

Istvan Szunyogh *

Eric J. Kostelich,[†] Gyorgyi Gyarmati, Brian R. Hunt, Eugenia Kalnay, Edward Ott,
Dhanurjay Patil, and James A. Yorke
University of Maryland, College Park, MD 20742

1. Introduction

The time has come when ensemble-based Kalman filter data assimilation schemes can be considered for implementation on operational weather forecast systems in the foreseeable future. For the first time, an ensemble Kalman filter has been reported to break even with a sophisticated operational 3D-Var system (Houtekamer et al 2004), to outperform the NCEP 3D-Var in reconstructing the state of the mid-troposphere from surface pressure observations (Whitaker et al. 2004), and to be efficient in assimilating simulated and real Doppler-radar observations of convective systems (Snyder and Zhang 2003; Zhang et al 2004; Dowell et al 2004).

The present paper reports on the current status of the 4-dimensional Local Ensemble Kalman Filter data assimilation system (4D LEKF) developed by our interdisciplinary team at the University of Maryland. Our plan has been to develop a largely model independent analysis system through completion of the following tasks:

- Derivation of the basic LEKF algorithm, and extensive testing on the Lorenz-96 model (using from 40 to 120 dynamical variables). This step has been completed and the results are reported in Ott et al. (2002 and 2004).
- Extending the theory, and the low-order model experiments, to the case when asynchronous observations are assimilated at a fixed assimilation time. The resulting scheme, called 4D LEKF, and the results for low-order model experiments are described in Sauer et al. (2004) and Hunt et al. (2004).
- Testing the largely model independent computer code of the LEKF on an operational global numerical weather forecast model. For this purpose we selected the Global Forecast System (GFS) of the National Centers for Environmental Prediction (NCEP). While preliminary

results are shown in Szunyogh et al. (2004), the present paper summarizes the final results of this recently completed task.

- Implementation of the LEKF on the Regional Spectral Model (RSM) of NCEP. The RSM has been selected since it has the most consistent dynamics with the GFS among all regional models. In our formulation, the regional analysis is essentially a refinement of the global analysis prepared for the GFS. This implementation of the LEKF is currently under testing, and we hope to report on some preliminary results soon.
- Development and testing of methods to compensate for model errors via modification of the data assimilation technique.
- Assimilating real observations (including a large number of remotely sensed wind observations) with the 4D LEKF into both the GFS and RSM forecast systems. This step is in progress, and we hope to report on some results in the coming year.

The LEKF scheme is an ensemble square-root filter (Tippett et al. 2002): one first obtains an estimate of the most likely state of the atmosphere and an analysis error covariance matrix that describes the uncertainty in the best estimate. Then an ensemble of analyses is generated that is centered on the most likely state and is representative of the uncertainty reflected by the analysis error covariance matrix. A distinguishing feature of the LEKF is that it solves the Kalman filter equations *locally in model grid space*; other square-root filters solve the Kalman filter equation locally in observation space (Bishop et al. 2001; Anderson 2001, Whitaker and Hamill 2002). More precisely, the LEKF obtains the analysis at the different grid points independently, using all observations that are thought to improve the analysis at the individual grid points. In this scheme, the same observation may be used to obtain the analysis at multiple grid points. On the other hand, sequential schemes assimilate the observations one by one (or by small

*Corresponding author address: University of Maryland, IPST, CSS Building, College Park, MD 20742-2431, USA

[†]Current address: Arizona State University, Tempe, AZ 85287, USA

groups when the errors between the observations are correlated), iteratively updating the state estimate at those grid-points, where the accuracy of the analysis is thought to be positively affected by a given observation (or group of observations).

We speculate that solving the Kalman filter equations locally in grid space may be computationally advantageous and, at the same time, may not noticeably degrade the accuracy of the assimilation. The local analyses can be processed in parallel, involve relatively small matrices, and treat all data simultaneously. These features suggest the potential for the LEKF method in rapidly and efficiently assimilating large amounts of data (e.g., as will become available from future satellite observing systems.)

In what follows, we first provide a short summary of the LEKF algorithm (Section 2), then explain the implementation of the scheme on the NCEP GFS (Section 3). This implementation is tested under the *perfect model hypothesis*, i.e., by assuming that a model run provides a perfect representation of the true evolution of the atmosphere, making possible the generation of simulated observations (with known error statistics) and the exact computation of analysis and forecast errors (Section 4). Experiments are carried out for different size ensembles and for varying observational data coverage (Section 5). The accuracy of the analysis scheme is measured by the root-mean-square distance between the true states and the analyses. The computational efficiency is measured by the wall-clock time needed to complete the analysis. The results of this experiment indicate that the LEKF may become an operationally feasible scheme (Section 6).

2. Local ensemble Kalman filter

A detailed description and mathematical justification of the different components of the LEKF scheme can be found in Ott et al. (2002 and 2004). Here we provide only a brief algorithmic summary needed to understand the implementation of the scheme on the NCEP GFS. The version of the scheme that we describe assumes that the rank of the background and analysis covariance matrices is k when the ensemble has $k + 1$ members. Ott et al. (2004) describes a more general formulation that allows for a reduction of the rank. We note, that while we use one particular form of the matrix square root, other square roots can also be considered (Tippett et al 2002; Ott et al. 2004). Finally, we consider the case where all observations collected for the current analysis are taken at the same time.

[A simple technique to extend the scheme to the assimilation of asynchronous observations is presented in (Hunt et al. 2004).]

2.1 Global and local background vectors

A $k + 1$ -member ensemble ($k \geq 1$) of global background state vectors, $\mathbf{x}_g^{b(i)}$, $i = 1, 2, \dots, k + 1$, is obtained by integrating the forecast model started from a $k + 1$ -member ensemble of analysis fields created in the previous analysis cycle.

For each grid point p , we define a corresponding *local region* that consists of all grid points within a suitably prescribed neighborhood of p . Let $\mathbf{x}_l(m, n, o)$ be the d -dimensional local vector representing the model state within the local region centered at the grid point (m, n, o) . The construction of this local vector is a linear mapping \mathbf{L} of the vector \mathbf{x}_g that represents the state of the model in the space defined by the global three-dimensional grid space. Since all the analysis operations take place at a fixed time t and are repeated for all local regions, henceforth we suppress the dependence of all vectors and matrices on t, m, n , and o .

