Maryland, USA*

⁴*Air Quality Research Branch, Meteorological Service of Canada*

⁵*Goddard Earth Science and Technology Center, University of Maryland Baltimore County, USA*

⁶*Department of Computer Science and Engineering, University of North Texas, USA*

(Received 30 December 2002; revised 22 December 2003)

SUMMARY

We have developed a new numerical algorithm, the Lagrangian filter, for solving the Kalman filter equations for assimilation of constituent observations directly on trajectories that propagate with the flow. This is a finite-dimensional approximation of the solution of the state estimation problem by characteristics, and may be thought of as an extension of methods such as trajectory mapping. The Lagrangian filter provides a natural framework for the study and solution of the constituent data assimilation problem because of the conservative properties of the state and its estimation-error variance and covariance along trajectories. Considerable insight into the behaviour of the filter is gained as a result of these properties. The Lagrangian filter also requires significantly fewer floating point operations than the Eulerian Kalman filter because of its simple error covariance propagation step. We implemented it for two-dimensional flow on isentropes in the stratosphere and assimilated methane observations from the Upper Atmosphere Research Satellite. Results are validated against those of the Eulerian filter.

KEYWORDS: Characteristics Estimation theory Kalman filter

1. INTRODUCTION

Lyster *et al.* (1997), Ménard *et al.* (2000, hereafter MCCL) and Ménard and Chang (2000, hereafter MC) developed a Kalman filter for the assimilation of chemical constituent observations based on a Eulerian numerical model for the dynamics. Following the work of Cohn (1993, 1997), the continuum dynamics of both the constituent field and its error covariance can be formulated using the method of characteristics. An approximate numerical implementation of this formulation uses Lagrangian trajectories determined by the wind field. Here we develop a Lagrangian Kalman filter algorithm based on this formulation. The analysis update is performed on an unstructured mesh that is generated using a Delaunay triangulation of the trajectory locations. The mapping from the trajectory locations to the observation locations, i.e. the observation operator, uses this triangulation for linear or higher-order interpolation.

Earlier work that uses Lagrangian trajectories to study stratospheric processes includes trajectory mapping (Sutton *et al.* 1994; Morris *et al.* 1995). In that approach, field values on trajectories are initialized by interpolating once from nearby observations. The value on each trajectory is not updated sequentially using subsequent observations, as in the Lagrangian filter, but rather trajectories are deemed unreliable after some period of time. The methods of contour advection and forward/reverse domain filling (Dritschel 1988; Fisher *et al.* 1993; O'Neill *et al.* 1994; Waugh and Plumb 1994; Schoeberl and Newman 1995) also use Lagrangian trajectories. In contour advection, specific contours of a field quantity that is conserved with the flow are followed using trajectories. In forward/reverse domain filling, clusters of space-filling trajectories are used to reveal the structure of the evolving fields. Fisher and Lary (1995) and Khattatov *et al.* (1999)

* Corresponding author, present affiliation: Code 900.3, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA. e-mail: lysterp@mail.nih.gov

use trajectories with photochemical box models in four-dimensional variational formulations. These works do not fully account for the developing errors in the fields. The Lagrangian filter is like domain-filling methods in that trajectories fill the volume of the simulation, and it is like trajectory mapping in that observations are used to adjust the value of the mixing ratio on individual trajectories. The Lagrangian filter is also a sequential estimation algorithm that accounts for errors in the dynamical model, the initial field, and the observations. The mixing ratio on trajectories is thereby updated appropriately as new observations become available, along with an estimate of the accuracy of the result. We carry the full spatial-error covariance of the mixing-ratio field.

Our development of the Lagrangian filter was motivated in part by the finding of MCCL that Eulerian filtering gives rise to a large, spurious loss of variance that must be corrected explicitly to conform to the continuum theory of error covariance propagation. The Lagrangian filter by its nature does not have this problem. The present work also indicates that the loss of variance results primarily from the cascade of correlation structure to fine scales, due to shear of the advecting wind, that are not resolved by a fixed Eulerian grid. This structure is dissipated away by the Eulerian filter, but is represented directly by the Lagrangian filter. We also found that the Lagrangian state and error covariance propagation, being inherently non-diffusive, give rise to sharper gradients in the mixing-ratio field than those of the Eulerian filter, at the cost of some graininess. The Lagrangian filter is also significantly less expensive computationally than the corresponding Eulerian filter, because the forecast step is achieved simply by propagating trajectories and adding the model-error covariance matrix to the analysis-error covariance matrix.

Section 2 introduces the Lagrangian filter and shows how the numerical algorithm is formulated in terms of a finite number of trajectories. We also describe the representativeness error that arises in the finite-dimensional approximation. Section 3 discusses the numerical implementation including the Delaunay triangulation algorithm for implementing interpolation between the unstructured grid of trajectory locations and arbitrary observation points on the sphere. Section 4 discusses the experimental results. We present tests of the Lagrangian filter including comparisons with the Eulerian filter using a baseline set of data, viz. methane retrieved from the Cryogenic Limb Array Etalon Spectrometer (CLAES) on board the Upper Atmosphere Research Satellite (UARS) and winds from the Goddard Earth Observing System Version 1 Data Assimilation System (GEOS-1 DAS) (GMAO 2000). Section 5 summarizes the results.

2. THE LAGRANGIAN FILTER

(a) *The continuum Kalman filter*

The Kalman filter can be derived in a number of contexts (e.g. Ghil *et al.* 1981; Jazwinski 1970). Cohn (1997) derived equations for the filter where the state and error covariances are considered to be functions defined on the continuum, rather than as discrete vectors. We use this formulation and also some of the ideas of Cohn (1993), which approach the solution of both the state equation for the advection of the mixing ratio and the corresponding state estimation-error covariance evolution equation by using the method of characteristics, and include the model-error covariance function in a consistent manner. The numerical method that we have developed following this approach uses Lagrangian trajectories. This is a discrete approximation to the solution by characteristics, in the sense of having a finite number N of trajectories, thus enabling the algorithm to be implementable on a computer.

The continuum formulation of the Kalman filter attempts to estimate the constituent mixing ratio $w(\mathbf{x}, t)$ whose dynamics are governed by

$$\left(\frac{\partial}{\partial t} + \mathbf{v}(\mathbf{x}, t) \cdot \nabla \right) w(\mathbf{x}, t) = f(\mathbf{x}, t), \tag{1}$$

where $\mathbf{v}(\mathbf{x}, t)$ denotes a prescribed deterministic velocity field, and $f(\mathbf{x}, t)$ is a random function that accounts for unknown forcing which we call the model error, assumed to be Gaussian-distributed in the space coordinates, with zero mean. In general, both \mathbf{x} and \mathbf{v} are three-dimensional. However, Lyster *et al.* (1997) and subsequently MCCL and MC considered the dynamics to be restricted to two-dimensional isentropic surfaces in the stratosphere with no sources or sinks of the chemical constituent; see MCCL for a discussion of the limitations of this approximation and the resulting model error. We have developed the Lagrangian filter with the same approximation.

The observations are taken at discrete instants of time $t_k, k = 1, 2, 3, \dots$, and there are p_k observations at each time, described by the vector

$$\mathbf{w}_k^o = \mathbf{H}_k^c w(\cdot, t_k) + \boldsymbol{\epsilon}_k^m, \tag{2}$$

where \mathbf{H}_k^c is a p_k -vector of linear operators on the continuum state $w(\mathbf{x}, t)$, and $\boldsymbol{\epsilon}_k^m$ is the p_k -vector measurement error. The measurement error is assumed to be Gaussian, white in time, and independent of the state $w(\mathbf{x}, t)$ and its initial condition $w_0(\mathbf{x})$; the latter is assumed to be a mean-zero Gaussian random field. The measurement error is assumed to be unbiased, $\langle \boldsymbol{\epsilon}_k^m \rangle = 0$, where $\langle \cdot \rangle$ denotes the expectation. The $p_k \times p_k$ measurement-error covariance matrix,

$$\mathbf{R}_k^m \equiv \langle \boldsymbol{\epsilon}_k^m \boldsymbol{\epsilon}_k^{mT} \rangle, \tag{3}$$

is assumed to be known. The superscript T denotes the transpose.

The filter is a time-sequential algorithm that seeks the best estimate of the state, called the *analysis*, and its error covariance. The analysis $w_{k-1}^a(\mathbf{x})$ at time t_{k-1} is defined as the expected value of the state conditioned on all observations up to time t_{k-1} :

$$w_{k-1}^a(\mathbf{x}) \equiv \langle w(\mathbf{x}, t_{k-1}) | \mathcal{W}_{k-1}^o \rangle, \tag{4}$$

where $\mathcal{W}_{k-1}^o \equiv \{\mathbf{w}_1^o, \mathbf{w}_2^o, \dots, \mathbf{w}_{k-1}^o\}$. The analysis at the next time t_k uses the observations \mathbf{w}_k^o and a forecast $w_k^f(\mathbf{x})$ generated using the known dynamics. Defining the conditional mean forecast

$$\widehat{w}(\mathbf{x}, t) \equiv \langle w(\mathbf{x}, t) | \mathcal{W}_{k-1}^o \rangle, \quad t_{k-1} < t \leq t_k, \tag{5}$$

its dynamics are described by taking the conditional mean of Eq. (1),

$$\left(\frac{\partial}{\partial t} + \mathbf{v}(\mathbf{x}, t) \cdot \nabla \right) \widehat{w}(\mathbf{x}, t) = 0. \tag{6}$$

The forecast $w_k^f(\mathbf{x}) \equiv \widehat{w}(\mathbf{x}, t_k)$ at time t_k is therefore the solution of Eq. (6) with $w_{k-1}^a(\mathbf{x})$ as the initial condition.

