QUARTERLY JOURNAL OF THE ROYAL METEOROLOGICAL SOCIETY
Q. J. R. Meteorol. Soc. **133**: 1281–1294 (2007)
Published online 13 July 2007 in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/qj.96

RMetS
ROYAL METEOROLOGICAL SOCIETY

Adaptive ensemble reduction and inflation

B. Uzunoglu,[a]* S. J. Fletcher,[b] M. Zupanski[b] and I. M. Navon[a]
[a] *School of Computational Science and Information Technology, Florida State University, USA*
[b] *Cooperative Institute for Research in the Atmosphere, Colorado State University, USA*

ABSTRACT: In this paper we address the question of whether it is possible consistently to reduce the number of ensemble members at a late stage in the assimilation cycle. As an extension, we consider the question: given this reduction, is it possible to reintroduce ensemble members at a later time, if the accuracy is decreasing significantly? To address these questions, we present an adaptive methodology for reducing and inflating an ensemble by projecting the ensemble onto a limited number of its leading empirical orthogonal functions, through a proper orthogonal decomposition. We then apply this methodology with a global shallow-water-equations model on the sphere in conjunction with an ensemble filter developed at Florida State University and the Cooperative Institute for Research in the Atmosphere at Colorado State University. An adaptive methodology for reducing and inflating ensembles is successfully applied in two contrasting test cases with the shallow-water-equations model. It typically results in a reduction in the number of ensemble members required for successful implementation, by a factor of up to two. Copyright © 2007 Royal Meteorological Society

KEY WORDS empirical orthogonal functions; ensemble data assimilation; proper orthogonal decomposition; square-root filter; Shannon entropy; preconditioning

Received 14 December 2006; Revised 10 April 2007; Accepted 13 April 2007

1. Introduction

The use of data, together with a model, to assess the current state of a system and to make predictions, is called 'data assimilation'. The algorithms for doing data assimilation are referred to as 'filters', 'smoothers', or 'predictors' (Jazwinski, 1970). The Kalman filter (KF) (Kalman and Bucy, 1961) is an optimal data-assimilation method for linear dynamics with an additive Gaussian model and observation errors (Jazwinski, 1970). However, the error-covariance calculations associated with the KF are difficult to implement for realistic systems, because of the computational cost or the nonlinearity of the dynamics. The ensemble Kalman filter (EnKF) (Evensen, 1994) addresses some of these issues by using ensemble representations for the forecast- and analysis-error covariances through pseudo-random Monte Carlo perturbations. Ensemble-filter algorithms are of special interest because of their simplicity of implementation: no adjoint operators are required (Evensen, 2003). Another advantage is their potential for efficient use on parallel computers with large-scale geophysical models.

There is a family of ensemble methods, including EnKF, that use matrix square roots of the error-covariance matrices. These methods are called square-root filters (SRFs) (Bierman, 1977). SRFs are not unique, since the square root of the covariance is not unique (Tippett *et al*., 2003). Ensemble assimilation methods can also be considered as low-rank approximations. A good summary of ensemble SRFs that perform a transformation of the ensemble after updating the ensemble mean can be found in (Tippett *et al*., 2003).

Examples of low-rank reduced filters include: the reduced-rank Kalman filter (RRKF) (Fisher and Andersson, 2001); the singular evolutive interpolated Kalman filter (Pham *et al*., 1998); the reduced-rank square-root filter (RRSQRT); the partially-orthogonal ensemble Kalman filter (POEnKF) and the complementary-orthogonal-subspace filter for efficient ensembles (COFFEE) (Heemink *et al*., 2001); and the singular evolutive extended Kalman filter (Pham *et al*., 1998). All these algorithms use a low-rank representation of the covariance matrix, either by using a pseudo-random ensemble, or explicitly. Thus, the filter analyses operate only in a low-dimensional error subspace that approximates the full error space.

The maximum-likelihood ensemble filter (MLEF) (Zupanski, 2005) differs from most others in that it uses the most likely state, rather than the ensemble mean, to generate the statistics. It is similar to the ensemble transform Kalman filter (ETKF) (Bishop *et al*., 2001), and can be regarded as a maximum-likelihood approximation to it. The main difference is that the ETKF uses the ensemble mean to approximate the true statistics, whereas the MLEF minimizes a nonlinear cost function similar to that used in three-dimensional variational data assimilation (3D-Var) (Lorenc, 1986) and uses the solution to approximate the square root of the analysis covariance matrix.

* Correspondence to: B. Uzunoglu, School of Computational Science and Information Technology, Florida State University, Tallahassee, FL 32306-4120, USA. E-mail: uzunoglu@csit.fsu.edu

The MLEF is different from the filter of (Hamill and Snyder, 2000), which is a hybrid of 3D-Var and the EnKF. Hamill and Snyder (2000) use a weighted sum of a climatological and a flow-dependent ensemble covariance for their forecast-error covariance term. The most important difference between this filter and the MLEF (and all other EnKFs) is in the analysis-correction subspace. In the case of the hybrid filter, this is the full error space. In the case of ECMWF's RRKF (Fisher and Andersson, 2001), it is a subspace consisting of some ensemble subspace plus the remaining orthogonal complement from the climatological covariance. The MLEF uses the ensemble-spanned subspace. We discuss the MLEF in more detail in the next section. However, the primary purpose of this paper is not to describe the MLEF, but rather to use it as an example. We will focus on the impact of ensemble reduction and inflation on this filter.

A question that is important in ensemble data assimilation is: how many ensemble members are required to achieve a desired order of accuracy? We do not tackle this question directly in this paper, but rather address the related question: given an ensemble of a certain size, is there a way consistently to reduce the number of ensemble members at a late stage in the assimilation cycle, after the errors from the initiation have been compensated for? An extension to this question is: given such a reduction, is it possible to reintroduce the ensemble members at a later time if the order of accuracy is decreasing?

The answer to both these questions is: yes. In this paper, we show this for the MLEF, but the approach is valid for any forecast-error covariance present in any ensemble filter. The method that we present here makes use of a Karhunen–Loève expansion with the full ensembles to generate specific model evaluations, or snapshots, from which the empirical orthogonal functions (EOFs) are extracted. This is a technique based on the proper orthogonal decomposition (POD) (Rowley, 2005). A detailed comparison between the method employed here and other methods is given in Section 3.6.

In this paper we apply the technique for ensemble reduction with Colorado State University's two-dimensional shallow-water-equations model on the sphere (Heikes and Randall, 1995a,b; Ringler and Randall, 2002), with two sets of initial conditions taken from the test suite described in (Williamson *et al.*, 1992): one representing a zonal flow around a mountain, and the other representing a Rossby–Haurwitz wave, which is known to generate certain types of flow depending on the choice of parameters (Wlasak, 2002; Fletcher, 2004). It is well known that the shallow-water-equations model captures some of the more important flows of the full atmosphere (Daley, 1996), and this model is therefore regarded as a good first test for any new methodology.

The remainder of this paper is organized as follows. The next section provides a brief summary of the MLEF, based on (Zupanski, 2005). Section 3 describes EOFs and the POD and their application to ensemble reduction and inflation. Section 4 describes the shallow-water-equations model that we use to test our method, and the initial conditions we use, followed by the experimental results with different initial ensemble sizes and scales of reduction and inflation. The paper ends with a summary of conclusions and further work.

2. The MLEF

The following summary of the MLEF is based on (Zupanski, 2005).

2.1. Forecast step

The MLEF consists of two steps. The first, the forecast step, is concerned with the evolution of the error covariances associated with the discrete Kalman filter (Jazwinski, 1970). We use the following notation: k is the time index, $\mathbf{P}_\mathrm{f}(k)$ is the forecast-error covariance matrix, $\mathbf{P}_\mathrm{a}(k-1)$ is the analysis-error covariance matrix, and M_e is the nonlinear model evolution. The model error is taken to be zero.

We define factorizations of $\mathbf{P}_\mathrm{a}(k-1)$ and $\mathbf{P}_\mathrm{f}(k)$ into their square-root forms as follows:

$$\left.\begin{aligned}\mathbf{P}_\mathrm{a}(k-1) &= \mathbf{P}_\mathrm{a}^{\frac{1}{2}}(k-1)\mathbf{P}_\mathrm{a}^{\frac{\mathrm{T}}{2}}(k-1)\\ \mathbf{P}_\mathrm{f}(k) &= \mathbf{P}_\mathrm{f}^{\frac{1}{2}}(k)\mathbf{P}_\mathrm{f}^{\frac{\mathrm{T}}{2}}(k)\end{aligned}\right\}. \tag{1}$$

The structure of the square-root analysis-error covariance matrix $\mathbf{P}_\mathrm{a}^{\frac{1}{2}}$ is:

$$\mathbf{P}_\mathrm{a}^{\frac{1}{2}}(k-1) = (\mathbf{p}_1 \quad \dots \quad \mathbf{p}_N), \tag{2}$$

where

$$\mathbf{p}_i = \begin{pmatrix} p_{1i} \\ \vdots \\ p_{Mi} \end{pmatrix},$$

M is the total number of state variables and N is the ensemble size. It is assumed that $N \ll M$.

Expanding this definition, we can express the square-root forecast-error covariance matrix $\mathbf{P}_\mathrm{f}^{\frac{1}{2}}(k)$ as

$$\mathbf{P}_\mathrm{f}^{\frac{1}{2}}(k) = (\mathbf{b}_1 \quad \cdots \quad \mathbf{b}_N),$$

where

$$\mathbf{b}_i = M_\mathrm{e}(\mathbf{x}_\mathrm{a}(k-1) + \mathbf{p}_i) - M_\mathrm{e}(\mathbf{x}_\mathrm{a}(k-1)), \tag{3}$$

$\mathbf{x}_\mathrm{a}(k-1)$ being the analysis of the most likely state from the previous cycle (the analysis part of the MLEF, described below), found from the posterior analysis probability density function (Lorenc, 1986). $\mathbf{P}_\mathrm{f}^{\frac{1}{2}}(k)$ can be obtained from N nonlinear ensemble forecasts plus one

DOI: 10.1002/qj

control forecast. Equation (3) implies the use of a control (deterministic) forecast, instead of an ensemble mean, commonly used in other ensemble data-assimilation methods.

