Accelerating the spin-up of Ensemble Kalman Filtering

Eugenia Kalnay* and Shu-Chih Yang
University of Maryland

Revised 26 August 2009

Abstract

Ensemble Kalman Filter (EnKF) may have a longer spin-up time to reach its asymptotic level of accuracy than the corresponding spin-up time in variational methods (3D-Var or 4D-Var). During the spin-up EnKF has to fulfill two independent requirements, namely that the ensemble mean be close to the true state, and that the ensemble perturbations represent the "errors of the day". As a result, there are cases, such as radar observations of a severe storm, or regional forecast of a hurricane, where EnKF may spin-up too slowly to be useful. A scheme is proposed to accelerate the spin-up of EnKF by applying a no-cost Ensemble Kalman Smoother, and using the observations more than once in each assimilation window during spin-up in order to maximize the initial extraction of information. The performance of this scheme is tested with the Local Ensemble Transform Kalman Filter (LETKF) implemented in a Quasi-geostrophic model, which requires a very long spin-up time when initialized from random initial perturbations from a uniform distribution. Results show that with the new "running in place" (RIP) scheme the LETKF spins-up and converges to the optimal level of error faster than 3D-Var or 4D-Var even in the absence of any prior information. Additional computations (2-12 iterations for each window) are only required during the initial spin-up, since the scheme naturally returns to the original LETKF after spin-up is achieved. RIP also accelerates spin-up when the initial perturbations are drawn from a well-tuned 3D-Var background error covariance, rather than being uniform noise, but the maximum number of iterations and RIP cycles required is reduced compared to the case without such prior information.

1. Introduction

The relative advantages and disadvantages of 4-dimensional Variational Data Assimilation (4D-Var), already operational in several numerical forecasting centers, and Ensemble Kalman Filter (EnKF), a newer approach that does not require the adjoint of the model, are the focus of considerable current research (e.g., Lorenc, 2003, Kalnay et al, 2007a, Gustafsson, 2007, Kalnay et al., 2007b, Miyoshi and Yamane, 2007, WWRP/THORPEX Workshop, 2008).

One area where 4D-Var may have an advantage over EnKF is in the initial spin-up, since there is evidence that 4D-Var converges faster than EnKF to its asymptotic level of accuracy. For example, Caya et al. (2005) compared 4D-Var and EnKF for a storm

* Corresponding author: 3431 Computer and Space Sciences Bldg., College Park, MD, 20742-2425. ekalnay@atmos.umd.edu

simulating the development in a sounding corresponding to 00UTC 25 May 1999. They found that "Overall, both assimilation schemes perform well and are able to recover the supercell with comparable accuracy, given radial-velocity and reflectivity observations where rain was present. 4DVAR produces generally better analyses than the EnKF given observations limited to a period of 10 min (or three volume scans), particularly for the wind components. In contrast, the EnKF typically produces better analyses than 4DVAR after several assimilation cycles, especially for model variables not functionally related to the observations." In other words, for the severe storm problem the EnKF eventually yields better results than 4D-Var, presumably because of the assumptions made in the 4D-Var background error covariance, but during the crucial *initial time of storm development*, when radar data starts to become available, EnKF provides a worse analysis. For a global shallow water model, which is only mildly chaotic, Zupanski et al. (2006) found that initial perturbations that had horizontally correlated errors converged faster and to a lower level of error than perturbations created with white noise. In agreement with these results, Liu (2007) found using the SPEEDY global primitive equations model that perturbations obtained from differences between randomly chosen states (which are naturally balanced and have horizontal correlations of the order of the Rossby radius of deformation) spun-up faster than white noise perturbations.

Yang et al (2009a) compared 4D-Var and the Local Ensemble Transform Kalman Filter (LETKF, Hunt et al., 2007) within a quasi-geostrophic channel model (Rotunno and Bao, 1996). They found that if the LETKF is initialized from randomly chosen fields with uniform distribution perturbations, it takes more than 100 days before it converges to the optimal level of error. If, on the other hand, the ensemble mean is initialized from an existent 3D-Var analysis, which is already close to the true state, using the same random perturbations, the LETKF converges to its optimal level very quickly, within about 3-5 days. However, with a well-tuned background error covariance, 3D-Var and 4D-Var converge fast without needing a good initial guess. This has also been observed for severe storm simulations (Caya et al., 2005), especially when using real radar observations (Jidong Gao, 2008, personal communication). It is not surprising that EnKF spins-up more slowly than 3D-Var or 4D-Var because in order to be optimal the ensemble has to satisfy two independent requirements, namely that the mean be close to the true state of the system, and that the ensemble perturbations represent the characteristics of the "errors of the day" in order to estimate the evolving background error covariance $\mathbf{B}$. In both 3D-Var and 4D-Var, by contrast, $\mathbf{B}$ is tuned and assumed to be constant.

Within a global operational system it is feasible to initialize the EnKF from a state close enough to the optimal analysis, such an existent 3D-Var analysis, with balanced perturbations drawn from a 3D-Var error covariance, so that spin-up may not a serious problem. However, there are other situations, such as the storm development discussed above, where radar information is not available before the storm starts, so that no information is available to guide the EnKF in the spin-up towards the optimal analysis. The system may start from an unperturbed state without precipitation, and if a severe storm develops within a few minutes and the EnKF takes considerable real time to spin-up from the observations, it will not give reliable forecasts until later in the storm evolution, and thus give results that are less useful for severe storm forecasting than 4D-Var or even 3D-Var. Similarly, a regional model initialized from a global analysis at lower resolution may take too long to spin-up when confronted with mesoscale observations.