The solution by the method of characteristics recognizes that the operator $(\partial/\partial t) + \mathbf{v}(\mathbf{x}, t) \cdot \nabla$ in Eq. (6) is the total time derivative along a trajectory, which can be parametrized by its initial position at time t_{k-1} . The trajectory is given by

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}(\mathbf{x}, t). \tag{7}$$

Hence the forecast at time t_k is

$$w_k^f(\mathbf{x}) = w_{k-1}^a(\mathbf{y}_{k-1}(\mathbf{x}, t_k)), \quad (8)$$

where $\mathbf{y}_{k-1}(\mathbf{x}, t_k)$ is the departure point at time t_{k-1} of the trajectory that ends at \mathbf{x} at time t_k . Cohn (1993) points out that \mathbf{y} may be evaluated by solving the advection equation with initial condition $\mathbf{y} = \mathbf{x}$. For discrete trajectories, where the solution is required at a set of points in \mathbf{x} at time t_k , \mathbf{y} can be calculated by backwards tracing in the given velocity field. Alternatively, and this is what we will adopt for the Lagrangian filter, one can choose the departure points of discrete trajectories and then the solution $w_k^f(\mathbf{x})$ will be available at the locations of the endpoints of the trajectories.

Along with the conditional mean state, we evolve the conditional estimation-error covariance function. Define the state estimation (forecast) error

$$e(\mathbf{x}, t) \equiv w(\mathbf{x}, t) - \widehat{w}(\mathbf{x}, t), \quad t_{k-1} < t \leq t_k. \quad (9)$$

The conditional forecast-error covariance function is

$$P(\mathbf{x}_1, \mathbf{x}_2, t) \equiv \langle e(\mathbf{x}_1, t)e(\mathbf{x}_2, t) | \mathcal{W}_{k-1}^0 \rangle, \quad t_{k-1} < t \leq t_k. \quad (10)$$

Using Eqs. (1), (6), and (9) it can be shown (Cohn 1993, appendix A) that

$$\left(\frac{\partial}{\partial t} + \mathbf{v}(\mathbf{x}_1, t) \cdot \nabla_1 + \mathbf{v}(\mathbf{x}_2, t) \cdot \nabla_2 \right) P(\mathbf{x}_1, \mathbf{x}_2, t) = Q(\mathbf{x}_1, \mathbf{x}_2, t), \quad (11)$$

where the subscript on ∇ indicates the set of spatial coordinates, \mathbf{x}_1 or \mathbf{x}_2 , on which the ∇ operator acts, and where we assume that the model error $f(\mathbf{x}, t)$ is uncorrelated with the initial error and is white in time,

$$\langle f(\mathbf{x}_1, t)f(\mathbf{x}_2, \tau) \rangle = Q(\mathbf{x}_1, \mathbf{x}_2, t)\delta(t - \tau), \quad t_0 \leq \tau < t. \quad (12)$$

Here $Q(\mathbf{x}_1, \mathbf{x}_2, t)$ is the model-error covariance function. Similarly to Eq. (4), the conditional analysis-error covariance function at time t_{k-1} is

$$P_{k-1}^a(\mathbf{x}_1, \mathbf{x}_2) \equiv \langle e_{k-1}^a(\mathbf{x}_1, t)e_{k-1}^a(\mathbf{x}_2, t) | \mathcal{W}_{k-1}^0 \rangle, \quad (13)$$

where

$$e_{k-1}^a(\mathbf{x}) \equiv w(\mathbf{x}, t_{k-1}) - w_{k-1}^a(\mathbf{x}). \quad (14)$$

The solution of Eq. (11) at time t_k with initial condition $P_{k-1}^a(\mathbf{x}_1, \mathbf{x}_2)$ is the conditional forecast-error covariance function $P_k^f(\mathbf{x}_1, \mathbf{x}_2)$. Once again, the operator on the left side of Eq. (11) is a total time derivative, and the solution by the method of characteristics is

$$P_k^f(\mathbf{x}_1, \mathbf{x}_2) = P_{k-1}^a(\mathbf{y}_{k-1}(\mathbf{x}_1, t_k), \mathbf{y}_{k-1}(\mathbf{x}_2, t_k)) + \int_{t_{k-1}}^{t_k} Q\{\mathbf{x}(\tau; \mathbf{y}_{k-1}(\mathbf{x}_1, t_k)), \mathbf{x}(\tau; \mathbf{y}_{k-1}(\mathbf{x}_2, t_k)), \tau\} d\tau, \quad (15)$$

where $\mathbf{x}(\tau; \mathbf{y}_{k-1}(\mathbf{x}_1, t_k))$ denotes the location at time τ of the trajectory starting at \mathbf{y}_{k-1} at time t_{k-1} and ending at \mathbf{x}_1 at time t_k (Cohn 1993, appendix B). Although the error covariance function involves more independent variables ($\mathbf{x}_1, \mathbf{x}_2, t$) than the state (\mathbf{x}, t), the auxiliary (characteristic) equations for its approximate solution use the same trajectories (Cohn 1993, section 3).

The analysis at the new time step t_k (Cohn 1997) is

$$w_k^a(\mathbf{x}) = w_k^f(\mathbf{x}) + \mathbf{K}_k(\mathbf{x})(\mathbf{w}_k^o - \mathbf{H}_k^c w_k^f(\cdot)). \quad (16)$$

The innovation $\mathbf{w}_k^o - \mathbf{H}_k^c w_k^f(\cdot)$ is a p_k -vector, and the Kalman gain \mathbf{K}_k is a $1 \times p_k$ row vector function of \mathbf{x} :

$$\mathbf{K}_k(\mathbf{x}) = (\mathbf{H}_{2k}^c P_k^f(\mathbf{x}, \cdot))^T \{ \mathbf{H}_{1k}^c (\mathbf{H}_{2k}^c P_k^f(\cdot, \cdot))^T + \mathbf{R}_k^m \}^{-1}, \tag{17}$$

where \mathbf{R}_k^m is the $p_k \times p_k$ measurement-error covariance matrix, Eq. (3), and \mathbf{H}_{1k}^c and \mathbf{H}_{2k}^c denote the action of the observation operator \mathbf{H}_k^c on the arguments \mathbf{x}_1 and \mathbf{x}_2 , respectively, of the forecast-error covariance function $P_k^f(\mathbf{x}_1, \mathbf{x}_2)$. Combining Eqs. (2), (16), and (13) gives the analysis-error covariance function

$$P_k^a(\mathbf{x}_1, \mathbf{x}_2) = (1 - \mathbf{K}_k(\mathbf{x}_1)\mathbf{H}_{1k}^c)\{(1 - \mathbf{K}_k(\mathbf{x}_2)\mathbf{H}_{2k}^c)P_k^f(\cdot, \cdot)\} + \mathbf{K}_k(\mathbf{x}_1)\mathbf{R}_k^m\mathbf{K}_k^T(\mathbf{x}_2). \tag{18}$$

This is called the *Joseph form*, while if Eq. (17) is substituted into Eq. (18) then we obtain the simpler formula

$$P_k^a(\mathbf{x}_1, \mathbf{x}_2) = (1 - \mathbf{K}_k(\mathbf{x}_1)\mathbf{H}_{1k}^c)P_k^f(\cdot, \mathbf{x}_2), \tag{19}$$

which is called the *optimal form* of the continuum analysis-error covariance function update equation.

Equations (8) and (15) constitute the forecast equations of the continuum Kalman filter for our constituent assimilation problem, and Eqs. (16), (17), and (19) constitute the corresponding continuum analysis update equations. The initial conditions for the whole data assimilation cycle are $w_0^a(\mathbf{x}) = \langle w_0(\mathbf{x}) \rangle$, and

$$P_0^a(\mathbf{x}_1, \mathbf{x}_2) = \langle (w_0(\mathbf{x}_1) - \langle w_0(\mathbf{x}_1) \rangle)(w_0(\mathbf{x}_2) - \langle w_0(\mathbf{x}_2) \rangle) \rangle. \tag{20}$$

(b) *The Lagrangian discretization*

The continuum Kalman filter equations require discretization to be implemented on a computer. Our implementation estimates the continuum field $w(\mathbf{x}, t)$ at the locations of N trajectories propagated according to Eq. (7), that is,

$$\frac{d\mathbf{x}(i)}{dt} = \mathbf{v}(\mathbf{x}(i), t), \quad i = 1, \dots, N, \tag{21}$$

where $\mathbf{x}(i), i = 1, \dots, N$, are the locations of the trajectories. Following Cohn (1997), the truth vector at time t_k is

$$\mathbf{w}_k^t \equiv \begin{bmatrix} w(\mathbf{x}(1), t_k) \\ w(\mathbf{x}(2), t_k) \\ \vdots \\ w(\mathbf{x}(N), t_k) \end{bmatrix}. \tag{22}$$

