Ensemble clustering in deterministic ensemble Kalman filters

Article

Published Version

CentAUR

Ensemble clustering in deterministic ensemble Kalman filters

By JAVIER AMEZCUA[1]*†, KAYO IDE[1,2,3,4], CRAIG H. BISHOP[5] and
EUGENIA KALNAY[1,2,3], [1]*Department of Atmospheric and Oceanic Science, University of Maryland,
College Park, MD, USA*; [2]*Earth System Science Interdisciplinary Center, University of Maryland, College Park,
MD, USA*; [3]*Institute for Physical Science and Technology, University of Maryland, College Park, MD, USA*;
[4]*Center for Scientific Computation and Mathematical Modeling, University of Maryland, College Park, MD,
USA*; [5]*Naval Research Laboratory, Monterey, CA, USA*

ABSTRACT

Ensemble clustering (EC) can arise in data assimilation with ensemble square root filters (EnSRFs) using
non-linear models: an M-member ensemble splits into a single outlier and a cluster of $M-1$ members. The
stochastic Ensemble Kalman Filter does not present this problem. Modifications to the EnSRFs by a periodic
resampling of the ensemble through random rotations have been proposed to address it. We introduce a metric
to quantify the presence of EC and present evidence to dispel the notion that EC leads to filter failure. Starting
from a univariate model, we show that EC is not a permanent but transient phenomenon; it occurs
intermittently in non-linear models. We perform a series of data assimilation experiments using a standard
EnSRF and a modified EnSRF by a resampling though random rotations. The modified EnSRF thus alleviates
issues associated with EC at the cost of traceability of individual ensemble trajectories and cannot use some
of algorithms that enhance performance of standard EnSRF. In the non-linear regimes of low-dimensional
models, the analysis root mean square error of the standard EnSRF slowly grows with ensemble size if the size
is larger than the dimension of the model state. However, we do not observe this problem in a more complex
model that uses an ensemble size much smaller than the dimension of the model state, along with inflation and
localisation. Overall, we find that transient EC does not handicap the performance of the standard EnSRF.

Keywords: non-linear models, ensemble square root filters, non-Gaussianity

1. Introduction

The ensemble Kalman filter (EnKF) is a Monte Carlo
implementation of the Kalman filter (KF, Kalman, 1960),
which relies on the evolution of a family (ensemble) of
trajectories in the forecast and performs the KF analysis
using the sample estimators for mean and covariance when
observations become available. Depending on the way the
analysis is carried out, EnKF algorithms can be classified
into stochastic and deterministic. The stochastic (perturbed
observations) EnKF makes use of random number realisa-
tions to update individual ensemble members in the
analysis (Burgers et al., 1998; Houtekamer and Mitchell,

1998). These random realisations can be a source of
sampling error that decreases as the sample size increases
(Whitaker and Hamill, 2002); in fact, the KF covariance
equation is satisfied only in a statistical sense. In contrast,
the ensemble square root filters (EnSRF) use algorithms
for transforming the forecast ensemble into analysis
ensemble with consistent statistics (Tippett et al., 2003).
The EnSRFs include ensemble adjustment Kalman filter
(EAKF, Anderson, 2001), serial EnSRF (Whitaker and
Hamill, 2002), ensemble transform Kalman filter (ETKF;
one-sided: Bishop et al., 2001; spherical simplex or sym-
metric: Wang et al., 2004), local ensemble Kalman filter
(LEKF; Ott et al., 2004) and local ensemble transform
Kalman filter (LETKF; Hunt et al., 2007).

As in the KF, the optimality of any EnKF is not
guaranteed when the non-linear error-growth in the fore-
cast becomes significant and the distribution of the
ensemble members is no longer Gaussian. In non-linear

*Corresponding author.
email: jamezcua@atmos.umd.edu
†Now at: Department of Meteorology, University of Reading,
Reading, UK

1

forecast models, the departure from linear error-growth can depend upon the frequency of observations, the length of the assimilation window and the magnitude of the observational error covariance (Lawson and Hansen, 2004; Kalnay et al., 2007).

In a seminal study, Lawson and Hansen (2004) analysed the update mechanisms of the stochastic EnKF and the serial EnSRF and compared their performance in linear and non-linear regimes of the forecast. Their analysis showed that the serial EnSRF is better than the stochastic EnKF at retaining higher order moments of the background distribution. This implies, however, that any departure from Gaussianity in the background ensemble is likely retained in the analysis and propagated forward in the forecast. An important finding was that, in non-linear regimes, the serial EnSRF could lead to the emergence of outliers from the main ensemble that are mostly responsible for keeping the variance predicted by the KF. The higher order moments of the serial EnSRF ensembles presented non-Gaussian values, and the rank histograms for the verification of the truth were U-shaped, implying that the truth and the analysis ensemble members could not be considered statistically indistinguishable. Nonetheless, there was no difference reported in the analysis root mean squared error (RMSE) between the stochastic and the deterministic EnKFs.

Using the EAKF in highly non-linear scenarios, Anderson (2010) observed ensemble clustering (EC), a phenomenon in which a M-member ensemble splits in an outlier and a tight cluster of $M-1$ members. He found that the EC may occur as a result of the disparity between the non-linear expansion of the ensemble spread in the forecast and the linear contraction of the ensemble spread in the analysis. More precisely, non-linear expansion in the forecast may push the outermost member further outwards while a linear compaction in the analysis that is sufficient to constrain the outlier may be larger than required for the remaining members. Using other models, this study also showed that the analysis RMSE of the EAKF increased with ensemble size M due to the non-linear expansion in the forecast. To address the EC, Anderson (2010) proposed a rank histogram filter.

Another approach to alleviate the EC in EnSRF is a periodic resampling of the ensemble (e.g. via bootstrapping) as suggested by Lawson and Hansen (2004). Leeuwenburgh et al. (2005) implemented the resampling in the one-sided ETKF (Bishop et al., 2001) using a random rotation of the transform matrix and called it the EnSRF+. For temperature assimilation in an ocean model, Leeuwenburgh et al. (2005) showed that the EnSRF+ outperformed the one-sided ETKF in terms of the RMSE, with the higher order moments of the ensemble closer to the Gaussian. There was a caveat, however,

in EnSRF+. The one-sided ETKF is not an unbiased square root filter (Livings et al., 2008; Sakov and Oke, 2008) and neither is the EnSRF+. The problems associated with the one-sided ETKF are illustrated in Sakov and Oke (2008), who compared the performance of the (unbiased) spherical-simplex ETKF (Wang et al., 2004) and an unbiased randomly rotated ETKF. Using the 40-variable Lorenz 1996 model (L96; Lorenz and Emanuel, 1998) with different ensemble sizes and multiplicative covariance inflation factors, Sakov and Oke (2008) found similar performance for the both filters in terms of analysis RMSE (see their Fig. 3). Their rotated ETKF, however, produced ensembles with more Gaussian-like characteristics in terms of higher order moments and flatter rank histograms in the verification of the truth.

The objective of this study is to shed light to the issues associated with EC in the EnSRFs. A particular focus is placed on the notion that, once EC occurs, it sets in and severely handicaps the performance of the EnSRF. We dispel this notion by showing that EC is in general a transient phenomenon and the EnSRFs quickly recover after individual events of EC. We base our EnSRFs on the unbiased ETKF (spherical-simplex ETKF: Wang et al., 2004; LETKF: Hunt et al., 2007) because of its susceptibility to EC by the choice of analysis ensemble spread to be the closest to the background analysis ensemble spread (Section 2). We introduce a second type of EnSRF that aims to alleviate EC by adopting random rotations in the ensemble transform matrix. To quantify degree of EC evolving in time, we introduce a metric called clustering degree (CD).