The evolution of the covariance matrix from $\mathbf{P}_{\mathrm{f}}^{\frac{1}{2}}(k)$ to $\mathbf{P}_{\mathrm{a}}^{\frac{1}{2}}(k)$ is defined by:

$$\mathbf{P}_{\mathrm{a}}^{\frac{1}{2}}(k) = \mathbf{P}_{\mathrm{f}}^{\frac{1}{2}}(k)\mathbf{T}(k),$$

where $\mathbf{T}(k)$ is a transformation that we define below. Thus, we have:

$$\mathbf{P}_{\mathrm{a}}(k) = \mathbf{P}_{\mathrm{f}}^{\frac{1}{2}}(k)\mathbf{T}(k)\mathbf{T}^{\mathrm{T}}(k)\mathbf{P}_{\mathrm{f}}^{\frac{\mathrm{T}}{2}}(k). \tag{4}$$

This matrix transform is equivalent to the matrix transform used in the ETKF algorithm (Bishop *et al*., 2001).

2.2. Analysis step

The discrete time index k will be omitted for most of the remainder of this paper. The analysis step for the MLEF involves solving a nonlinear cost function similar to that of (Lorenc, 1986), which is based on a Gaussian assumption for the background and observation variables.

The cost function is defined in terms of $\mathbf{P}_{\mathrm{f}}$, although this matrix is never calculated in the process of the filter. The cost function is of the form:

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x}-\mathbf{x}_{\mathrm{b}})^{\mathrm{T}}\mathbf{P}_{\mathrm{f}}^{-1}(\mathbf{x}-\mathbf{x}_{\mathrm{b}}) + \\ + \frac{1}{2}(\mathbf{y}-\mathbf{h}(\mathbf{x}))^{\mathrm{T}}\mathbf{R}^{-1}(\mathbf{y}-\mathbf{h}(\mathbf{x})), \tag{5}$$

where $\mathbf{y}$ is the vector of observations, $\mathbf{h}$ is the nonlinear observation operator, $\mathbf{R}$ is the observation covariance matrix, and $\mathbf{x}_{\mathrm{b}}$ is a background state. For our purposes, we take

$$\mathbf{x}_{\mathrm{b}} = M_{\mathrm{e}}(\mathbf{x}_{\mathrm{a}}(k-1)).$$

Note that if $\mathbf{P}_{\mathrm{f}}$ is defined in an ensemble subspace only, then its inverse is uniquely defined on its range.

To minimize the expression in Equation (5), a form of Hessian preconditioning is introduced, through the following change of variable:

$$\mathbf{x}-\mathbf{x}_{\mathrm{b}} = \mathbf{P}_{\mathrm{f}}^{\frac{1}{2}}(\mathbf{I}+\mathbf{C}_{\mathrm{e}})^{-\frac{\mathrm{T}}{2}}\boldsymbol{\xi}, \tag{6}$$

where $\boldsymbol{\xi}$ is a vector of control variables defined in the ensemble subspace, and

$$\mathbf{C}_{\mathrm{e}} = \mathbf{P}_{\mathrm{f}}^{\frac{\mathrm{T}}{2}}\mathbf{H}^{\mathrm{T}}\mathbf{R}^{-1}\mathbf{H}\mathbf{P}_{\mathrm{f}}^{\frac{1}{2}} \\ = (\mathbf{R}^{-\frac{1}{2}}\mathbf{H}\mathbf{P}_{\mathrm{f}}^{\frac{1}{2}})^{\mathrm{T}}(\mathbf{R}^{-\frac{1}{2}}\mathbf{H}\mathbf{P}_{\mathrm{f}}^{\frac{1}{2}}), \tag{7}$$

$\mathbf{H}$ being the Jacobian matrix of $\mathbf{h}$ evaluated at $\mathbf{x}$.

It may be the case that the observation operators are highly nonlinear, and difficult to differentiate analytically, or even discontinuous. Therefore we use information about the forecast-error covariance matrix to approximate the square root of $\mathbf{C}_{\mathrm{e}}$ componentwise as

$$\mathbf{z}_i = (\mathbf{R}^{-\frac{1}{2}}\mathbf{H}\mathbf{P}_{\mathrm{f}}^{\frac{1}{2}})_i \\ = \mathbf{R}^{-\frac{1}{2}}\mathbf{H}\mathbf{b}_i \\ \approx \mathbf{R}^{-\frac{1}{2}}(\mathbf{h}(\mathbf{x}+\mathbf{b}_i)-\mathbf{h}(\mathbf{x})). \tag{8}$$

The mean of $\mathbf{x}$ is not used in the MLEF formalism: $\mathbf{x}$ can only be related to the mode for the Gaussian probability density function. The MLEF directly minimizes the cost function with nonlinear observation operators, whereas other EnKFs insert nonlinearities into the original linear KF solution. Many practical EnKF algorithms need an additional covariance inflation to account for nonlinearities. The variable $\mathbf{x}$ in $\mathbf{h}(\mathbf{x}+\mathbf{b}_i)$ is the current iterative solution of the control state vector, so $\mathbf{x} = \mathbf{x}_{\mathrm{b}}$ at the first iteration. At a particular iteration, it is not necessarily either the mean or the mode: it is just the result of the minimization. The final result of the minimization approximates the mode.

A matrix

$$\mathbf{Z} = (\,\mathbf{z}_1 \quad \cdots \quad \mathbf{z}_N\,), \tag{9}$$

is defined so that

$$\mathbf{C}_{\mathrm{e}} = \mathbf{Z}^{\mathrm{T}}\mathbf{Z}.$$

The square-root analysis-error covariance matrix is updated as:

$$\mathbf{P}_{\mathrm{a}}^{\frac{1}{2}} = \mathbf{P}_{\mathrm{f}}^{\frac{1}{2}}(\mathbf{I}+\mathbf{C}_{\mathrm{e}}(\mathbf{x}_{\mathrm{a}}))^{-\frac{\mathrm{T}}{2}}, \tag{10}$$

where $\mathbf{x}_{\mathrm{a}}$ minimizes $J(\mathbf{x})$ (Equation (5)).

To generate the perturbation for the ensemble, we simply use the columns of the $\mathbf{P}_{\mathrm{a}}^{\frac{1}{2}}$ matrix. This represents a square-root inverse Hessian estimated at the minimum, and so it is a good approximation of the square-root analysis-error covariance. The transformation $\mathbf{T}$ defined above can be written as

$$\mathbf{T} = (\mathbf{I}+\mathbf{C}_{\mathrm{e}}(\mathbf{x}_{\mathrm{a}}))^{-\frac{\mathrm{T}}{2}}.$$

Then Equation (10) becomes:

$$\mathbf{P}_{\mathrm{a}}^{\frac{1}{2}} = \mathbf{P}_{\mathrm{f}}^{\frac{1}{2}}\mathbf{T}. \tag{11}$$

The final algorithmic detail that we give here concerns the inversion of the matrix $\mathbf{C}_{\mathrm{e}}$. This is accomplished through an eigenvalue decomposition

$$\mathbf{C}_{\mathrm{e}} = \mathbf{V}\boldsymbol{\Lambda}_{\mathrm{e}}\mathbf{V}^{\mathrm{T}},$$

DOI: 10.1002/qj

where $\mathbf{V}$ is the vector of orthogonal eigenvectors and $\mathbf{\Lambda}_e$ is a diagonal matrix containing the eigenvalues of $\mathbf{C}_e$. Then the transformation $\mathbf{T}$ can be rewritten as

$$(\mathbf{I} + \mathbf{C}_e(\mathbf{x}_a))^{-\frac{T}{2}} = \mathbf{V}(\mathbf{I} + \mathbf{\Lambda}_e)^{-\frac{1}{2}}\mathbf{V}^T. \tag{12}$$

For more details of the MLEF algorithm, see (Zupanski, 2005). We do not know of any biased analyses in the MLEF. This may be related to the fact that there is no factor $1/(N-1)$ in Equation (10), so that the MLEF can create larger, and in general different estimates for ensembles, compared to ETKF.

3. Adaptive ensemble reduction and inflation

In this section we describe the criteria that we use for ensemble reduction and inflation. We start with a brief overview of the theory of EOFs and the POD. This leads us to the Shannon entropy and the variance percentage, which are the two measures that we use to assess the need for a reduction in the number of ensemble members. In Section 3.1 we show how we can use the Shannon entropy to ascertain how many ensemble members we need to keep. In Section 3.3 we explain the procedure for inflating the ensemble. Lastly, in Section 3.6 we explain the differences between the method derived here and those described in (Oczkowski *et al*., 2005) and (Heemink *et al*., 2001).

3.1. POD and Shannon entropy

As we have already mentioned, an important feature of the MLEF is its ability to define the full Hessian preconditioner in ensemble space, Equation (6). The preconditioner is based on the ensemble correlation matrix $\mathbf{C}_e$, Equation (7). For the remainder of the paper we will drop the notation $\mathbf{x}_a$ for the minimum of $\mathbf{C}_e(\mathbf{x})$. We can write

$$\begin{aligned}\mathbf{C}_e &= \mathbf{Z}^T\mathbf{Z} \\ &= \begin{pmatrix} \mathbf{z}_1^T\mathbf{z}_1 & \cdots & \mathbf{z}_1^T\mathbf{z}_N \\ \vdots & \ddots & \vdots \\ \mathbf{z}_N^T\mathbf{z}_1 & \cdots & \mathbf{z}_N^T\mathbf{z}_N \end{pmatrix}.\end{aligned} \tag{13}$$

The Hessian preconditioner transformation $\mathbf{T}$ could also have employed a covariance matrix defined in observation space, $\mathbf{C}_o = \mathbf{Z}\mathbf{Z}^T$, as discussed in (Bishop *et al*., 2001). However, in order to calculate $\mathbf{C}_o$ we need to determine an $M \times M$ covariance matrix defined in observation space. The dual approach involves the ensemble subspace, with an $N \times N$ covariance matrix $\mathbf{C}_e = \mathbf{Z}^T\mathbf{Z}$. An advantage of this is that generally $N \ll M$. We also know that the leading eigenvalues of the transformation $\mathbf{T}$ form an incomplete subset of the leading eigenvalues of $\mathbf{C}_o$. Therefore the two matrices have the same leading-eigenvalue spectrum (Sirovich, 1987; Bishop *et al*., 2001; Antoulas, 2005). We try to adaptively control the size of the ensemble by employing the transformation $\mathbf{T}$.

