

Spatially and temporally varying adaptive covariance inflation for ensemble filters

By JEFFREY L. ANDERSON*, *NCAR Data Assimilation Research Section PO Box 3000 Boulder, CO 80307-3000, USA*

(Manuscript received 7 January 2008; in final form 16 July 2008)

ABSTRACT

Ensemble filters are used in many data assimilation applications in geophysics. Basic implementations of ensemble filters are trivial but are susceptible to errors from many sources. Model error, sampling error and fundamental inconsistencies between the filter assumptions and reality combine to produce assimilations that are suboptimal or suffer from filter divergence. Several auxiliary algorithms have been developed to help filters tolerate these errors. For instance, covariance inflation combats the tendency of ensembles to have insufficient variance by increasing the variance during the assimilation. The amount of inflation is usually determined by trial and error. It is possible, however, to design Bayesian algorithms that determine the inflation adaptively. A spatially and temporally varying adaptive inflation algorithm is described. A normally distributed inflation random variable is associated with each element of the model state vector. Adaptive inflation is demonstrated in two low-order model experiments. In the first, the dominant error source is small ensemble sampling error. In the second, the model error is dominant. The adaptive inflation assimilations have better mean and variance estimates than other inflation methods.

1. Introduction

Because of their ease of implementation and power, ensemble filters (Evensen, 2006) are being used for data assimilation in a growing number of geophysical applications (Kalnay et al., 2007). Ensemble filters use a sample of model state vectors to estimate the relation between observations and model state variables. Basic ensemble methods (Evensen, 1994) are trivial to implement but are subject to errors from many sources. Some of these, like model error (Buizza et al., 1999; Dee and Todling, 2000; Hansen, 2002), error in the description of the observing system (Eyre et al., 1993), representativeness error (Daley, 1993) and gross observation error (Kistler et al., 2001; Uppala et al., 2005), are shared with all data assimilation techniques. Sampling error due to a finite ensemble is unique to ensemble methods, whereas errors due to assumptions of linearity and Gaussianity (Harlim and Hunt, 2007) are shared by filters and most other assimilation methods.

Errors in ensemble filters have been addressed by a number of algorithmic enhancements, including covariance inflation (Anderson and Anderson, 1999) and localization (Mitchell and Houtekamer, 2000; Hamill et al., 2001; Houtekamer and Mitchell, 2001; Mitchell et al., 2002), to deal with sampling

error when computing the statistical relation between an observation and a state variable. Almost all sampling-error sources, along with model and observing system systematic errors, lead to insufficient variance in ensemble state estimates (Furrer and Bengtsson, 2007). Covariance inflation was originally developed to deal with the loss of variance due to sampling errors, but it can help filters tolerate all types of errors. Inflation simply increases the prior variance estimate to counter a spurious variance deficiency.

Covariance inflation has been effective for a variety of ensemble filter algorithms and geophysical applications, but it has several problems (Hamill and Whitaker, 2005). First, it can be costly to find inflation values that perform well for a given application. For large models with large observation sets, the cost of performing ‘tuning’ experiments to find appropriate inflation values can be prohibitive. Recently, ensemble algorithms that adaptively determine appropriate values of inflation have been developed (Anderson, 2007a, hereafter A07). These algorithms are related to work on correcting model bias in deterministic assimilation (Dee, 1995; Dee and da Silva, 1999; Dee et al., 1999). An alternative approach that deals with the loss of variance by adaptively modifying the ensemble size is described in Uzunoglu et al. (2007).

A more serious problem occurs when a single value of inflation is not appropriate for all state variables. Assimilation of in situ observations, like radiosonde and aircraft observations, in a global numerical weather prediction model provides an

*Correspondence.

e-mail: jla@ucar.edu

DOI: 10.1111/j.1600-0870.2008.00361.x

example. In densely observed regions like the upper troposphere over North America, ensemble variance can be inappropriately small due to model bias and sampling error. Inflation can reduce this problem. However, over the Southern Ocean, there are very few observations to constrain the model. Repeated application of inflation values large enough to correct problems over North America can systematically increase the variance of the ensemble over the Southern Ocean. Eventually, this can lead to values that are inconsistent with climatological values, and in the worst case, incompatible with the model's numerical methods. The result is ridiculous solutions, at best, and model failure, at worst.

The adaptive inflation algorithm developed here avoids this problem by associating a different inflation value with each model state-vector component. Inflation might be large in the upper troposphere over North America and small over the Southern Ocean. Heuristic tuning for a spatially varying inflation algorithm would be prohibitively expensive even in small model applications, so an adaptive algorithm is required to obtain appropriate values of the inflation as a function of time. This extends the method described in A07 which used a single temporally adaptive covariance inflation for all state variables.

Section 2 reviews ensemble filter methods and Section 3 discusses adaptive inflation. Section 4 presents results from a low-order model assimilation while Section 5 provides discussion and conclusions.

2. Overview of ensemble filtering

Ensemble filter data assimilation methods can be derived as Monte Carlo approximations to Bayes theorem under the assumption of Gaussian prior and likelihood distributions. Derivations of various ensemble filter algorithms can be found in Evensen (2006) and Anderson (2001). An ensemble filter represents the probability distribution of a model's state with a set of M -dimensional state vectors,

$$\mathbf{x}_n, \quad n = 1, \dots, N, \quad (1)$$

where N is the ensemble size and M is the model size.

If the observational error distributions associated with each observation are independent, implementations of many ensemble filters (Burgers et al., 1998; Pham, 2001; Whitaker and Hamill, 2002; Tippett et al., 2003; Ott et al., 2004) can be described as follows.

(1) Each ensemble member is advanced from the time of the most recently used observation to the time of the next observation using the model.

(2) An ensemble of prior estimates of the observation is computed by applying a forward observation operator, h , to each sample of the state vector.

(3) The prior ensemble estimate of the observation $y_n = h(\mathbf{x}_n)$, $n = 1, \dots, N$, the observed value y^o and the observational error variance σ_o^2 (a function of the observing system) are combined

using Bayes rule to compute increments Δy_n to the prior ensemble estimates. Different flavours of ensemble filters differ primarily in how this step is done (Anderson, 2003), but the details do not matter to the discussion here.