The algorithm calculates the analysis vector,

$$\mathbf{w}_k^a \equiv \begin{bmatrix} \langle w(\mathbf{x}(1), t_k) | \mathcal{W}_k^o \rangle \\ \langle w(\mathbf{x}(2), t_k) | \mathcal{W}_k^o \rangle \\ \vdots \\ \langle w(\mathbf{x}(N), t_k) | \mathcal{W}_k^o \rangle \end{bmatrix}, \tag{23}$$

along with the $N \times N$ analysis-error covariance matrix \mathbf{P}_k^a , whose (i, j) th element is the covariance between locations $(\mathbf{x}(i), \mathbf{x}(j))$ of trajectories i and j . The ensuing *forecast* equations for the *Lagrangian filter* are given by Eqs. (8) and (15) in discrete, or vector,

form:

$$\mathbf{w}_k^f = \mathbf{w}_{k-1}^a, \quad (24)$$

$$\mathbf{P}_k^f = \mathbf{P}_{k-1}^a + \mathbf{Q}_{k-1}. \quad (25)$$

The propagating parts of the state and error covariance matrix involve trivial equality in Eqs. (24) and (25) because, apart from the model-error covariance term \mathbf{Q}_{k-1} , their values are fixed along characteristics. The computational cost is the advection of N trajectories. Thus the $\mathcal{O}(N^2)$ cost of covariance propagation in conventional discrete Kalman filtering is significantly reduced. The non-propagating part of the forecast-error covariance matrix arises from the model-error covariance term in Eq. (15), which is approximated to first order in time in the Lagrangian filter by

$$\begin{aligned} \mathbf{Q}_{k-1}(i, j) &= \int_{t_{k-1}}^{t_k} \mathcal{Q}[\mathbf{x}\{\tau; \mathbf{y}_{k-1}(\mathbf{x}(i), t_k)\}, \mathbf{x}\{\tau; \mathbf{y}_{k-1}(\mathbf{x}(j), t_k)\}, \tau] d\tau \\ &\approx (t_k - t_{k-1}) \mathcal{Q}(\mathbf{x}_{k-1}(i), \mathbf{x}_{k-1}(j), t_{k-1}), \end{aligned} \quad (26)$$

where $\mathbf{x}_{k-1}(i)$ is the location of the i th trajectory at time t_{k-1} . We also use a fixed time step $\Delta t \equiv t_k - t_{k-1}$.

The *analysis* equations at the new time step t_k are given by Eqs. (16), (17), and (19) in discrete or vector form:

$$\mathbf{w}_k^a = \mathbf{w}_k^f + \mathbf{K}_k(\mathbf{w}_k^o - \mathbf{H}_k \mathbf{w}_k^f), \quad (27)$$

$$\mathbf{K}_k = \mathbf{P}_k^f \mathbf{H}_k^T (\mathbf{H}_k \mathbf{P}_k^f \mathbf{H}_k^T + \mathbf{R}_k^m)^{-1}, \quad (28)$$

$$\mathbf{P}_k^a = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_k^f, \quad (29)$$

where $\mathbf{w}^{f,a}$ are N -vectors, $\mathbf{P}^{f,a}$ and \mathbf{Q} are $N \times N$ matrices, \mathbf{K} is a $N \times p_k$ matrix, \mathbf{H}_k is a $p_k \times N$ interpolation matrix, and \mathbf{I} is the identity matrix. At time t_0 we initialize N trajectories on a regular two-dimensional latitude–longitude grid. Hence the forecast equations are solved by forward Lagrangian trajectories. We developed a suite of trajectory generators for the sphere and these will be discussed in the next section.

In general the point observations do not occur exactly at locations of the trajectories. Thus the \mathbf{H}^c operator in Eqs. (16)–(19) must be approximated by an interpolation or smoothing matrix \mathbf{H} that maps a field defined at the locations of the trajectories to the locations of the observations. The experiments carried out for the Eulerian filter by MCCL and MC used observations of chemical constituents retrieved from UARS instruments. \mathbf{H} was modelled as bi-linear interpolation on a regular latitude–longitude grid. Here we make a similar approximation, using linear interpolation on spherical triangular elements of the trajectory grid that are generated using a Delaunay triangulation on the sphere at each time step. These algorithms will be described in the next section. The approximation of \mathbf{H}^c by \mathbf{H} gives rise to representativeness error in the Lagrangian filter. With the truth field defined by Eq. (22), we may rewrite Eq. (2) as

$$\mathbf{w}_k^o = \mathbf{H}_k \mathbf{w}_k^t + \boldsymbol{\epsilon}_k^r + \boldsymbol{\epsilon}_k^m, \quad (30)$$

where

$$\boldsymbol{\epsilon}_k^r = \mathbf{H}_k^c w_k(\cdot, t_k) - \mathbf{H}_k \mathbf{w}_k^t \quad (31)$$

is called the representativeness error. Part of this error arises from the interpolation error itself, and another part arises from the component of the continuum state which is not resolved by our model discretization (cf. Cohn 1997).

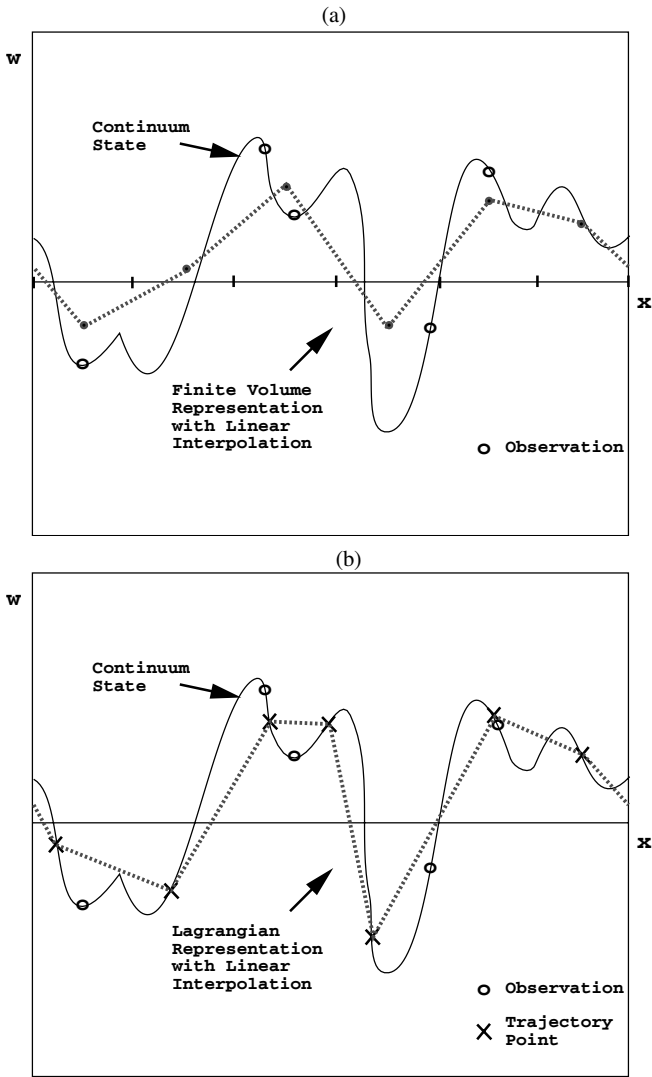


Figure 1. Schematics in one spatial dimension of the representativeness error for the cases of (a) an Eulerian filter and (b) a Lagrangian filter. The underlying continuum state $w_k(\mathbf{x})$ is the solid curve which is the same in both cases. The point observations, indicated by the circles, are also the same in both cases. The Lagrangian representation $\mathbf{H}_k \mathbf{w}_k^t$ is shown by the dotted lines. In (a) there are six grid points indicated by the segments on the \mathbf{x} -axis and the values of the volume average are indicated by the dots at the nodes of the linear interpolation function. In (b) the crosses indicate the values of the underlying continuum state at the locations of the trajectories. See text for explanation of symbols.

Figures 1(a) and (b) show schematics of the representativeness error in one spatial dimension for the cases of a Eulerian filter and a Lagrangian filter, respectively. For the Eulerian filter the grid intervals are bounded by the tick marks on the \mathbf{x} axis and the discretization uses volume averages. The dotted lines show possible values of $\mathbf{H}_k \mathbf{w}_k^t$ where linear interpolation is used in both cases. Therefore, for point observations, the vertical difference between the solid and dotted curves represents the values of ϵ_k^t : for observations shown by circles in Fig. 1, the difference between the curves at those

locations is the representativeness error. This error may be reduced by increasing the resolution of the dynamical models (i.e. by adding more grid points or trajectories). Also, for the same number of grid points and trajectories (i.e. for Eulerian and Lagrangian filters of approximately the same resolution), the representativeness error has similar values in a global sense for both Figs. 1(a) and (b). However, in the degenerate case where trajectories are chosen to pass through the locations of point observations, the representativeness error for the Lagrangian filter is zero.

The error ϵ_k^r in Eq. (30) adds to the measurement error. If the serial (time) correlation and state dependence of this error is neglected, an additional error covariance ($\langle \epsilon_k^r \epsilon_k^{rT} \rangle$) appears inside the parentheses of Eq. (28) (Cohn 1997). For this reason the representativeness error is included in data assimilation formulations by augmenting the measurement-error covariance matrix with a representativeness-error covariance matrix (Lorenç 1986; Daley 1993; Mitchell and Daley 1997). This term diminishes the contribution of the innovation to the analysed state by an amount depending on the presumed representativeness-error covariance. A parametrized representativeness-error covariance matrix was formulated for the Eulerian filter in our earlier work (MCCL; MC). In section 4 we discuss a similar implementation for the Lagrangian filter.