The basic problem that we address here is how to identify bases $\mathbf{U}$ that represent the coherent structures, or patterns, in a random vector field

$$\mathbf{Z} = (\,\mathbf{z}_1 \quad \cdots \quad \mathbf{z}_N\,). \tag{14}$$

Given an ensemble of random vector fields $\mathbf{Z}$, we seek bases that have a structure typical of the members of the ensemble. Specifically, we consider the optimization problem over the orthogonal transformations. This is because the orthogonal transformations are linear, so that the optimization problem amounts to finding best orthogonal bases or coherent structures (Sirovich, 1987).

Thus, given a field $\mathbf{Z}$, we seek the orthogonal matrix $\mathbf{U}$ for which the inner product $\langle \mathbf{z}_i, \mathbf{U}\rangle$ is the standard Euclidean inner product. Each column $\mathbf{z}_i$ in $\mathbf{Z}$ may be expressed in terms of the basis vectors as

$$\mathbf{z}_i = a_1\mathbf{U}_1 + \cdots + a_M\mathbf{U}_M, \tag{15}$$

where $a_k = \langle \mathbf{z}_k, \mathbf{U}_i\rangle$. Here k and i range from 1 to M, and $\mathbf{U}$ defines the orthogonal basis. The inner-product norm of $\langle \mathbf{z}_i, \mathbf{U}\rangle$ here is a Frobenius norm, which is a discrete version of the Euclidean norm (Sirovich, 1987; Preisendorfer, 1988).

The POD method is one way to find the structures that we seek. This method is able directly to compute the coherent structures in a random field of ensemble members, in an optimal-linear sense (Sirovich, 1987). This type of decomposition has many names, including 'principal components analysis', 'Karhunen–Loève decomposition', and 'total-least-squares estimation'. It leads to two dual eigenvalue problems for the covariance matrix. The detailed derivations of the eigenvalue problems used to obtain these bases can be found in (Sirovich, 1987) or (Berkooz *et al*., 1993).

Two corresponding dual eigenvalue problems for the covariance matrix are:

$$\left.\begin{aligned}\mathbf{C}_o\mathbf{U} &= \mathbf{\Lambda}_o\mathbf{U} \\ \mathbf{C}_e\mathbf{V} &= \mathbf{\Lambda}_e\mathbf{V}\end{aligned}\right\}, \tag{16}$$

where the leading eigenvalues in $\mathbf{\Lambda}_e$ of the symmetric $N \times N$ matrix $\mathbf{C}_e = \mathbf{Z}^T\mathbf{Z}$, and the leading eigenvalues in $\mathbf{\Lambda}_o$ of the symmetric $M \times M$ matrix $\mathbf{C}_o = \mathbf{Z}^T\mathbf{Z}$, are equal.

There is also a direct relationship between the singular-value decomposition (SVD) and the dual eigenvalue problems of POD. If we compute the SVD of $\mathbf{Z}$,

$$\mathbf{Z} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T,$$

then $\mathbf{U}$ corresponds to the matrix of left-singular vectors, $\mathbf{V}$ is the matrix of right-singular vectors, and $\mathbf{\Sigma}$ is the matrix containing the singular values of $\mathbf{Z}$. If we let λ_i be the ith-largest eigenvalue of $\mathbf{\Lambda}_o$, then $\sigma_i^2 = \lambda_i$ for $i = 1, \ldots, N$, where σ_i the ith-largest singular value of $\mathbf{\Sigma}$. See (Volkwein, 2004).

DOI: 10.1002/qj

The eigenvalue problem employed in the transformation $\mathbf{T}$ described in Section 2 is to find the eigenvalues and eigenvectors of a symmetric $N \times N$ correlation matrix defined by $\mathbf{C}_e = \mathbf{Z}^T\mathbf{Z}$. The POD basis for this eigenvalue problem is orthonormal, i.e.

$$\mathbf{V}_i^T\mathbf{V}_j = \begin{cases} 0 & i \neq j \\ 1 & i = j \end{cases}. \tag{17}$$

The trace of the matrix $\mathbf{C}_e$ represents the total variance in the ensembles,

$$E = \sum_{i=1}^{N} \lambda_i. \tag{18}$$

A *variance percentage* can then be assigned to each eigenvector:

$$E_k = \frac{\lambda_k}{E}. \tag{19}$$

A standard measure of the information content in these eigenvalues of the covariance matrix is the *Shannon entropy*, which is defined as

$$H_s = -\sum_{k=1}^{N} E_k \ln E_k. \tag{20}$$

Here, N will be defined by the inflated or reduced size of the ensemble from the previous cycle. We can interpret the scaled eigenvalues E_k of the covariance matrix as probabilities, since they sum to one. It is then possible to show that the POD eigenvectors are optimal in an information-theoretic sense (Watanabe, 1965). In the ensemble reduction, we would like to retain as much of the original information as possible. This is equivalent to maximizing H_s (Watanabe, 1965). After addition of all ordered scaled eigenvalues E_k, H_s will be flattened:

$$\frac{\partial H_s}{\partial N} \approx 0.$$

This can be interpreted as giving a truncation degree for the size of the ensemble, beyond which adding further eigenvalues adds very little information. The significance of the Shannon entropy in this context is that it provides a measure of the distribution of the magnitudes of the eigenvalues. The eigenvalues correspond to percentages of energy, if energy scaling, as defined in (Oczkowski *et al*., 2005), is employed.

3.2. Application to adaptive ensemble reduction

In POD, since the eigenvalues are arranged in descending order of size, the basis is arranged in descending order of variance. After addition of all ordered scaled eigenvalues E_k, H_s is flattened. This provides us with a criterion for determining the ensemble size adaptively, according to the variance percentage retained, the limit being expressed by:

$$\frac{\partial H_s}{\partial N} \approx 0.$$

Equivalently, this can be thought of as a damped SVD solution in which filter factors of zero are used for basis vectors associated with smaller singular values (Hansen, 1998). The most common approach to regularization of rank-deficient problems is to consider the given matrix $\mathbf{C}_e$ as a noisy representation of a rank-deficient matrix, and to replace $\mathbf{C}_e$ by a matrix that is close to $\mathbf{C}_e$ and rank-efficient (Hansen, 1998). Using the relation between POD and SVD defined above, one can relate 'truncated SVD' and regularization to the POD eigenvalue truncation employed in this paper.

3.3. Application to adaptive ensemble inflation

The algorithm must also allow for ensemble enrichment, in case the reduction in ensemble size is not consistent with the decrease in the root mean square (RMS) error. One option is to increase the correlation length of the matrix $\mathbf{C}_e$ with a new update of the system by ensemble initialization (Zupanski *et al*., 2006).

In this study we have adopted the following approach. For ensemble inflation, the computed eigensets derived from the snapshots depend on the initial and final portions of the ensemble set under consideration. For this purpose, the POD method can be extended, as described in (Glezer *et al*., 1989; Uzunoglu, 2001; Uzunoglu and Nair, 2001). For statistically stationary data, the extended POD (EPOD) is equivalent to the classical POD. The corresponding eigenvalue problem is then to find the eigenvalues and eigenfunctions of a symmetric $2N \times 2N$ correlation matrix, defined by $\mathbf{C}_e$, which inherits the current $\mathbf{z}_N(k)$ and the previous $\mathbf{z}_N(k-1)$. For ensemble enrichment by inflation, the reduced ensemble can be inflated using the ensemble data generated by the filter at previous cycles. This means that we can define a matrix

$$\widetilde{\mathbf{Z}} = (\mathbf{z}_1(k-1) \cdots \mathbf{z}_N(k-1)\ \mathbf{z}_1(k) \cdots \mathbf{z}_N(k)), \tag{21}$$

that enables us to write

$$\widetilde{\mathbf{C}}_e = \widetilde{\mathbf{Z}}^T\widetilde{\mathbf{Z}},$$

which is our new $\mathbf{C}_e$ matrix. This idea can be extended to any number of previous cycles of data assimilation. This approach can also be employed to enrich the current ensembles from previous ensembles while keeping the ensemble size the same. Enrichment by EPOD might also lead to greater efficiency in the reduction strategy used here, since we would have richer ensembles. Since our problem is non-stationary and nonlinear, new directions will be probably introduced by EPOD; these may improve the results, and at least should not make them deteriorate (Glezer *et al*., 1989).

3.4. Measures of performance

We use the RMS error as our measure of the performance of the ensemble filter. This is simply

$$e = \sqrt{\frac{1}{M}\sum_{i=1}^{M}(x_i - x_t)^2}, \tag{22}$$

DOI: 10.1002/qj

where x is one of the state variables of the model, M is the number of instances of the variable, and x_t is the 'true' solution, which is obtained from a run of the model without any errors added.

When the RMS error is not available, we must employ other tools to judge the quality of a reduced or inflated assimilation cycle. Our aim is to have an online tool for measuring the quality of the adaptive data assimilation. The quality of a given (reduced- or inflated-ensemble) assimilation cycle can be judged by the extent to which it preserves the value of the ensemble covariance matrix structures, from the previous cycle $\mathbf{Z}(k-1)$ to the current cycle $\mathbf{Z}(k)$. These structures are expected to be different, since our assimilation is evolving; however, we would not expect a significant difference between the error covariance matrices of two successive cycles. They should be similar, and they should evolve to be more and more similar if the data assimilation is successful. This approach can be used to assess the quality of reduction or inflation over a cycle.