Our analysis starts with a univariate quadratic model (Section 3) and moves onto more complex models: the three-variable Lorenz 1963 (L63) model (Section 4) and a medium complexity atmospheric general circulation model known as SPEEDY (Molteni, 2003) where we discuss localisation (Section 5). Through analysis and experiments, we show advantages and disadvantages in both types of the EnSRFs with or without random rotations; each appeals to different situations and specific assimilation goals. Our conclusions and summary are presented in Section 6.

2. Metric for ensemble clustering and choice of the EnSRF methods

Starting with a univariate ensemble, we define clustering degree (CD) as:

$$CD = \frac{\sigma^2_{M-1}}{\sigma^2_M}. \qquad (1)$$

The denominator of eq. (1) is the variance of the M-member ensemble, while the numerator is the variance

of the $M-1$-member ensemble that remains after removing the outermost member of the original. By this definition, CD spans from zero to one. If EC is present, most of the variance comes from the outermost member (which for this case can be labelled as an outlier) and hence CD will tend towards zero. The average CD value of a 'healthy' (unclustered) ensemble depends on the ensemble size, and it has less variation as this ensemble size increases. Nonetheless, as our experiments suggest that small CD robustly detect EC for any M.

For a multivariate case, this metric can be generalised to:

$$\mathrm{CD} = \frac{\mathrm{Trace}(P_{M-1})}{\mathrm{Trace}(P_M)}. \qquad (2)$$

The denominator is the trace of the M-member-ensemble covariance matrix, while the numerator is the trace of the $M-1$-member-ensemble covariance matrix after removing the outermost member. Equation (2) is adequate in the multivariate case only when the variables have the same units. If this condition is violated, one can use a proper norm (e.g. an energy norm) when summing the variances or one can perform the analysis separately for different sets of state variables.

Having defined the metric for EC, we lay out our strategy using two types of EnSRFs. The first type is the one prone to EC. We choose symmetric ETKF (S-ETKF), which is equivalent to the spherical-simplex ETKF (Wang et al., 2004) as well as the LETKF (Hunt et al., 2007) without the application of localisation. This form of ETKF is unbiased and preserves the mean (Ott et al., 2004; Wang et al., 2004). Moreover, S-ETKF retains the analysis ensemble spread the closest to the original background ensemble spread (Hunt et al., 2007). This property makes S-ETKF suitable for this study because of its susceptibility to EC. It also makes S-ETKF desirable because of its ability to capture the 'errors of the day' through the KF analysis. The second type is based on S-ETKF but designed to alleviate EC. We thus apply a constrained random rotation to the ensemble transform matrix in S-ETKF, in such a way that ensemble mean is preserved and thus address the bias issues (Livings et al., 2008; Sakov and Oke, 2008). The resulting method is non-symmetric ETKF (NS-ETKF). An essential difference between the two types of ETKFs is that NS-ETKF cannot trace individual ensemble trajectories to the past and thus loses one of the benefits of EnSRFs (Anderson, 2001). In the case of ETKFs, this means loss of traceability of ensemble weights. Another significant difference is that NS-ETKF also loses the 'errors of the day' information by resampling. Technical details of S-ETKF and NS-ETKF are provided in the Appendix. To perform detailed analysis and compare the results, all data assimilation experiments are carried out in the identical twin setting. In following section, these methods are compared through the background (i.e. forecast) and analysis RMSE, the third-order moment of the ensemble (sample skewness) as defined in the Appendix of Lawson and Hansen (2004) and the time evolution of the CD.

3. Ensemble clustering in a simple non-linear model

Following Anderson (2010), we consider the univariate quadratic ordinary differential equation (ODE) $dx/dt = x + b|x|x$. A forecast model based on the Euler forward discretisation of this ODE is

$$x_{t+1} = x_t + \Delta(x_t + b|x_t|x_t), \qquad (3)$$

where $\Delta = 0.05$ is chosen as the time step. This model exhibits the necessary non-linear expansion described in Section 1 through the non-linear coefficient b. The system described by eq. (3) has an unstable fixed point at the origin, which we use as the truth, i.e. $x^t = 0$. Observations are made every two model steps unless otherwise noted, by adding random Gaussian noise $\sim N(0, 1)$. We assimilate every time we observe. We vary the ensemble size, $M = \{10, 20, 100\}$ and the non-linearity coefficient, $b \in [0, 0.2]$, where Anderson (2010) used $b = 0.2$. The initial ensemble members are drawn uniformly from the interval $[-1, 1]$.

Figure 1 shows the time evolution of the analysis ensemble for the case $b = 0.1$ and $M = 10$ with S-ETKF (a) and NS-ETKF (b). Panel c shows CD: for S-ETKF, CD smoothly decreases towards zero indicating EC, while, for NS-ETKF, it changes abruptly at every analysis, but the variation remains around a mean value.

Using S-ETKF, we study the effects of ensemble sizes $M = \{10, 15, \ldots, 100\}$ and non-linearity coefficients $b = \{0.02, 0.025, \ldots, 0.7\}$. The results are presented in Fig. 2. In panel a, we present the time evolution of CD for combinations of three different values of both M and b. EC occurs for any $b > 0$; its emergence sets in earlier as b increases. Moreover, EC takes place more gradually for smaller M and more abruptly for larger M. In all cases, EC seems to occur at the same time that depends solely on b, indicating that this phenomenon is related to the intrinsic non-linearity in the model dynamics. Moreover, for a given b, the curves for different M seem to come together around a small CD value, which we have empirically estimated to be $\mathrm{CD}_c \approx 0.04$ (indicated by a horizontal line in the panel). For each combination of M and b, we measure the time at which CD crosses below the CD_c threshold (we call this time t_c). In panel b, we plot t_c as a function of b, with a line for each value of M. All lines
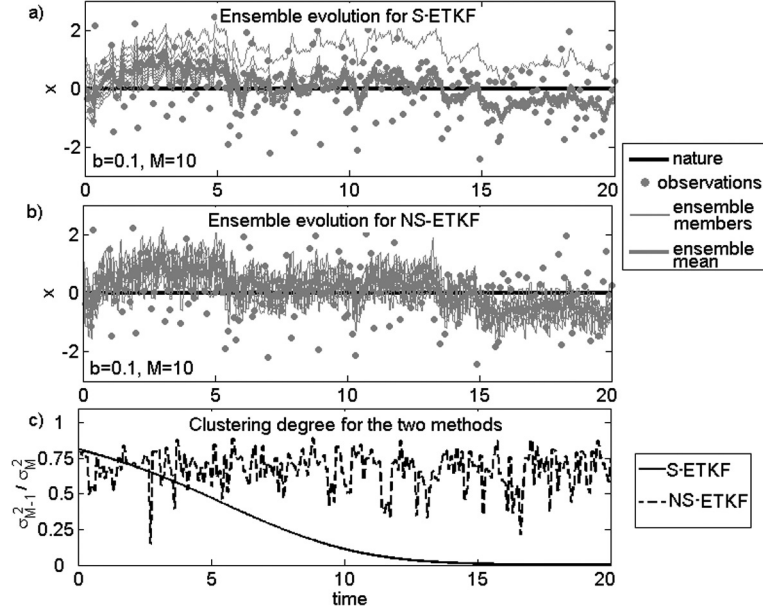
Fig. 1. Data assimilation experiment with the model $x_{t+1} = x_t + 0.05(x_t + b|x_t|x_t)$, observations every two model steps, and $M = 10$ and $b = 0.1$. S-ETKF (panel a) presents EC soon after five time units while the NS-ETKF (panel b) does not. In panel (c), we quantify the clustering degree (as defined in Section 2) of the ensemble obtained for both assimilation methods as time advances.

are almost indistinguishable using log–log plot, while y-intercepts slightly increase with M. This indicates that the convergence rate to a clustered ensemble in this simple prototypical model follows a power law: $t_c = Ab^{-\gamma}$. The coefficient A is related to the y-intercept of the log–log plot and has a slight dependence on M. The exponent is

estimated to be $\gamma \approx -0.92$. Further investigation of the properties associated with this figure is beyond the scope of this work.