The local background error covariance matrix \mathbf{P}_l^b and the most probable local background state $\bar{\mathbf{x}}_l^b$ are derived from the $k + 1$ -member ensemble of global state field vectors $\mathbf{x}_g^{b(i)}$, $i = 1, 2, \dots, k + 1$. The most probable local state is estimated by

$$\bar{\mathbf{x}}_l^b = \mathbf{L} \left[(k + 1)^{-1} \sum_{i=1}^{k+1} \mathbf{x}_g^{b(i)} \right], \quad (1)$$

while the $d \times d$ local background error covariance matrix \mathbf{P}_l^b is estimated by

$$\mathbf{P}_l^b = k^{-1} \sum_{i=1}^{k+1} \delta \mathbf{x}_l^{b(i)} \left(\delta \mathbf{x}_l^{b(i)} \right)^T, \quad (2)$$

where the superscript T denotes transpose and

$$\delta \mathbf{x}_l^{b(i)} = \mathbf{L} \mathbf{x}_g^{b(i)} - \bar{\mathbf{x}}_l^b. \quad (3)$$

We can express \mathbf{P}_l^b in terms of the $d \times (k + 1)$ matrix,

$$\mathbf{X}_l^b = (k)^{-1/2} \left[\delta \mathbf{x}_l^{b(1)} \mid \delta \mathbf{x}_l^{b(2)} \mid \dots \mid \delta \mathbf{x}_l^{b(k+1)} \right], \quad (4)$$

as

$$\mathbf{P}_l^b = \mathbf{X}_l^b \mathbf{X}_l^{bT}. \quad (5)$$

2.2 Projection onto the k -dimensional analysis space

By using a $k + 1$ -member ensemble, we assume that an estimate of the background covariance matrix of rank k is sufficient to obtain accurate analyses. Experience accumulated by others (Houtekamer and

Mitchell 2000; Keppenne and Rienecker 2002) suggest that $k + 1$ may be as small as 40. For the purpose of subsequent computations, we consider the coordinate system of the k -dimensional space determined by the k orthonormal eigenvectors $\{\mathbf{u}^{(j)}\}$ of \mathbf{P}_l^b , which we use to form the *internal coordinate system* for the k -dimensional *local analysis space*. Since \mathbf{P}_l^b has rank k , it has k positive eigenvalues

$$\lambda^{(1)} \geq \lambda^{(2)} \geq \dots \geq \lambda^{(r)} \geq \dots \geq \lambda^{(k)} > 0. \quad (6)$$

Thus,

$$\mathbf{P}_l^b = \sum_{j=1}^k \lambda^{(j)} \mathbf{u}^{(j)} (\mathbf{u}^{(j)})^T. \quad (7)$$

Since the size of the ensemble ($k + 1$) is envisioned to be much less than the dimension d of \mathbf{x}_l^b , the computation of the basis vectors $\{\mathbf{u}^{(j)}\}$ is most efficiently done in the basis of the ensemble vectors. That is, we consider the eigenvalue problem for the $(k+1) \times (k+1)$ matrix $\mathbf{X}^{bT} \mathbf{X}^b$, whose nonzero eigenvalues are those of \mathbf{P}_l^b and whose corresponding eigenvectors left-multiplied by \mathbf{X}^b are the k eigenvectors $\mathbf{u}^{(j)}$ of \mathbf{P}_l^b .

We denote the projection of vectors into the k -dimensional space and the restriction of matrices to the same space by a superscribed circumflex (hat). The operator of this projection is

$$\mathbf{Q} = \{\mathbf{u}^{(1)} \mid \mathbf{u}^{(2)} \mid \dots \mid \mathbf{u}^{(k)}\}. \quad (8)$$

For instance, for the d -dimensional local background vector \mathbf{x}_l^b , the vector $\hat{\mathbf{x}}_l^b$ is a k -dimensional column vector given by

$$\hat{\mathbf{x}}_l^b = \mathbf{Q}^T \mathbf{x}_l^b. \quad (9)$$

Similarly, for a $d \times d$ matrix, such as the local background covariance matrix \mathbf{P}_l^b , the matrix $\hat{\mathbf{P}}_l^b$ is $k \times k$ and is given by

$$\hat{\mathbf{P}}_l^b = \mathbf{Q}^T \mathbf{P}_l^b \mathbf{Q}. \quad (10)$$

We also note that, in the internal coordinate system, $\hat{\mathbf{P}}_l^b$ is diagonal:

$$\hat{\mathbf{P}}_l^b = \text{diag} \left(\lambda^{(1)}, \lambda^{(2)}, \dots, \lambda^{(k)} \right), \quad (11)$$

and so is trivial to invert.

2.3 Local analysis

Let \mathbf{y}_l^o be the vector of current observations within the local region. We solve the Kalman filter equation in the local low-dimensional subspaces. Let

$$\Delta \hat{\mathbf{x}}_l^a = \mathbf{x}_l^a - \bar{\mathbf{x}}_l^b. \quad (12)$$

The most probable value of $\Delta \hat{\mathbf{x}}_l^a$ is

$$\Delta \hat{\mathbf{x}}_l^a = \hat{\mathbf{P}}_l^a \hat{\mathbf{H}}_l^T \mathbf{R}^{-1} \left(\mathbf{y}_l^o - \mathcal{H}(\bar{\mathbf{x}}_l^b) \right). \quad (13)$$

Here $\hat{\mathbf{H}}_l = \mathbf{H}_l \mathbf{Q}$ is the Jacobian matrix of partial derivatives of the observation operator \mathcal{H}_l (evaluated at $\bar{\mathbf{x}}_l^b$); thus, $\hat{\mathbf{H}}_l = \mathbf{H}_l \mathbf{Q}$ maps variables from the k -dimensional representation of the analysis analysis to the observation space. $\Delta \hat{\mathbf{x}}_l^a = \mathbf{Q}^T \Delta \mathbf{x}_l^a$ is the analysis increment in the k -dimensional analysis space, and $\hat{\mathbf{P}}_l^a$ is the analysis error covariance matrix in the same k -dimensional space. (If there are s scalar observations in the local region at the analysis time, $\bar{\mathbf{y}}_l^o$ is s -dimensional and the rectangular matrix \mathbf{H}_l is $s \times d$). In Eq. (13), $\hat{\mathbf{P}}_l^a$ is determined from the usual Kalman filter equations (e.g., Kalnay 2003), but restricted to the k -dimensional internal coordinate system, we have

$$\hat{\mathbf{P}}_l^a = \hat{\mathbf{P}}_l^b \left[\mathbf{I} + \hat{\mathbf{H}}_l^T \mathbf{R}_l^{-1} \hat{\mathbf{H}}_l \hat{\mathbf{P}}_l^b \right]^{-1}. \quad (14)$$

Finally, going back to the local space representation, we have

$$\bar{\mathbf{x}}_l^a = \mathbf{Q} \Delta \hat{\mathbf{x}}_l^a + \bar{\mathbf{x}}_l^b. \quad (15)$$

2.4 Ensemble of local analyses

The ensemble of local analysis fields $\{\mathbf{x}_l^{a(i)}\}$, $i = 1, 2, \dots, k + 1$ is obtained by first finding the $(k + 1)$ local analysis perturbations $\delta \mathbf{x}_l^{a(i)}$,

$$\delta \mathbf{x}_l^{a(i)} = \mathbf{Q} \delta \hat{\mathbf{x}}_l^{a(i)}, \quad (16)$$

then forming the local analysis ensemble

$$\mathbf{x}_l^{a(i)} = \bar{\mathbf{x}}_l^a + \delta \mathbf{x}_l^{a(i)}. \quad (17)$$

The local analysis perturbations $\delta \hat{\mathbf{x}}_l^{a(i)}$ are a linear combination of the local background perturbations in the k -dimensional analysis space:

$$\hat{\mathbf{X}}_l^a = \hat{\mathbf{X}}_l^b \mathbf{Y}_l, \quad (18)$$

where

$$\hat{\mathbf{X}}_l^{a,b} = k^{-1/2} \left\{ \delta \hat{\mathbf{x}}_l^{a,b(1)} \mid \delta \hat{\mathbf{x}}_l^{a,b(2)} \mid \dots \mid \delta \hat{\mathbf{x}}_l^{a,b(k+1)} \right\} \quad (19)$$

and

$$\mathbf{Y}_l = \left[\mathbf{I} + \hat{\mathbf{X}}_l^{bT} \left(\hat{\mathbf{P}}_l^b \right)^{-1} \left(\hat{\mathbf{P}}_l^a - \hat{\mathbf{P}}_l^b \right) \left(\hat{\mathbf{P}}_l^b \right)^{-1} \hat{\mathbf{X}}_l^b \right]^{1/2}. \quad (20)$$

This construction of the local analysis perturbations has the desirable properties that it does not distort the mean of the analysis ensemble, it correctly represents the analysis uncertainty, and it preserves the smoothness of the background ensemble fields as closely as possible (see Ott et al. 2004 for details).