3. NUMERICAL IMPLEMENTATION

For the implementation of \mathbf{H} in the Lagrangian filter, Eqs. (27)–(29), we use packages developed by Renka (1997a,b) for generating a Delaunay triangulation of trajectory locations on the sphere (STRIPACK) and for the interpolation or smoothing of the trajectory fields to an arbitrary point (SSRFPACK). The Delaunay triangulation is a list of spherical triangles whose vertices are composed of $\mathbf{x}(i)$. The algorithm generates reasonably uniform triangles by guaranteeing that no circumcircle of a triangle contains a node in its interior. The algorithm has $N \log N$ computational complexity. For the results of this paper we used continuous (C^0) interpolation (i.e. point evaluation of a piecewise linear interpolation function).

The Lagrangian filter was configured to perform comparisons with our earlier work on the Eulerian filter (MCCL; MC). That work focused on a baseline run over 8 days in September 1992. Winds were obtained from the GEOS-1 Data Assimilation System. The Eulerian transport algorithm of Lin and Rood (1996) was used with $N_x = 72$ grid points in longitude and $N_y = 46$ grid points in latitude, i.e. $5^\circ \times 4^\circ$ horizontal resolution. For the Lagrangian filter we initialize the trajectories at the centres of grid cells of the Lin–Rood model, so that the average resolution of the trajectories is approximately the same as that of the Eulerian transport model. Although $N_x \times N_y = 3312$ trajectories are carried by the filter, they are thinned for the purpose of the triangulation subroutines in order to guarantee an approximately equal-area coverage of the sphere. The results are not sensitive to the method of thinning. For the present work, trajectories initialized within 11 zonal circles of grid cells closest to the poles are thinned. Starting out at the poles and moving towards the equator, the number of initially equally spaced trajectories is 1, 2, 3, 7, 9, 12, 14, 18, 36, 36 and 36 (all the other grid circles have 72 initial trajectories). Thus the triangulation and interpolation subroutines use 2076 trajectories.

We implemented algorithms that generate accurate trajectories based on the works of Ritchie (1987), Côté and Staniforth (1988) and McDonald and Bates (1989). These three algorithms explicitly avoid the singularity in coordinates at the poles. They were originally intended for semi-Lagrangian applications, but were easily modified for the fully Lagrangian use here. Our implementation uses velocity fields specified at




Figure 2. Delaunay triangulation for the baseline run for (a) the initial distribution of trajectories on a uniform latitude–longitude grid with thinning near the poles, and (b) the distribution of trajectories at end of the run (8 days).

the centres of grid volumes defined by the Lin–Rood transport model. Figure 2(a) shows the Delaunay triangulation generated by STRIPACK for the initial distribution of trajectories. Figure 2(b) shows the Delaunay triangulation after 8 days of simulation time using the GEOS-1 winds interpolated to the 1100 K isentrope. The time step of the trajectory generator is 15 minutes. The GEOS-1 archives provide winds at six-hour intervals, and the instantaneous velocity field at each grid volume centre is specified by linear interpolation to the time step between the six-hour winds. These plots and all the remaining results of this paper use the McDonald–Bates trajectory generator. The other two trajectory generators gave almost identical results.

Equations (21), (24), (25) and (27)–(29) constitute the Lagrangian filter. These correspond to Eqs. (9)–(12) and (14) of the Eulerian filter of Lyster *et al.* (1997). Other than the significant differences between the grids as described previously, Eqs. (21) and (25) replace the corresponding term $\mathbf{M}^m \mathbf{P}^a \mathbf{M}^{mT}$ (Lyster *et al.* 1997, Eq. (10)), where \mathbf{M}^m is the model matrix for the Eulerian filter. Equation (25) takes $\mathcal{O}(N^2)$ floating point operations per time step to add the two matrices, and Eq. (21) involves N trajectory calculations. This represents considerably fewer calculations than for the corresponding term $\mathbf{M}^m \mathbf{P}^a \mathbf{M}^{mT}$. Lyster *et al.* (1997) estimate the cost of calculating this term to be hn^2 floating point operations per time step, where $h \approx 20$ depends on the sparsity of the transport model matrix and n is the number of grid points (we take $n \approx N$ for the Eulerian and Lagrangian filters of approximately equal resolution). Furthermore, the parallel implementation of the Eulerian filter using the message-passing interface library requires a global transpose of the covariance matrix distributed column-wise across processors for each time step—the covariance decomposition shown in Lyster *et al.* (1997), Fig. 1(b)—but the parallel implementation of the Lagrangian filter using the same covariance decomposition requires no global transpose. We have achieved on an SGI/Cray T3E-1200 computer a 200-fold speed-up for the Lagrangian filter over the same problem solved using the Eulerian filter on a single processor of a Cray C90.

Equation (29) requires $\mathcal{O}(pN^2)$ operations. Lyster *et al.* (1997) describe how this work is parallelized for the Eulerian filter by distributing columns of \mathbf{P}^f and \mathbf{P}^a

across processors. For the Lagrangian filter the analysis equations, Eqs. (27)–(29), are parallelized in the same way as the Eulerian filter, but the FORTRAN loops indexed over the regular latitude–longitude static grid in the Eulerian filter are replaced by loops over the mesh of trajectory locations. Loops over arrays dimensioned by the number of observations are treated the same way for both filters. The method of distributing columns of \mathbf{P}^f and \mathbf{P}^a across processors also naturally parallelizes the $\mathcal{O}(N^2)$ floating point operations in Eq. (25).

The trajectories require only $\mathcal{O}(N)$ floating point operations, so we decided to evaluate trajectories redundantly on all processors. The Kalman gain is also evaluated on all processors (i.e. \mathbf{K} is not distributed) for both filters. The main limitation on scalability of the Lagrangian filter is then Amdahl's (1967) bottleneck, because the floating point operations to evaluate trajectories and the Kalman gain are not distributed across processors.

4. EXPERIMENTAL RESULTS

To validate and assess the Eulerian filter, MCCL and MC used the time period 6–14 September 1992, with GEOS-1 DAS stratospheric winds interpolated to the 1100 K isentrope (about 38 km) and observations of methane mixing ratio retrieved from the UARS CLAES instrument. This period encompasses a southern hemisphere wave-breaking event which proved useful for the assessment of the Eulerian filter. For the CLAES retrievals there are about 1300 observations per day. For the assimilation experiments of MCCL, MC and the present work these observations are assumed to be point measurements with a relative measurement error of 7.3%.

The key parametrizations and parameter values of the Lagrangian filter are chosen to be identical to those of the Eulerian filter for purposes of comparison. Both the Eulerian model and the trajectory generator use a 15 min time step. We use the same 'spun up' initial field \mathbf{w}_0 as MCCL. The initial-error covariance matrix is derived from the initial field \mathbf{w}_0 and the first-order autoregressive correlation function:

$$\mathbf{P}_0(i, j) = \gamma^2 \mathbf{w}_0(i) \mathbf{w}_0(j) \mathbf{C}(i, j), \quad (32)$$

with

$$\mathbf{C}(i, j) = \exp\left(-\frac{|\mathbf{x}_0(i) - \mathbf{x}_0(j)|}{L}\right), \quad (33)$$

where $\mathbf{x}_0(i)$ and $\mathbf{x}_0(j)$ are the initial locations of trajectories i and j , $|\cdot|$ denotes Euclidean distance on \mathbb{R}^3 (chordal distance on the sphere), and the length-scale $L = 3600$ km. The representativeness error is modelled by adding to \mathbf{R}_k^m in Eq. (28) a term \mathbf{R}_k^r ,

$$\mathbf{R}_k^r = \beta^2 \mathbf{R}_k, \quad (34)$$

where

$$\mathbf{R}_k = \text{diag}\{(\mathbf{H}_k \mathbf{w}_k^f(1))^2, \dots, (\mathbf{H}_k \mathbf{w}_k^f(p_k))^2\}. \quad (35)$$

The model-error covariance matrix is assumed to have the form

$$\mathbf{Q}_k(i, j) = \delta^2 \mathbf{w}_k^a(i) \mathbf{w}_k^a(j) \mathbf{C}^q(i, j), \quad (36)$$

where the model-error correlation model \mathbf{C}^q is chosen to be the same as the initial-error correlation model, Eq. (33), with the same length-scale parameter $L = 3600$ km.