As shown in Fig. 1, NS-ETKF successfully avoids EC by resampling. To illustrate the point, we depict the update process of both filters from background to analysis for each
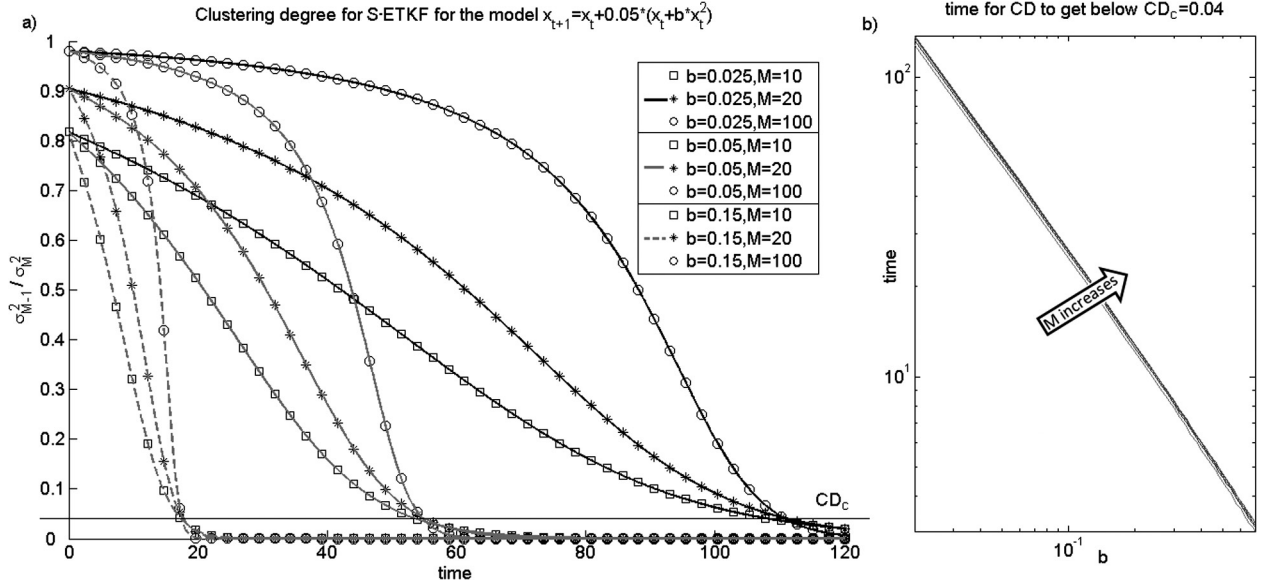

Fig. 2. Data assimilation experiment using S-ETKF with the model $x_{t+1} = x_t + 0.05(x_t + b|x_t|x_t)$ and observations every two model steps. The effect of different values of non-linearity b and ensemble size M are explored. In panel (a), the time evolution of the CD metric is shown. As expected, EC appears faster as the non-linearity increases, and this appears to be independent of the ensemble size. In panel (b), we measure the time it takes for CD to get below CD_c, this relationship follows a power law.

one of the $M = 10$ members in Fig. 3. To accelerate the occurrence of EC, we take observations every five model steps and $b = 0.2$. For S-ETKF (panel a), the analysis ensemble is chosen to be the closest to the background ensemble (Hunt et al., 2007). Therefore, any deformation introduced by the non-linear expansion in the forecast will remain in the analysis; the separation of the outliner member from the cluster cannot be stopped once it starts in this simple model. By contrast, NS-ETKF (panel b) effectively erases any deformation occurred during the forecast by resampling at each analysis.

Rank histograms for the verification of the truth with respect to the analysis ensemble were computed for both methods (not shown). For S-ETKF, the truth very often falls either outside the ensemble or between the outlier and the cluster. For NS-ETKF, the histograms are generally flat, evidence that the truth is statistically undistinguishable for the ensemble. However, both methods estimate very similar analysis means, leading to indistinguishable performances in terms of RMSE.

The experiments presented so far seem to suggest that, once EC sets in, it is irreversible and can imply a major obstacle for an EnSRF. It is crucial to realise that non-linearity in the simple univariate model (3) has been maintained constant by b that is fixed in time. In higher-dimensional models, intrinsic non-linearity is spatially and temporally variable as the trajectory may visit different regimes of the phase space. In the reminder of this paper, we demonstrate that this variability of the non-linearity can help revert the EC and thus the EC is a transient phenomenon that occurs intermittently.

To introduce the variability of non-linearity in the univariate model as simple as eq. (3), we let b change every T model time steps, where T comes from a uniform distribution $\sim U(T_0, 2T_0)$. Every time a 'cycle' of length T completes, a new b is drawn from $N(0, 0.1)$; in this way,

$\sim 95\%$ of the cycles have $|b| < 0.2$. Hence, forecast model dynamics experience different dynamical regimes for ensemble spread near the truth: unstable expansion ($b > 0$) or stable contraction ($b < 0$). Dynamics is quasi-linear for $|b| \sim 0$.

In panels a and b of Fig. 4, we show b in grey line (right vertical axis) along with the CD of the S-ETKF in black line (left vertical axis) for the interval $t \in [200, 900]$. We show the results for the cases $T_0 = \{50, 500\}$ model steps in these panels. By introducing the variability in non-linearity of the forecast model, the S-ETKF no longer suffers from irreversible EC. In panel a, around $t \approx 200$ EC sets in due to large positive values of b. EC persists until $t \approx 600$ but decays as the outlier returns to the rest of the ensemble as shown in panel c. EC reemerges subsequently (not shown in this panel), but it is again transient and subsides. In general, it is an intermittent phenomenon.

In this simple model, EC can persistent over a long period, although the introduction of the artificial variability in non-linearity eventually resolves the EC. In higher-dimensional models with natural variability in non-linearity, EC is less persistent as we demonstrate in the next two sections.

4. Experiments with L63

L63 is a non-linear 3-variable model widely used to test data assimilation schemes because of its challenging properties near regime changes (e.g. Miller et al., 1994; Evans et al., 2004; Kalnay et al., 2007). The system of non-linear coupled ODEs describing its evolution is

$$\dot{x}^{(1)} = p\left(x^{(2)} - x^{(1)}\right)$$
$$\dot{x}^{(2)} = x^{(1)}\left(r - x^{(3)}\right) - x^{(2)} . \qquad (4)$$
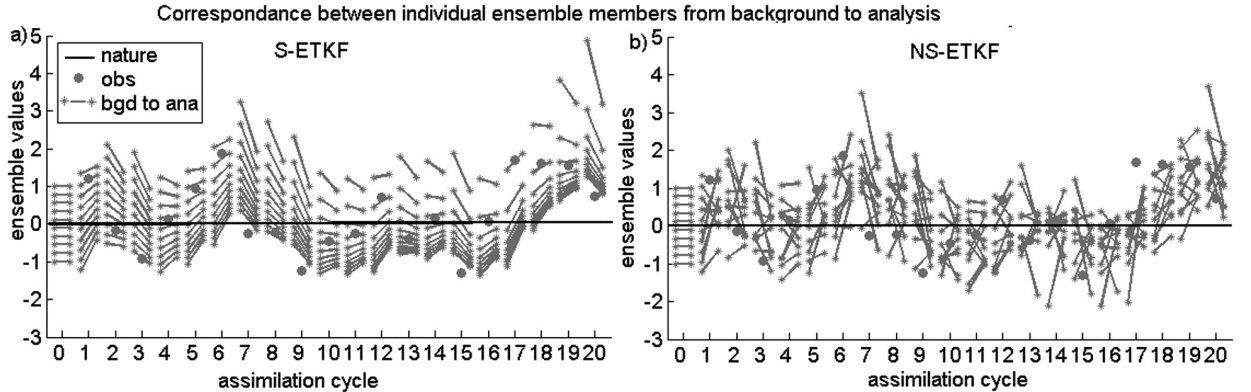$$\dot{x}^{(3)} = x^{(1)}x^{(2)} - bx^{(3)}$$