(4) Increments to the prior ensemble of each state-vector element are computed by linearly regressing the observation increments onto each state-vector component independently using the prior joint sample statistics. Increments for the m th state-vector component are

$$\Delta x_{m,n} = \left(\sigma_{p,m} / \sigma_p^2 \right) \Delta y_n, \quad m = 1, \dots, M, \quad n = 1, \dots, N, \quad (2)$$

where $\sigma_{p,m}$ is the prior sample covariance of the observed variable and the m th element of the state vector and σ_p^2 is the prior sample variance of the observed variable.

Details of this implementation of the ensemble filter can be found in A07.

3. Adaptive inflation

To combat a persistent loss of variance, improve filter performance and avoid filter divergence, ad hoc methods to restore or avoid a loss of variance have been used in filters. One common method is covariance inflation (Anderson and Anderson, 1999) in which the ensemble estimate of each state-vector component has its variance increased as

$$x_{m,n}^{\text{inf}} = \sqrt{\lambda}(x_{m,n} - \bar{x}_m) + \bar{x}_m, \quad m = 1, \dots, M; n = 1, \dots, N. \quad (3)$$

The overbar is an ensemble mean and λ is called an (covariance) inflation factor. Inflation is most commonly applied to the prior ensemble just before forward operators are computed. The observational error distributions associated with each observation are assumed independent here as was required for the derivation of the basic filter in the previous section.

To deal with the challenges outlined in the introduction, a separate value of inflation is associated with each state vector element here. The inflation associated with each element is represented by a random variable, λ_m , $m = 1, \dots, M$. Covariance inflation can be thought of as a simple model to correct unknown deficiencies in an ensemble filter that lead to underestimates of prior variance. The inflation factor vector, $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \dots, \lambda_M]$, is the state for this simple model. Observations used via Bayes theorem can incrementally improve the estimate of $\boldsymbol{\lambda}$. The prior ensemble estimate of an observation, y_n , $n = 1, \dots, N$, the observation, y^o , and the observational error variance, σ_o^2 , can be used to improve the estimate of $\boldsymbol{\lambda}$.

In the algorithm presented here, a Bayesian estimator is used to approximate the values of $\boldsymbol{\lambda}$. The prior probability density function is a multivariate normal. The prior marginal distribution for λ_m at time t is described by its mean $\bar{\lambda}_{m,p}$ and variance $\sigma_{\lambda,m,p}^2$, where the subscript p indicates a prior and the subscript λ on the variance distinguishes this from the prior sample variance of

an observation. The inflation factors are required to be positive definite, so, a Gaussian with its unbounded tails is probably not the most appropriate choice of distribution. However, as shown below, choosing a Gaussian led to a computationally efficient approximate algorithm for the Bayesian update that produced satisfactory results for a number of different applications. Future research should explore the use of distributions that formally restrict the inflation to positive values. The correlation between any two elements of the inflation is assumed to be the same as the prior sample correlation between the corresponding elements of the model state vector $\mathbf{x}$,

$$r_{\lambda i, \lambda j} = r_{x i, x j}, \quad 1 \leq i, j \leq M. \quad (4)$$

This choice allows the inflation algorithm developed here to be applied in a sequential fashion, similar to that described for the basic ensemble filter above. It is also the only available estimate of how inflation state vector elements are related.

Suppose the posterior distribution for λ after assimilating an observation at time t_{a-1} is $p(\lambda, t_{a-1} | Y_{a-1})$, where Y_{a-1} is the set of all observations taken at times $t \leq t_{a-1}$ that have been assimilated. A filter algorithm for estimating the inflation vector has two parts. First, a method for computing the prior distribution at the time, t_a , of the next observation from the posterior distribution at t_{a-1} is required. This is a model for the time evolution of λ . It is not obvious what such a model should look like. For geophysical models with sensitive dependence on initial conditions and varying observational density and quality, it seems likely that the inflation variance, $\sigma_{\lambda, m}^2$, would increase with time whereas the mean, $\bar{\lambda}_m$, might relax towards some climatological value. Here, the prior at t_a is assumed to be identical to the posterior at t_{a-1} so that the only changes to the distribution for λ are caused by the assimilation of observations. Developing a better model for the time evolution of λ could lead to improved adaptive filter performance (see Section 5).

Second, a method for computing updated estimates of λ given a prior estimate and a set of observations at t_a is required. Following A07, the same observations used to update $\mathbf{x}$ in the standard ensemble filter are used to update the estimate of λ . The observations are assumed to have error distributions that permit sequential assimilation of the observations. Because the inflation prior is assumed Gaussian, marginal distributions for each λ_m can be updated sequentially and independently (Anderson, 2003). Examining the impact of a single observation on a single state variable is therefore sufficient to describe the inflation update algorithm and simplifies notation.

The prior marginal distribution for an element of the inflation vector is

$$p(\lambda, t_a | Y_{a-1}) = \text{Normal}(\bar{\lambda}_p, \sigma_{\lambda, p}^2), \quad (5)$$

with ensemble mean $\bar{\lambda}_p$ and variance $\sigma_{\lambda, p}^2$. The subscript indexing the element of the inflation vector has been dropped for simplicity. By applying Bayes theorem, as in A07, assimilating

an observation at t_a gives

$$p(\lambda, t_a | Y_a) = p(y^o | \lambda) p(\lambda, t_a | Y_{a-1}) / \text{norm}, \quad (6)$$

where the left hand side (l.h.s.) is the posterior distribution, the first term in the numerator is the observation likelihood and the second term in the numerator represents the prior distribution at t_a .

An expression for the observation likelihood in the numerator of (6) is required. Suppose the prior observation ensemble is given by

$$y_{p, n} = h(\mathbf{x}_n), \quad n = 1, \dots, N, \quad (7)$$

with sample mean and variance $\bar{y}_p$ and σ_p^2 respectively. The observation has value y^o and error variance σ_o^2 . The observation likelihood in (6) expresses the probability that y^o would be observed given a value of the inflation component, λ .