2.5 Ensemble of global analyses

The components of the most probable global analysis field \bar{x}_g^a at the grid point (m, n, o) are obtained by first selecting the local analysis vector $\bar{x}_l^a(m, n, o)$ associated with the local region centered at (m, n, o) , then copying the components of $\bar{x}_l^a(m, n, o)$ at its central grid point. The same strategy is used to obtain the members of the global analysis ensemble $\{x_g^{a(i)}\}$, $i = 1, 2, \dots, k + 1$.

3. Implementation on the NCEP GFS

3.1 The forecast model

The NCEP GFS is a spectral model, which means that the model state variables are coefficients of a spherical harmonic expansion on the globe. We use a version of the NCEP GFS that was in operational use at the beginning of 2001. In this version the model variables are spectral coefficients of the two-dimensional vorticity and divergence, virtual temperature, logarithm of the surface pressure, specific humidity, and ozone mixing ratio. The only difference between our version and the operational one is in the resolution, which we have reduced to T62 in the horizontal direction and to 28 level in the vertical direction. This resolution is well tested insofar as many operational forecast products of NCEP have been obtained at this resolution for more than a decade. This resolution allows for a large number of experiments with the computational resources available for us.

3.2 Definition of the local regions

The nonlinear and physical parameterization terms in the T62 resolution NCEP GFS are computed on a 192×94 Gaussian longitude-latitude grid. We utilize this grid to implement the LEKF, which is formulated in model grid space. This means that our three-dimensional global grid has $194 \times 92 \times 28$ points. Since all variables are defined on the global grid, except for the logarithm of the surface pressure, the total number of grid point variables is 2 544 768. (The number of model variables—1 137 024 spectral coefficients—is about 44% of the number of grid point variables.)

In our experiments, the observed variables are the two horizontal components of the wind, the virtual temperature, and the surface pressure. We use this information to obtain an analysis of all grid point variables. Then the spectral transform is applied to the grid point variables of different types to obtain an analysis in spectral space at each level. Finally,

the spectral coefficients of the two-dimensional vorticity and divergence are computed from the spectral coefficients of the two horizontal components of the wind vector. (We note that the spectral transform from grid to spectral space has a smoothing effect, which means that the final analysis used as model initial condition is smoother than the original analysis prepared in grid point space.)

Our implementation of the LEKF scheme involves both horizontal and vertical localizations of the $192 \times 94 \times 28$ model grid. The horizontal localization is done on the 192×94 longitude-latitude grid. To avoid introducing artificial singularities, extra latitudes are also added at the poles to make possible the proper definition of the local regions on the entire globe. Visually, one may think of our approach as defining the local regions on a polar stereographic map projection near the poles. In all experiments reported here, the number of grid points in the zonal and meridional directions are equal, while the vertical localization is done by centering narrow layers around all 28 model levels.

3.3 Algorithmic complexity

An algorithm is $O(f(n))$ if, given input data of length n , the required number of machine instructions is bounded by $Cf(n)$ for some constant C when n is sufficiently large. This so-called *order notation* describes the computational complexity of a given algorithm and provides a rough measure of the computing time on a single processor. For instance, multiplying an n -vector by a scalar is an $O(n)$ procedure. The product of an $n \times k$ matrix with a k -vector is $O(nk)$, insofar as one must form a linear combination of k vectors, each of length n ; the classical algorithm for computing the product of an $n \times k$ matrix with a $k \times m$ matrix is $O(nkm)$.

Table 1 summarizes the computational complexity of each step in the LEKF algorithm for a single local region containing n observations together with d dynamical variables in each of k ensemble solutions. In typical applications, d is expected to be one to two orders of magnitude larger than k . Therefore, the most expensive steps in the linear algebra are the change of coordinates between the model space and the local coordinates (i.e., forming $\hat{\mathbf{X}}^b$ and \mathbf{X}^a), the computation of the nonzero eigenvalues and associated eigenvectors of \mathbf{P}_l^b , and the computation of \mathbf{Y}_l . In the idealized experiments described here, the observation operator \mathcal{H} is trivial, multiplication by $\hat{\mathbf{H}}_l$ is simply a gather-scatter operation, and each covariance matrix \mathbf{R}_l of local measurement errors is constant and diagonal. In practice, of course, \mathcal{H} and $\hat{\mathbf{H}}_l$ are more complex,

Eq. (1)	$O(kd)$
Eq. (4)	$O(kd)$
$\lambda^{(j)}, \mathbf{u}^{(j)}$	varies; $O(k^3)$
$\hat{\mathbf{X}}^b$	$O(k^2 d)$
$\hat{\mathbf{P}}^a$	up to $O(k^3)$
$\mathbf{y}_i^o - \mathcal{H}(\bar{\mathbf{x}}_i^b)$	varies; $O(n)$ here
Eq. (13)	varies; $O(kn)$ here
Eq. (20)	varies; $O(k^3)$
\mathbf{X}^a	$O(k^2 d)$

Table 1: The computational complexity of the main steps in the LEKF algorithm for each local region: k denotes the size of the ensemble; d , the dimension of the local atmospheric vector; n , the number of observations in the local region.

depending on the nature of the data. For conventional observations, such as radiosonde data, \mathcal{H} may involve only simple linear interpolation between model grid points, \mathbf{R}_i may be banded or even diagonal, and the evaluation of Eq. (13) may still be $O(kn)$. Remotely sensed data, such as satellite measurements, may involve the evaluation of a nonlinear observation operator and, hence, Eq. (13) may be more expensive to compute. However, in either case, if \mathbf{R}_i is approximately constant, then \mathbf{R}_i^{-1} may be precomputed, thus reducing the cost of evaluating Eq. (13).

Similar comments apply to the evaluation of the analysis covariance matrix in Eq. (14). The computational complexity may vary from $O(k)$, as in the experiments described here, to $O(k^3)$, depending on the structure of \mathbf{R}_i and $\hat{\mathbf{H}}_i$.