To assess the quality of the $\mathbf{C}_e(k-1)$ and $\mathbf{C}_e(k)$ matrices, which are computed and saved at the end of each data-assimilation cycle, we define a matrix $\mathbf{S}$ (Buizza, 1994), whose (i, j) entry is

$$s_{ij} = \langle \mathbf{V}_i, \widetilde{\mathbf{V}}_j \mathbf{\Lambda}_e^j \rangle, \tag{23}$$

the squared scalar product of basis $\mathbf{V}_i(k-1)$ of $\mathbf{C}_e(k-1)$ and basis $\widetilde{\mathbf{V}}_j(k)$ of the reduced or inflated model $k-1$. Thus, $\mathbf{S}$ represents the amount of the variance of the ith eigenvector of the covariance matrix $\mathbf{C}_e(k-1)$ that is 'explained by' the jth eigenvector of the covariance matrix $\mathbf{C}_e(k)$. The $\mathbf{\Lambda}_e$ contains the eigenvalues of $\mathbf{C}_e(k)$, representing the weight factors that are the eigenvalues of the current cycle (Buizza, 1994). $\mathbf{C}_e(k)$ belongs to the current cycle while $\mathbf{C}_e(k-1)$ belongs to the previous cycle. Based on the matrix $\mathbf{S}$, a 'similarity index' is defined as

$$SM = \frac{1}{N} \sum_{i,j=1}^{N} s_{ij}. \tag{24}$$

This takes values between 0 and 1. For identical models, $SM = 1$, since their bases will be parallel. As SM approaches the value 1, we would expect the RMS error to decrease. Since our system is evolving continuously through each cycle, we would not expect $SM = 1$.

3.5. Summary of the algorithm

We can summarize our algorithm for adaptive ensemble reduction and inflation, and the preconditioning algorithm, as follows.

1. Generate the forecast ensemble

$$\mathbf{Z} = \mathbf{R}^{-\frac{1}{2}} \mathbf{H} \mathbf{P}_f^{\frac{1}{2}},$$

where

$$\mathbf{P}_f^{\frac{1}{2}}(k) = (\mathbf{b}_1 \quad \cdots \quad \mathbf{b}_M)$$

and

$$\mathbf{b}_i = M_e(\mathbf{x}_a(k-1) + \mathbf{p}_i) - M_e(\mathbf{x}_a(k-1)).$$

2. Minimize the cost function. Generate the analysis-error covariance matrix

$$\mathbf{C}_e(\mathbf{x}_a) = \mathbf{Z}^T \mathbf{Z}.$$

3. Solve the eigenvalue problem

$$\mathbf{C}_e(\mathbf{x}_a)\mathbf{V} = \lambda \mathbf{V}.$$

4. Calculate total variance

$$E = \sum_i \lambda_i.$$

5. Decide on a new ensemble size corresponding to the number of eigenvectors for which

$$\frac{\partial H_s}{\partial N} \approx 0.$$

After addition of all ordered normalized eigenvalues E_k, H_s will be flattened, as expressed by the above equation. This can be interpreted as giving the truncation degree for the ensemble, since adding further eigenvalues adds very little information. Note that N represents the ensemble size of the previous cycle.

6. Check the performance of ensemble reduction by computing

$$SM = \frac{1}{N} \sum_{i,j=1}^{N} s_{ij}$$

and using the bases of both cycles, as already computed. If $SM(k-1) < SM(k)$ and H_s is becoming flatter, then continue reduction; otherwise (if SM is decreasing or H_s is not flattening), inflate the ensemble using EPOD. Determine the new ensemble size by using ensemble members from the previous cycle, continuing until H_s becomes flat.

7. Generate the square-root analysis-error covariance matrix

$$\mathbf{P}_a^{\frac{1}{2}} = \mathbf{P}_f^{\frac{1}{2}} (\mathbf{I} + \mathbf{C}_e(\mathbf{x}_a)^{-\frac{T}{2}})$$

(Zupanski, 2005).

3.6. Comparison with other methods

The method proposed in this paper gives a global reduction of the ensemble. There is research by the Chaos Group at the University of Maryland that is based on a local ensemble Kalman filter (Patil *et al.*, 2001), now called the local ensemble transform Kalman filter (LETKF) (Oczkowski *et al.*, 2005). The important feature of the LETKF is the use of a quantity known as the 'ensemble dimension' (Patil *et al.*, 2001), which measures the variance associated with each ensemble

DOI: 10.1002/qj

member in terms of the eigenvalues of the forecast-error covariance matrix. The LETKF employs spatial localizations of the total model domain. The measure is applied only in the local domains, and has no interaction with the surrounding domains. In (Patil *et al*., 2001), no ensemble transformation is used to define the new ensemble size for the evolving cycles.

The approach to adaptive ensemble reduction and inflation presented in this paper is different from that of the variance-reduced ensemble Kalman filters, such as RRSQRT, POEnKF and COFFEE, presented in (Heemink *et al*., 2001). In our approach, the error covariance matrix evolves in the ensemble subspace. Heemink *et al*. (2001) discuss matrix-inversion problems caused by ill-conditioning of the matrix, and filter-divergence problems, but do not provide systematic quantitative results on adaptive ensemble reduction and inflation.

The covariance matrix $\mathbf{C}_e$ is similar to the matrix defined in (Oczkowski *et al*., 2005): both are defined in the ensemble subspace. When Oczkowski *et al*. (2005) calculate the ensemble size they calculate the elements of the observed variable vectors using an energy scaling. They define this scaling so that the Euclidean norm of the vectors formed by the different variables has the dimensions of energy. We perform no such scaling. Although this may be a good approach to relating energetic coherent structures to eigenvalues of such a covariance matrix (Sirovich, 1987), it requires definition of a relevant $\mathbf{T}$ transformation that is different from that defined in (Zupanski, 2005). In such a (continuous, infinite-dimensional) case, the L^2 energy norm of the Euclidean norms would have to be used, and then a transformation like $\mathbf{T}$ would have to be defined for such a norm. Oczkowski *et al*. (2005) do not address the question of how to transform the reduced-ensemble members from one cycle to another.

Oczkowski *et al*. (2005) use a regional covariance matrix, whereas we use a global covariance matrix. Furthermore, they do not normalize the covariance matrix $\mathbf{R}^{-\frac{1}{2}}$ by the observations, as in Equation (8). This normalization avoids problems of ill-conditioning (Bishop *et al*., 2001).

The MLEF estimates the mode, without requiring a reference to the mean. It just happens that these two estimates (of the mean and mode) are identical for the Gaussian distribution and some other symmetric probability distributions. Nor does the MLEF try to estimate the mean using the mode: it simply estimates the mode by maximizing the posterior probability density function. The covariances in the MLEF are defined as symmetric positive-semidefinite matrices (Gaspari and Cohn, 1999), so the expectation (mean) is not explicitly used.

4. Experiments and results

In this section we outline the model and the initial conditions that we use for the experiments, and then present the results.

4.1. Model, initial conditions and observations

The shallow-water equations have often been used in the development of new numerical methods for atmospheric models, because they exhibit the same wave behaviour as the more complex baroclinic equations governing the motion of the atmosphere (Daley, 1996).

The nonlinear shallow-water equations on the sphere that we use are of the form given in (Ringler and Randall, 2002). The prognostic variables of vorticity, divergence and mass are all defined at the grid-cell centres. We use the 'Z' grid introduced by Heikes and Randall (1995a, 1995b): a twisted icosahedral geodesic grid, composed of a tessellation of hexagons and pentagons. The resolution of the grid enables us to have approximately 2600 height points and 5100 points for wind components.

As initial conditions we use two of the test cases from (Williamson *et al*., 1992):

- Test case 1 is a geostrophically-balanced zonal flow over an isolated conical mountain. This induces flows with strong nonlinearity in the vicinity of the mountain. This set of conditions is characterized by the excitation of Rossby and gravity waves. The initial zonal flow is 20 ms^{-1}. The mountain is centred at (30 °N, 90 °W), and is 2000 m tall.
- Test case 2 is a Rossby–Haurwitz wave. This is the analytic solution of the nonlinear barotropic-vorticity equation on the sphere, and in this model it is non-divergent. This property does not appear in the shallow-water model.

The observations are created by adding random perturbations from a normal distribution, $\mathcal{N}(\mathbf{0}, \mathbf{R})$, where $\mathbf{R}$ is the covariance matrix, assumed to be diagonal. Its variables are uncorrelated with a model forecast, which we refer to as the 'truth'.

There are 1025 observations defined in each analysis cycle, uniformly distributed with respect to the Z grid around the globe. These observations consist of 513 height observations and 512 wind observations in our assimilation observation space. Since the two wind components (east–west and north–south) are co-located, there are 256 observation points of the wind as a vector. In both of our test cases the geopotential and velocity observations are uncorrelated.

Numerical experiments are conducted to compare the impacts on the RMS height analysis (as well as the u and v RMS, not shown), and the similarity index SM, for the adaptively reduced and inflated and the full-size ensembles.

Two different criteria have been tested: a variance-percentage criterion and an entropy criterion. As we will see below, the entropy criterion appears to be much more reliable than the variance-percentage criterion.

DOI: 10.1002/qj

4.2. Test case 1

4.2.1. Ensemble reduction

The assimilation for ensemble reduction is performed over an interval of 120 h, which corresponds to 40 cycles of data assimilation, each representing a 3 h interval. The initial conditions for the experimental run are as defined above, but shifted to −3 h in order to create an unbalanced set of initial conditions. The observation standard deviation for the height is 5 m, and for the wind is 0.5 ms^{-1}. The background standard deviation for the height is 2 m, and for the wind is 0.2 ms^{-1}.

The algorithm allows the adaptive reduction to be started from any cycle of the full-ensemble run. We perform various numerical experiments, starting the adaptive ensemble reduction after 1, 20, 30 or 40 cycles. This provides a rigorous test of the procedure and allows for a fuller understanding and validation of the procedure.

The figures that follow show:

- the cycle number of the full-ensemble run at which the adaptive reduction is initiated;
- the variance percentage used as a criterion (cases of 95% and 99% explained variance have been chosen);
- the resulting reduction in the number of ensemble members.