In this note we propose a new method to accelerate the spin-up of the EnKF by "running in place" (RIP) during the spin-up phase and using the observations more than once in order to extract maximum initial information. We find that it is possible to accelerate the convergence of the EnKF so that (in real time) it spins-up even faster than 3D or 4D-Var without losing accuracy after spin-up and without requiring prior information such as the initial background error covariance. Section 2 contains a brief theoretical motivation and discussion of the method, results are presented in Section 3 and a discussion is given in Section 4.

2. Spin-up, no-cost smoothing and "running in place" in EnKF

In this section we briefly review the conditions that justify the rule that in Kalman Filter data should be used once and then discarded. We then suggest that this rule is not strictly valid during spin-up, when the initial covariance is still influencing the results, or when the statistics of the "errors of the day" suddenly change, as during the initial development of a severe storm. During these transition periods, the ensemble perturbations are not representative of the "errors of the day" and extracting information from observations using them only once is not efficient.

Hunt et al. (2007) provide a derivation of the linear Kalman Filter equations by showing that in the cost function

$$J(\mathbf{x}) = \left[\mathbf{x} - \overline{\mathbf{x}}_n^b\right]^T \left(\mathbf{P}_n^b\right)^{-1} \left[\mathbf{x} - \overline{\mathbf{x}}_n^b\right] + \left[\mathbf{y}_n^o - \mathbf{H}_n \mathbf{x}\right]^T \left(\mathbf{R}_n^{-1}\right) \left[\mathbf{y}_n^o - \mathbf{H}_n \mathbf{x}\right], \qquad (1)$$

the background term represents the Gaussian distribution of a state with the maximum likelihood trajectory (history), i.e., the analysis/forecast trajectory that best fits the data from $t = t_1, \ldots, t_{n-1}$. This state is obtained by using the forecast model $\mathbf{M}_{t_{n-1}, t_n}$ to advance the previous maximum likelihood analysis $\overline{\mathbf{x}}_{n-1}^a$ and the corresponding analysis error covariance $\mathbf{P}_{n-1}^a$ to the new analysis time t_n. Taking logarithms of the Gaussian distribution this means that for some constant *c*,

$$\sum_{j=1}^{n-1} \left[\mathbf{y}_j^o - \mathbf{H}_j \mathbf{M}_{t_n, t_j} \mathbf{x}\right]^T \mathbf{R}_j^{-1} \left[\mathbf{y}_j^o - \mathbf{H}_j \mathbf{M}_{t_n, t_j} \mathbf{x}\right] = \left[\mathbf{x} - \overline{\mathbf{x}}_n^b\right]^T \left(\mathbf{P}_n^b\right)^{-1} \left[\mathbf{x} - \overline{\mathbf{x}}_n^b\right] + c \qquad (2)$$

After the cost function in (1) is minimized finding the analysis $\overline{\mathbf{x}}_n^a$ and its corresponding covariance $\mathbf{P}_n^a$, a similar relationship holds for the analysis at t_n, and another constant *c*':

$$\left[\mathbf{x}-\overline{\mathbf{x}}_n^b\right]^T\left(\mathbf{P}_n^b\right)^{-1}\left[\mathbf{x}-\overline{\mathbf{x}}_n^b\right]+\left[\mathbf{y}_n^o-\mathbf{H}_n\mathbf{x}\right]^T\left(\mathbf{R}_n^{-1}\right)\left[\mathbf{y}_n^o-\mathbf{H}_n\mathbf{x}\right]=\left[\mathbf{x}-\overline{\mathbf{x}}_n^a\right]^T\left(\mathbf{P}_n^a\right)^{-1}\left[\mathbf{x}-\overline{\mathbf{x}}_n^a\right]+c' \quad (3)$$

Equating the terms in (3) that are linear and quadratic in $\mathbf{x}$, the linear Kalman Filter equations for a perfect model are obtained.

This derivation makes clear that Kalman Filter yields the maximum likelihood estimate $\overline{\mathbf{x}}_n^a$ with the corresponding error covariance $\mathbf{P}_n^a$ at time t_n if the model is linear and perfect and if the previous analysis $\overline{\mathbf{x}}_{n-1}^a$ at t_{n-1} is also the maximum likelihood state estimate at the previous analysis time. Hunt et al. (2007) also address the problem of initialization: a system can be initialized assuming a prior background distribution at the initial time t_0 such that the initial background error covariance $\mathbf{P}_0^b$ is large but not infinitely large. This introduces into the cost function an additional quadratic term, but Hunt et al. (2007) point out that "with sufficient observations over time, the effect of this term [on the background error covariance] at time t_n decreases in significance as n increases". In other words, with sufficient observations, the Kalman Filter spins-up and eventually converges and yields the maximum likelihood solution and its error covariance independently from the initial conditions. During spin-up, however, or when the statistical properties of the dynamical system suddenly change, the background may be a very unlikely state, and it may be desirable to use the observations more than once in order to extract maximum information from them.

The EnKF, like the Kalman Filter, also provides a maximum likelihood analysis, except that the background and analysis error covariances are estimated from an ensemble of K generally nonlinear forecasts:

$$\mathbf{P}_n^b \approx \frac{1}{K-1}\mathbf{X}_n^{bT}\mathbf{X}_n^b, \quad (4)$$

where $\mathbf{X}_n^b$ is a perturbation matrix whose k^{th} column is the background (forecast) perturbation $\mathbf{x}_{n,k}^b-\overline{\mathbf{x}}_n^b$ and $\overline{\mathbf{x}}_n^b=\frac{1}{K}\sum_{k=1}^{K}\mathbf{x}_n^b$ is the most likely forecast state, i.e., the ensemble average. Similar equations are valid for the analysis mean $\overline{\mathbf{x}}_n^a$ and the analysis error covariance $\mathbf{P}_n^a$. Thus, EnKF, like the original Kalman Filter, is a sequential data assimilation system where, after the new data is used at the analysis time it should be discarded (Ide et al., 1997), but this is true only if the previous analysis and the new background are not only the most likely states given the past observations, but they also have already "forgotten" the choice of initial ensemble. In other words, *if the system has converged after the initial spin-up, all the information from past observations is already included in the background* and the data can be discarded after the new analysis is computed. In contrast, 4D-Var is a smoother that best fits all the observations (even asynoptic data) within an assimilation window. We note that EnKF can be extended to 4-dimensions as in 4D-Var, allowing for the assimilation at the right time of asynoptic observations made between two analyses (e.g., Hunt et al., 2004, 2007), but, being a filter, the EnKF analysis is only obtained at the end of the assimilation window.