Fig. 3. Update mechanisms for S-ETKF (panel a) and NS-ETKF (panel b) for the individual ensemble members. S-ETKF preserves the structure from the background ensemble into the analysis ensemble. The NS-ETKF effectively scrambles the ensemble every time an assimilation occurs. The model is $x_{t+1} = x_t + 0.05(x_t + b|x_t|x_t)$, $b = 0.2$, observations every five model steps, and $M = 10$.
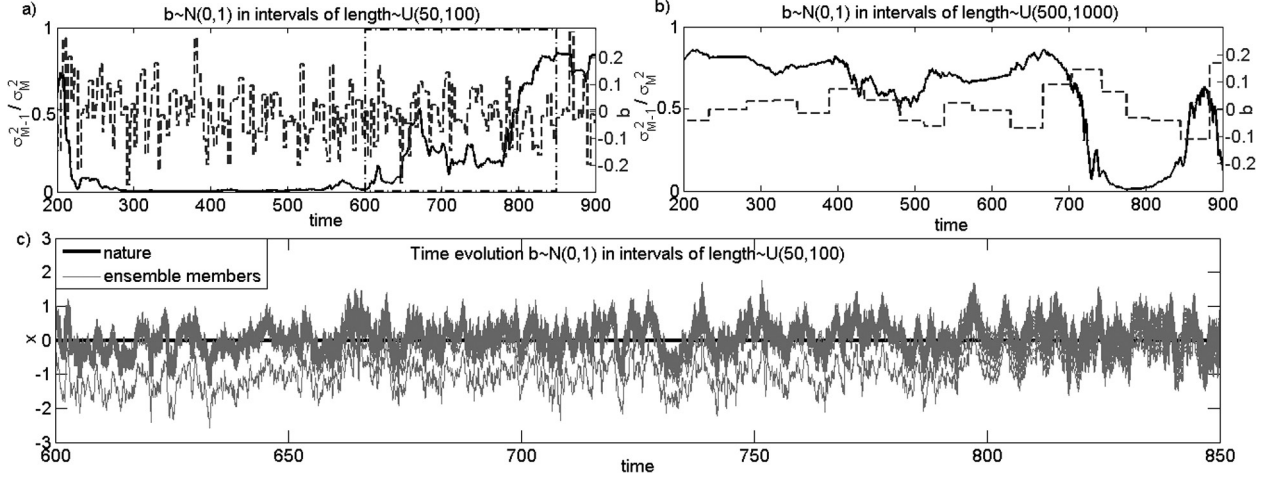
Fig. 4. Assimilation experiments with the model $x_{t+1} = x_t + 0.05(x_t + b_t |x_t| x_t)$. We allow the non-linear coefficient b_t to vary as a piecewise function of time (grey line, right vertical axes). CD is represented by the black line and left vertical axes. The time intervals in which b_t is fixed are different for each panels: $T_0 = 50$ for (a) and $T_0 = 500$ for (b). Panel (c) shows the ensemble evolution for the time interval $t \in [600, 850]$ of case (a); the reattachment of the outlier occurs in a natural way.

The standard values for the parameters are $p = 10$, $r = 28$ and $b = 8/3$, which result in a chaotic behaviour with two regimes in a very well-known butterfly-shaped fractal attractor in the phase space. To generate the nature run, the model is integrated with the Runge–Kutta fourth-order method (RK4) using a time step $\Delta t = 0.01$ for 10^6 steps after a short spin-up to put the trajectory on the attractor. The observations are generated by adding a random noise $N(\mathbf{0}, \mathbf{R} = 2\mathbf{I})$ to the nature run; all variables are observed. Two types of observing systems are used by varying the frequency of the observations: one with a short assimilation window using frequent observations at every eight model steps and the other with a longer window sing infrequent observations at every 24 model steps. The short and long assimilation windows correspond to the linear and non-linear regimes for ensemble spread in the mode forecast (Kalnay et al., 2007).

Using L63, we study the effect of the ensemble size M with respect to the dimension of the model state N as well as that of non-linearity in the forecast model on the background ensemble spread by changing the assimilation window length. We present results for two ensemble sizes $M = \{3, 20\}$. For $M = 3$ with the rank-deficient background covariances P_M and P_{M-1} in eq. (2) due to $M - 1 < N$, multiplicative covariance inflation $\mathbf{X}^b \rightarrow \mathbf{X}^b(1 + \delta)$ is applied with $\delta = 0.04$ for the short assimilation window and $\delta = 0.04$ for the long assimilation window. These values are close to the optimal values obtained in Kalnay et al. (2007) and Amezcua (2012).

Figure 5 shows the CD for $t \in [1525, 1550]$. The top row illustrates the cases for the linear regime, while the bottom row represents the cases within the non-linear

regime. For $M = 3$ (left column), we observe very rapid variations in CD for both S-ETKF (black line) and NS-ETKF (grey line). Still, some instances of clustering (e.g. $t \approx 1548$) emerge in the non-linear regime for S-ETKF.

With large M (only the case $M = 20$ is shown), there is a clear difference in the CD between S- and NS-ETKFs. For NS-ETKF, it varies abruptly (but around a mean value) every time the assimilation is performed, but the variation is smaller as M increases. For S-ETKF, the variations in CD are slower and smoother; CD can reach low values in both the linear and non-linear regimes, but it does so more often in the non-linear regime. There are no cases of irreversible collapse of the ensemble; when EC occurs, it is only transient and not as persistent as with the simple quadratic model. Figure 6 illustrates EC in the non-linear regime with $M = 20$ using S-ETKF. The top panel shows the CD evolution for a longer time period $t \in [165, 200]$. There is an indication of EC around $t = 190$. The three panels in the bottom row of this figure show the trajectories for the truth (black line) and the analysis ensemble members (grey lines) at three different instants with different CD values. The middle panels show the case with EC, being evident in what seems to be a two-member ensemble. This, however, does not prevent the ensemble to revert the EC afterwards.

Why is EC less persistent in this model? In the univariate quadratic model, EC occurs and decays with the varying magnitude of the non-linear expansion and contraction of the ensemble spread. In higher dimensional models, not only the magnitude but also the direction changes temporally and spatially. A way to study the characteristics in the local perturbation growth is by using bred vectors
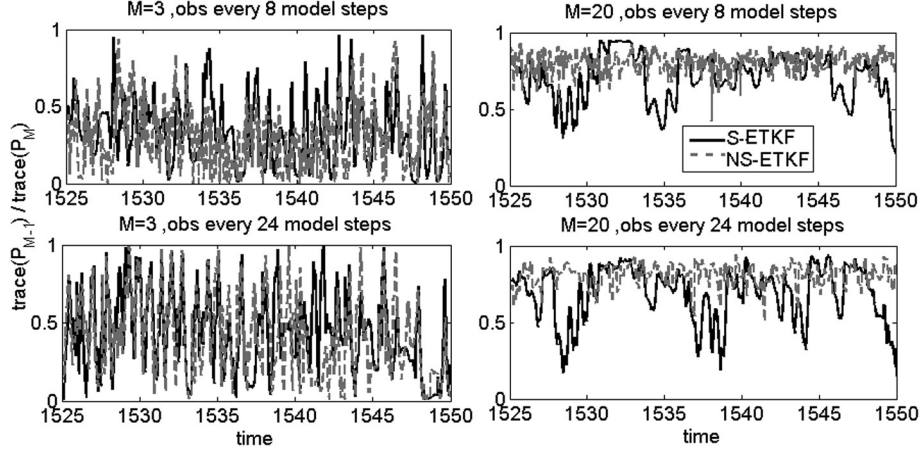
Fig. 5. Time evolution of the CD for S-ETKF (black solid line) and NS-ETKF (grey-dashed line) from an assimilation experiment with the L63 model. Two ensemble sizes (columns) are used in a linear regime (top row) and a non-linear regime (bottom row).