Some of the graphs show the evolution of the RMS for the full ensemble compared to the RMS for the

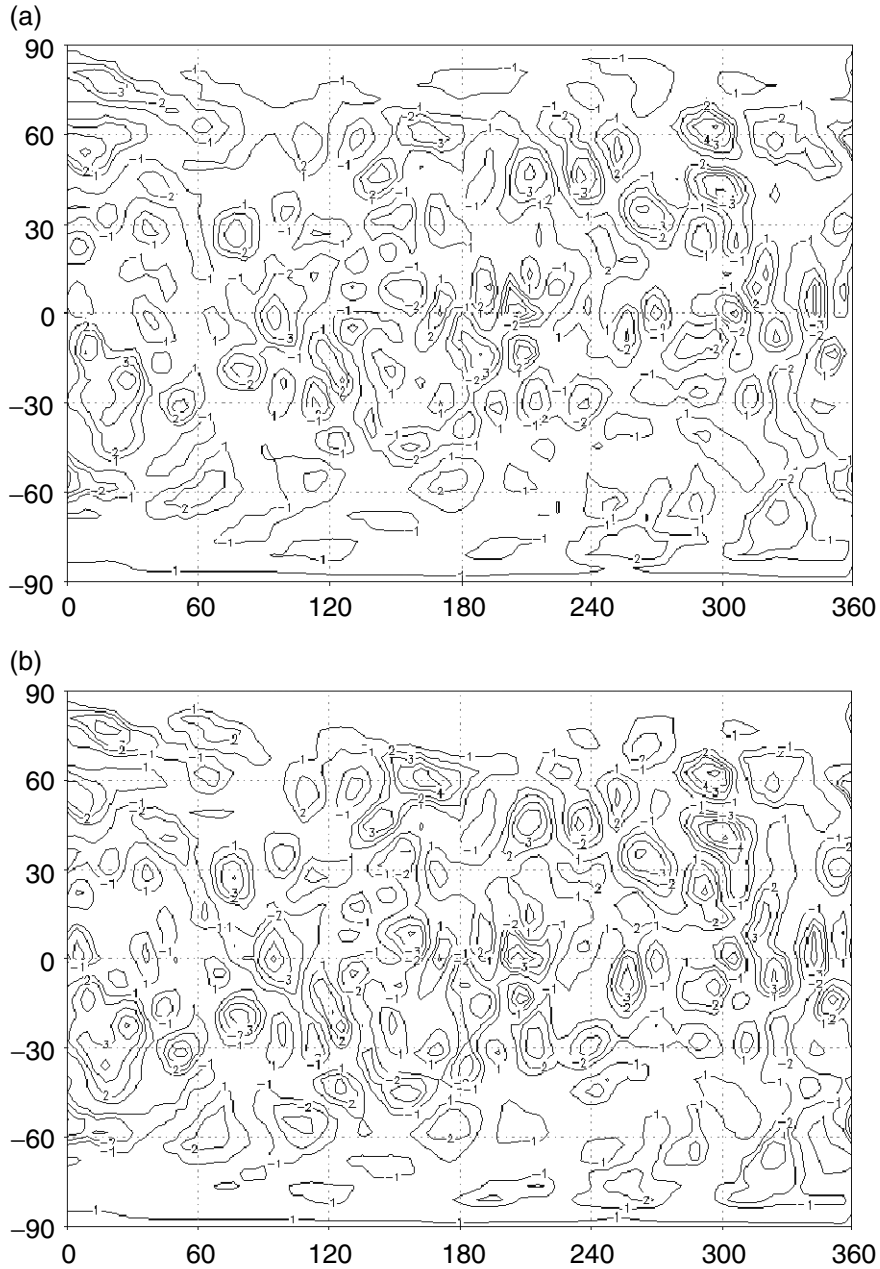

Figure 1. Error isolines of height analysis (m) for test case 1 at cycle 26. (a) Full ensemble (1000 members). (b) Reduced ensemble (280 members) with 95% variance retained.

DOI: 10.1002/qj

adaptively-reduced ensemble; others show the evolution of the similarity index for the same two ensembles.

Figure 1 consists of two plots. Figure 1(a) shows the height analysis error field for an ensemble of size 1000; Figure 1(b) shows the same field but for a reduced ensemble size of 280, adaptively reduced at cycle 26; both plots show the field at cycle 26. The reduction criterion was for a variance percentage of 95%. It is clear that the height analysis error increases only slightly despite a substantial reduction in ensemble size.

Figure 2 shows the height analysis results for test case 1 at cycle 26. Figure 2(a) shows the height field for an ensemble of size 1000, while Figure 2(b) shows the height field with adaptive reduction having been applied at cycle 20, retaining 95% variance at that time. The ensemble size is reduced to 280. It is again clear that the analysis height fields are similar even with a reduced ensemble size.

Figure 3 shows results for various times at which the reduction criterion is introduced. Here we are retaining 99% of the variance, but reducing in accordance with the Shannon-entropy criterion. The left-hand plots are of RMS error for the full (1000-member) and reduced ensemble. The right-hand plots are of the ensemble size.

We see that with the variance criterion for reduction at cycle 1 we can reduce the ensemble size by nearly 300, but that we degrade the solution by doing so. If we use this criterion at cycle 20, we can reduce the ensemble size by a factor of more than two, while still obtaining results consistent with the full field, although there are slight differences in the later cycles. If we reduce at cycle 30, then again we can reduce the ensemble size by a factor of more than two, and the results are still more consistent with the full ensemble. We could improve the RMS height analysis results for the reduced ensemble by increasing the variance percentage used a criterion for

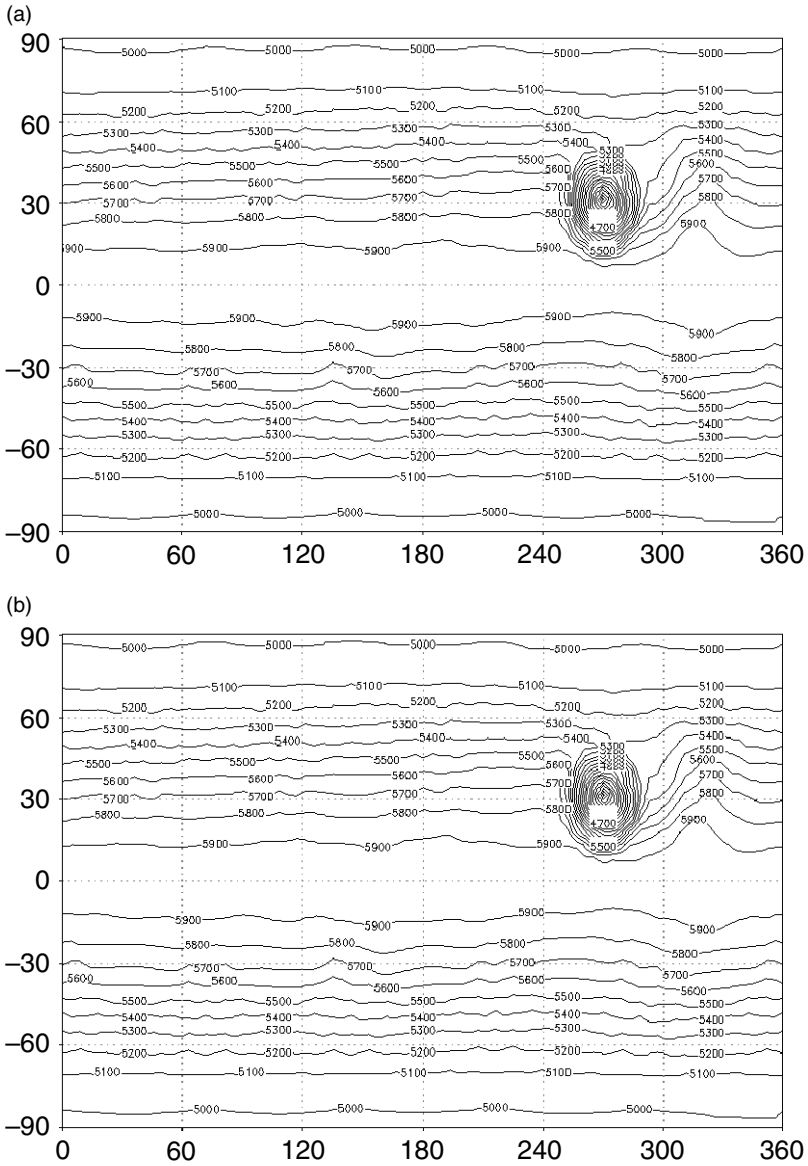

Figure 2. Isolines of height analysis (m) for test case 1 at cycle 26. (a) Full ensemble (1000 members). (b) Reduced ensemble (280 members) Adaptive reduction was introduced at 20.