In summary, after the initial spin-up, all the information from past observations is already included in the background field, so that the observations should be used only once and then discarded. However, there is no theoretical reason why this constraint should also be applied when EnKF is "cold-started", and the initial ensemble is not representative of the most likely state and its uncertainty, since during spin-up the background term still "remembers" the arbitrarily chosen initial ensemble. In practical applications, the rule of using the data only once is usually applied even during spin-up (e.g., Zupanski et al. 2006), and depending on the initial ensemble, a slow EnKF spin-up can then be observed.

In this note we suggest that when a quick EnKF spin-up (in real time) is needed in order to make useful short-range forecasts for fast weather instabilities, the initial observations can be used more than once in order to extract more initial information from them, and that this procedure leads to a much faster spin-up of the initial ensemble. This "running in place" (RIP) algorithm is made possible by the use of a "no-cost" Ensemble Kalman Smoother (EnKS) (Kalnay et al., 2007b, Yang et al., 2009a).

The no-cost EnKS is based on the property that in EnKF the ensemble analysis is computed as a linear combination (weighted average) of the ensemble forecasts at the end of the assimilation window (see schematic Figure 1). Since a linear combination of ensemble trajectories within an assimilation window is also a model trajectory, a linear combination that is close to the truth at one time within the window should remain close to the truth over the entire window (at least as close as model errors allow). This argument (B. Hunt, personal comm., 2009) indicates that the weights used in constructing the analysis ensemble mean, although determined at the end of the assimilation window, should be valid throughout the window. A similar argument suggests that the ensemble analysis perturbation weights obtained using Bayes' theorem are also valid throughout the assimilation window $\left[t_{n-1}, t_n\right]$ (Hunt et al., 2007).

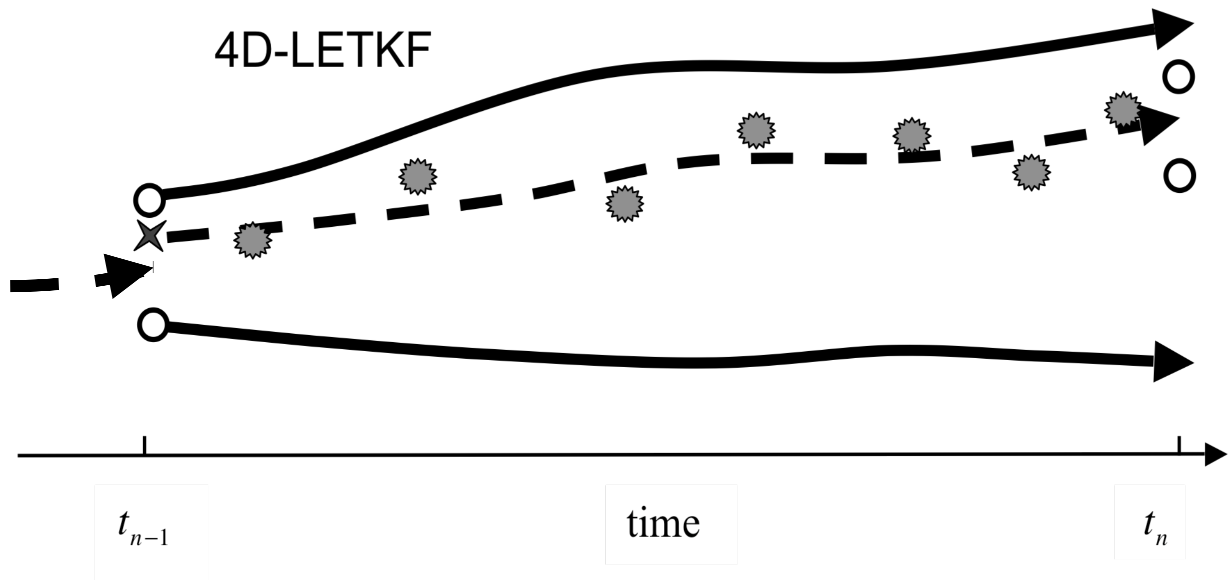

Figure 1: Schematic of an EnKF assimilation window with a two-member ensemble of forecasts (full lines) started from the analysis ensemble (circles). An arrow points to the analysis ensemble mean, and the dashed line is the model trajectory that goes through the analysis mean at the end of the trajectory. The cross is the analysis smoother valid at the beginning of the window, obtained applying the same weights to the forecasts as obtained at the end of the assimilation window. The smoother is more accurate than the analysis mean because it has information about the "future" observations within the assimilation window.

The no-cost EnKS is easy to implement if the weights that transform the ensemble forecasts into the ensemble analysis are explicitly computed and available, as is the case in the LETKF. The analysis ensemble members at time t_n are each a weighted average (linear combination) of the ensemble forecasts valid at t_n (Hunt et al., 2007). Since the ensemble analysis estimates the linear combination of the trajectories that best fits the observations within an assimilation window, not just at the end of the interval, the no-cost EnKS valid at the beginning of the window is obtained by simply applying the same weights obtained at analysis time t_n to the initial ensemble at t_{n-1}.