Our initial implementation of the LEKF is in Fortran 95, which provides a simple, portable, and efficient notation for handling dense matrices. We have used version 3 of the LAPACK library (Anderson et al. 1999) to compute matrix inverses, eigenvalues, and eigenvectors, because it is numerically robust, thoroughly tested, and widely available; most computer vendors provide optimized implementations of the Basic Linear Algebra Subroutines (BLAS) (Lawson et al. 1979; Dongarra 1988 and 1990) upon which the LAPACK library is built. The LAPACK routine DSYEVR implements the algorithm of choice for finding all eigenvalues and eigenvectors of a symmetric $k \times k$ matrix. [However, for maximum efficiency, DSYEVR requires IEEE-754 (IEEE, 1985) infinity arithmetic to be implemented without trapping (Anderson et al. 1999, p. 146)]. This requirement can be problematic, depending on the processor and Fortran compiler, although we have had no difficulty with version 6.1 of the Lahey Fortran compiler

(Lahey Computer Systems, www.lahey.com) on Intel Pentium Xeon processors running Red Hat Linux (Red Hat, Inc. www.redhat.com). Because the underlying algorithm is iterative, the operation count associated with DSYEVR is not fixed: the convergence rate depends on the data, but in general is expected to be $O(k^3)$.

The overall efficiency of the LEKF algorithm also is influenced by the quality of the Fortran intrinsic function MATMUL, which we use heavily to multiply matrices. The level-3 BLAS routines DGEMM and DGEMV (Dongarra 1990) may be substituted to yield faster code, depending on the Fortran implementation, though with some loss of clarity and simplicity.

The total time required to complete the LEKF is proportional to the total number of local regions. However, because the assimilation is performed on each region independently, the algorithm is amenable to efficient implementation on parallel computer architectures, which substantially reduces the wall-clock time.

4. Experimental design

We assume that the NCEP GFS provides a perfect representation of the true atmosphere, an approach frequently called the *perfect model hypothesis*. Under this assumption, forecast errors arise and grow exclusively due to uncertainties in the initial conditions and the sensitivity of the model solutions to these uncertainties. In other words, the model is a chaotic system, in the sense that uncertainties in the initial conditions are more frequently amplified than damped during the forecast phase of the analysis cycle. The role of the data assimilation system, on the one hand, is to use the information contained in the observations to remove the growing component of the errors from the background. On the other hand, the data assimilation must use the information contained in the background to filter the observational noise and to spread the information to unobserved locations. We have designed a series of experiments that measure the efficiency of the LEKF in all of these three areas (removing growing errors, reducing noise, and spreading information).

4.1 Observations

First, a time series of “true” states, $\mathbf{x}^t(t)$, was generated by a 60-day integration of the T62 GFS model, started from the operational NCEP analysis at 0000 UTC on 1 January 2000. Then, simulated observations were prepared at each grid point by adding zero-mean Gaussian random noise (simulated observational error) to the true states ev-

ery six hours (at 0000 UTC, 0600 UTC, 1200 UTC, 1800 UTC). The standard deviation of the assumed observational errors is 1 K, 1.1 m/s, and 1 hPa for the virtual temperature, horizontal wind components, and surface pressure, respectively. The humidity and ozone variables and the physical parameters describing the conditions of the underlying surface (e.g., sea surface temperature, albedo, snow and ice coverage, soil type, etc.) are not observed. In most of the experiments, to simulate reduced observational networks, only subsets of the observations are assimilated. These subsets have been created in a systematic manner, gradually removing observational locations to obtain sparser observational data sets.

4.2 Diagnostics

The accuracy of an analysis field at a given time is measured by the root-mean-square (rms) distance between the analyzed and “true” meteorological fields. The rms error at a given model level is computed by taking the mean over all horizontal grid points in the verification area. Time-mean results are then obtained by averaging the rms values over time. When horizontal distributions of time-mean errors are shown, the fields are obtained by averaging the absolute value of the error at each grid point. (Following the conventions of numerical weather prediction, the rms is never taken over both space and time.)

5. Numerical experiments

Our strategy to validate and tune our implementation of the LEKF on the NCEP GFS is based on first designing a *base experiment*, then exploring changes in the behavior of the data assimilation system under gradual changes to selected parameters of the scheme. In the base experiment, the surface pressure, horizontal wind components, and virtual temperature are observed at 2000 randomly selected geographical locations, and the parameters of the LEKF are the following: the number of ensemble members is 40, the horizontal size of the local region is 7×7 , the depth of the local layers varies with altitude (see Fig. 1). The scheme also applies a uniform 4% multiplicative inflation (Anderson and Anderson 1999) to the background ensemble to compensate for the variance lost to nonlinearities and the limited sample size. According to the results of the numerical experiments described below, the base configuration is a reasonable, though not optimal, configuration of the LEKF. As the results show, this configuration is a good starting point

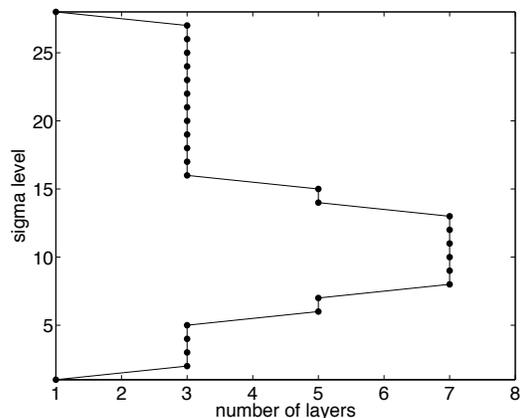


FIG. 1: Number of model levels in the local regions (x axis) centered at the different levels (y axis)

to explore the sensitivity of the scheme to changes in the various assimilation parameters.

5.1 Temporal evolution of errors

With the rms analysis error rapidly tends to a level that is much smaller than the rms error of the observations. We have found this to be a robust property: it is characteristic of all experiments reported in this paper. While the speed of convergence is rapid, it is slightly different for the different variables: fastest for the temperature and slowest for the wind components. Fig. 2 shows an example of the time evolution of the rms error for the surface pressure of the base experiment. In this example, the rms error tends to a level that is about 40% of the observational error within a couple of days. While some slow temporal fluctuations of the error can be observed, the efficient filtering of the observational noise by the data assimilation scheme is evident. Since observations are taken at only 2000 locations (or about 11% of all the grid points), the results also indicate that the LEKF efficiently propagates information to the unobserved locations.

Since the temporal fluctuation of the errors is modest, the different configurations of the LEKF can be meaningfully compared by the examination of time-mean results. To ensure that the time means are not affected by the initial transient, the time averages are computed for the last 45 days (the last 180 analysis cycles) of each experiment, i.e., the first 15 days (60 analysis cycles) are a “spin-up” period whose results are ignored.

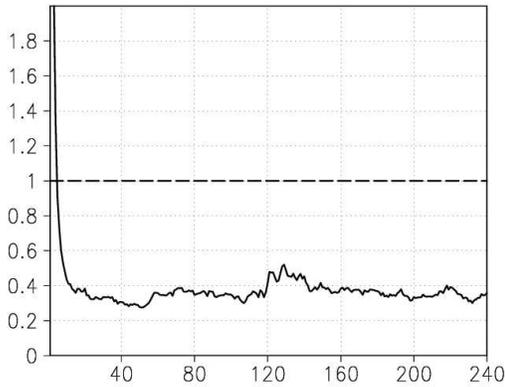


FIG. 2: Time evolution of the surface pressure rms error for the base configuration (solid line). The rms error of the observations is shown by dashes.