(Toth and Kalnay, 1997). Evans et al. (2004) applied this technique and showed different magnitudes of growth for different regions of the attractor. Zhang et al. (in preparation) have recently extended this study and have illustrated the change in direction as well.

Figure 7a shows statistical measures of both filters for the linear (left column) and non-linear (right column) regimes depending on the observation frequency, with boxplots for the CD (top row) and analysis RMSE (bottom row) for the two ensemble sizes. The black dots accompanying the boxplots represent the mean for each metric; these values are also displayed in the figure. For a small ensemble size, performance of S-ETKF and NS-ETKF is practically the same (see Fig. 5). For a larger ensemble size, differences arise. S-ETKF in general presents smaller CD values, a sign that it is more prone to EC. For the linear regime, both the background and analysis RMSEs have a similar distribution with little difference in the mean for S- and NS-ETKFs. For the non-linear regime, NS-ETKF has less outliers in the ensemble spread, leading to smaller mean RMSE. This is consistent with the finding by Anderson (2010) that the mean analysis RMSE of the EAKF increased for the larger ensemble size. One can hypothesise if there is any relationship between the mean CD value in the forecast and the analysis RMSE at the end of that window.
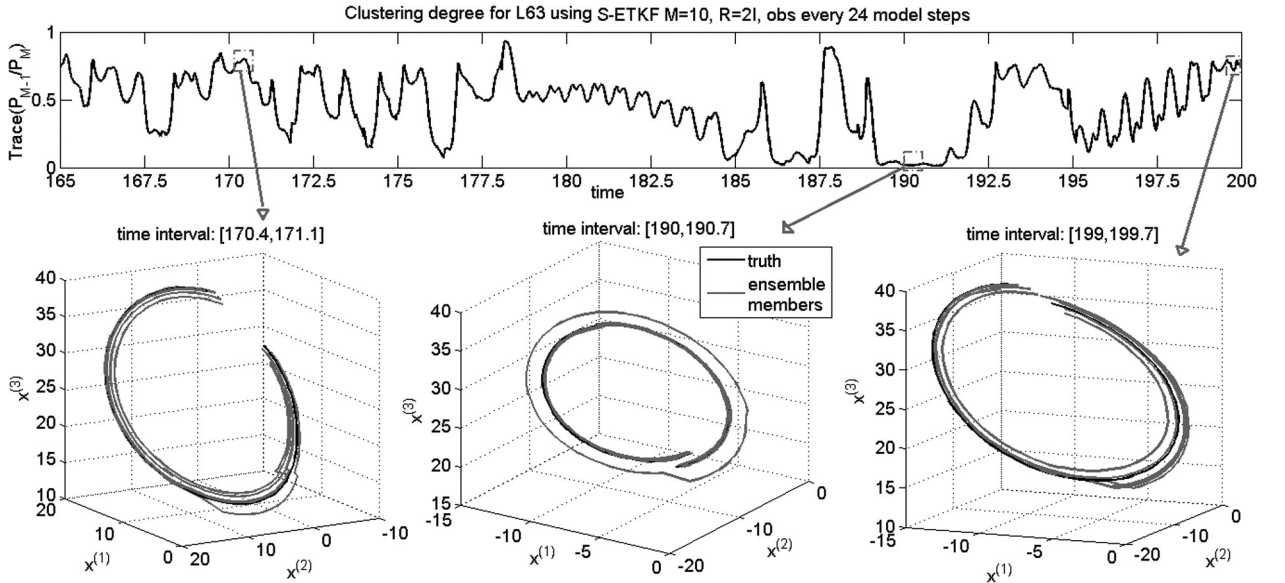

Fig. 6. Experiments with L63, observations every 24 model steps and **R = 2I**. The evolution of the CD is shown in the top. Snapshots of the phase space are presented for three time intervals with contrasting CD values, the one in the middle shows EC occurring.
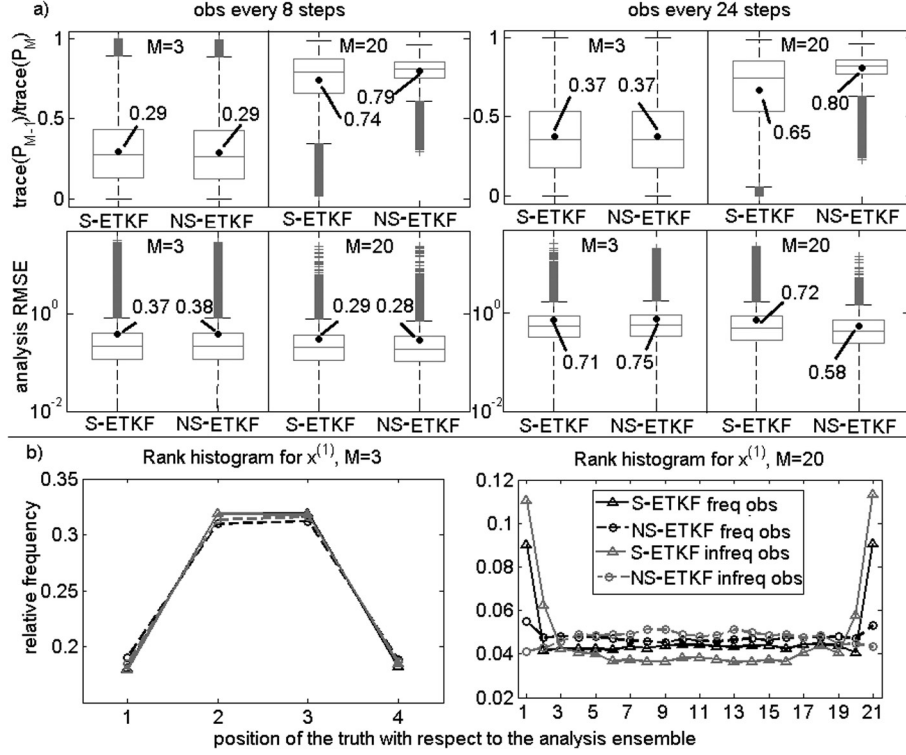
Fig. 7. Statistical summary of the experiment with L63. In part (a), the left column shows the results in the linear regime and the right column in the non-linear regime. Boxplots for CD (top row) and analysis RMSE (bottom row) are shown for both S-ETKF and NS-ETKF. The dots inside the boxplots represent the mean for each metric; the actual values displayed. In part (b), rank histograms for the verification of the truth with respect to the analysis ensemble are presented for variable $x^{(1)}$.

Figure 7b presents the rank histograms for the verification of the truth with respect to the analysis ensemble for variable $x^{(1)}$. For $M = 3$, there is no difference between S- and NS-ETKF: all ensembles are over-dispersive. This may be a result of the use of inflation. For $M = 20$, the S-ETKF has a U-shaped histograms, especially in the non-linear regime. Using NS-ETKF, on the other hand, produces flat rank histograms.

The results in this section show that NS-ETKF has a better performance in the non-linear regime when $M > N$, and this difference is more evident as M grows (experiments were performed with larger ensemble sizes, and these results are not shown). In practical applications, however, usually $M < < N$ and techniques such as localisation and covariance inflation are needed to compensate for the limited ensemble size. This is the focus of next section.

5. Results with more complex models and the use of R-localisation

To investigate EC in more complex model, localisation must be incorporated into the S-ETKF and NS-ETKF. R-localisation – introduced for LETKF (Hunt et al., 2007) – is

a natural choice of the localisation scheme for the post-multiplicative EnSRFs where transformation from background to analysis is performed in the ensemble space. In this scheme, an independent analysis is carried out for every single grid point using observations within a certain distance and assuming that the observation error increases with the distance to the grid point (see Greybush et al., 2011, for details). For stability in the model forecast, it is important that the analyses obtained in neighbouring grid points vary smoothly. This was one of the reasons for the use of symmetric square root in LEKF (Ott et al., 2004) and LETKF (Hunt et al., 2007).