The no-cost EnKS was tested by Yang et al. (2009a) on the quasi-geostrophic model of Rotunno and Bao (1996). Figure 2 compares the analysis error of the LETKF with that obtained using the no-cost EnKF, and shows that indeed, the no-cost ensemble Kalman smoother at t_{n-1} is more accurate than the analysis ensemble valid at t_{n-1}, as could be expected from the fact that the smoothed ensemble at the beginning of the window has benefited from the information provided by the "future" observations within the window $\left[t_{n-1}, t_n\right]$. Although the no-cost smoothing improves the accuracy of the initial analysis at t_{n-1}, it does not improve the final analysis at t_n, since the forecasts started from the new initial analysis ensemble will end as the final analysis ensemble (at least in a linear sense, see Figure 1).

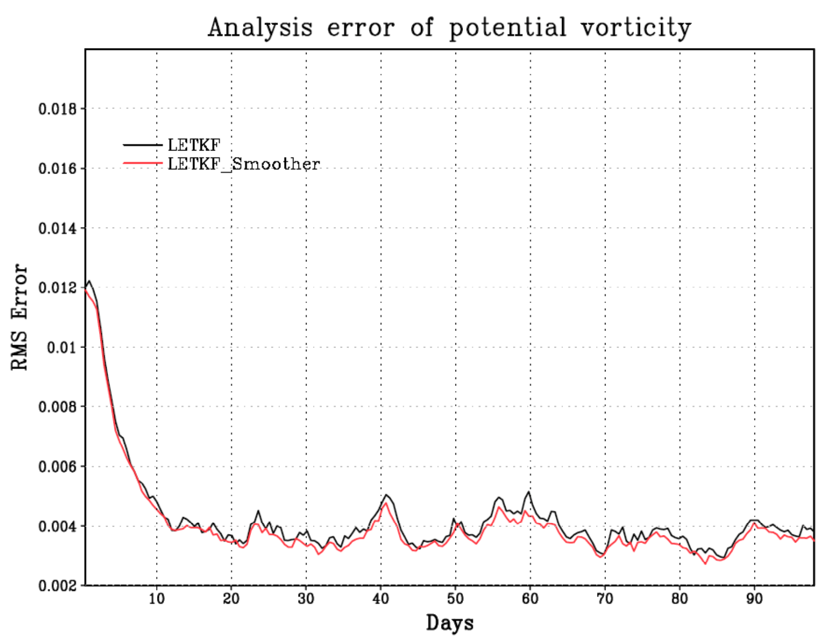

Figure 2: Comparison of the RMS domain averaged error of the LETKF analysis (black line) computed at the end of each assimilation window, and the no-cost LETKF smoother (red line) computed by combining the forecasts valid at the beginning of the window with weights obtained at the end of the window.

With the no-cost EnKS it is thus possible to go backwards in time within an assimilation window, and then advance with the regular EnKF using the initial observations repeatedly in order to extract maximum information from them. During spin-up this improves the quality (likelihood) of the initial ensemble mean faster, and leads the ensemble-based background error covariance to be more representative of the true forecast error statistics.

As indicated above, the EnKF requires the choice of an initial prior ensemble at t_0 with covariance $\mathbf{P}_0^b$. We have tested the RIP algorithm using three different initial ensembles, all with the same randomly chosen ensemble mean but with different distributions of the random perturbations: 1) a uniform distribution; 2) a Gaussian distribution and 3) perturbations drawn from a carefully optimized 3D-Var error covariance. Cases 1) and 2) include no prior information, case 3) contains the same prior information used for 3D-Var, and an optimal rescaling, in 4D-Var.

The RIP algorithm that we have tested (not necessarily optimal) is as follows: at t_0 we integrate the initial ensemble to t_1. Then the RIP loop with $n=1$, is:

a) Perform a standard EnKF analysis and obtain the analysis weights at t_n, saving the mean square observations minus forecast, $OMF^2 = \left[\mathbf{y}_n^o - \mathbf{H}_n\mathbf{x}\right]^T \left(\mathbf{R}_n^{-1}\right)\left[\mathbf{y}_n^o - \mathbf{H}_n\mathbf{x}\right]$, computed by the EnKF.

b) Apply the no-cost smoother to obtain the smoothed analysis ensemble at t_{n-1} by using the same forecast weights obtained at t_n.

c) Perturb the smoothed analysis ensemble with a small amount of random perturbations, a method similar to additive inflation. These added perturbations have two purposes: they avoid the problem of otherwise reaching the same final analysis at t_n as in the previous iteration (Figure 1), and they allow the ensemble perturbations to evolve into fast growing directions that may not have been included in the unperturbed ensemble subspace[1].

d) Integrate the perturbed smoothed ensemble to t_n. If the forecast fit to the observations is smaller than in the previous iteration according to a criterion such as

$$\frac{OMF^2(iter) - OMF^2(iter+1)}{OMF^2(iter)} > \varepsilon, \qquad (5)$$

go to a) and perform another iteration. If not, let $t_{n-1} \leftarrow t_n$ and proceed to the next assimilation window. In the results presented here we have used $\varepsilon = 0.05$ as the criterion for relative improvement.

---

[1] We have tested the use of additive perturbations with a Gaussian and with a uniform distribution, and both worked well. The results presented here have Gaussian additive perturbations.

e) If no additional iteration beyond the first one is needed, the RIP analysis is the same as the standard EnKF. When the system converges, no additional iterations are needed, so that if several assimilation cycles take place without invoking a second iteration, the RIP can be switched off and the system returns to a normal EnKF. In the results presented here we switched off RIP after 5 cycles without invoking a second iteration.