Let θ be the expected value of the distance between the prior ensemble mean for the observed quantity, $\bar{y}_p$, and the observed value, y^o . If the ensemble for the ‘observed’ variable y is inflated by a factor λ^o and assuming that the inflated prior observation distribution is unbiased,

$$\theta = \sqrt{\lambda^o \sigma_p^2 + \sigma_o^2}. \quad (8)$$

To find the observation likelihood, λ^o must be related to the state space inflation vector component, λ . All that is known about this relationship comes from the prior joint ensemble distributions of x and y . Since the prior distribution of the inflation values is assumed to be multivariate normal, the impact of changes in λ^o on each state space inflation component can be computed independently (Anderson, 2003). Since x and y are assumed to be binormally distributed with a sample correlation of r , and the prior values of λ and λ^o are assumed to have a binormal distribution with the same correlation, inflating the variance of x_m by a factor of $\bar{\lambda}_p$ leads to an expected variance inflation for y of

$$E[\lambda^o] = [1 + r(\sqrt{\bar{\lambda}_p} - 1)]^2. \quad (9)$$

If localization and/or a hierarchical filter (Anderson, 2007b) are being used to reduce sampling error, it is necessary to introduce a further factor

$$E[\lambda^o] = [1 + \gamma(\sqrt{\bar{\lambda}_p} - 1)]^2, \quad (10)$$

where $\gamma = \alpha r$ and α is the product of any localization factor and any hierarchical filter factor that are being applied for the impact of observation y on state variable x_m . Obviously, errors in the prior sample state covariance will lead to corresponding errors in the adaptive inflation update step.

The actual distance between the prior ensemble mean and the observed value is $D \equiv |\bar{y}_p - y^o|$. The assumption that the prior sample and observation are unbiased would imply that D is drawn from a distribution that is Normal $(0, \theta)$. By the definition

of the normal,

$$p(y|\lambda) = (\sqrt{2\pi}\theta)^{-1} \exp(-D^2/2\theta^2) \quad (11)$$

is the observation likelihood term and (6) becomes

$$p(\lambda, t_a|Y_a) = (\sqrt{2\pi}\theta)^{-1} \exp(-D^2/2\theta^2) \text{Normal}(\bar{\lambda}_p, \sigma_{\lambda,p}^2) / \text{norm}. \quad (12)$$

Both the observation likelihood (11) and the prior (5) are Gaussian but in different variables. The former has a variance that is a function of λ_p whereas the latter has mean $\bar{\lambda}_p$, so their product is not Gaussian. Nevertheless, the algorithm here implicitly assumes a Gaussian for the posterior since it is taken to be equivalent to the next prior which is assumed Gaussian. A more detailed discussion of the prior and the likelihood along with a figure depicting the functions is found in A07.

In general, the prior probability distribution function (pdf) will be very strongly peaked compared to the likelihood term. This indicates that there is very little information about λ available from a single observation. Since the prior is assumed Gaussian, and the likelihood is usually nearly flat over many standard deviations of the prior, the posterior (6) is very similar to the prior and is nearly Gaussian. A posterior (updated) mean and variance must be found so that $\text{Normal}(\bar{\lambda}_u, \sigma_{\lambda,u}^2)$ closely approximates the result of (12). One could try to find the mean of the product through analytic integration, but this is difficult because λ is associated with the mean in the prior and the variance in the likelihood terms in a complicated fashion. One could also try to find the mode of the posterior and use it as a surrogate for the mean. This computation can be performed analytically by solving a cubic equation, when $\gamma = 1$. This approach is taken in A07 while developing a spatially constant adaptive inflation algorithm. When $0 < \gamma < 1$, finding the mode requires the solution of a sixth-order equation; this is not analytically feasible.

However, an approximate mode of (6) can be obtained by doing a truncated Taylor expansion of the likelihood around the mean value of the prior inflation,

$$p(y|\lambda) \cong p(y|\bar{\lambda}_p) + \frac{\partial}{\partial \lambda} [p(y|\lambda)]_{\bar{\lambda}_p} (\lambda - \bar{\lambda}_p) + O(\lambda - \bar{\lambda}_p)^2. \quad (13)$$

Appendix A gives details of this computation and its associated cost. The estimated mode is taken as the mean of the updated inflation distribution.

The variance of the posterior inflation is found by the same naïve numerical technique used in A07. The ratio, R , of the value of the numerator of (12) at $\bar{\lambda}_u + \sigma_{\lambda,p}$ to the value at $\bar{\lambda}_u$ is computed by evaluating the numerator at both points. It is assumed that the posterior is normal and its variance is computed as

$$\sigma_{\lambda,u}^2 = -\sigma_{\lambda,p}^2 / 2 \ln(R). \quad (14)$$

Other methods could be used to approximate the updated variance. One could compute ratios as just outlined for a large num-

ber of points and use the mean value of the approximate variances to account for the impact of higher order moments in the posterior. Another method uses quadrature and an optimization to find the value of the variance that minimizes the integrated area between the exact posterior and the Gaussian approximation. The approximate method used here is biased to produce slightly larger values of $\sigma_{\lambda,u}^2$ in the mean over many observations. This is a convenient property as noted in the discussion of filter divergence for the λ filter in A07.

The ensemble filter algorithm with spatially and temporally varying adaptive inflation proceeds as follows.

- (1) The ensemble estimate of the model state $\mathbf{x}$ is advanced to the time of the next set of observations using the model.
- (2) The prior estimates for λ are assumed to be equal to the posteriors at the previous observation time.
- (3) For each observation y available at this time, do the following iteratively:

- (a) Compute the prior ensemble distribution of y by applying the forward operator to the state vector ensemble.
- (b) Compute the prior mean and variance of y and the distance D .
- (c) For each state variable component in $\mathbf{x}$:
 - (i) compute the prior sample correlation between the observation and x_m ;
 - (ii) compute any localization factors;
 - (iii) compute γ and then the expected value of λ^o using (10);
 - (iv) compute θ using the expected value of λ^o and (8);
 - (v) use (12), (13) and the numerical method in Appendix A to compute an updated value for the mean of λ , corresponding to this state-vector component and
 - (vi) use (14) to compute the updated value for the variance of λ as an algorithmic variant; this step can be skipped and the variance of λ can be kept at a constant value.