DOI: 10.1002/qj

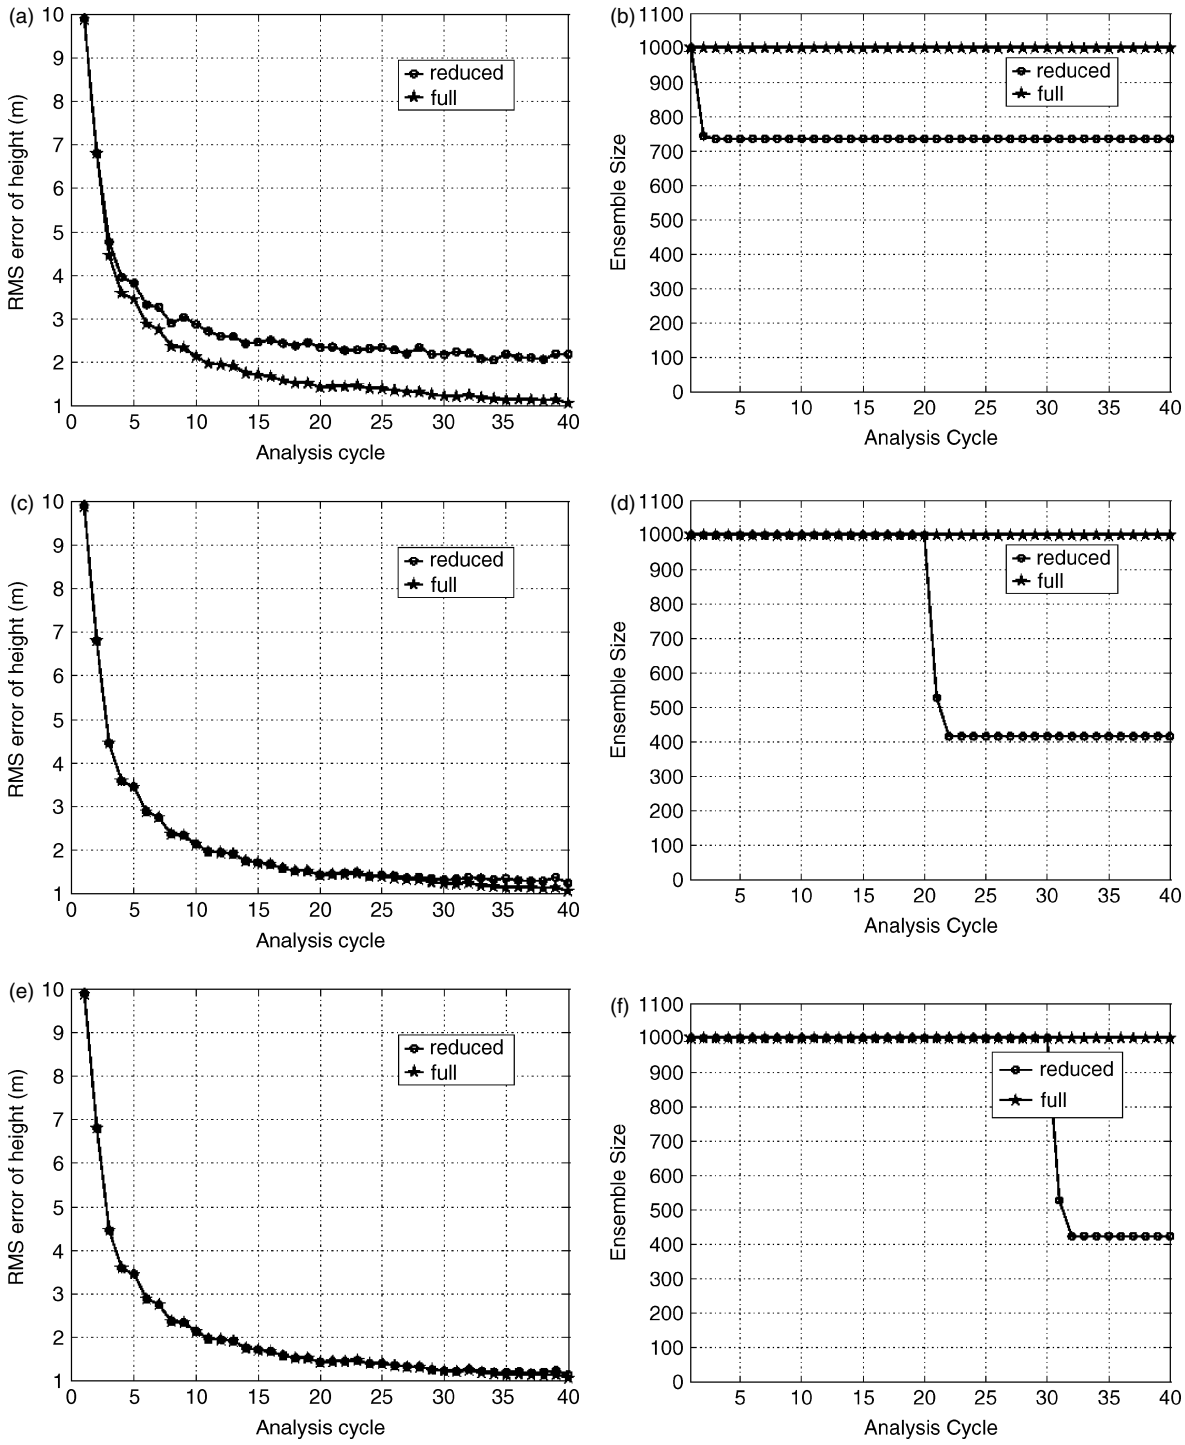

Figure 3. Left panels: evolution of RMS height error for full and reduced ensembles, for adaptive reduction beginning (from top to bottom) at cycle 1, cycle 20 and cycle 30, with a retained variance of 99%. Right panels: evolution of ensemble size for the same experiments.

the reduction – for example, by requiring an explained variance of 99.5% instead of 99%.

Since the variance-percentage criterion does not perform well for the initial cycles of the ensemble data assimilation, the remainder of the results that we show here are based on the Shannon-entropy criterion. Nevertheless, the variance-percentage criterion might perform well in the later cycles.

Figure 4(a) illustrates a decrease in the eigenvalue spectrum, scaled by the total variance E, for the full 1000-member ensemble. As the assimilation proceeds, the gradients of the normalized spectra become steeper, while the scaled magnitudes of the eigenvalues of $\mathbf{C}_e$ decrease. This signifies a reduction in the rank of the matrix $\mathbf{C}_e$ as the assimilation proceeds.

Figure 4(b) shows the Shannon entropy for cycles 1, 3, 10 and 30. For cycle 30 we have

$$\frac{\partial H_s}{\partial N} \approx 0,$$

DOI: 10.1002/qj

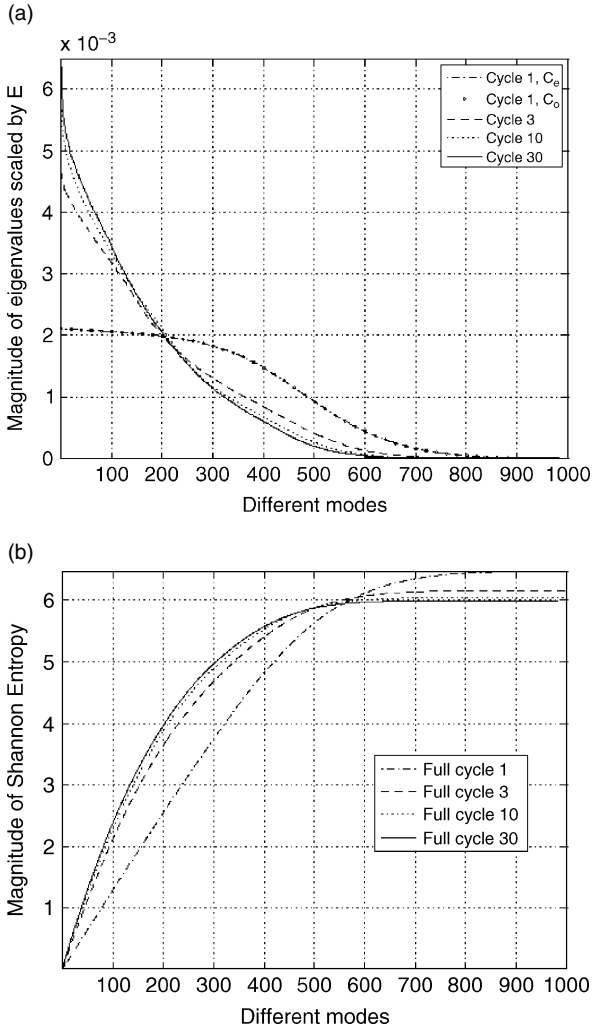

Figure 4. (a) Eigenvalue spectrum, for the full 1000-member ensemble, at cycles 1, 3, 10 and 30, of the matrix C_e, and the associated eigenvalue spectrum at cycle 1 of C_o. (b) The corresponding Shannon entropy, for the same cycles, of C_e.

which may indicate that a reduction can occur at this point, as the change in entropy between two modes is very small, indicating that no further information is being added.

Figure 5 illustrates a successful application of our adaptive ensemble reduction algorithm. The four panels show the evolution of RMS height error, similarity index, ensemble size, and Shannon entropy, for the same experiments as in Figure 3.

In Figure 5(a), the RMS error of the reduced-ensemble runs is compared with that of the full-ensemble run. The Shannon-entropy criterion gives better results than the variance-percentage criterion, as the RMS is nearer to that of the full-ensemble run.

Figure 5(b) shows the similarity index between two successive cycles $k-1$ and k, for the full-ensemble run and the two adaptively-reduced runs. A similarity index of one would indicate equivalence between the current cycle and the previous cycle. When the adaptive algorithm based on the variance-percentage criterion is activated at a given cycle of the full-ensemble run, we observe an immediate dip in the graph, which recovers after three or four cycles. A deeper dip may be taken as an indication of less successful reduction. We do not observe this dip with the Shannon-entropy criterion.

The similarity index seems to be less sensitive in recovery than the RMS error. However, careful investigation shows the same convergence behaviour as for RMS error.

Figure 5(c) shows the evolution of ensemble size for the percentage-variance and Shannon-entropy criteria. Figure 5(d) shows the evolution of the Shannon-entropy spectrum for several experiments. The behaviour of the reduced ensemble is comparable to that of the full ensemble. This provides further insight into how the ensemble-reduction algorithm works in practice.

Using adaptive reduction for cases where the RMS height analysis error is small gives a factor-of-two economy in CPU time (corresponding roughly to a factor-of-two reduction in the number of ensemble members), compared to the full-ensemble run.

In all the experiments with both full-ensemble and reduced-ensemble runs, the number of observations is identical. There is no distance-dependent filtering in the MLEF.

4.2.2. *Adaptive algorithm*

The assimilation with the adaptive algorithm is performed over an interval of 45 hours, corresponding to 15 cycles of data assimilation, each representing a 3 h interval. The initial conditions for the experimental run are as defined above, but shifted to -3 h in order to create an unbalanced set of initial conditions. The background and observation standard deviation for the height is 5 m, and for the wind is 0.5 ms^{-1}.

Figure 6(a) shows the efficiency of the inflation strategy that has been adopted. The RMS error for the full 1000-member ensemble is plotted as a benchmark for comparison. There are no spin-up cycles. The RMS error of the full ensemble remains marginally lower than that of the adaptive ensemble. This is expected, since the algorithm is trying to compensate for the ensemble members that have been lost, and it appears to be quite efficient at this, with a much smaller start-up ensemble (400 members, in comparison to the full ensemble size of 1000). Figure 6(b) shows similar behaviour for RMS height background errors.

Figure 6(c) illustrates the adaptive algorithm, compared to the full-ensemble run, in terms of the ensemble size. It demonstrates successful adaptive reduction for the test case. Figure 6(d) shows the evolution of the Shannon-entropy spectrum for the adaptive experiments. It demonstrates that ensemble inflation introduces new directions in the subspace spanned by the ensemble.