5.2 Spatial distribution of errors

The geographical distribution of the analysis errors is strongly zonal (Figs. 3 and 4); the largest errors are in the tropics and over the polar regions; the errors in the mid-latitudes are the smallest. The time-mean error in the temperature analysis is small compared to the rms observational errors (1 K), even at the locations of the largest errors. The errors for the wind are the largest in the same regions as for the temperature. Despite this similarity, there is an important difference between the errors for the two variables: compared to the observational errors, the relative errors are clearly larger for the wind (Figs. 3 and 4). (We recall that the observational wind error is 1.1 m/s .)

Fig. 4 indicates that the analysis of the wind is the most difficult in the region of ascending motions in the Hadley cells. The error quickly decreases toward the poles in the neighboring regions of descending motions. A comparison of Fig. 4 and Fig. 14.3 of Emanuel (1994) suggests that the errors are the largest in the layers, where the convective available potential energy (CAPE) is the largest. A picture emerges in which the largest errors in the tropical wind analysis are associated with deep convective processes. This conclusion was also well supported by an inspection of the geographical (latitude-longitude) distribution of the errors (not shown). This revealed that the locations of the largest errors are sandwiched between the regions of easterly trade winds to the north and the south. These regions are associated with pools of

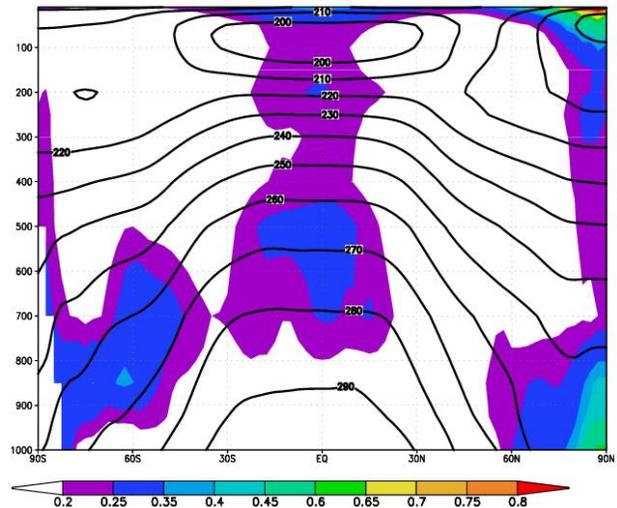


FIG. 3: Longitude-averaged time-mean rms error in the temperature analysis for the base configuration (shades). The time-mean of the “true” temperature is also shown (contours).

warm air located in the southwest Pacific/northeast Indian region and off the northeast coast of South America. The sources of CAPE are air parcels lifted from the surface. Because the analysis is very accurate near the surface, the temperature flux associated with the ascending warm air parcels may be relatively well analyzed in the subcloud layer. This indicates that the relatively poor analysis in the deep convective clouds originates from a poor analysis of processes within the clouds themselves. Also, because the wind analyses are of poorer quality than the temperature analysis, it may be considerably more difficult to analyze the momentum fluxes in deep convective clouds than the associated temperature fluxes. These difficulties, encountered in the regions of deep convection, cannot be explained by the inadequate parameterization of convection; in our experiments, the true state has also been generated by the model.

The version of the NCEP GFS that we use employs a modified Arakawa-Schubert scheme (Pan and Wu 1995) for the parameterization of deep convection. On the one hand, the scheme is simpler than the original Arakawa-Schubert scheme (Arakawa and Schubert 1974); it assumes that the deep convection is associated with one type of cloud (the deepest cloud) instead of a spectrum of clouds. On the other hand, the scheme corrects one important deficiency of the original scheme: it allows for the transport of momentum by down-

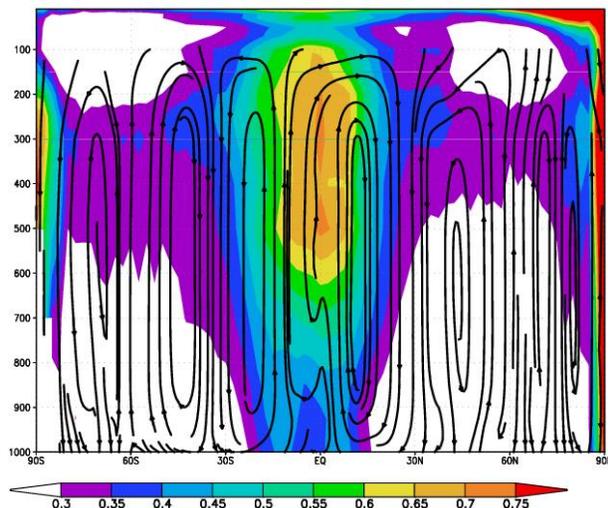


FIG. 4: Longitude-averaged time-mean rms error in the analysis of the zonal component of the wind for the base configuration (shades). The time-mean of the “true” circulation in the latitude-height plane is shown by streamlines (contours with small arrowheads showing the direction of the flow).

drafts. Otherwise, the scheme retains the central hypothesis of the Arakawa-Schubert scheme: convection is essentially a rapid-response mechanism to neutralize the destabilizing effects of such large-scale processes as surface fluxes and radiation. More precisely, the scheme assumes that the consumption of CAPE, by an ever present ensemble of deep convective clouds, is in a statistical equilibrium with the CAPE generated by the large scale processes.

One may ask whether our results have any significance in the case where real observations, collected in real deep convective clouds, are assimilated. As Emanuel (1994) pointed out, the observational evidence to support the central hypothesis of Arakawa and Schubert (1974) is striking. He also pointed out, however, that the *entraining plume model*, on which the Arakawa-Schubert scheme is based, is a poor representation of the individual clouds. Thus, in the regions of deep convection, we can expect large model errors to occur, which makes a good analysis even more difficult to obtain than we have found here.

This example shows the importance of carrying out experiments under the perfect model hypothesis. Had we started assimilating real observations first, we might attributed all difficulties to the inadequate parameterization of deep convection in the

model. Our results show that preparing the analysis for the regions of deep convection can be challenging even in the absence of model errors.

We later show that in the perfect model set-up, the analysis error can be efficiently reduced in the regions of deep convection by increasing the ensemble size, reducing the depth of the local layers, increasing the number of observations, and using information about the humidity. We also note that since 2001, NCEP has made several important upgrades to the parameterization of deep convection in the operational GFS. We would not be surprised if an implementation of the LEKF on the current operational model behaved somewhat differently than that reported here.

Another important feature is the good quality of the analysis in the mid-latitude regions. Fig. 3 shows that, in the lower troposphere (below the 50-kPa level), the errors are practically negligible. This is most remarkable, since these are the regions where baroclinic instabilities—the most energetic instabilities in the Earth’s atmosphere—convert the available potential energy to kinetic energy. (The main zones of baroclinic energy conversion can be recognized in Fig. 4 by the on-average upward motions that they generate in the mid-latitudes.) This finding shows the efficiency of the Kalman filter in correcting fast-growing background errors associated with processes initiated by baroclinic instabilities. The largest, but still modest, analysis errors in the mid-latitudes occur in and right below the jet layer. These error patterns are directly connected to the patterns of large errors in the tropics. We have prepared animations of the error propagation which clearly show that errors in the regions of deep convection frequently propagate to the mid-latitudes. Thus, we conjecture that errors in the mid-latitude upper troposphere can be reduced by improving the analysis of the deep convection in the tropics.