- (4) Inflate the prior ensemble for each state vector component by the mean of the corresponding updated inflation distribution.
- (5) Complete a standard ensemble filter assimilation to update the state vector for the observations at this time.

Modifications to this algorithm can enhance computational efficiency for large model applications (Keppenne and Rienecker, 2002). For instance, it may be costly to compute forward observation operators twice as required in the algorithm above (once in step 3a, once in the standard filter assimilation in step 5). For parallel implementations of filters like the one in the Data Assimilation Research Testbed at NCAR (www.IMAGe.ucar.edu/DAReS/DART), expensive data transpose operations are required to transform from a situation where a complete model state vector is available on a processor to a situation where all ensemble elements of a particular state-vector component are available (Anderson and Collins, 2007). In such

cases, repeatedly computing forward observation operators can be even more expensive.

An approximation to the algorithm above that only computes forward observation operators once can be designed. Steps 1 and 2 are unchanged, but subsequent steps proceed as follows.

(3) Inflate the prior ensemble for each state-vector component by the mean of the corresponding updated inflation distribution.

(4) Compute the forward observation operator for all observations; compute and retain the ensemble mean and variance for each.

(5) Do the following for each observation in turn:

(a) Compute the observation increments using the most current estimate of the observation ensemble mean and variance, the observation and the observation error variance.

(b) Do the following for each state variable component in turn:

(i) Regress the observation increments onto the state variable component increments (2),

(ii) Compute γ ,

(iii) Using the original values of observation ensemble variance saved in step 4, remove the estimated impact of inflating this state variable from the variance. The resulting deflated observation variance is equal to the stored value divided by the right-hand side (r.h.s) of (10). In steps (iv)–(vi), use this deflated variance.

(iv) Compute θ using the expected value of λ^o and (8);

(v) use (12), (13) and the numerical method in Appendix A to compute an updated value for the mean of λ corresponding to this state-vector component and

(vi) use (14) to compute the updated value for the variance of λ .

For linear forward observation operators and Gaussian prior distributions with no sampling error, this second algorithm is equivalent to the first (Anderson and Collins, 2007).

4. Low-order model results

The spatially and temporally varying adaptive inflation algorithm is applied to the 40-variable model of Lorenz (1996; see Appendix B) that has been used in many assimilation studies (Lorenz and Emanuel, 1998). Forty state variables are defined as being equally spaced on a $[0, 1]$ periodic domain. A set of 20 observing stations at locations

$$z_i = 0.395 + 0.01i, i = 1, \dots, 20, \quad (14)$$

equally spaced between model state vector locations is used. The forward operator is linear interpolation between the two closest state vector locations. This creates a densely observed region between 0.4 and 0.6 and a sparsely observed region. Observations are generated by applying the forward operator to a free integration of the Lorenz-96 model started from a point

on the model's attractor. A random draw from $\text{Normal}(0, 1)$ is added to each observation to simulate observational error. Observations are taken every time step for a total of 4000 steps but results use only the last 2000 steps. Neglecting the initial 2000 steps appears to be sufficient to avoid significant transient behaviour. The initial ensemble conditions are random draws from the model climatology created by taking every 10000th step in a long free integration of the model. The forcing F in (B.1) is set to 8, unless otherwise noted.

When the regression of increments from an observation located at z_o onto a state variable at z_s is performed with (2), the regression coefficient is multiplied by $\zeta(d, c)$, where

$$d = \min(|z_o - z_s|, 1 - |z_o - z_s|) \quad (14)$$

is the distance between the observation and state variable and ζ is the fifth-order polynomial function of Gaspari and Cohn (1999), with half-width c . For $d \geq 2c$, the observation has no impact on the state variable. For $d < 2c$, ζ approximates a Gaussian. The half-width is $c = 0.20$ for all experiments.

4.1. Perfect model experiments

The first results are for a perfect model case in which assimilation error is due to sampling error from small ensembles, inadequate localization functions, and the assumption of a linear relation between the observation and state variable priors. A very small ensemble size, $N = 6$, is chosen to maximize the relative impact of sampling error. Four assimilations are performed: (P-No) has no inflation; (P-ST) has spatially and temporally varying adaptive inflation as described in section 3; (P-Con) has a specified constant inflation and (P-T) has temporally varying but spatially fixed adaptive inflation, as described in A07. The initial means of the inflation components are $\bar{\lambda} = 1$ for P-ST and P-T and the standard deviations of the inflation are held constant at $\sigma_\lambda = 0.1$ throughout the assimilation (eq. 14 is not used). Results show small sensitivity to the fixed value of the inflation standard deviation between 0.01 and 0.5. A complete discussion of the use of a fixed versus adaptive variance for the adaptive inflation is found in A07. The constant value of inflation in P-Con was 1.1, a value that minimized the ensemble mean prior rms error and was also the largest value that ran without 'blowing up'. The prior rms error and spread, as a function of state variable location, are used to compare these assimilations.

Figure 1 shows the rms error and spread as a function of state variable location for P-No. The error is smallest in the densely observed region with smaller values extending further to the right- than to the left-hand side. This is consistent with the group velocity of the model carrying information from the observations downstream (to the right-hand side). The error is the largest just upstream of the densely observed region. The spread has a similar shape but is significantly smaller than the error for all locations. If the prior ensemble from the filter is to be regarded as a random sample of the analysis distribution, the




Fig. 1. Ensemble mean rms error and ensemble spread averaged over last 2000 steps of a Lorenz-96 model assimilation with a perfect model and no inflation (case P-No). Twenty observations with error variance 1.0 are equally-spaced in the region between 0.405 and 0.595 and assimilated every time step.




Fig. 3. As in Fig. 1 for case with a perfect model and spatially and temporally varying adaptive inflation (case P-ST).




Fig. 2. As in Fig. 1 for case with a perfect model and constant 1.1 inflation (case P-Con).

prior spread and the prior rms error should have the same time mean values; this is not the case here.