MC describe a χ^2 -validation procedure for tuning the parameters γ , β , and δ . The running time average value of χ^2 is

$$\overline{\chi_k^2} = \frac{1}{P_k} \sum_{i=k-\Delta k/2}^{k+\Delta k/2} \mathbf{v}_i^T \mathbf{S}_i^{-1} \mathbf{v}_i, \quad (37)$$

where the innovation vector is

$$\mathbf{v}_k \equiv \mathbf{w}_k^o - \mathbf{H}_k \mathbf{w}_k^f, \quad (38)$$

the innovation covariance matrix is

$$\mathbf{S}_k \equiv \mathbf{H}_k \mathbf{P}_k^f \mathbf{H}_k^T + \mathbf{R}_k^r + \mathbf{R}_k^m, \quad (39)$$

the moving time window is $\Delta k = 96$ (i.e. one day), and the total number of observations in the window centred around time t_k is

$$P_k \equiv \sum_{i=k-\Delta k/2}^{k+\Delta k/2} p_k. \quad (40)$$

MC show that the running $\overline{\chi^2}$ should be approximately equal to one when the input-error covariances of the filter algorithm are specified consistently. This condition must hold, for instance, to guarantee accurate estimates of the analysis-error variance, one of the most basic attributes of a useful filter. For the baseline CLAES dataset of September 1992, $P_k \approx 1300$, and the validated (MCCL) parameters for the Eulerian filter are $\gamma = 0.11$, $\beta = 0.10$ and $\delta = 0.003$.

Both the state estimates $\mathbf{w}^{f,a}$ and the estimation-error covariance matrices $\mathbf{P}^{f,a}$ for the Lagrangian filter represent point values on unstructured meshes of trajectory locations. The Delaunay triangulation may be used to remap the fields to a regular latitude–longitude grid (Renka 1997a,b) for the purposes of post processing or visualization of the results. Since we often wish to compare results with those of the Eulerian filter, we define the remapping, $\mathbf{M}^f \mathbf{w}^{f,a}$, which interpolates from the Lagrangian grid to centres of grid volumes of the Lin–Rood transport model at the $5^\circ \times 4^\circ$ resolution. This \mathbf{M}^f represents a 3312×2076 matrix. The corresponding remapping of the covariance matrices to the $5^\circ \times 4^\circ$ grid in both pairs of coordinates is $\mathbf{M}^f \mathbf{P}^{f,a} \mathbf{M}^{fT}$. For the purpose of later discussion we refer to $\text{diag}(\mathbf{M}^f \mathbf{P}^{f,a} \mathbf{M}^{fT})$ as the indirect-remapped variance, and to $\mathbf{M}^f \text{diag}(\mathbf{P}^{f,a})$ as the direct-remapped variance.

MCCL described a spurious loss of variance and related increase in correlation length-scales in the Eulerian filter associated with the finite-difference implementation of the error covariance matrix propagation. This was corrected by separately advecting the variance and rescaling the correlation matrix by the corrected variance. A second correction was then applied to the covariance matrix to reduce the spatial extent of the correlations. This is referred to here as the shape correction. Figure 3 shows colour contour plots of variance for the Eulerian filter and Lagrangian filter after 2 days of simulation in pure forecast mode, i.e. with no observations assimilated and with the model-error parameter δ set to zero. The direct-remapped variance illustrates the variance field actually calculated by the cycled Lagrangian filter algorithm. In fact, in pure forecast mode the variance is explicitly conserved along the flow. By their similarity, Figs. 3(b) and (d) show that the Lagrangian filter is a direct algorithm that performs variance evolution accurately, without the need for a variance correction. Interestingly,

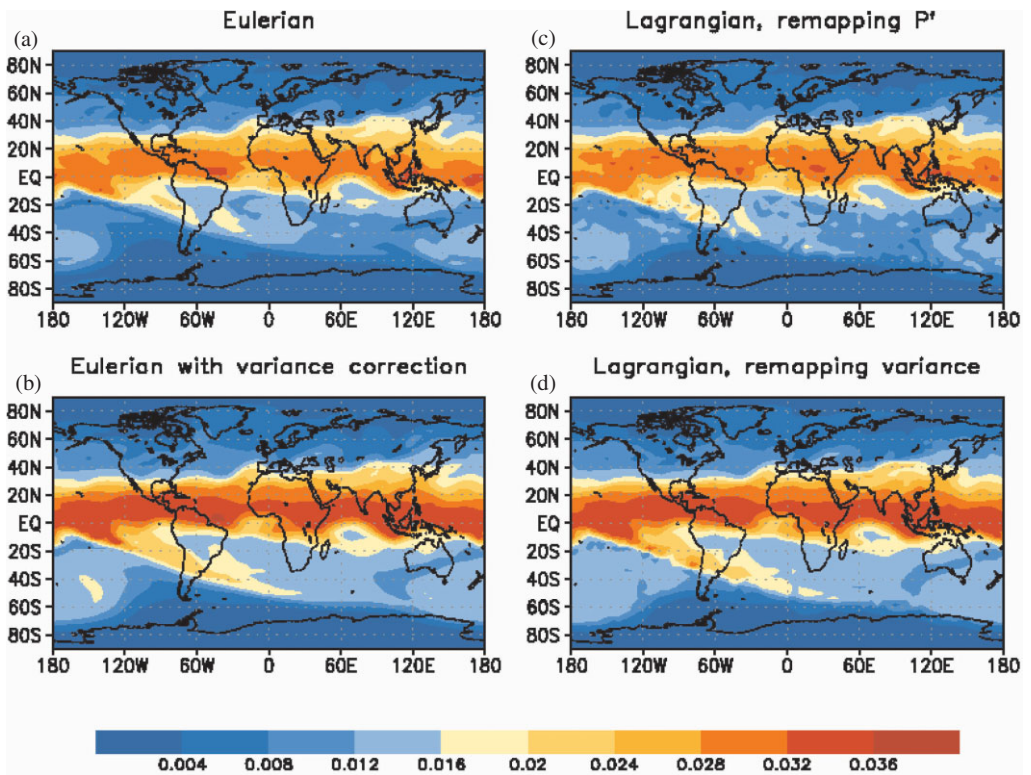


Figure 3. Colour contour plots of variance for (a)–(b) the Eulerian filter and (c)–(d) the Lagrangian filter in pure forecast mode after 2 days of simulation time with identical initial conditions. (a) and (b) show the variance of the Eulerian filter before and after, respectively, variance and shape correction. (c) and (d) show the indirect- and direct-remapped variance, respectively, of the Lagrangian filter.

Figs. 3(a) and (c) show that the indirect-remapped variance for the Lagrangian filter is problematic, even though the filter itself calculates the variance accurately. It yields a visualization of the variance field similar to the erroneous field of the uncorrected filter.

Figure 4 shows the total variance (TV), $\int \mathbf{P}^f(\mathbf{x}, \mathbf{x}) \, d\mathbf{x}$, evaluated numerically on the regular latitude–longitude grid, plotted versus time up to 8 simulation days. The Eulerian and Lagrangian filters were run in pure forecast mode where the driving winds have been filtered to make them divergence free. Under these conditions the total variance should be constant in time (cf. MCCL, Eq. (3.11)). The TV for the direct-remapped variance of the Lagrangian filter at intervals of 1 day is nearly constant as expected. The TV of the Eulerian filter without variance and shape correction illustrates the steady variance loss. Comparing the direct-remapped and indirect-remapped variances for the Lagrangian filter shows that, while the loss of variance due to the indirect remapping is only a visualization artifact and not cumulative, it nonetheless grows steadily worse with time. Figure 4 shows also that the cumulative variance loss of the (uncorrected) Eulerian filter exceeds that due to the indirect remapping in the Lagrangian filter after about a day and a half, increasing further thereafter. In fact MC (see MC, Fig. 11) found that in assimilation mode the uncorrected Eulerian filter performs as well as the corrected one for about a day and a half, but deteriorates thereafter.

To gain a quantitative measure of the stretching and contraction of the error covariance structures, for the Lagrangian filter we define the gridded maximum and




Figure 4. The total variance (TV) plotted versus simulation time. The Eulerian and Lagrangian filters were run in pure forecast mode where the driving winds have been filtered to make them divergence-free. The circles show the TV for the direct-remapped variance of the Lagrangian filter at intervals of 1 day. The solid curve shows the TV of the Eulerian filter without variance and shape correction. The squares (including one at time zero) show the TV for the indirect-remapped variance of the Lagrangian filter.

minimum correlation lengths as follows. First the Lagrangian covariance matrix is remapped to a regular grid, $\mathbf{P}_g^a = \mathbf{M}^T \mathbf{P}^a \mathbf{M}^{rT}$. Then a correlation matrix is evaluated, $\mathbf{P}_c^a = \mathbf{D}^{-1/2} \mathbf{P}_g^a \mathbf{D}^{-1/2}$, where \mathbf{D} is the diagonal matrix $\mathbf{D} = \text{diag } \mathbf{P}_g^a$. Then for each grid point the maximum correlation length is defined to be the great circle distance to the furthest grid point where the correlation is reduced by $1/e$; similarly the minimum correlation length is the great circle distance to the closest grid point where the correlation is reduced by $1/e$. Figure 5(a) shows a colour contour plot of the gridded ratio of the minimum to maximum correlation lengths for the Lagrangian filter in pure forecast mode after 4 days of simulation time. This may be thought of as the aspect ratio of the flow-dependent error covariance structures. It overestimates the actual aspect ratio because, as we have seen, the operation $\mathbf{P}_g^a = \mathbf{M}^T \mathbf{P}^a \mathbf{M}^{rT}$ is a smoothing operation. Figure 5(b) shows the ratio of the indirect-remapped variance to the direct-remapped variance for the same run, which according to Fig. 3 is approximately the variance loss that would occur for an equivalent Eulerian filter run. The light blue areas indicate where stretching and contraction of the error covariance structures has occurred in the vicinity of a tongue of tropical air that penetrated into the southern midlatitudes (Fig. 7), where the shear in the flow is at a maximum. The similarity of the plots in Fig. 5 illustrates the close association between this stretching and contraction (Fig. 5(a)) and the variance loss (Fig. 5(b)).