Finally, we note the narrow regions of relatively large errors at the top of the atmosphere and over the polar regions. While the origin of these errors is not clear to us, we suspect that minor flaws in our implementation of the LEKF may play some role. For instance, we hope that using a reduced grid near the poles would take into account the decreasing distance between grid points and would help to reduce the errors. Similarly, we may find a better way to prepare the analysis near the upper boundary of the model atmosphere. While we expect that refinements of our implementation would somewhat reduce the errors in these regions, we do not expect them to completely eliminate the elevated errors. The polar regions are similar to the regions of deep

convection in that physical parameterizations play an important role, with the important difference that very strong downdrafts characterize the motions instead of the updrafts in deep convection. Also, the artificial upper boundary condition (the real atmosphere has no upper boundary) is known to generate strong artificial instabilities (Kalnay and Toth 1996; Hartman et al. 1997) that may not be efficiently corrected by the data assimilation. Finally, it is important to point out that these errors are still smaller than the observational errors (except for the wind error at the top of the atmosphere), and we have not found indications that these errors propagate deep into other regions (mid-latitudes and lower layers of the atmosphere).

In what follows, we present a series of experiments. In each of these experiments, a parameter is selected and gradually altered, starting from the value used in the base experiment. Since we have found that the major patterns in the error distribution change only a little in the tested parameter ranges, we describe the effects of changing the selected parameter by showing average errors over large regions. Errors in the wind analysis are computed separately for the tropics (30° S – 30° N) and for the two extratropical regions (30° N – 90° N and 30° S – 90° S), but due to the strong similarities between the hemispheres, results are not shown for the extratropics in the Southern Hemisphere. Also, only global averages are shown for the temperature, since the relatively small differences between the tropics and the extratropics have made a separation of the two regions unnecessary.

5.3 Sensitivity to the size of the local regions

The size of the local regions is a crucial parameter of the LEKF scheme. On the one hand, the local region must be sufficiently large to allow nearby observations, which contain useful information, to affect the analysis at the center of the region. On the other hand, the region must be sufficiently small to permit an efficient filtering of statistical fluctuations, which occur due to the relatively small ensemble size. (The computational cost increases as much as cubically with the number of ensemble members.) The best way to find a nearly optimal region size is through numerical experimentation.

First we vary the horizontal size of the local regions, using 3×3 , 5×5 , 7×7 , 9×9 , and 11×11 patches of grid points. The dimension d of the local vectors for the different-sized regions is shown in Fig. 5. The horizontal size of the local regions has no significant effect on the speed of the convergence, and the 15 days that we allow for the tran-

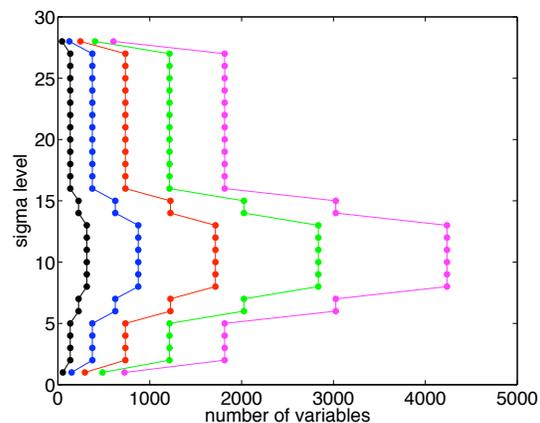


FIG. 5: Number of model variables in the local regions (x axis) centered at the different model levels (y axis). The number of variables is shown for local regions consisting of 3×3 (black), 5×5 (blue), 7×7 (base experiment, red), 9×9 (green) and 11×11 (purple) grid points at each level.

sients to settle is plenty in all cases (results are not shown). While the rms error of the surface pressure analysis is essentially the same for all experiments (including the base experiment shown in Fig. 2) once the transient dies out, the patch size has a noticeable impact on the accuracy of the temperature and wind analyses. The vertical distributions of the time-mean rms error differ with region size for the temperature over the globe (Fig. 6) and the zonal component of the wind vector in the Northern Hemisphere extratropics (not shown) and the tropics (Fig. 7), albeit not dramatically. The results are the best for the 5×5 horizontal regions, providing uniformly good results over the different variables and regions. In the extra-tropics, the 7×7 local regions provide good performance similar to the 5×5 local regions, while in the tropics, the 5×5 local regions have a clear advantage over the 7×7 and the similarly well-performing 3×3 regions. Increasing the size of the horizontal regions to 9×9 , and especially to 11×11 , clearly degrades the performance of the scheme.

In summary, we conclude that the rms error of the analyses is about 30% of the observational errors, when the ensemble has $k = 40$ members and the patch-size is selected from a reasonable range (from 3×3 to 7×7 horizontal grid points at the number of vertical levels shown in Fig. 1).

It is reasonable to assume that when the ensemble size is increased, detection of the (typically

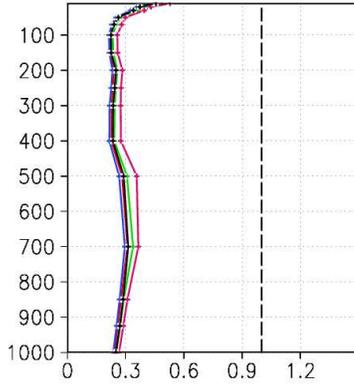


FIG. 6: The time mean of the rms error of the temperature analysis over the globe (x axis) as a function of height, using pressure as the vertical coordinate (y axis). The color scheme is the same as in Fig. 5. The order of the errors, from small to begin, is 5×5 (blue), 7×7 (base experiment, red), 3×3 (black), 9×9 (green) and 11×11 (purple).

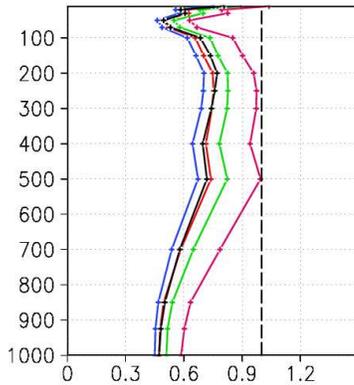


FIG. 7: The time mean of the rms error of the zonal wind component analysis in the tropics (x axis) as a function of height, using pressure as the vertical coordinate (y axis). The color scheme is the same as in Fig. 5. The rms error of the observations is shown by dashes.