Figure 2 shows error and spread for P-Con. The error is similar to that for P-No. The spread is much closer to the error in the vicinity of the observations, although it is still slightly smaller. Outside of the densely observed region, the spread is much larger than the error. An inflation value larger than 1.1 would further improve the agreement between spread and error in the observed region but leads to unconstrained growth of ensemble variance in the unobserved region. This is consistent with P-T, which quickly ‘blew up’ as it tried to increase its spatially constant inflation to

make the prior ensemble and the observations consistent. The P-T inflation grew to values greater than 1.1, which led to rapid unbounded variance growth in the unobserved region. This is analogous to cases in which GCM ensemble assimilations with fixed (or temporally varying but spatially fixed) inflation blow up by generating too much variance in poorly observed regions.

Figure 3 shows error and spread for P-ST. The error pattern is qualitatively similar to the other cases. In the observed region, the spread and error agree very well (the curves cannot be distinguished in the plot), but the spread is slightly smaller than the error outside the observed region.

Figure 4 plots the error and spread for the three successful cases in the observed region. The error from P-ST is the smallest everywhere, although P-Con is nearly as small in the upstream portion. The error for P-No is much larger. Even at this enhanced resolution, the error and spread for P-ST are virtually indistinguishable, indicating that the adaptive inflation is working correctly in the observed region. The P-Con spread is not much larger than the P-No spread, and both are substantially less than the corresponding errors.

Figure 5 shows the time mean inflation values for the last 2000 assimilation steps for P-ST (solid). Maximum values are close to 1.5 whereas the minima are 1.0 (no inflation) far from the observed region. In this perfect model case, the main source of erroneously small variance is assimilation sampling error due to the small ensemble size, so it is expected that the need for inflation will be concentrated near the observed region. In this region, P-ST produces spread and error that are nearly identical. In regions remote from the observations, there should be no need for inflation and the spread and error should both be at the climatological values. However, Fig. 3 shows that the P-ST spread is too small away from the observations. Experimentation indicates that this is due to a second type of systematic sampling error.




Fig. 4. rms error and spread from perfect model cases in Figs. 1–3, displayed only over the region where observations are available. Thick lines are rms and thin lines are spread. Solid lines are no inflation case (P-No), solid with ‘x’ are constant inflation (P-Con) and dashed are spatially varying adaptive inflation (P-ST).




Fig. 5. Values of the adaptive inflation as a function of state variable location averaged over last 2000 time steps of perfect model (P-ST, solid) and imperfect model (I-ST, dashed) assimilations.

Since the model will not generate spread greater than the climatological spread in the time mean, occasional chance correlations between observations and state variables that are relatively far apart lead to erroneous reductions in variance far from the observations. Fixing this problem would require a more appropriate localization like that produced by hierarchical ensemble filters (Anderson, 2007b).

4.2. Imperfect model experiments

A second set of experiments, I-No, I-Con, I-T and I-ST, explore assimilations when model error is the dominant source of error




Fig. 6. Ensemble mean rms error and ensemble spread averaged over last 2000 steps of a Lorenz-96 model assimilation with an imperfect model and no inflation (case I-No). The model integration used to generate the observations had $F = 8$ while the assimilating model had $F = 5$. Twenty observations with error variance 1.0 are equally spaced in the region between 0.405 and 0.595 and assimilated every time step.

in the assimilation. Model error is simulated by changing the value of the forcing, F , in the Lorenz-96 model (eq. B.1) to 5 in the assimilating models. The control integration that generates the observations continues to have $F = 8$. A free integration of the model is periodic for $F = 5$ but chaotic for $F = 8$. To reduce the impact of sampling error, 20-member ensembles are used. Results for 40-, 80-, 160- and 320-member ensembles were quantitatively similar, indicating that model error is dominating in the 20 member cases. All other details of the four cases are identical to the corresponding perfect model experiments.

Figure 6 shows error and spread for I-No. The rms error and spread are smaller in the observed region, but the spread is nearly 10 times smaller than the error over most of the domain. The assimilation is unaware that the model time tendency has large errors, so the truth generally lies far outside the prior ensemble. Rank histogram plots (Anderson, 1996, not shown) indicate that the truth lies in one of the outer bins in nearly every case.

Figure 7 shows error and spread for I-Con with a fixed inflation $\lambda = 1.09$. This is the largest value of inflation that did not ‘blow up’. Smaller values of λ produced less spread and greater rms errors in the observed region. The rms error is much smaller in the observed region than for I-No, and the spread is more consistent with the rms. Nevertheless, the spread is still about a factor of 4 too small in the observed region. Inflating enough to account for the insufficient spread due to model error in the observed region leads to uncontrolled growth of ensemble spread in the unobserved region.

Figure 8 shows error and spread for the I-ST case. The error in the observed region is smaller than in the other imperfect




Fig. 7. As in Fig. 6 for case with an imperfect model and constant 1.09 inflation (case I-Con).




Fig. 8. As in Fig. 6 for case with an imperfect model and spatially and temporally varying adaptive inflation (case I-ST).

model cases, and the spread is nearly indistinguishable from the error. Outside of the observed region, the rms error is about 4, but the spread is only about 2. Since there are no observations in this region, the assimilation cannot determine that the model has error. There is no way to address this without providing additional prior information about model deficiency to the assimilation system. The resulting spread is approximately equal to the 'climatological' spread from the assimilating model with $F = 5$. Between the unobserved and observed regions, especially on the upstream side, the spread is nearly as large as the error and is much larger than the climatological spread. There is enough information from the nearby observations to identify the presence of large model error, but not enough information to significantly reduce the rms error. Fig. 5 shows the time mean




Fig. 9. rms error and spread from imperfect model cases in Figs. 6–8, displayed only over the region where observations are available. Thick lines are rms and thin lines are spread. Solid lines are no inflation case (I-No), solid with 'x' are constant inflation (I-Con) and dashed are spatially varying adaptive inflation (I-ST).




Fig. 10. Values of spatially varying inflation over last 500 assimilation times for state variables at location 0.25 (thick solid), 0.375 (thick dashed), 0.5 (thin solid) and 0.6 (thin dashed).

of the mean inflation over the last 2000 assimilation steps. The value is roughly uniform at about 1.7 in the observed region and 1 (no inflation) in the middle of the unobserved region. The values greater than 1 on the flanks of the observed region are responsible for the spread being larger than the climatological value in these regions.