These results lead us to conclude that the variance loss of the (uncorrected) Eulerian filter is due to the generation of correlation structure at increasingly fine scale that is always resolved by the Lagrangian algorithm, but not by the Eulerian algorithm. Figures 3 and 4 indicate that one can consider the Eulerian filter covariance propagation to be, in part, an indirect variance remapping algorithm, with local remapping error after two days essentially identical to that of the indirect-remapping algorithm we introduced for the Lagrangian filter. Figure 5 illustrates a relationship between local remapping error (variance loss) and the local aspect ratio of the correlation length-scales. This

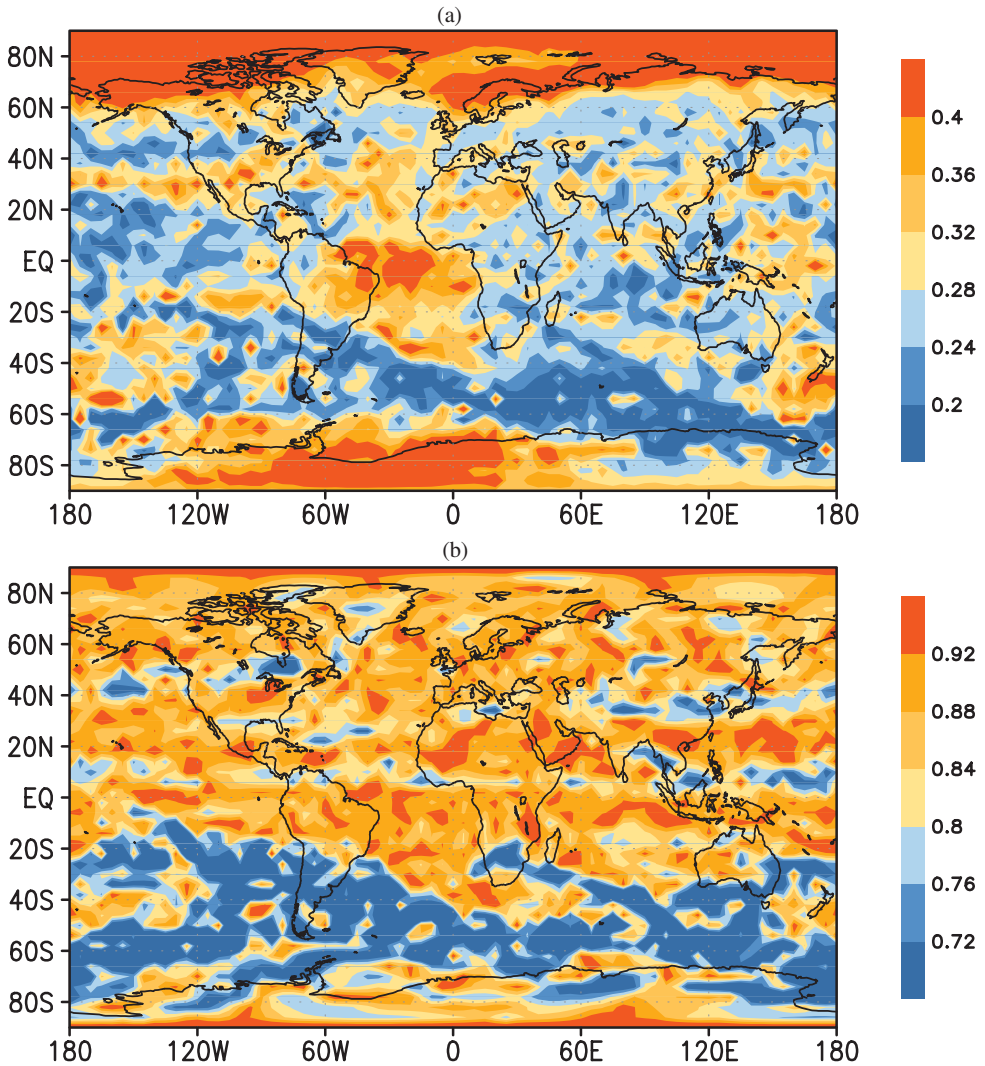


Figure 5. (a) Ratio of the shortest to the longest correlation length, and (b) ratio of the indirect-remapped variance to the direct-remapped variance, for the Lagrangian filter in pure forecast mode after 4 days of simulation time.

aspect ratio is initially 1 because the initial correlation is isotropic, but decreases as the wind shear creates a cascade of correlation structure to smaller scales. The indirect-remapping error for the Lagrangian filter (Fig. 4) is due solely to this cascade. Figure 4 suggests that the variance loss of the Eulerian filter is due mostly to this cascade, and only partly compounded by the fact that the loss is cumulative for the Eulerian filter. The compounding effect would be present even for a solid-body rotation wind field, which preserves isotropic correlation functions.

Next we repeat the baseline assimilation experiment of MCCL and MC using the CLAES dataset. Figure 6 shows first the running mean $\overline{\chi^2}$ for the Eulerian and Lagrangian filters. As described earlier, the value should be about 1 for a valid filter. The result for the Lagrangian filter, shows that it maintains $\overline{\chi^2}$ close to 1, perhaps

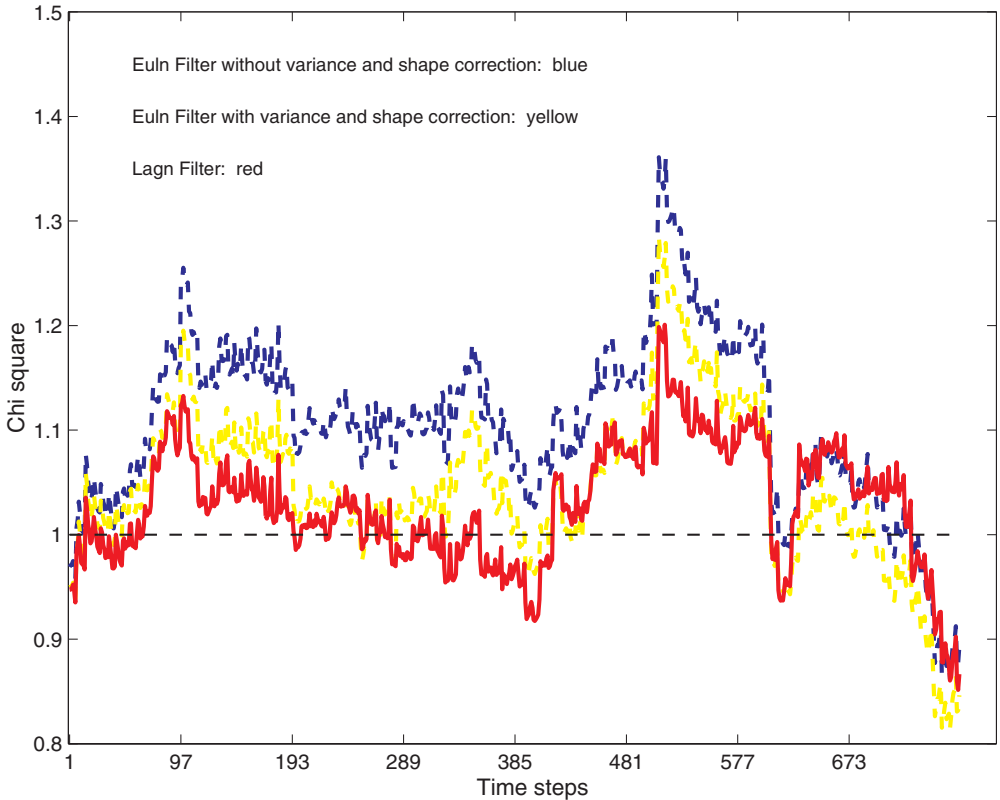


Figure 6. The running mean $\overline{\chi^2}$ plots for the Eulerian filter and the Lagrangian filter for the baseline dataset and parameters of MCCL. As discussed in the text, the value should be about 1 for a valid filter. The dashed yellow curve is the same as Fig. 1 of MC for the Eulerian filter for baseline values of $\gamma = 0.11$, $\beta = 0.10$, and $\delta = 0.003$ with adjustments for variance loss and shape correction. The solid red curve is the result for the Lagrangian filter for the same parameters. The dashed blue curve is the Eulerian filter without variance or shape correction.

closer than for the (corrected) Eulerian filter. For reference, results for the Eulerian filter without variance and shape correction are also shown. That the Lagrangian filter should not require significant retuning of the parameter values based on $\overline{\chi^2}$ might have been expected. We have already discussed in section 2 why the representativeness error (parametrized by β) should be similar in an average or global sense for the two filters. The model error (parametrized by δ) arises from wind errors, missing sources or sinks in the chemistry, diabatic mixing across isentropes, and numerical discretization error in the transport schemes. Only the last error source is different between the Eulerian and Lagrangian filters, and apparently this difference does not substantially modify the value of the running mean $\overline{\chi^2}$. We accept the Lagrangian filter as statistically valid.