The results described above are obtained using information from previous cycle perturbations, as defined in Section 3.3. Given the inflation strategy, it would be possible, for this example, to start with a much smaller ensemble, which could then increase to capture the correct covariances, and later decrease.

DOI: 10.1002/qj

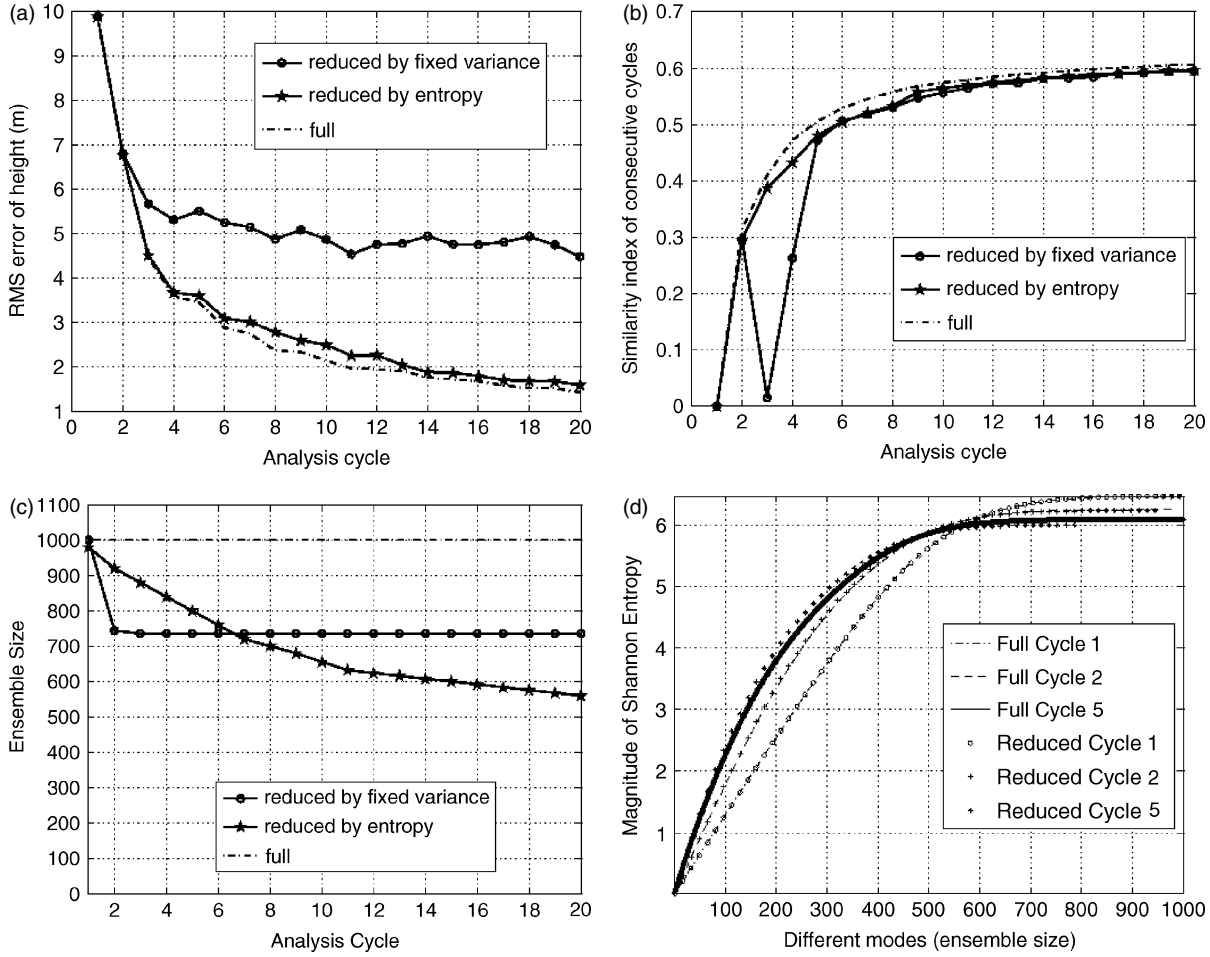

Figure 5. (a) Evolution of RMS height error, for cycles 1–20, for the full ensemble and for the variance-percentage and Shannon-entropy criteria. (b) The corresponding similarity indices. (c) The corresponding ensemble sizes. (d) Evolution of Shannon entropy for several full- and reduced-ensemble experiments.

Adaptive inflation of ensembles in later cycles might lead to economies in computational time, compared to the full-ensemble run. The results for test case 1 show that the inflation-and-reduction technique introduced in Section 3.3 exhibits adaptive behaviour.

4.3. Test case 2

The choices of parameters for the wave (Williamson *et al.*, 1992) are as follows. The wave number is $K = 4$, which is known to be the highest stable wave number in this model (Hoskins, 1973). We set the ω and R parameters to 7.848×10^{-6} s^{-1}, and the minimum height at the pole, h_0, to 8000 m. These parameters generate a fast tall wave, with eddies forming in the troughs by 120 h (Wlasak, 2002). The observation standard deviations for the height and wind are taken as 150 m and 3.0 ms^{-1} respectively. The background standard deviations for the height and wind are 50 m and 5.0 ms^{-1} respectively. The observations are assimilated every 6 h.

Two ensemble runs are chosen for test case 2. The results are displayed in Figure 7, which shows the RMS height error for ensemble sizes of 40 and 1000, each using a 6 h interval. The figure shows convergence. The ensemble sizes are 40 and 1000 for the last 10 cycles. This test case gives a low-dimensional ensemble subspace with correlation matrix $\mathbf{C}_o$ of dimensions 1025×1025, so the ensemble subspace can be much smaller than the observation space. This is the point that we wish to emphasize in this figure.

In a full three-dimensional model with flows that are geostrophically balanced, as in test case 1, we may require more ensembles initially to capture the important parts of the flow. For test case 2, we require fewer ensembles. This may imply (although more research is needed here) that there are only a few directions that dominate the ensemble space, so that it is wasteful to use a very large ensemble.

5. Summary and conclusions

A novel adaptive methodology for reducing and inflating an ensemble by projecting it onto a limited number of its leading EOFs has been successfully applied in the framework of the MLEF and ETKF for data assimilation. The methodology is general enough to be applied and

DOI: 10.1002/qj

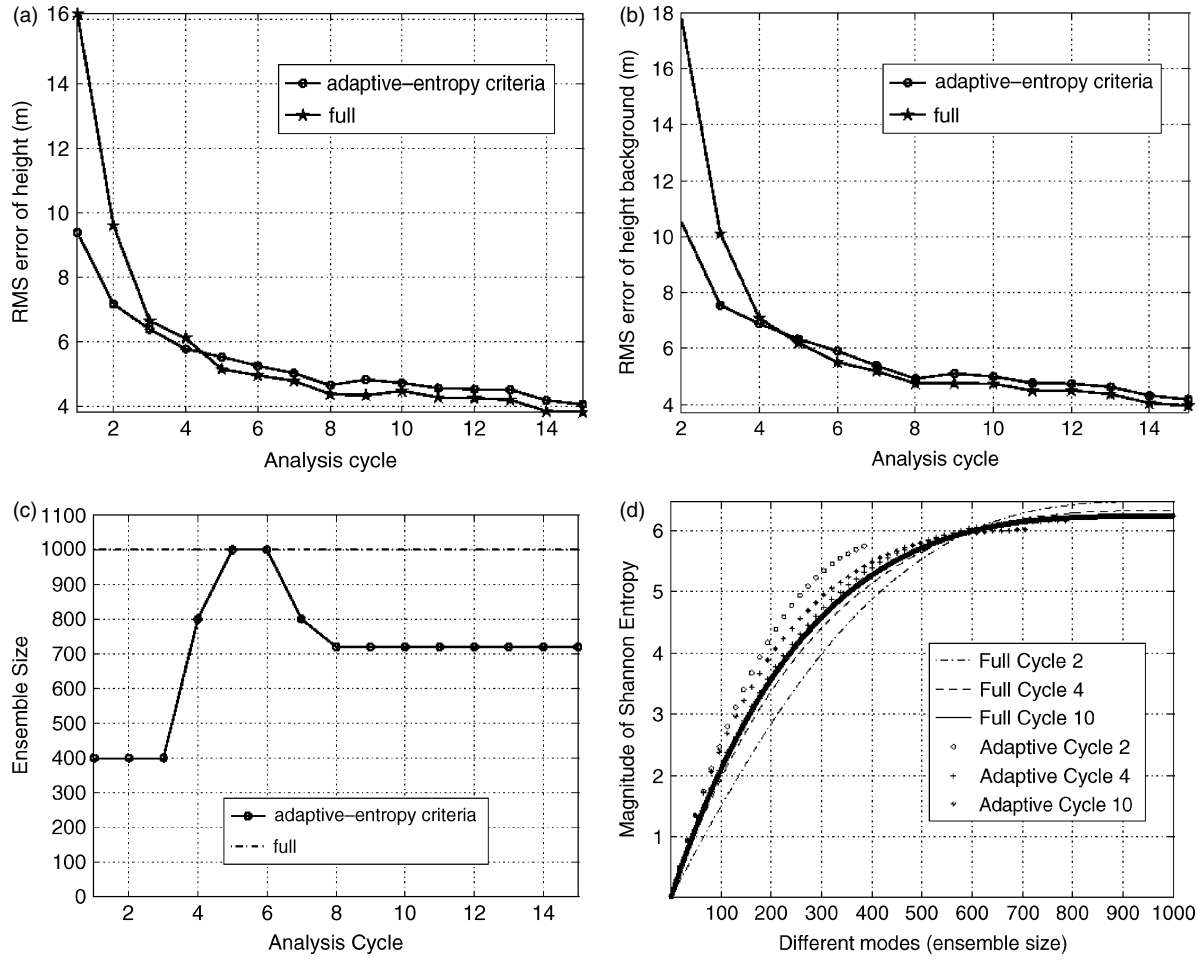

Figure 6. (a) Evolution of RMS height error, for cycles 1–15, for the full ensemble and for the adaptive entropy criterion. (b) The corresponding RMS height background error. (c) The corresponding ensemble size. (d) Evolution of the Shannon entropy.