With R-localisation, S-ETKF becomes LETKF (Hunt et al., 2007) and is denoted as S-LETKF in this study for consistency. To implement R-localisation in NS-ETKF while guaranteeing the smoothness among from one grid to neighbouring grid, we impose a single, global random transition matrix on S-LETKF performed at individual grid points. We call this implementation as NS-LETKF. While NS-LETKF can benefit from the resampling, there are disadvantages associated with the loss of traceability of the weight matrix (see also Section 2). In particular, a variety of the schemes that help the S-LETKF enhance its performance is not applicable to NS-LETKF. Such schemes

include quasi-outer-loop and running-in-place that improve the accuracy of ensemble Kalman filtering under non-linear, non-Gaussian perturbation growth (Yang et al., 2009; Kalnay and Yang, 2010; Yang et al., in press).

Having localisation implemented, we carry out data assimilation experiments using a model that is more representative of those used in operational numerical weather prediction, known as simplified parameterizations, primitive-equation dynamics (SPEEDY), developed by Molteni (2003) and adapted for data assimilation by Miyoshi (2005). SPEEDY has a spectral primitive-equation dynamic core and a set of simplified physical parameter-isation schemes with horizontal spectral truncation at 30 wave numbers and seven vertical levels. The model is formulated in σ coordinates and calculates five field variables: zonal wind u, meridional wind v, temperature T, relative humidity q and surface pressure ps. The geopotential height z for different pressure levels may be obtained by interpolation.

The nature run for our experiments starts after a 1-yr spin-up from state of rest and lasts 2 months (January and February). The observations are taken every 6 h at all seven vertical levels at horizontal positions that resemble a realistic radiosonde observational network (Fig. 9b of Miyoshi, 2011) by adding Gaussian random perturbations to every variable with the following SDs: 1 m/s for u and v, 1 K for T, 10^{-3} Kg$_{water}$/Kg$_{air}$ for q and 1 hPa for ps. The observation density is higher over continents than over the oceans, and the Northern Hemisphere is better observed than the Southern Hemisphere. Both EnSRFs use an ensemble size $M = 20$ and the R-localisation length scales of $\lambda = 500$ km in the horizontal and $\lambda_v = 0.1 \ln p$ in the vertical. We use the adaptive multiplicative covariance inflation (Miyoshi, 2011), tailored for R-localisation by estimating the time-evolving inflation parameter at each gridpoint.

For the variables $\{u, v, T, q, z\}$, we compute two latitude-weighted metrics separately at each vertical level both globally and per region: NH (25N–75N), tropics (25S–25N) and SH (75S–25S). These metrics are: analysis RMSE and sample skewness of the analysis ensemble. No noticeable difference in skewness or RMSE values is observed for the variables $\{u, v, z\}$; the variables $\{T, q\}$ do present differences. The results for T are shown in Fig. 8. The right panel shows the mean skewness value (along with bars indicating 1 SD) for the ensembles generated by S- and NS-LETKFs for the four geographic regions previously indicated (rows) and for three levels of the atmosphere (columns). S-LETKF tends to create asymmetric ensembles in the tropics and the SH. These are poorly observed regions in which non-linear behaviour can arise. Nonetheless, in spite of these non-Gaussian ensembles, the analysis RMSE values (left part of the plot) show no difference between S- and NS-LETKFs. The same is true for q (not shown). We plot rank histograms for the verification of the truth with respect to the analysis ensemble for the variables at different
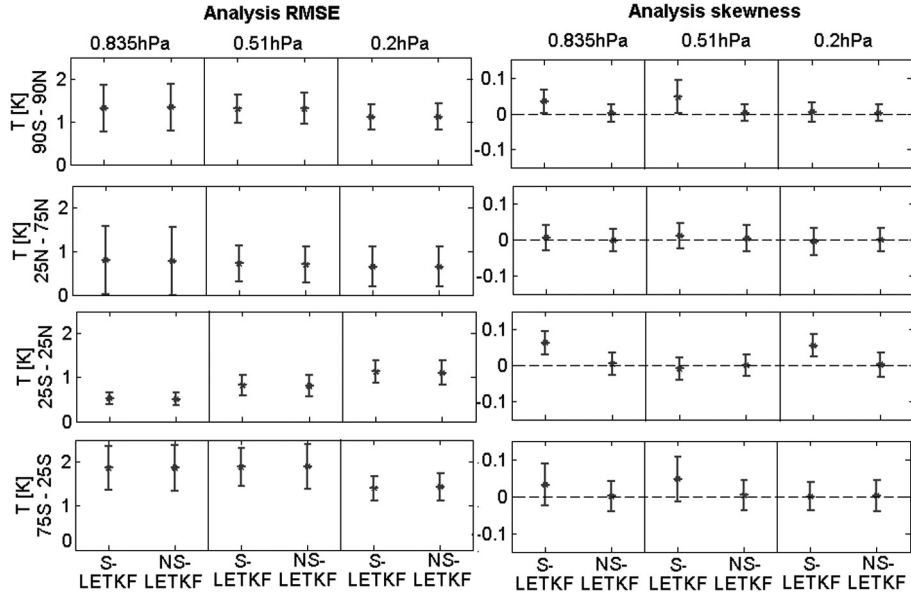

Fig. 8. Latitude weighted analysis RMSE (left) and analysis skewness (right) for the variable T computed per region (rows) for three vertical levels (columns) in the SPEEDY model. The bars represent 1 SD of the metric around its mean. S-(L)ETKF can lead to asymmetric ensembles (e.g. in the tropics in the lower and upper atmosphere), but there is no noticeable difference in its performance with respect to NS-(L)ETKF in terms of analysis RMSE.

pressure levels and for different regions (not shown). We do not observe differences between the two methods, both lead to under-dispersive ensembles.

6. Conclusions and discussion

In this work, we have studied EC that arises when using EnSRFs in non-linear forecast models. It results from the consecutive iteration of the non-linear expansion of the ensemble spread in the forecast and the linear contraction in the analysis. We have introduced a metric, CD (2), to quantify and follow the behaviour of EC in time. Very small values of this metric help detect the presence of EC. One of our main goals has been to dispel the notion that EC is an irreversible phenomenon that severely handicaps EnSRFs.

We have based our EnSRFs on the unbiased ETKF (spherical-simplex ETKF: Wang et al., 2004; LETKF: Hunt et al., 2007) as a standard EnSF because of its susceptibility to EC; the analysis ensemble spread is the closest possible to the background ensemble spread, due to the use of symmetric square root for the transform matrix in the ensemble space. We call this EnSRF as S-ETKF in this study. For comparison, we have introduced a second type of EnSRF that aims to alleviate EC by resampling the ensemble through random rotations of the transform matrix while preserving the ensemble mean (i.e. unbiased). It is called the NS-ETKF due to lack of symmetry in the transition matrix. Extensions of S- and NS-ETKFs with R-localisation are called S- and NS-LETKFs, respectively.

Using models of increasing complexity, we have assessed the performance of S-(L)ETKF and that of NS-(L)ETKF in the following aspects: (a) the accuracy of the ensemble mean as best estimator of the truth in terms of background and analysis RMSE, (b) the behaviour of higher order moments of the ensemble, in particular sample skewness, and (c) the statistical reliability of the ensemble with respect to the truth as measured by rank histograms as function of ensemble size M and the dimension of the system N. In the linear regimes, the two filters have indistinguishable performances. It is in the non-linear regimes, differences arise as expected; the remaining of the text refers to this case.