3. Results

The LETKF with the RIP method was implemented in the Rotunno and Bao (1996) quasigeostrophic model. The data assimilation experiments are performed with a 6-hour analysis cycle. The analysis is validated every 6-hour against the truth simulation, a long nature run of this QG model. The validation is done through the RMS analysis error, defined as the domain-averaged RMS difference of the model variables (potential vorticity and temperature) between the analysis and truth. The observations are zonal and meridional wind components and temperature, generated by adding random Gaussian errors on the truth. Details of the QG DA setup can be found in Yang et al. (2009a).

In this section we compare several data assimilation methods started from the same randomly chosen mean. We define the (real time) spin-up as the number of cycles it takes to reduce the RMS error in potential temperature, which starts from a nondimensional value of 0.76, to a level of 0.038 (i.e., 5% of the initial analysis error). The results, including both spin-up time and asymptotic level of analysis error are also summarized in Table 1.

Figures 3a, b and c show the RMS error of the analysis obtained during spin-up, using several methods over 200 analysis cycles of 12 hours each (corresponding to a total of 100 days). In Figure 3a we compare the number of cycles required for spin-up for the LETKF with initial random perturbations uniformly distributed, with and without using RIP (black), with 3D-Var and 4D-Var (grey). As indicated before, all the experiments started from the same a randomly chosen mean state. 3D-Var (dashed grey line) takes about 60 cycles to spin-up, and 4D-var (full grey line) takes about 80 cycles, but converges to a much lower RMS error than 3D-Var. The standard LETKF (full black line) using the observations once and discarding them takes much longer, a total of 170 cycles. It is interesting the LETKF devotes the first 120 cycles essentially to create ensemble perturbations representative of the “errors of the day”, with little reduction in the analysis mean error, and only then, between 120 and 170 cycles, the LETKF converges rather quickly to the optimal level of error. After spin-up, the LETKF and 4D-Var have a similar asymptotic RMS error but significant day-to-day differences.

	LETKF (1) Random Uniform Initial Ensemble		**LETKF** (2) Random Gaussian Initial Ensemble		**LETKF** (3) B3DVar Initial Ensemble		**LETKF** Random Initial Ensemble	**Variational**	
	No RIP	With RIP	No RIP	With RIP	No RIP	With RIP	Fixed 10 RIP iterations	3D-Var/ B3DVar	4D-Var/ 0.05B3 D-Var
Spin-up: DA cycles needed to reduce error to 5%	141	47	57	37	53	39	37	44	54
RMS error ($x10^{-2}$)	.51	.51/.53	0.51	.51/.56	.50	.50/.56	1.26	1.24	.53

Table 1: Comparison of the spin-up time (number of cycles to reduce the initial RMS error in potential temperature to 5% of the original value) and the asymptotic RMS error for LETKF ensembles with and without RIP, and fixing the number of RIP iterations to 10 rather than determining them adaptively. In the runs with RIP the first value of the RMS error (averaged over 120 cycles after spin-up) corresponds to the case RIP is switched off after 5 cycles with only one iteration (step e of the RIP algorithm). The second value is the average error if the RIP continues to be executed even after convergence, so that an occasional second iteration is executed and slightly degrades the results. Variational methods starting from the same initial state as the ensemble mean are also compared. The variational error covariances have been optimally tuned for both 3D and 4D-Var.

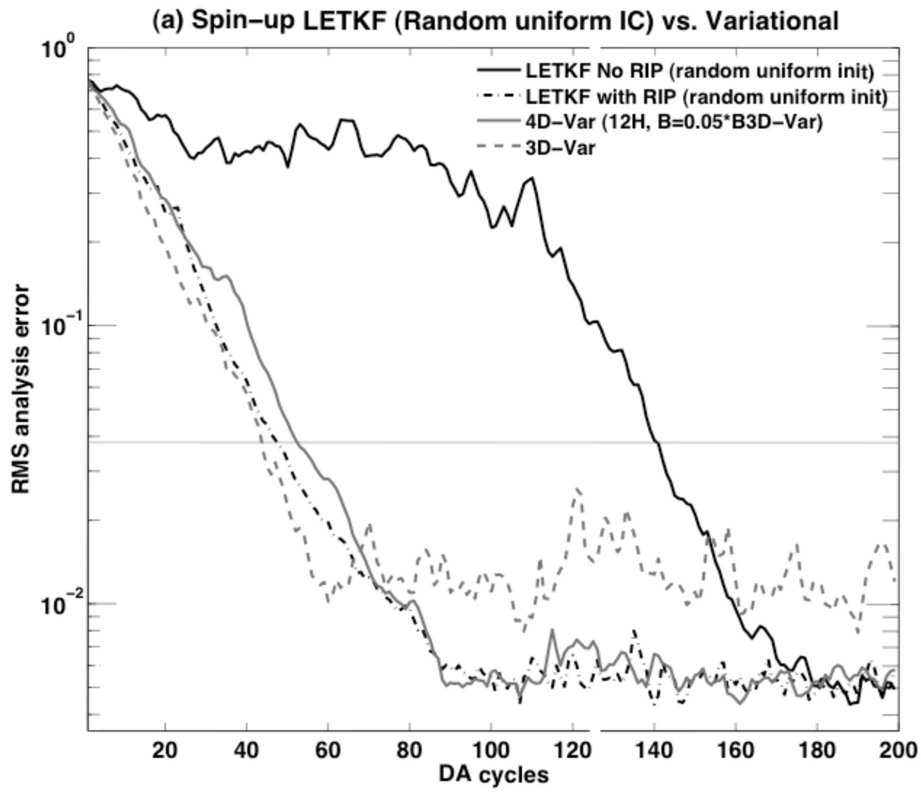
(a) Spin−up LETKF (Random uniform IC) vs. Variational
LETKF No RIP (random uniform init)
LETKF with RIP (random uniform init)
4D−Var (12H, B=0.05*B3D−Var)
3D−Var
RMS analysis error
DA cycles
20
40
60
80
100
120
140
160
180
200