Figure 9 compares the rms error and spread from Figs. 6–8 over the observed region. The rms error for I-ST is much smaller than for the other two cases, and the spread is very close to the rms. The spread for the other cases is much smaller than the error.

Figure 10 plots the temporal evolution of the mean inflation from I-ST over the last 500 observing times for four state

variable locations. The inflation has relatively small variation with time as one would expect with the model and observation network fixed in time. If the value of the inflation variance had been allowed to vary in time in the assimilation, this variance would be a non-increasing function of time and would asymptote to 0. As the inflation variance became small, the inflation mean would be increasingly resistant to modification by observations. Eventually, the inflation would become a constant as a function of time (see A07 for a detailed discussion of these issues). Error and spread results from an experiment in which the inflation variance is adaptive are quantitatively similar to those from I-ST. An analogous plot to Fig. 5 (not shown) would not show any discernible differences whereas an analogy to Fig. 10 would show no discernible time variability in the inflation values.

5. Discussion

Covariance inflation is a common method for dealing with error in ensemble filters. Use of a single spatially constant but temporally varying value of inflation for all state variables is known to be problematic in applications with spatially heterogeneous error characteristics. The most common cause of spatially varying prior error is a spatially heterogeneous observing system. A spatially and temporally varying adaptive inflation algorithm has been shown to work better than existing spatially and temporally constant algorithms and spatially constant but temporally varying adaptive algorithms in the Lorenz-96 model with a densely observed region and an unobserved region. The algorithm works better than the optimal spatially and temporally constant inflation, whether the dominant source of insufficient ensemble variance is small ensemble sampling error or model error. Many additional cases have been explored in the Lorenz-96 model, and the new algorithm is always found to produce lower rms error and more consistent prior variance than the best fixed spatially and temporally constant inflation assimilations. These cases have included a variety of heterogeneous observation distributions, different observation frequencies and a variety of nonlinear forward observation operators.

Many large geophysical ensemble data assimilation applications require error tolerant algorithms like inflation to produce small rms and consistent error variance. A good example is found for global numerical weather prediction using in situ observations. The spatial density of the observing system is heterogeneous with Northern-Hemisphere continental regions being much better observed than the remote Southern Ocean. Computational cost necessitates the use of ensembles that are tiny compared with the model state vector, and even the best global atmospheric models are known to have significant prediction errors. Tuning the assimilation by trying a range of constant inflation values can be costly, whereas an attempt to find appropriate spatially varying distributions of inflation would be prohibitively expensive.

The algorithm developed here is routinely applied to global assimilations using several versions of the National Center for Atmospheric Research Community Atmospheric Model (CAM). The algorithm produces much lower 6-hour forecast rms errors and more consistent ensemble variance than assimilations with fixed inflation. Assimilations have been cycled for several months with the adaptive inflation, but fixed inflation runs with small rms errors during the first week blow up when the assimilation period is extended. The spatial distribution of inflation after a month ranges from 1 (no inflation) over some areas of the Southern Ocean to more than 10 in Northern-Hemisphere upper troposphere regions that are heavily trafficked by Aircraft Communications and Reporting System (ACARS) equipped aircraft. This disparity between the adaptive inflation values emphasizes the need for spatially varying inflation to obtain good results over the whole globe. The largest values of inflation occur in densely observed regions where the ensemble filter produces posterior distributions with very small variances. These variances would be correct if there were no error in the model, the filter or the observations, but this is not the case. Instead, even a small amount of error, in particular model error, means that the posterior variance and the subsequent forecast variance are far too small to be consistent with the observations. Very large inflation values are needed to make the priors consistent in these cases. These assimilation results in CAM will be discussed in detail in a future report.

Despite the success of the spatially and temporally varying adaptive inflation in Lorenz-96 and CAM, challenges remain. As noted in Section 3, individual observations generally provide only a small amount of information about the inflation distribution. The inflation adapts slowly as a function of time unless observations are dense in space at all times. If the distribution of observations (or model error) varies rapidly as a function of time, the estimated inflation may not be able to keep up. The in situ global observing system has significant variations in observation density during the diurnal cycle. Radiosonde observations are predominantly available at 0 and 12 UTC, whereas aircraft observations peak at different times for different routes. In the CAM assimilations, observations are assimilated every 6 h and the adaptive inflation values over the Northern-Hemisphere land tend to be somewhat too large at 0 and 12 UTC and too small at 6 and 18 UTC.

Larger temporal heterogeneity of observations is found when assimilating Doppler radar observations for storm-scale prediction. Observations of a wind component are only available when there is sufficient reflectivity. There may be no observations available until a convective cell begins to form, then there may be frequent spatially dense observations. It is challenging for adaptive inflation to adjust quickly enough to minimize the threat of filter divergence in situations like this.

A related problem occurs when observations that have been regularly available in time suddenly disappear. When there are no observations available, the algorithm described in Section 3

does not change the inflation. If observations have been available in a certain region, the algorithm may produce inflations larger than 1. When the observations disappear, this inflation will continue to be applied at every time that observations are available anywhere. An example occurs when doing global assimilation of in situ observations. During the Southern-Hemisphere summer, there are many more observations from ships and aircraft over the Southern Ocean than during winter. During summer, an assimilation may require inflation to deal with model error in this region. When winter comes, the inflation continues to be applied, but there are no observations to constrain the growth of ensemble variance. The result may be catastrophic. A solution is to include a model for the time variation of the inflation estimate. The simplest possibility is to damp the inflation value towards 1 as a function of time. If observations are no longer available, the inflation will gradually disappear. This problem may be much more serious in the Doppler radar case. During a convective event, large inflation may be required. When the convection ends, Doppler wind observations disappear and the inflation would need to be eliminated. A similar situation occurs for tropical storm prediction; observations in a region can become dense as the storm passes and then revert to a sparse background level once the storm has moved on. Appropriate models of the time evolution of both the mean and variance of the inflation distribution could improve assimilation performance in cases like this.