In terms of RMSE, the results of experiments with the L63 show that differences are noticeable only when the ensemble size becomes much larger than the number of variables. The NS-ETKF has a lower mean RMSE because a smaller number of cycles with very large RMSE appear, but the general distribution of the RMSE is not very different from that of S-ETKF as shown by boxplots. For the SPEEDY model, the RMSE values obtained by the two methods are indistinguishable for all variables even in the poorly observed regions of the globe.

Our experiments show that S-(L)ETKF tends to shift ensembles away from the Gaussian distribution in non-linear regimes, consistent with the finding by Anderson (2010). For the SPEEDY model using S-LETKF, this tendency is clear especially for the tracer variables $\{T, q\}$ in the tropics and the SH, where the sample skewness values are clearly different from zero. This, however, does not lead to higher RMSE of S-LETKF values with respect to that of NS-LETKF.

When verifying the truth against the analysis ensemble in L63 when $M > N$, the rank histograms obtained from NS-ETKF tend to be flat, while those obtained from S-ETKF are not. While the truth tends to be statistically indistinguishable from the NS-ETKF ensemble, it is not the case for the S-ETKF ensemble. This is not the case when $M = N$; in this case, we get over-dispersive ensembles for both filters. For SPEEDY, using $M < < N$ with localisation and adaptive multiplicative inflation, the rank histograms obtained by both filters have the same behaviour, viz. they show an under-dispersive ensemble. Moreover, persistence time in individual EC events becomes less persevering for $M < < N$, i.e. the reversal occurs before EC severely hinders S-LETKF performance.

While NS-(L)ETKF can benefit from the resampling to stay away from EC and the ability to sustain the ensemble spread statistically closer to Gaussian distribution, it lacks advantages of S-(L)ETKF, i.e. access to useful algorithms that enhance performance and ability to capture the 'errors of the day'. Overall our experiments show that EC is a transient phenomenon and does not severely handicap performance of S-(L)ETKF, a representative of standard EnSRFs. We do not intend to assert that one filter is better than the other; as a matter of fact, the conclusion would be different depending on the choice of the focus. We end this work echoing a conclusion from Lawson and Hansen (2004), namely, that the key to handle different filters is to understand their mechanisms, implications and limitations.

7. Acknowledgements

References

Amezcua, J. 2012. *Advances in Sequential Data Assimilation and Numerical Weather Forecasting: An Ensemble Transform Kalman–Bucy Filter, a Study on Clustering in Deterministic Ensemble Square Root Filters, and a Test of a New Time Stepping Scheme in an Atmospheric Model*. PhD Dissertation. University of Maryland, College Park.

Anderson, J. L. 2001. An ensemble adjustment filter for data assimilation. *Mon. Weather Rev.* **129**, 2884–2903.

Anderson, J. L. 2010. A non-Gaussian ensemble filter update for data assimilation. *Mon. Weather Rev.* **138**, 4186–4198.

Bishop, C. H., Etherton, B. J. and Majumdar, S. J. 2001. Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Mon. Weather Rev.* **129**, 420–436.

Burgers, G., van Leeuwen, P. J. and Evensen, G. 1998. Analysis scheme in the ensemble Kalman filter. *Mon. Weather Rev.* **126**, 1719–1724.

Evans, E., Bhatti, N., Kinney, J., Pann, L., Peña, M. and co-authors. 2004. RISE undergraduates find that regime changes in Lorenz's model are predictable. *Bull. Am. Meteorol. Soc.* **85**, 520–524.

Greybush, S. J., Kalnay, E., Miyoshi, T., Ide, K. and Hunt, B. R. 2011. Balance and ensemble Kalman filter localization techniques. *Mon. Weather Rev.* **139**, 511–522.

Houtekamer, P. I. and Mitchell, H. I. 1998. Data assimilation using an ensemble Kalman filter technique. *Mon. Weather Rev.* **126**, 796–811.

Hunt, B. R., Kostelich, E. J. and Szunyogh, I. 2007. Efficient data assimilation for spatiotemporal chaos: a local ensemble transform Kalman filter. *Phys. D* **230**, 112–126.

Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Trans. ASME J. Basic Eng. D* **82**, 35–45.

Kalnay, E., Li, H., Miyoshi, T., Yang, S. and Ballabrera-Poy, J. 2007. 4D-Var or ensemble Kalman filter? *Tellus A* **59**, 758–773.

Kalnay, E. and Yang, S. 2010. Accelerating the spin-up of ensemble Kalman filtering. *Q. J. Roy. Meteorol. Soc.* **136B**, 1644–1651.

Lawson, W. G. and Hansen, J. A. 2004. Implications of stochastic and deterministic filters as ensemble-based data assimilation methods in varying regimes of error growth. Mon. *Weather Rev.* **132**, 1966–1981.

Leeuwenburgh, O., Evensen, G. and Bertino, L. 2005. The impact of ensemble filter definition on the assimilation of temperature profiles in the tropical Pacific. *Q. J. Roy. Meteorol. Soc.* **131**, 3291–3300.

Livings, D. M., Dance, S. L. and Nichols, N. K. 2008. Unbiased ensemble square root filters. *Phys. D* **237**, 1021–1028.

Lorenz, E. N. 1963. Deterministic non-periodic flow. *J. Atmos. Sci.* **20**, 130–141.

Lorenz, E. N. 1996. Predictability: A problem partly solved. *Proc. Semina ron Predictability*. vol. 1, ECMWF, Reading, Berkshire, UK, 1–18.

Lorenz, E. N. and Emanuel, K. A. 1998. Optimal sites for supplementary weather observations: simulations with a small model. *J. Atmos. Sci.* **55**, 399–414.

Miller, R. N., Ghil, M. and Gauthiez, F. 1994. Advanced data assimilation in strongly nonlinear dynamical systems. *J. Atmos. Sci.* **51**, 1037–1056.

Miyoshi, T. 2005. *Ensemble Kalman Filter Experiments with a Primitive-equation Global Model*. PhD Thesis, University of Maryland, College Park, 226 pp.

Miyoshi, T. 2011. The Gaussian approach to adaptive co-variance inflation and its implementation with the local ensemble transform Kalman filter. *Mon. Weather Rev.* **139**, 1519–1535.

Molteni, F. 2003. Atmospheric simulations using a GCM with simplified physical parametrizations. I: Model climatology and variability in multi-decadal experiments. *Clim. Dyn.* **20**, 175–191.

Ott, E., Hunt, B. R., Szunyogh, I., Zimin, A. V., Kostelich, E. J. and co-authors. 2004. A local ensemble Kalman filter for atmospheric data assimilation. *Tellus A* **56**, 415–428.

Sakov, P. and Oke, P. R. 2008. Implications of the form of the ensemble transformation in the ensemble square root filter. *Mon. Weather. Rev.* **136**, 1042–1053.

Tippett, M. K., Anderson, J. L., Bishop, C. H., Hamill, T. M. and Whitaker, J. S. 2003. Ensemble square-root filters. *Mon. Weather Rev.* **131**, 1485–1490.

Toth, Z. and Kalnay, E. 1997. Ensemble forecasting at NCEP: the breeding method. *Mon. Weather Rev.* **125**, 3297–3318.

Wang, X., Bishop, C. H. and Julier, S. 2004. Which is better, and ensemble of positive–negative pairs or a centered spherical ensemble? *Mon. Weather Rev.* **132**, 1590–1605.

Whitaker, J. S. and Hamill, T. M. 2002. Ensemble data assimilation without perturbed observations. *Mon. Weather Rev.* **130**, 1913–1924.

Yang, S., Kalnay, E., Hunt, B. R. and Bowler, N. E. 2009. Weight interpolation for efficient data assimilation with the local ensemble transform Kalman filter. *Q. J. Roy. Meteorol. Soc.* **135**, 251–262.