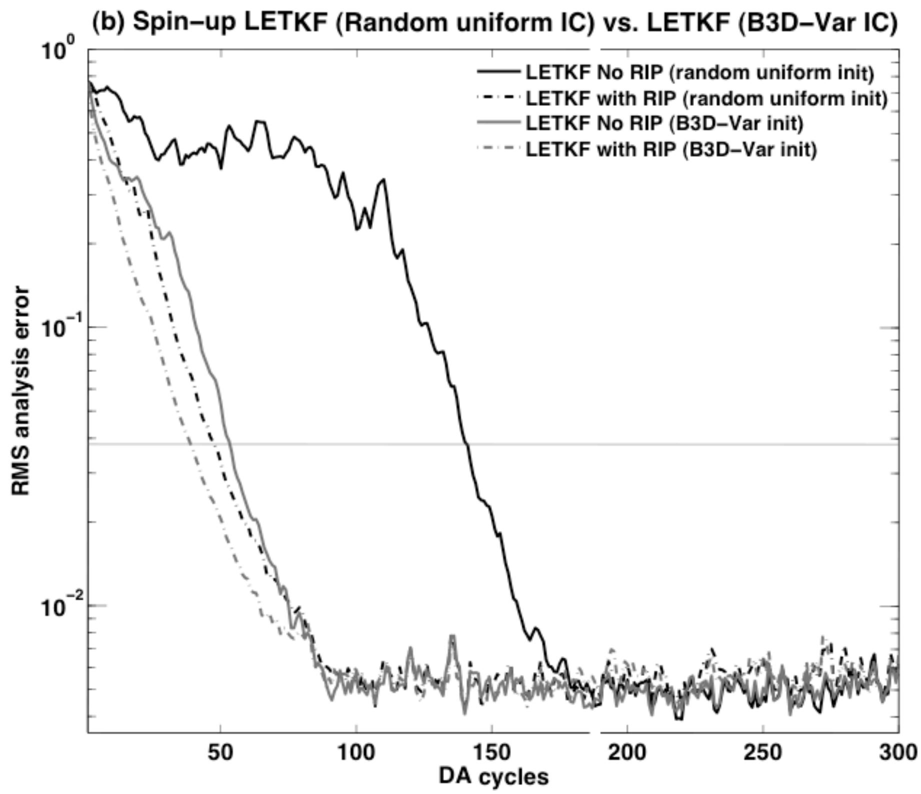
(b) Spin−up LETKF (Random uniform IC) vs. LETKF (B3D−Var IC)
LETKF No RIP (random uniform init)
LETKF with RIP (random uniform init)
LETKF No RIP (B3D−Var init)
LETKF with RIP (B3D−Var init)
RMS analysis error
DA cycles
10^0
10^{-1}
10^{-2}
50
100
150
200
250
300

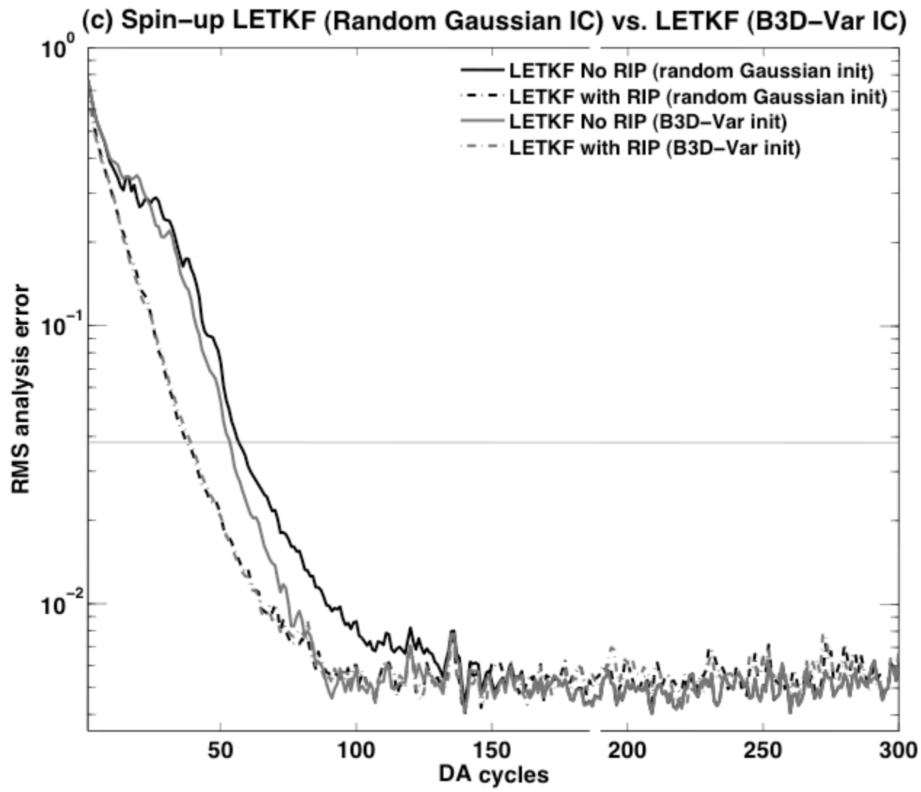

Figure 3: a) Comparison of the analysis RMS error in potential temperature versus the number of analysis cycles for the LETKF started from a random ensemble (1), i.e., perturbations from a uniform distribution (black full line), LETKF with RIP (black dot-dashed line), 4D-Var (grey full line) and 3D-Var (grey dot-dashed line), both with tuned initial error covariances. b) Comparison of the LETKF analysis errors starting from random uniform perturbations with and without RIP (black lines as in 3a), and LETKF starting from 3D-Var random perturbations with and without RIP (full and dot-dashed grey lines). c) Comparison of the LETKF analysis errors starting from random Gaussian perturbations with and without RIP (full and dot-dashed black lines), and LETKF starting from 3D-Var random perturbations with and without RIP (grey lines as in 3b).