Figure 7 shows colour contour plots of the analysis methane mixing ratio for the MCCL baseline data set and parameters after 2, 4 and 8 days. For the Lagrangian filter the plots were produced by offline remapping. For the full 8-day length of the dataset we did not, as we had originally anticipated, find the need to remap the fields or covariance matrices online (i.e. during the course of the run). For the set of experiments performed in this work, gaps or clumping of trajectories do not affect adversely the results of the Lagrangian filter. Gaps do appear (Fig. 2), but represent areas where gradients in the

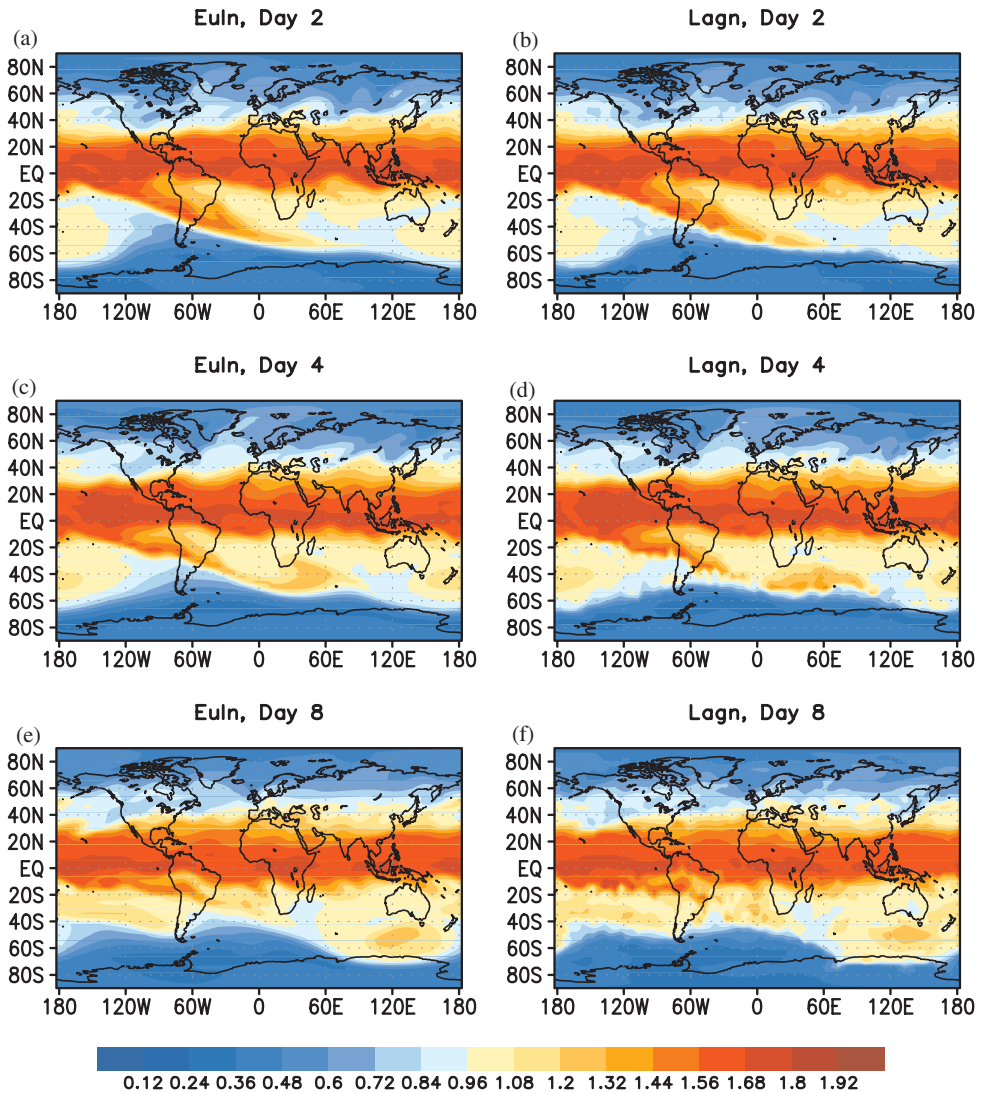


Figure 7. Colour contour plots of the analysis methane mixing ratio w^a for the MCCL baseline dataset and parameters after (a) and (b) 2 days, (c) and (d) 4 days, and (e) and (f) 8 days: for the Eulerian filter (as shown in Ménard and Chang (2000), Fig. 2) with variance and shape correction (left column) and for the Lagrangian filter (right column).

fields are small. The assimilation process also should have a healing effect on the fields: the observational data force the gradients in the fields to resemble the physical world more closely, thus washing out any numerically induced artifacts of the highly irregular meshes. The methane analyses produced by the two filters are quite similar, and both appear to capture well the tongue of air with high methane mixing ratio that protrudes into southern midlatitudes and then breaks away. The mixing-ratio gradients appear to be somewhat sharper for the Lagrangian filter, indicative of the reduced dissipation of its state and error covariance transport scheme. There is a corresponding graininess due to the finite number of trajectories representing the field and the associated remapping required to visualize the field.

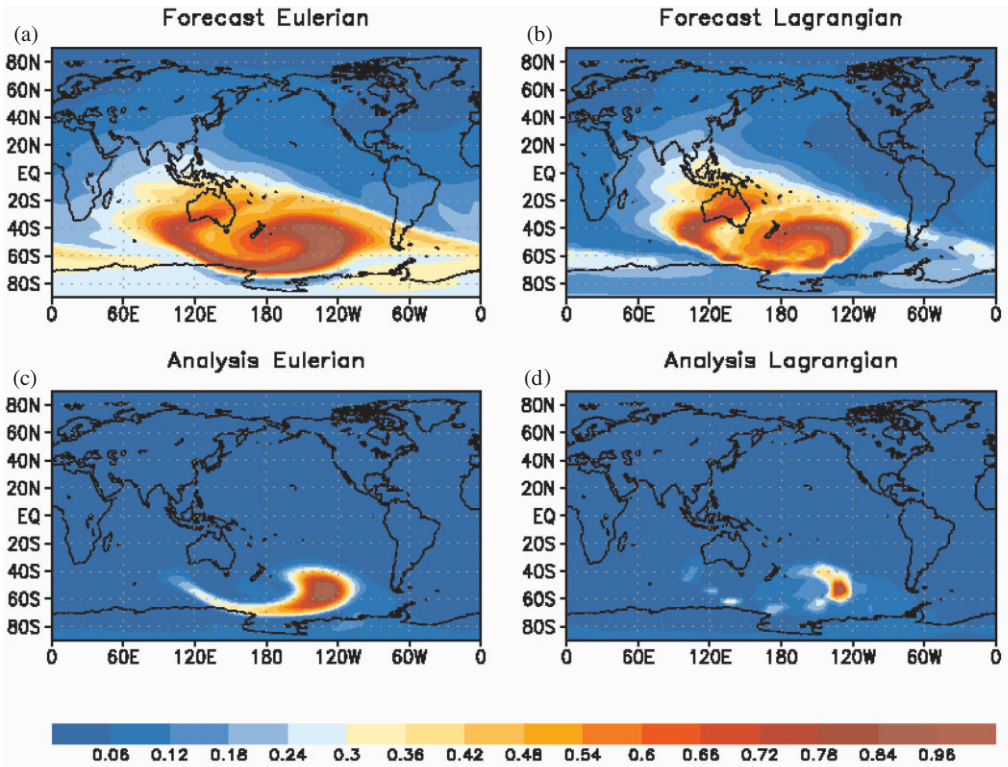


Figure 8. Forecast- and analysis-error correlation structures for the Eulerian and Lagrangian filters after 2 days with respect to a base trajectory that started out at (32°S, 90°E) at day 0: (a)–(b) for the runs in pure forecast mode, and (c)–(d) for the baseline assimilation runs. (a) and (c) Correlation $\text{corr}_E^{f,a}$ from the Eulerian filter; (b) and (d) $\text{corr}_L^{f,a}$ from the Lagrangian filter. See text for further explanation.

Finally we compare the correlation structures of the two filters by evaluating and plotting the correlation between a specific trajectory location $\mathbf{x}(j)$ of the Lagrangian filter and the grid points (i') at the centres of the grid volumes of the Lin–Rood transport model. For the Eulerian filter, both dimensions of the covariance matrix $\mathbf{P}_E^{f,a}$ are indexed by the coordinates of the Lin–Rood model grid. Hence, we evaluate the covariance $\mathbf{P}_E^{f,a}(i', j)$ using bilinear interpolation of the four columns of the covariance matrix corresponding to the four grid points nearest the trajectory location $\mathbf{x}(j)$. The correlation is then defined as:

$$\text{corr}_E^{f,a}(i', j) = \mathbf{P}_E^{f,a}(i', j) / \sqrt{\mathbf{P}_E^{f,a}(i', i') \mathbf{P}_E^{f,a}(j, j)}, \quad (41)$$

where $\mathbf{P}_E^{f,a}(i', i')$ is the variance at the grid point i' of the Eulerian filter, and $\mathbf{P}_E^{f,a}(j, j)$ is evaluated using bilinear interpolation of the four values of $\mathbf{P}_E^{f,a}(i', j)$ corresponding to the four grid points i' nearest the trajectory location $\mathbf{x}(j)$. For the Lagrangian filter, both dimensions of $\mathbf{P}_L^{f,a}$ are indexed by the locations of the trajectories. Hence, we evaluate the covariance $\mathbf{P}_L^{f,a}(i', j) = \mathbf{M}^T \mathbf{P}_L^{f,a}(\cdot, j)(i')$, i.e. the remapped j th column of $\mathbf{P}_L^{f,a}$. The correlation is then defined as:

$$\text{corr}_L^{f,a}(i', j) = \mathbf{M}^T \mathbf{P}_L^{f,a}(\cdot, j)(i') / \sqrt{\mathbf{M}^T \text{diag } \mathbf{P}_L^{f,a}(i') \mathbf{P}_L^{f,a}(j, j)}, \quad (42)$$

where $\mathbf{M}^f \text{diag } \mathbf{P}_L^{f,a}(i')$ is the direct-remapped variance at grid point i' , and $\mathbf{P}_L^{f,a}(j, j)$ is the variance at the trajectory location (j) of the Lagrangian filter.