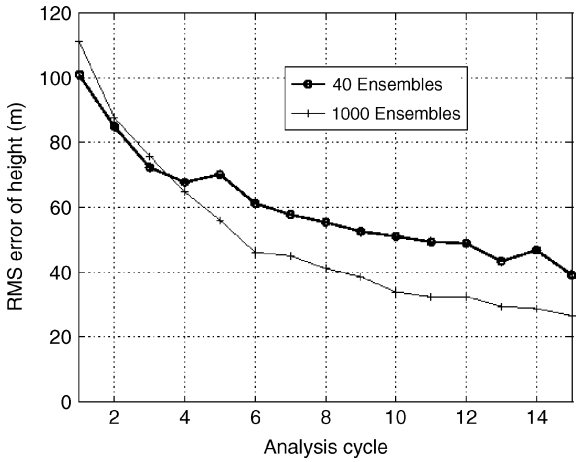

Figure 7. Evolution of RMS height error for Rossby–Haurwitz wave ensembles. Convergence is obtained with a much lower dimension of ensemble space.

extended to most ensemble-filter techniques, such as the EnKF and its various flavours.

Twin experiments have been run for a global shallow-water-equations model (Heikes and Randall, 1995a, 1995b), using test cases taken from the suite of (Williamson *et al*., 1992).

The idea used here originates in reduced-order modelling theory as well as in regularization theory (Hansen, 1998). It allows significant economies in the application of ensemble approaches to data assimilation, typically reducing the ensemble size required for successful implementation by a factor of up to two. The idea is extended to the case where inflation, rather then reduction, of the ensemble might be necessary. This idea is successfully tested using the same test cases.

Numerical experiments conducted with the MLEF in the twin-experiments framework show that the adaptive ensemble reduction and inflation perform well, compared to the full-ensemble run, in terms of the RMS error and similarity index. These two indicators reveal marginal effects on the reduced-ensemble results, compared to the full-ensemble run. This is further illustrated by considering the impact of adaptive reduction on geopotential and velocity fields, and by considering the spectrum of eigenvalues of the matrix $\mathbf{C}_e$ in the MLEF ensemble filter. Similar experiments have been conducted with inflation of the ensemble. These lead to a decrease in RMS error, thus validating the inflation approach.

The goal is not necessarily to reduce, or increase, the ensemble size, but rather to find an optimal utilization of the available ensembles given the computing limitations

DOI: 10.1002/qj

and the complexity of the problem. From this perspective, reduction is an important part of the picture. Once we know when and how to do that, we can come up with a strategy for enrichment (e.g. resampling), if this is deemed relevant to the modelling–observation system we are using.

Crommelin and Majda (2004) highlight some limitations of the EOF approach, particularly for systems that exhibit sudden transitions between different states, e.g. bursting behaviour. Our test cases, and short-period atmospheric circulations in general, do not exhibit such behaviour, but this limitation should be borne in mind, and other optimal bases (such as principal interaction patterns or bred vectors) should be considered in such cases.

Future research will be aimed at applying adaptive ensemble reduction and inflation to non-Gaussian ensembles, as well as extending the present methodology to typical EnKF data assimilation with realistic models.

Acknowledgements

The research of Dr Bahri Uzunoglu and Dr Michael Navon was sponsored by the National Science Foundation Collaboration in Mathematical Geosciences Grant ATM-0327818. That of Dr Milija Zupanski and Dr Steven Fletcher was supported by the National Science Foundation Collaboration in Mathematical Geosciences Grant ATM-0327651. Our gratitude is extended to the National Center for Atmospheric Research, which is sponsored by the National Science Foundation, for the computing time used in this research. The authors would also like to thank Prof. Gordon Erlebacher for his valuable discussions.

References

Antoulas CA. 2005. *Approximation of Large-Scale Dynamical Systems*. SIAM.

Berkooz G, Holmes P, Lumley JL. 1993. The generalized inverse of a nonlinear quasi-geostrophic ocean circulation model. *Ann. Rev. Fluid Mech.* **25**: 539–575.

Bierman G-J. 1977. *Factorization Methods for Discrete Sequential Estimation*. Academic Press.

Bishop CH, Etherton BJ, Majumdar SJ. 2001. Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Mon. Weather Rev.* **129**: 420–436.

Buizza R. 1994. Sensitivity of optimal unstable structures. *Q. J. R. Meteorol. Soc.* **120**: 429–451.

Crommelin DT, Majda AJ. 2004. Strategies for model reduction: Comparing different optimal bases. *J. Atmos. Sci.* **61**: 2206–2217.

Daley R. 1996. *Atmospheric Data Analysis*. Cambridge University Press.

Evensen G. 1994. Sequential data assimilation with a non-linear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.* **99**: 10143–10162.

Evensen G. 2003. The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynam.* **53**: 343–367.

Fletcher SJ. 2004. *Higher Order Balance Conditions Using Hamiltonian Dynamics for Numerical Weather Prediction*. PhD thesis, Department of Mathematics, University of Reading.

Fisher M, Andersson E. 2001. 'Developments in 4D-Var and Kalman filtering'. ECMWF Tech. Memo. 347.

Gaspari G, Cohn SE. 1999. Construction of correlation functions in two and three dimensions. *Q. J. R. Meteorol. Soc.* **125**: 723–757.

Glezer A, Kadioglu Z, Pearlstein A. 1989. Development of an extended proper orthogonal decomposition and its application to a time periodically forced plane mixing layer. *Phys. Fluids* **1**: 1363–1373.

Hamill TM, Snyder C. 2000. A hybrid ensemble Kalman filter 3d variational analysis system. *Mon. Weather Rev.* **128**: 2905–2919.

Hansen P-C. 1998. *Rank-Deficient and Discrete Ill-Posed Problems*. SIAM.

Heemink AW, Verlaan M, Segers JA. 2001. Variance reduced ensemble Kalman filtering. *Mon. Weather Rev.* **129**: 1718–1728.

Heikes R, Randall DA. 1995a. Numerical integration of the shallow-water equations on a twisted icosahedral grid. Part I: Basic design and results of tests. *Mon. Weather Rev.* **123**: 1862–1880.

Heikes R, Randall DA. 1995b. Numerical integration of the shallow-water equations on a twisted icosahedral grid. Part II: A detailed description of the grid and an analysis of numerical accuracy. *Mon. Weather Rev.* **123**: 1881–1887.

Hoskins BJ. 1973. Stability of the Rossby–Haurwitz wave. *Q. J. R. Meteorol. Soc.* **99**: 723–745.

Jazwinski AH. 1970. *Stochastic Processes and Filtering Theory*. Academic Press.

Kalman R, Bucy R. 1961. New results in linear prediction and filtering theory. *J. Basic Eng. – T. AMSE* **83D**: 95–108.

Lorenc AC. 1986. Analysis methods for numerical weather prediction. *Q. J. R. Meteorol. Soc.* **112**: 1177–1194.

Oczkowski M, Szunyogh I, Patil DJ. 2005. Mechanisms for the development of locally low-dimensional atmospheric dynamics. *J. Atmos. Sci.* **62–4**: 1135–1156.

Patil DJ, Hunt BR, Kalnay E, Yorke JA, Ott E. 2001. Local low-dimensionality of atmospheric dynamics. *Phys. Rev. Lett.* **86**: 5878–5881.

Pham DT, Verron J, Roubaud MC. 1998. A singular evolutive extended Kalman filter for data assimilation in oceanography. *J. Marine Syst.* **16**: 323–340.

Preisendorfer R. 1988. *Principal Component Analysis in Meteorology and Oceanography*. Elsevier.

Ringler TD, Randall DA. 2002. A potential enstrophy and energy conserving numerical scheme for solution of the shallow-water equations on a geodesic grid. *Mon. Weather Rev.* **130**: 1397–1410.

Rowley CW. 2005. Model reduction for fluids, using balanced proper orthogonal decomposition. *Int. J. Bifurcat. Chaos* **15**: 997–1103.

Sirovich L. 1987. Turbulence and the dynamics of coherent structures. Parts I–III. *Q. Appl. Math.* **45**(3): 561–590.

Tippett MK, Anderson JL, Bishop CH, Hamill TM, Whitaker JS. 2003. Ensemble square-root filters. *Mon. Weather Rev.* **131**: 1485–1490.

Uzunoglu B. 2001. *Reduced Modeling and Controller Design for Boundary Element Flow Around Cylinders*. PhD thesis, University of Southampton.

Uzunoglu B, Nair PE. 2001. 'Optimal flow control framework for adaptive variable fidelity-reduced modelling'. In: *Global Flow Instability and Control Symposium*, Crete, Greece, 11–13 June 2001.

Volkwein S. 2004. 'Interpretation of proper orthogonal decomposition as singular value decomposition and HJB-based feedback design'. In: *Proceedings of the Sixteenth International Symposium on Mathematical Theory of Networks and Systems (MTNS)*, Leuven, Belgium, 5–9 July 2004.

Watanabe S. 1965. 'Karhunen–Loeve expansion and factor analysis'. In: *Trans. 4th Prague Conf. on Inf. Theory, Statist. Decision Functions, and Random Proc.*, Prague, 1965: 635–660.

Williamson DL, Drake JB, Hack JJ, Jakob R, Swarztrauber PN. 1992. A standard test set for numerical approximations to the shallow water equations in spherical geometry. *J. Comput. Phys.* **102**: 211–224.

Wlasak MA. 2002. *The Examination of Balanced and Unbalanced Flow Using Potential Vorticity in Atmospheric Modelling*. PhD thesis, Department of Mathematics, University of Reading.

Zupanski M. 2005. Maximum likelihood ensemble filter: theoretical aspects. *Mon. Weather Rev.* **133**: 1710–1726.

Zupanski M, Fletcher SJ, Navon IM, Uzunoglu B, Heikes RP, Randall DA, Ringler TD, Daescu DN. 2006. Initiation of ensemble data assimilation. *Tellus* **58A**: 159–170.

DOI: 10.1002/qj