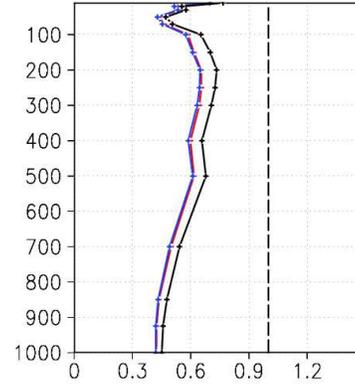


FIG. 8: Same as Fig. 7, but for an 80-member ensemble. [The results are virtually the same for the 5×5 (blue) and 7×7 (base experiment, red) regions and are not shown for the 9×9 (green) and 11×11 (purple) horizontal regions.]

smaller) longer-distance correlations in the background uncertainties becomes less sensitive to statistical fluctuations. In other words, we can expect that, for a larger ensemble, the optimal region size also increases and the better utilization of more observed information leads to a reduction of the analysis errors. To assess the effect of ensemble size on the optimal region size and the analysis errors, experiments with $k = 80$ ensemble members have also been carried out. We find that the increased ensemble size has the largest positive effect on the configuration using 7×7 horizontal grid points. This configuration now breaks even with the configuration using 5×5 horizontal grid points in all regions and for all variables. This occurs as the performance of the latter configuration is only slightly affected by the increased ensemble size. These results indicate that a 40-member ensemble suffices to resolve the analysis uncertainty in the 5×5 local horizontal regions. Since the overall improvement from increase of the ensemble size from 40 to 80 is very modest for the extratropics, the results are illustrated only for the tropics (Fig. 8).

5.4 Sensitivity to the density of observations

Earlier experiments with simple models show that ensemble Kalman filters have a growing advantage over 3D-Var-type schemes, which use static estimates of the background covariance matrix, as the density of the observations decreases (e.g., Hamill

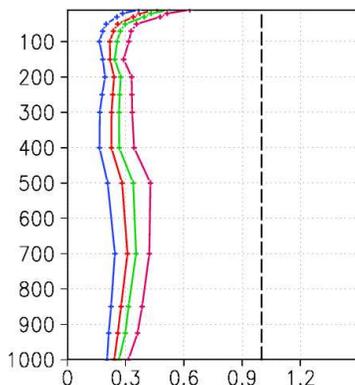


FIG. 9: The time mean of the rms error of the temperature analysis over the globe (x axis) as function of height using pressure as vertical coordinate (y axis). The analysis error is shown for observing networks consisting of 18 048 (blue), 2000 (base experiment, red), 1000 (green) and 500 (purple) observational locations. This corresponds to observing all grid points (blue), and 11% (red), 5.5% (green), and 2% (purple) of all grid points. The rms error of the observations is shown by dashes.

et al. 2001; Ott et al 2004). More precisely, while the performance of the 3D-Var-type schemes degrades dramatically with decreasing observational density, the accuracy of the Kalman filter is only slightly affected until the density of observations reaches a critically-low value. As the density of observations decreases, the correct estimation of the flow-dependent covariance becomes ever more reliant on locations to which information has been propagated from elsewhere. To test whether our implementation of the LEKF on the NCEP GFS retains this important property, further experiments with differing number of observations were also carried out. The results are compared for four different observational networks: observing all, 2000, 1000, and 500 locations around the globe. While the observational density has only a modest impact on the quality of the temperature analysis (Fig. 9), the accuracy of the wind analysis, especially in the tropics, degrades more dramatically as the density of the observations is reduced (Fig. 10). Most importantly, Fig. 10 shows that increasing the number of observations is a very efficient way to reduce the analysis errors in the tropics (in regions of deep convection).

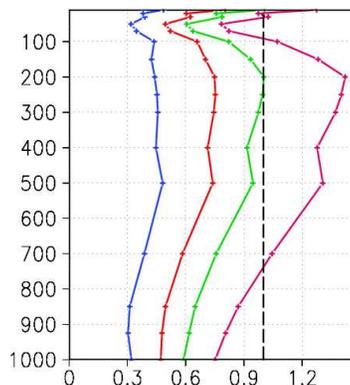


FIG. 10: The time mean of the rms error of the zonal wind component analysis in the tropics (x axis) as a function of height using pressure as the vertical coordinate (y axis). The color scheme is the same as in Fig. 5.4.

5.5 Timing results

The timing results described here are from our initial implementation of the LEKF, which is on a relatively modest Beowulf cluster consisting of 25 dual-processor nodes, each with 2 GB of random access memory and connected by a 1-gigabit Ethernet; each processor is a 2.8-GHz Intel Pentium Xeon with hyperthreading disabled. Most runs use 40 processors on the cluster.

Typically, one complete cycle of the algorithm takes 15 minutes of wall-clock time for an ensemble of 40 solutions when observations are available at every model grid point using the GFS at T62/L28 resolution and vertical localization, as described below. This computation involves $192 \times 94 \times 28 = 505\,344$ local regions and slightly more than 1.5 million observations; each local region is a cube of $7 \times 7 \times \nu$ model grid points, where $\nu = 1, 3, 5$ or 7 , depending on altitude. The time includes that spent computing the transforms from spectral space to physical space and back (which could be parallelized but which we have implemented only on a single processor); the i/o and network overhead to transport the appropriate model grid and observation data to each processor; one step of the LEKF algorithm to each local region; and 40 six-hour forecasts from the resulting analysis. Excluding the spectral transforms and forecasts, the wall-clock time is about 505 seconds for the above parameters. The average time needed to process one local region, where the ensemble size is 40

parameters $k + 1, N$	wall-clock time (min)	LEKF step only (sec)	patch mean (10^{-3} sec)
40, all	15	505	31
40, 2000	14	447	28
80, all	45	1973	122
80, 2000	42	1682	104

Table 2: Time needed to run the LEKF algorithm.

and the patch size is 490, is about 31 milliseconds. Table 2 shows timing results for local regions that consist of $7 \times 7 \times v$ cubes ($v = 1, 3$) as described above, with either 40 or 80 ensemble members. (For this calculations the $v=5,7$ layers were replaced by layers of $v=3$.) The observing network consists of observations of temperature and wind speed at all 28 vertical levels, plus the surface pressure, at each of N points. The notation $N = \text{all}$ refers to the case where observations are available at each model grid point (1.5 million observations in total); $N = 2000$ refers to an observing network that consists of 2000 randomly-chosen longitude-latitude model grid points at which the observations are assumed to exist at each vertical level (i.e., a total of 170 000 observations). The wall-clock time includes the total time needed to perform all spectral transforms, the LEKF algorithm, and $k + 1$ six-hour forecasts from the resulting analysis. The column labeled “LEKF only” refers to the maximum amount of time that any given processor spends performing only the data assimilation step after the model grid is distributed more-or-less evenly across the cluster; the last column shows the average time spent processing a single local region.

These results suggest that, for a given observing network and patch size, the overall time required to perform the LEKF assimilation step grows roughly quadratically with the number of ensemble solutions and is relatively insensitive to the amount of data to be assimilated. In practice, of course, the work required to decode observations from data files and evaluate the observation operators may be substantially greater than what is needed here. Nevertheless, we are optimistic that, with further tuning and a somewhat larger computer, the LEKF data assimilation algorithm can be performed within the time constraints of a typical operational forecast center.

6. Conclusion

We have demonstrated, by an implementation on the NCEP GFS, that the LEKF is a highly accu-

rate and computationally efficient data assimilation scheme. Based on the results presented here, we believe that the LEKF is an operationally attainable. We hope to further test the potentials of the LEKF for operational purposes by an implementation on a state-of-the-art, high resolution, operational regional model and by assimilating real observations into both the global and the regional models. A relatively quick implementation is made possible by the important feature of the LEKF algorithm, that it is independent of the details of a given weather model. The code required to interface with the GFS model is localized into a module called “Grid Manager,” which handles the interface between the spectral transforms and the physical grid used for the data assimilation. The entire software suite, including code to implement the LEKF algorithm, spectral transforms, and diagnostics, comprises about 8500 lines of Fortran 95 code (including extensive comments). About half of the total is contained in the Grid Manager and the spectral transforms; only this portion of the code needs to be rewritten to apply the algorithm to a different forecast model.