Additive inflation (Sandu et al., 2007) is an alternative to the multiplicative inflation used here. Developing adaptive additive inflation is challenging because inflation values are dimensional, and computing the appropriate relations between observation space and physical space inflation for (9) is more difficult. Additive inflation is less likely to lead to rapid increases in ensemble variance in unobserved regions. Additive inflation may also be a better model of the appropriate correction for variance when the leading source of error is model deficiency. However, multiplicative inflation is probably more appropriate when insufficient ensemble variance is caused by sampling error in the assimilation step. A fully general adaptive inflation algorithm might address problems from different error sources in different ways.

The spatial distribution of inflation in applications to atmospheric models can develop structures with small spatial scale. It is possible that the small-scale structure introduced to the priors by the inflation could introduce spurious gradients or dynamic imbalances that would corrupt the assimilation, although this has not been evident in realistic assimilation performed to date. Using a spatially smoothed version of the prior sample covariance for the multivariate normal distribution of inflation could eliminate small-scale structures in the inflation.

Traditionally, inflation has been applied to the state variables just before forward observation operators are applied for assimilation. However, if the dominant source of insufficient variance is sampling error, the posterior ensemble and the resulting forecasts

will have too little variance. This could be remedied by applying the inflation immediately after the assimilation instead. Whereas this is trivial to implement for fixed inflation, an adaptive posterior inflation is more challenging. The available observations will already have been used to adjust the ensemble, so a simple analogy for (6) cannot be constructed. If a subset of observations were withheld from the ensemble assimilation, these could be used for the adaptive inflation assimilation. More elaborate methods can be developed if observations are not withheld and will be described in future work.

6. Acknowledgments

The author thanks Tim Hoar, Kevin Raeder, Nancy Collins and Hui Liu for developing and supporting DART and Ryan Torn, Chris Snyder and Altug Aksoy for their insights on filtering and inflation.

7. Appendix A: Computing the updated inflation distribution

The mean of the updated inflated distribution is estimated using an approximation of the mode of (6). One could also attempt to approximate the mean of (6) by evaluating the integral of the product of the prior and likelihood terms.

The likelihood (13) is first approximated by a Taylor expansion truncated at first order:

$$p(y^o|\lambda) \cong \bar{l} + l'(\lambda - \bar{\lambda}_p), \quad (\text{A.1})$$

where

$$\bar{l} = p(y^o|\bar{\lambda}_p) = (\sqrt{2\pi\bar{\theta}})^{-1} \exp(-\frac{1}{2}D^2\bar{\theta}^{-2}), \quad (\text{A.2})$$

$$l' = \frac{\partial}{\partial \lambda} [p(y^o|\lambda)]_{\bar{\lambda}_p} = \bar{l}(D^2\bar{\theta}^{-2} - 1) \bar{\theta}^{-1}(\partial\theta/\partial\lambda)_{\bar{\lambda}_p}, \quad (\text{A.3})$$

$$\bar{\theta} \equiv \sqrt{\bar{\lambda}^o\sigma_p^2 + \sigma_o^2}, \quad (\text{A.4})$$

$$\bar{\lambda}^o \equiv [1 + \gamma(\sqrt{\bar{\lambda}_p} - 1)]^2 \quad (\text{A.5})$$

and

$$(\partial\theta/\partial\lambda)_{\bar{\lambda}_p} = \frac{1}{2}\sigma_p^2\gamma(1 - \gamma + \gamma\sqrt{\bar{\lambda}_p})/(\bar{\theta}\sqrt{\bar{\lambda}_p}). \quad (\text{A.6})$$

The approximate mode of the posterior (6) can be found by setting the derivative of the approximate posterior to 0 after removing the constant coefficient from the prior Gaussian

$$\frac{\partial}{\partial \lambda} \left\{ [\bar{l} + l'(\lambda - \bar{\lambda}_p)] \exp \left[-\frac{1}{2}(\lambda - \bar{\lambda}_p)^2 \sigma_\lambda^{-2} \right] \right\} = 0 \quad (\text{A.7})$$

and solving for λ . This results in a quadratic equation for λ :

$$\lambda^2 + (\bar{l}/l' - 2\bar{\lambda}_p)\lambda + (\bar{\lambda}_p^2 - \sigma_\lambda^2 - \bar{l}\bar{\lambda}_p/l') = 0. \quad (\text{A.8})$$

Solutions to (A.8) are two real roots, and the one closest to $\bar{\lambda}_p$ is selected.

Computational cost for each state variable and observation pair in the Data Assimilation Research Testbed implementation is dominated by five square roots and an exponential. For ensemble sizes greater than about 10, the cost of this computation is significantly smaller than the cost of computing the sample regression coefficient, so that the adaptive algorithm is not a major contributor to computational cost of the overall adaptive filter algorithm.

8. Appendix B: The Lorenz-96 model

The Lorenz-96 (Lorenz, 1996) model has M state variables, $X_1, X_2, \dots, X_M$ and is governed by the equation

$$dX_m/dt = (X_{m+1} - X_{m-2})X_{m-1} - X_m + F, \quad (\text{B.1})$$

where $m = 1, \dots, M$ with cyclic indices. Here, M is 40, $F = 8.0$ unless otherwise noted, and a fourth-order Runge–Kutta time step with $dt = 0.05$ is applied as in Lorenz and Emanuel (1998).