Figure 8 shows the gridded correlation plots after two days for the trajectory that was initialized at (32°S, 90°E). The plots show an extended comma-shaped structure that reflects the shear associated with the tongue of tropical air penetrating southern midlatitudes. The scale of this structure is reduced by the assimilation of observations assumed to be spatially uncorrelated. The structure is considerably narrower with the Lagrangian filter than with the Eulerian filter, particularly in assimilation mode. The Lagrangian filter also displays some graininess. However, because of its intrinsic ability to represent fine scales more accurately than the Eulerian filter, on the whole the Lagrangian filter probably gives the more accurate portrayal of the correlation structure.

5. CONCLUDING REMARKS

We have developed a new numerical algorithm, the Lagrangian filter, for solving the continuum Kalman filter equations for assimilation of constituent observations directly on particle trajectories. This is a finite-dimensional approximation of the solution of the state estimation problem by characteristics, and may be thought of as an extension of the method of trajectory mapping. We assert that the Lagrangian filter provides a more natural framework for studying and solving the constituent assimilation problem than the Eulerian filter. The Lagrangian filter does not require the variance and shape correction of the Eulerian filter, and it requires fewer floating point operations because of its simple estimation-error covariance propagation step. We implemented the Lagrangian filter for two-dimensional isentropic flow in the stratosphere and compared it with results of the Eulerian filter of Lyster *et al.* (1997), Ménard *et al.* (2000) and Ménard and Chang (2000). We chose parametrizations and parameter values for the initial error, model error and representativeness error identical to those of Ménard and Chang (2000), and accepted them as statistically valid for the Lagrangian filter.

We have found evidence that the Lagrangian filter representation of flow-dependent error correlation structures is more accurate than for the Eulerian filter. We have also found that the Lagrangian state and error covariance propagation, being inherently non-diffusive, gives rise to sharper gradients in the analysed fields, at the cost of some graininess. Finally, because of the conservative properties of the dynamics of the state, error variance, and error covariance along trajectories we were able to probe the behaviour of both the Lagrangian and Eulerian filters. Our experiments suggested that the spurious variance loss of the Eulerian filter is due mostly to the cascade of correlation structure to scales unresolved by the fixed Eulerian grid. The Lagrangian filter resolves these scales explicitly.

ACKNOWLEDGEMENTS

We would like to acknowledge valuable discussions with Stephen Bloom, Gregory Gaspari, Clark Mobarry, Richard Rood, Grace Wahba, and Steven Zalesak. Tom Clune helped us with the optimization of the code on the SGI/Cray T3E machines. Part of Peter Lyster's and Kevin Olson's contribution was funded by the NASA High Performance Computing and Communications Earth and Space Sciences (HPCC ESS) project. Other funding was provided by the NASA EOS Interdisciplinary Science Program, and by the NASA Research and Applications Program. The computing was performed on the SGI/Cray T3E at NASA Goddard Space Flight Center, while some timing runs were performed on the SGI/CRAY LC512 T3E-1200 at Chippewa Falls, MN.

REFERENCES

- Amdahl, G. M. 1967 'Validity of the single-processor approach to achieving large scale computing capabilities'. Pp. 483–485 in Proceedings of AFIPS conference 30, 18–20 April 1967, Atlantic City, N.J. AFIPS Press, Reston, Va, USA
- Cohn, S. E. 1993 Dynamics of short-term univariate forecast error covariances. *Mon. Weather Rev.*, **121**, 3123–3149
- 1997 An introduction to estimation theory. *J. Meteorol. Soc. Jpn.*, **75**, 257–288
- Côté, J. and Staniforth, A. 1988 A two-time-level semi-Lagrangian semi-implicit scheme for spectral models. *Mon. Weather Rev.*, **116**, 2003–2012
- Daley, R. 1993 Estimating observation error statistics for atmospheric data assimilation. *Annales Geophysicae*, **11**, 634–647
- Dritschel, D. G. 1988 Contour surgery: A topological reconnection scheme for extended integrations using contour dynamics. *J. Comput. Phys.*, **77**, 240–266
- Fisher, M. and Lary, D. J. 1995 Lagrangian four-dimensional variational data assimilation of chemical species. *Q. J. R. Meteorol. Soc.*, **121**, 1681–1704
- Fisher, M., O'Neill, A. and Sutton, R. 1993 Rapid descent of mesospheric air into the stratospheric polar vortex. *Geophys. Res. Lett.*, **20**, 1267–1270
- Gelb, A. 1974 *Applied optimal estimation*. M.I.T. Press, Cambridge, Massachusetts, USA
- Ghil, M., Cohn, S. E., Tavantzis, J., Bube, K. and Isaacson, E. 1981 'Applications of estimation theory to numerical weather prediction'. Pp. 139–224 in *Dynamic meteorology: Data assimilation methods*. Eds. L. Bengtsson, M. Ghil and E. Källén. Springer-Verlag, New York
- GMAO 2000 Algorithm theoretical basis document. Global Modeling and Assimilation Office, NASA Goddard Space Flight Center. Available at http://gmao.gsfc.nasa.gov/sci_research/atbd.php
- Jazwinski, A. H. 1970 *Stochastic processes and filtering theory*. Academic Press, New York
- Khattatov, B. V., Gille, J. C., Lyjak, L. V., Brasseur, G. P., Dvortsov, V. L., Roche, A. E. and Waters, J. W. 1999 Assimilation of photochemically active species and a case analysis of UARS data. *J. Geophys. Res.*, **104**, 18715–18737
- Lin, S.-J. and Rood, R. B. 1996 Multidimensional flux-form semi-Lagrangian transport schemes. *Mon. Weather Rev.*, **124**, 2046–2070
- Lorenc, A. C. 1986 Analysis methods for numerical weather prediction. *Q. J. R. Meteorol. Soc.*, **112**, 1177–1194
- Lyster, P. M., Cohn, S. E., Ménard, R., Chang, L.-P., Lin, S.-J. and Olsen, R. 1997 Parallel implementation of a Kalman filter for constituent data assimilation. *Mon. Weather Rev.*, **125**, 1674–1686
- McDonald, A. and Bates, J. R. 1989 Semi-Lagrangian integration of a gridpoint shallow water model on the sphere. *Mon. Weather Rev.*, **117**, 130–137
- Ménard, R. and Chang, L.-P. 2000 Assimilation of stratospheric chemical tracer observations using a Kalman filter. Part II: χ^2 -validated results and analysis of variance and correlation dynamics. *Mon. Weather Rev.*, **128**, 2672–2686
- Ménard, R., Cohn, S. E., Chang, L.-P. and Lyster, P. M. 2000 Assimilation of stratospheric chemical tracer observations using a Kalman filter. Part I: Formulation. *Mon. Weather Rev.*, **128**, 2654–2671
- Mitchell, H. L. and Daley, R. 1997 Discretization error and signal/error correlation in atmospheric data assimilation II: The effect of unresolved scales. *Tellus*, **49A**, 54–73
- O'Neill, A., Grose, W. L., Pope, V. D., McLean, H. and Swinbank, R. 1994 Evolution of the stratosphere during northern winter 1991–2 as diagnosed from the UKMO analysis. *J. Atmos. Sci.*, **51**, 2800–2817
- Renka, R. J. 1997a ALGORITHM 772. STRIPACK: Delaunay triangulation and Voronoi diagram on the surface of a sphere. *ACM Trans. Math. Software*, **23**, No. 3, 416–434
- 1997b ALGORITHM 773. SSRFPACK: Interpolation of scattered data on the surface of a sphere with a surface under tension. *ACM Trans. Math. Software*, **23**, No. 3, 435–442
- Ritchie, H. 1987 Semi-Lagrangian advection on a Gaussian grid. *Mon. Weather Rev.*, **115**, 608–619

- Schoeberl, M. R. and Newman, P. A. 1995 A multiple-level trajectory analysis of vortex filaments. *J. Geophys. Res.*, **100**, 25801–25815
- Sutton, R. T., Maclean, H., Swinbank, R., O'Neill, A. and Taylor, F. W. 1994 High-resolution stratospheric tracer fields estimated from satellite observations using Lagrangian trajectory calculations. *J. Atmos. Sci.*, **51**, 2995–3005
- Waugh, D. W. and Plumb, R. A. 1994 Contour advection with surgery: A new technique for investigation of fine scale structure in tracer transport. *J. Atmos. Sci.*, **51**, 530–540