Yang, S., Kalnay, E. and Hunt, B. R. In press. Handling nonlinearity in Ensemble Kalman Filter: Experiments with the three-variable Lorenz model. *Mon. Weather. Rev.*

8. Appendix

A.1. *S-ETKF and a NS-ETKF*

In this work, we use an unbiased (Livings et al., 2008) post-multiplicative EnSRF known as LETKF (Hunt et al., 2007), equivalent to the spherical-simplex ETKF (Wang et al., 2004) when no localisation is needed. Let $\mathbf{x} \in \Re^{\mathbf{N}}$ be a vector of state variables and $\mathbf{y} \in \Re^{\mathbf{L}}$ be a vector of observations; they are related through the observation equation:

$$\mathbf{y} = \mathbf{Hx} + \mathbf{v}, \tag{5}$$

where $\mathbf{H} \in \Re^{L \times N}$ is the observation matrix and $\mathbf{v} \in \Re^{L}$ represents the observational error which is assumed

$\sim N(\mathbf{0},\ \mathbf{R})$, and $\mathbf{R}$ is usually assumed to diagonal. An M-member ensemble is written as:

$$\overline{\overline{\mathbf{X}}} = [\mathbf{x}_1|\mathbf{x}_2|\cdots|\mathbf{x}_M] \in \Re^{N \times M}. \qquad (6)$$

The sample mean can be computed as:

$$\bar{\mathbf{x}} = \frac{1}{M}\sum_{m=1}^{M}\mathbf{x}_m \in \Re^{N}. \qquad (7)$$

The ensemble of perturbations can then be written as:

$$\mathbf{X} = [\mathbf{x}_1 - \bar{\mathbf{x}}|\mathbf{x}_2 - \bar{\mathbf{x}}|\cdots|\mathbf{x}_M - \bar{\mathbf{x}}] \in \Re^{N \times M}. \qquad (8)$$

And the sample covariance can be calculated as:

$$\mathbf{P} = (M-1)^{-1}\mathbf{X}\mathbf{X}^T \in \Re^{N \times N}. \qquad (9)$$

The S-ETKF obtains the analysis ensemble of perturbations $\mathbf{X}^a$ by a post-multiplication of the background ensemble of perturbations:

$$\mathbf{X}^a = \mathbf{X}^b\mathbf{W}^a. \qquad (10)$$

The matrix of weights $\mathbf{W}^a$ has to be obtained in a way such that $\mathbf{P}^a$ has the value prescribed by the KF. In particular, for the S-ETKF:

$$\mathbf{W}^a = \mathbf{C}(\mathbf{I}+\Gamma)^{-1/2}\mathbf{C}^{\mathbf{T}} = \left((M-1)\tilde{\mathbf{P}}^a\right)^{1/2}, \qquad (11)$$

where

$$\mathbf{C}\Gamma\mathbf{C}^{\mathbf{T}} = eig\left(\frac{\left(\mathbf{H}\mathbf{X}^{\mathbf{b}}\right)^{\mathbf{T}}\mathbf{R}^{-1}\left(\mathbf{H}\mathbf{X}^{\mathbf{b}}\right)}{M-1}\right), \qquad (12)$$

i.e. Γ is the matrix containing the eigenvalues of the multidimensional ratio of ensemble covariance (projected into observational space) and observational error covariance, while $\mathbf{C}$ is the matrix with the corresponding eigenvectors as columns. As indicated in eq. (11), $\mathbf{W}^a$ is proportional to the symmetric square root of the analysis covariance in ensemble space $\tilde{\mathbf{P}}^a$. This solution minimises the 'distance' between $\mathbf{W}^a$ and the identity matrix, thus getting an $\mathbf{X}^a$ as close as possible as $\mathbf{X}^b$ (Ott et al., 2004).

The factor $\mathbf{C}(\mathbf{I}+\Gamma)^{-1/2}$ in $\mathbf{W}^a$ is enough to guarantee the fulfillment of the KF covariance equation. As a matter of fact, the original (one-sided) ETKF (Bishop et al., 2001) used the matrix of weights as $\mathbf{W}^a = \mathbf{C}(\mathbf{I}+\Gamma)^{-1/2}$. Nonetheless, besides fulfilling the Kalman equation for covariance, the analysis perturbations must be unbiased, i.e.

$$\mathbf{X}^a\mathbf{1} = \mathbf{0}, \qquad (13)$$

where $\mathbf{1} \in \Re^M$. The one-sided ETKF in general does not produce analysis perturbations centered in zero, while the symmetric solution $\mathbf{W}^a = \mathbf{C}(\mathbf{I}+\Gamma)^{-1/2}\mathbf{C}^{\mathbf{T}}$ (used in the S-ETKF) does (Wang et al., 2004; Hunt et al., 2007)

(spherical-simplex ETKF: Wang et al., 2004; LETKF: Hunt et al., 2007). This formulation fulfills the covariance equation since $\mathbf{C}$ is orthonormal, i.e. $\mathbf{C}^{\mathbf{T}}\mathbf{C}=\mathbf{I}$

A general non-symmetric ETKF can be written as

$$\mathbf{W}^a = \mathbf{C}(\mathbf{I}+\Gamma)^{-1/2}\mathbf{S}^{\mathbf{T}}, \qquad (14)$$

where $\mathbf{S}$ can be any orthonormal matrix, but a non-symmetric solution will be unbiased if $\mathbf{S}$ is such that $\mathbf{W}$ contains $\mathbf{1}$ as an eigenvector. Instructions to construct this matrix are listed next:

Generate a matrix with random entries $\overline{\overline{G}} = [g_1, g_2, \cdots, g_M] \in \Re^{M \times M}$.

(1) Compute a matrix of perturbations $G = [g_1 - \bar{g}, g_2 - \bar{g}, \cdots g_M - \bar{g}]$, where $\bar{g} = M^{-1}\overline{\overline{G}}\mathbf{1}$. By construction $\mathbf{G1}=\mathbf{0}$.

(2) Perform the eigenvalue decomposition of the matrix $\mathbf{G^TG}=\mathbf{S\Lambda S^T}$. Since $\mathbf{G^TG}$ is symmetric (and therefore normal), $\mathbf{S}$ has orthonormal columns, i.e. $\mathbf{S^TS}=\mathbf{SS^T}=\mathbf{I}$. Moreover, all the eigenvalues in Λ are non-negative.

(3) Sort the eigenvalues by magnitude, and order the eigenvectors in $\mathbf{S}$ correspondingly. Since $\mathbf{G} \in \Re^{M \times M}$ and $\mathbf{G1}=\mathbf{0}$, $\lambda_M = 0$ and $\mathbf{S}_M = M^{-1/2}\mathbf{1}$. The elements of Λ and $\mathbf{S}$ are $\Lambda = diag([\lambda_1, \lambda_2, \ldots, \lambda_{M-1}, 0])$, $\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_{M-1}, M^{-1/2}\mathbf{1}]$.

A detailed proof of unbiasedness of this scheme both for $N > M$ (the usual case in applications) and $M > N$ (an unusual case, but the one we have used in the simple univariate quadratic model in this work) can be found in Chapter 3 of Amezcua (2012). We denominate a solution of this form a NS-ETKF.

A note on R-localisation (Hunt et al., 2007) is important. In this scheme, the filter independently calculates a local matrix of weights $\mathbf{w}_{\text{local}}^a \in \Re^{M \times M}$ for each gridpoint; it implies constructing $\mathbf{X}^a$ by sets of rows at a time, the size of each set corresponding to the number of variables in every gridpoint. For stability, it is important that the analyses obtained in neighbouring grid points vary smoothly. This is guaranteed automatically by S-LETKF (see the explanation in Hunt et al., 2007, as LETKF), but not by NS-LETKF. Hence, this method cannot be applied *directly* with R-localisation. Nonetheless, once $\mathbf{X}^a$ is completely calculated from local symmetric analyses (i.e. using S-LETKF), it can be globally rotated. This version of the NS-LETKF has no problems of divergence and can benefit from the ensemble resampling but requires an extra step.