References

- Anderson, J. L. 1996. A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *J. Clim.* **9**, 1518–1530.
- Anderson, J. L. 2001. An ensemble adjustment Kalman filter for data assimilation. *Mon. Wea. Rev.* **129**, 2894–2903.
- Anderson, J. L. 2003. A local least squares framework for ensemble filtering. *Mon. Wea. Rev.* **131**, 634–642.
- Anderson, J. L. 2007a. An adaptive covariance inflation error correction algorithm for ensemble filters. *Tellus* **59A**, 210–224.
- Anderson, J. L. 2007b. Exploring the need for localization in ensemble data assimilation using an hierarchical ensemble filter. *Physica D* **230**, 99–111.
- Anderson, J. L. and Anderson, S. L. 1999. A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Wea. Rev.* **127**, 2741–2758.
- Anderson, J. L. and Collins, N. 2007. Scalable implementations of ensemble filter algorithms for data assimilation. *J. Atmos. Ocean Tech.* **24A**, 1452–1463.
- Buizza, R., Miller, M. and Palmer, T. N. 1999. Stochastic representation of model uncertainties in the ECMWF ensemble prediction system. *Quart. J. Roy. Meteorol. Soc.* **125**, 2887–2908.
- Burgers, G., van Leeuwen, P. J. and Evensen, G. 1998. Analysis scheme in the ensemble Kalman filter. *Mon. Wea. Rev.* **126**, 1719–1724.
- Daley, R. 1993. Estimating the observation error statistics for atmospheric data assimilation. *Ann. Geophys.* **11**, 634–647.
- Dee, D. P. 1995. On-line estimation of error covariance parameters for atmospheric data assimilation. *Mon. Wea. Rev.* **123**, 1128–1145.
- Dee, D. P. and da Silva, A. M. 1999. Maximum-likelihood estimation of forecast and observation error covariance parameters, Part I: methodology. *Mon. Wea. Rev.* **127**, 1822–1834.
- Dee, D. P. and Todling, R. 2000. Data assimilation in the presence of forecast bias: the GEOS moisture analysis. *Mon. Wea. Rev.* **128**, 3268–3282.
- Dee, D. P., Gaspari, G., Redder, C., Rukhovets, L. and da Silva, A. M. 1999. Maximum-likelihood estimation of forecast and observation error covariance parameters, Part II: applications. *Mon. Wea. Rev.* **127**, 1835–1849.
- Evensen, G. 1994. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to do forecast error statistics. *J. Geophys. Res.* **99**(C5), 10143–10162.
- Evensen, G., 2006. *Data Assimilation: The Ensemble Kalman Filter*. Springer, 280 pp.
- Eyre, J. R., Kelly, G. A., McNally, A. P., Andersson, E. and Persson, A. 1993. Assimilation of TOVS radiance information through one-dimensional variational analysis. *Quart. J. Roy. Meteor. Soc.* **119**, 1427–1463.
- Furrer, R. and Bengtsson, T. 2007. Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *J. Multivariate Anal.* **98**(2), 227–255.
- Gaspari, G. and Cohn, S. E. 1999. Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteorol. Soc.* **125**, 723–757.
- Hamill, T. M. and Whitaker, J. S. 2005. Accounting for error due to unresolved scales in ensemble data assimilation: a comparison of different approaches. *Mon. Wea. Rev.* **133**, 3132–3147.
- Hamill, T. M., Whitaker, J. S. and Snyder, C. 2001. Distance-dependent filtering of background-error covariance estimates in an ensemble Kalman filter. *Mon. Wea. Rev.* **129**, 2776–2790.
- Hansen, J. A. 2002. Accounting for model error in ensemble-based state estimation and forecasting. *Mon. Wea. Rev.* **130**, 2373–2391.
- Harlim, J. and Hunt, B. R. 2007. A non-gaussian ensemble filter for assimilating infrequent noisy observations. *Tellus* **59A**, 225–237.
- Houtekamer, P. L. and Mitchell, H. L. 2001. A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Wea. Rev.* **129**, 123–137.
- Kalnay, E., Li, H., Miyoshi, T., Yang, S.-C. and Ballabrera-Poy, J. 2007. 4-D-Var or ensemble Kalman filter? *Tellus* **59A**, 758–773.
- Keppenne, C. L. and Rienecker, M. M. 2002. Initial testing of a massively parallel ensemble Kalman filter with the Poseidon isopycnal ocean general circulation model. *Mon. Wea. Rev.* **130**, 2951–2965.
- Kistler, R., Collins, W., Saha, S., White, G. and Woolen, J. 2001. The NCEP-NCAR 50-year reanalysis: monthly means CD-ROM and documentation. *Bull. Amer. Met. Soc.* **82**, 247–267.
- Lorenz, E. N. 1996. Predictability: a problem partly solved. In: *Proceedings of the ECMWF Seminar on Predictability*, Vol. I, ECMWF, Reading, United Kingdom, 1–18.
- Lorenz, E. N. and Emanuel, K. A. 1998. Optimal sites for supplementary weather observations: simulation with a small model. *J. Atmos. Sci.* **55**, 399–414.
- Mitchell, H. L. and Houtekamer, P. L. 2000. An adaptive ensemble Kalman filter. *Mon. Wea. Rev.* **128**, 416–433.
- Mitchell, H. L., Houtekamer, P. L. and Pellerin, G. 2002. Ensemble size, balance, and model-error representation in an ensemble Kalman filter. *Mon. Wea. Rev.* **130**, 2791–2808.
- Ott, E., Hunt, B., Szunyogh, I., Zimin, A., Kostelich, E. and co-authors 2004. A local ensemble Kalman filter for atmospheric data assimilation. *Tellus* **56A**, 415–428.
- Pham, D. T. 2001. Stochastic methods for sequential data assimilation in strongly non-linear systems. *Mon. Wea. Rev.* **129**, 1194–1207.

- Sandu, A., Constantinescu, E. M., Carmichael, G. R., Chai, T., Seinfeld, J. H., and co-authors 2007. Localized ensemble Kalman filter dynamic data assimilation for atmospheric chemistry. *Lecture Notes Comput. Sci.* **4487**, 1018–1025.
- Tippett, M. K., Anderson, J. L., Bishop, C. H., Hamill, T. M. and Whitaker, J. S. 2003. Ensemble square root filters. *Mon. Wea. Rev.* **131**, 1485–1490.
- Uppala, S. M., Kallber, P. W., Simmons, A. J., Andrae, U., Da Costa Bechtold, V. and 45 contributing authors, 2005. The ERA-40 re-analysis. *Quart. J. Roy. Meteor. Soc.* **131**, 2961–3012.
- Uzunoglu, B., Fletcher, S. J., Zupanski, M. and Navon, I. M. 2007. Adaptive ensemble reduction and inflation. *Quart. J. Roy. Meteor. Soc.* **133**, 1281–1294.
- Whitaker, J. S. and Hamill, T. M. 2002. Ensemble data assimilation without perturbed observations. *Mon. Wea. Rev.* **130**, 1913–1924.