Figure 3b compares the spin-up of the LETKF starting from a random uniform ensemble (Case 1), with and without RIP, as in Figure 3a, with the LETKF started from perturbations drawn from the 3D-Var error covariance (Case 3), i.e., where each ensemble perturbation is a column of the matrix $\sqrt{\mathbf{B}_{3D-Var}}\mathbf{E}$. Here $\mathbf{E}$ is an *MxK* matrix whose columns are random Gaussian numbers such that $\mathbf{E}\mathbf{E}^T \approx \mathbf{I}$, *M* is the dimension of the model and *K* the number of ensemble members. It is apparent from Figure 3b that, as suggested by both Anderson (2008, personal communication) and Zupanski et al. (2006), when the initial ensemble is drawn from the 3D-Var covariance matrix, the spin-up is much faster than when started from random, uniformly distributed perturbations. We can view Figure 3b as a comparison between the worst and best choices of the initial ensemble perturbations one could make, since in the first case we make no use of prior information and assume perturbations are non-Gaussian, and in the second we use best tuned 3D-Var prior information. Nevertheless it is remarkable that even in the case of faster spin-up, the application of the RIP algorithm is able to accelerate the spin-up even further.

Figure 3c compares the spin-up of the LETKF starting from 3D-Var covariance ensemble (Case 3) with and without RIP, as in Figure 3b, and Gaussian initial ensemble (Case 2), without any a-priori information. Initializing from perturbations with 3D-Var structures spins-up faster than using uncorrelated Gaussian perturbations, but such difference disappears when RIP is applied and similar results are obtained. Given that an optimally tuned 3D-Var covariance matrix may not be always available for ensemble-based data assimilation systems, the use of RIP appears to be an attractive alternative.

In an additional experiment in which the LETKF RIP algorithm was forced to always perform 10 iterations (not shown), the LETKF showed an even faster spin-down but it converged to a higher level of error, close to that of 3D-Var (Table 1). This is not surprising, since once the system is close to the maximum likelihood solution, as indicated by the theoretical arguments discussed above, observations should be used only once and then discarded. By performing 10 iterations even after the system spun-up, the EnKF analysis fits the data too closely and this increases the analysis errors.

Finally, Figure 4 shows the number of iterations needed to accelerate the spin-up in the RIP algorithm when started with random initial ensemble perturbations and with the 3D-Var initial perturbations. One iteration in the figure corresponds to the normal LETKF case, i.e., when a second iteration would give a relative improvement in the fit of the forecast to the observations of less than $\varepsilon = 0.05$ (equation 5), and thus it is not used. For the 3D-Var initial ensemble (Case 3) only 2-6 iterations are needed during the spin-up, but the other two ensembles without prior information need 11 iterations at cycle 19. The last second iteration is executed at cycle 65, 46 and 41 for ensembles (1), (2) and (3) respectively. After RIP is turned off, the analysis accuracy for the three ensembles is essentially identical (Table 1). We found that using a lower value of $\varepsilon = 0.01$ (not shown) leads to a faster initial reduction of errors but requires a large number of iterations.

Values of ε within a range of 0.02-0.05 gave optimal results, leading to a spin-down of the initial errors similar to 3D-Var and faster than 4D-Var, and converging to and error level at least as good as that of 4D-Var.

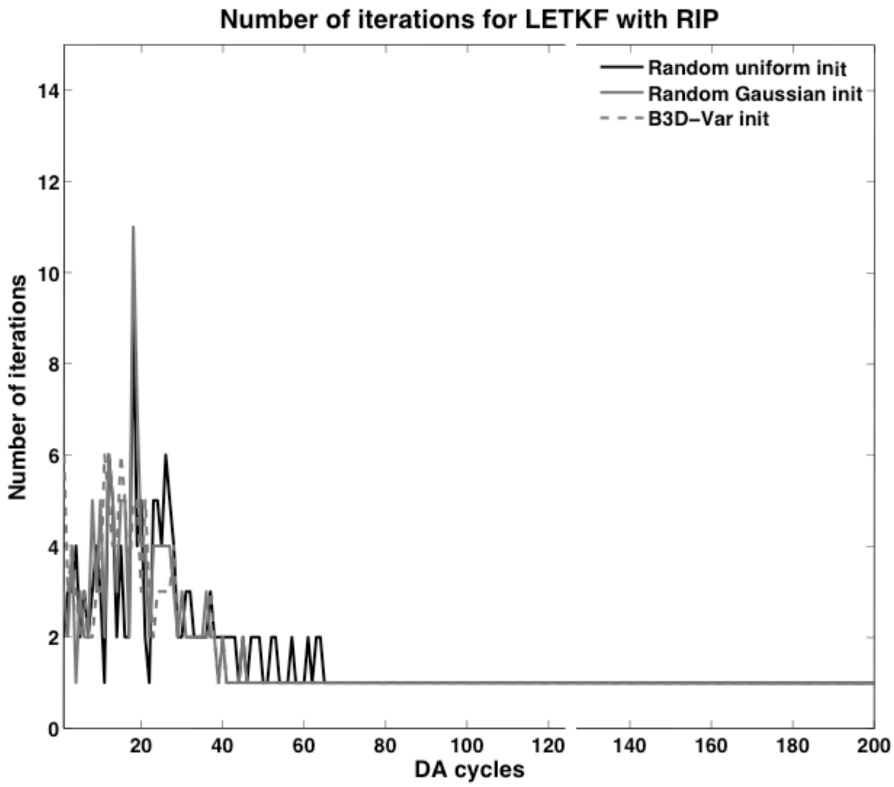

Figure 4: Number of iterations used by the RIP adaptive algorithm (5) with $\varepsilon = 0.05$, and one iteration corresponding to the regular LETKF system. After 5 cycles without requiring a second iteration RIP is switched off. Black: ensemble (1), uniform random initial ensemble perturbations. Blue: ensemble (2), Gaussian random initial ensemble perturbations. Red: ensemble (3), initial ensemble perturbations from the 3D-Var error covariance matrix.