7. Acknowledgments

This work was supported by the Army Research Office, a James S. McDonnell 21st Century Research Award, by the Office of Naval Research (Physics), and by the National Science Foundation (Grants #0104087 and PHYS 0098632).

8. References

- Anderson, E. et al., 1999: *LAPACK Users' Guide*, third ed. Society for Industrial and Applied Mathematics, Philadelphia.
- Anderson, J. L., 2001: An ensemble adjustment filter for data assimilation. *Mon. Wea. Rev.*, **129**, 2884-2903.
- Anderson, J. L., and S. L. Anderson, 1999: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Wea. Rev.*, **127**, 2741-2758.
- IEEE, 1985: *IEEE Standard for Binary Floating-Point Arithmetic*. ANSI/IEEE Std. 754-1985. American National Standards Institute.
- Anderson, E. et al., 1999: *LAPACK Users' Guide*, third ed. Society for Industrial and Applied Mathematics, Philadelphia.
- Arakawa, A., and W. H. Schubert, 1974: Interaction

- of a cumulus cloud ensemble with the large-scale environment. Part I. *J. Atmos. Sci.*, **31**, 674–701.
- Bishop, C. H., B. J. Etherton, and S. Majumdar, 2001: Adaptive sampling with the Ensemble Transform Kalman Filter. Part I: Theoretical aspects. *Mon. Wea. Rev.*, **129**, 420–436.
- Dongarra, J. J. et al., 1988: An extended set of FORTRAN Basic Linear Algebra Subprograms. *ACM Trans. Math. Soft.*, **14**, 1–17.
- Dongarra, J. J. et al., 1990: A set of Level 3 Basic Linear Algebra Subprograms. *ACM Trans. Math. Soft.* **16**, 1–17.
- Dowell, D. C., F. Zhang, L. J. Wicker, C. Snyder, and N. A. Crook, 2004: Wind and thermodynamic retrievals in the 17 May 1981 Arcadia, Oklahoma, supercell: Ensemble Kalman filter experiments. *Mon. Wea. Rev.*, **132**, (in press)
- Emanuel, K. A., 1994: *Atmospheric convection*. Oxford University Press, New York.
- Evensen, G. 2003: The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynamics*, **53**, 343–367.
- Hamill, T. M., J. Whitaker, and C. Snyder, 2001: Distance-dependent filtering of background error covariance estimates in an Ensemble Kalman Filter. *Mon. Wea. Rev.*, **129**, 2776–2790.
- Hartman, D. L., T. N. Palmer, and R. Buizza, 1995: Finite-time instabilities of lower-stratospheric flow. *J. Atmos. Sci.*, **53**, 2129–2143.
- Houtekamer, P. L., and H. L. Mitchell, 2001: A sequential ensemble Kalman Filter for atmospheric data assimilation. *Mon. Wea. Rev.*, **129**, 796–811.
- Houtekamer, P. L., H. L. Mitchell, G. Pellerin, M. Buehner, M. Charron, L. Spacek, and B. Hansen, 2003: Atmospheric data assimilation with the ensemble Kalman filter: Results with real observations. Submitted. Available from peter.houtekamer@ec.gc.ca.
- Hunt, B. R., E. Kalnay, E. J. Kostelich, E. Ott, and D. J. Patil et al., 2004: Four-dimensional ensemble Kalman filtering. *Tellus*, **56A**, (in press).
- Kalnay, E., 2003: *Atmospheric modeling, data assimilation, and predictability*. Cambridge University Press, Cambridge, 341 pp.
- Kalnay, E. and Z. Toth, 1996: The breeding method. *Proc. of the ECMWF Seminar on Predictability*, Reading, United Kingdom, ECMWF, 69–82.
- Keppenne, C., and H. Rienecker, 2002: Initial testing of a massively parallel ensemble Kalman Filter with the Poseidon Isopycnal Ocean General Circulation Model. *Mon. Wea. Rev.*, **130**, 2951–2965.
- Lahey Computer Systems. www.lahey.com.
- Lawson, C. L. et al., 1979: Basic Linear Algebra Subprograms for FORTRAN usage. *ACM Trans. Math. Soft.* **5**, 308–323.
- Lorenz, E. N., 1996: Predictability: A problem partly solved. *Proc. Seminar on Predictability*, Vol. 1, ECMWF, Reading, Berkshire, UK, 1–18.
- Ott, E., B. R. Hunt, I. Szunyogh, M. Corazza, E. Kalnay, D. J. Patil, J. A. Yorke, A. V. Zimin, and E. Kostelich, 2002: Exploiting local low dimensionality of the atmospheric dynamics for efficient Kalman filtering. <http://arxiv.org/abs/physics/0203058>.
- Ott, E., B. R. Hunt, I. Szunyogh, A. V. Zimin, E. J. Kostelich et al., 2004: A local ensemble Kalman filter for atmospheric data assimilation. *Tellus*, **56A**, (in press).
- Pan, H.-L. and W.-S. Wu, 1995: Implementing a Mass Flux Convection Package for the NMC Medium-Range Forecast model. *NMC Office Note 409*. Available from NCEP, 5200 Auth Road, Washington, 20233.
- Sauer, T., B. R. Hunt, J. A. Yorke, A. V. Zimin, and E. Ott et al., 2004: 4D Ensemble Kalman filtering for assimilation of asynchronous observations. Preprints of the 20th Conference on Weather Analysis and Forecasting/16th Conference on Numerical Weather Prediction. Seattle, WA, USA, January 11–15, 2004.
- Snyder, C. and F. Zhang, 2003: Assimilation of simulated Doppler radar observations with an Ensemble Kalman filter. *Mon. Wea. Rev.*, **131**, 1663–1677.
- Szunyogh, I., E. J. Kostelich, G. Gyarmati, B. Hunt, and E. Kalnay et al., 2004: A Local Ensemble Kalman Filter for the NCEP GFS. Preprints of the 20th Conference on Weather Analysis and Forecasting/16th Conference on Numerical Weather Prediction. Seattle, WA, USA, January 11–15, 2004.
- Tippett, M. K., J. L. Anderson, C. H. Bishop, T. M. Hamill, and J. S. Whitaker, 2003: Ensemble square-root filters. *Mon. Wea. Rev.*, **131**, 1485–1490.
- Whitaker, J. S., and T. H. Hamill, 2002: Ensemble

Data Assimilation without perturbed observations. *Mon. Wea. Rev.*, **130**, 1913-1924.

Whitaker, J. S., G. P. Compo, X. Wei, and T. M. Hamill, 2003: Reanalysis without radiosondes using ensemble data assimilation. *Mon. Wea. Rev.*, submitted. Available from jefrey.s.whitaker@noaa.gov.

Zhang, F., C. Snyder, and J. Sun, 2004: Impacts of initial estimate and observation availability on convective-scale data assimilation with an Ensemble Kalman filter. *Mon. Wea. Rev.*, **132**, in press.