4 Discussion

The "Running in Place" algorithm could be used to accelerate the spin-up of the EnKF whenever the background error statistics change suddenly, as in the case of a developing storm, or are otherwise not appropriate, as when a regional model with high resolution is started from initial conditions from a global model, or when no prior information is available to start the ensemble.

The results obtained with RIP are very encouraging: it is possible to significantly accelerate the spin-up of the LETKF (and other EnKF algorithms for which the weights of the ensemble forecasts are available) when fast convergence to the optimal level of error (in terms of real time) is required by simply using the initial observations several times rather than only once. The no-cost Ensemble Kalman Smoother, with the smoothed analysis ensemble at the beginning of an assimilation window given by using the analysis weights of the ensemble forecast at the end of the window enables this algorithm to extract the maximum information from the initial observations. It is necessary to add small perturbations to the ensemble, in a procedure akin to additive inflation. The number of iterations needed is estimated by checking whether the smoothed analysis reduces the forecast error, estimated from the innovation OMF. A level of relative reduction ε of about 2-5% was found to work well in this quasi-geostrophic model, leading to about 2-6 iterations during spin-up. After the system converges it naturally returns to the original LETKF.

Acknowledgements:

We are grateful to Jidong Gao who pointed out the slow spin-up of EnKF in severe storm prediction using radar data, and to Brian Hunt for explaining why forecast weights computed at the end of the LETKF assimilation window are valid throughout the window. Two reviewers of the original note made extremely constructive comments and one of them (Jeff Anderson) suggested starting the ensemble from perturbations drawn from the 3D-Var error covariance. This research was generously supported by NASA grants NNG06GB77G, NNX08AD90G and NNG06GB77G, and DOE grant DEFG0207ER64437.

5. References

Caya, C, J. Sun, and C. Snyder, 2005: A Comparison between the 4DVAR and the Ensemble Kalman Filter Techniques for Radar Data Assimilation, *Mon. Wea. Rev.*,

133, 3081-3094.
Hunt, B.R., E. Kalnay, E.J. Kostelich, E. Ott, D.J. Patil, T. Sauer, I. Szunyogh, J.A. Yorke, and A.V. Zimin, 2004: Four-Dimensional Ensemble Kalman Filtering. Tellus 56A, 273-277.
Hunt, B. R., E. Kostelich, I. Szunyogh, 2007: Efficient data assimilation for spatiotemporal chaos: a Local Ensemble Transform Kalman Filter. *Physica D,* **230***,* 112-126.
Kalnay, E., H. Li, T. Miyoshi, S.-C. Yang and J. Ballabrera, 2007a: 4D-Var or Ensemble Kalman Filter? *Tellus A*, **59**, 758–773.
--, 2007b: Response to Comments by N. Gustafsson. *Tellus A*, **59**, 778-780
Liu, J., 2008: Applications of the LETKF to adaptive observations, analysis sensitivity, observation impact and the assimilation of moisture. *PhD dissertation*, University of Maryland, 154 pages.
Lorenc, A. C., 2003: The potential of the ensemble Kalman filter for NWP – a comparison with 4D-Var. Quart. J. Roy. Meteor. Soc., 129, 3183-3203.
Ott, E., B. R. Hunt, I. Szunyogh, A. V. Zimin, E. J. Kostelich, M. Corazza, E. Kalnay, D. J. Patil, and J. A. Yorke, 2004: A local ensemble Kalman filter for atmospheric data assimilation. *Tellus*, **56A**, 415-428.
Parrish, D. and J. Derber, 1992: The National Meteorology Center's spectral statistical-interpolation analysis system. *Mon. Wea. Rev.*, **120**, 1747-1763.
Rotunno R, Bao JW, 1996: A case study of cyclogenesis using a model hierarchy. *Mon. Rev. Rev.*, **124**, 1051-1066.
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. M. Hamill, 2002: Ensemble data assimilation without perturbed observations, Mon. Wea. Rev. 130, 1913–1924.
WWRP/THORPEX Workshop on 4D-Var and Ensemble Kalman Filter Intercomparisons. Buenos Aires, 10-13 November 2008. Lectures and papers available at http://4dvarenkf.cima.fcen.uba.ar/
Yang, S-C, M. Corazza, A. Carrassi, E. Kalnay, and T. Miyoshi, 2009a: Comparison of ensemble-based and variational-based data assimilation schemes in a quasi-geostrophic model. *Mov. Wea. Rev.,* **137**, 639-709.
Yang, S-C, E. Kalnay, B. Hunt, N. Bowler, 2008: Weight interpolation for efficient data assimilation with the Local Ensemble Transform Kalman Filter, *Quart. J. Roy. Meteor. Soc.*, **135**, 252-262.
Zupanski, M., S. J. Fletcher, I. M. Navon, B. Uzunoglu, R. P. Heikes, D. A. Randall, T. D. Ringlee and D. Daescu, 2006: Initiation of ensemble data assimilation. Tellus, 58A, 159-170.