A Hierarchical Ensemble Filter for Data Assimilation

Jeffrey L. Anderson
NCAR Data Assimilation Initiative
Submitted to Monthly Weather Review, May, 2004

Abstract

Applying small ensemble filters to models with a large number of state variables has traditionally required the heuristic specification of functions that limit the impact of an observation to some set of state variables that is ‘close’ to the observation in a sense that is not rigorously defined. As a step toward the development of nearly generic filter assimilation systems, an algorithm is developed that precludes the need to specify ‘localization’ functions when using small ensemble filters in large models. Localization has been required to ameliorate sampling error that arises when small ensembles are used to sample the statistical relation between an observation and a state variable. This sampling error can be reduced more rigorously by using a Monte Carlo technique to detect and reduce the impact of spurious sample correlations between an observation and model state variables. A method referred to as a 'hierarchical ensemble filter' is applied, where groups of identical ensemble filters are used to minimize sampling error in the ensembles. Unlike traditional ensemble filters, hierarchical filters can adapt to a wide array of ensemble sizes and observational error characteristics without a need for heuristic tuning. Hierarchical filters also allow observations to efficiently impact state variables, even when the notion of ‘distance’ between the observation and the state variables cannot be easily defined. For instance, defining the distance between an observation of radar reflectivity from a particular radar and beam angle taken at 1133 GMT and a model temperature variable at 700 hPa some distance from the radar beam at 1200 GMT is a challenging task. The hierarchical filter evaluates the sampling error from a group of ensembles and computes a factor between 0 and 1, analogous to traditional distance dependent weightings, that accounts for this sampling error. There is no need to define some *a priori* notion of distance. Results are shown in both a low-order model and a simple atmospheric GCM. For simple models, the hierarchical filter produces ‘localization’ functions that are very similar to those already derived heuristically. However, as observation types become more complex or as observations are taken at different times from the state specification (as required for ensemble smoothers for instance), the localization functions become increasingly distinct from those derived heuristically. In the GCM, this complexity reaches a level that suggests that

hierarchical filter performance is unlikely to be matched by using traditional localizations. It is suggested that hierarchical filters can be used over a short training period to develop localization statistics that can then be used in a traditional ensemble filter to produce high quality assimilations at reasonable cost.

1. Introduction

Ensemble filter methods for data assimilation in the atmosphere and ocean have now been in existence for more than a decade. Through increased understanding of these methods, progressively more powerful and simpler implementations of ensemble filters have been applied to an ever increasing array of problems, ranging from low order idealized model studies, through applications in operational atmospheric prediction systems (Lorenc 2003, Keppenne and Rienecker 2002, Houtekamer et al 2004).

One goal of the Data Assimilation Initiative at NCAR is to develop simple, generic assimilation methods that can be applied by scientists with modeling and /or observational expertise but without corresponding knowledge of the details of assimilation methods. Ensemble filters described in the literature are becoming more generic, but still require the specification of numerous model and observation specific parameters for good performance. One of the most critical of these heuristic aspects of filters is the specification of functions that limit the impacts of observations to a subset of the model state variables, often a subset that is physically close in some sense to the observation (Hamill et al 2001). This localization is essential for small ensemble filters to provide high quality assimilations while avoiding problems of sampling error in large models.

In simple models, for instance univariate low-order models often used for testing ensemble filters (Lorenz 1996), it is not too difficult to define appropriate functions that localize the impact of observations. Making a physically motivated assumption that observation impacts are limited by a locally supported function that is similar to a Gaussian and tuning the width of this function works very well in many low-order model problems. Matters are much more complicated in large, multivariate, multidimensional models that are used for atmospheric and oceanic prediction. While many large ensemble filter applications have localized observation impact in the horizontal by a two-dimensional Gaussian-like function, the question of what to do in the vertical has been more troubling (Mitchell et al 2002). Limiting multivariate impacts, for

example, the impact of a temperature observation on a wind observation, has received even less attention. Observations taken at times different from that corresponding to the state variables being updated in an assimilation can also come into play. Questions like : “how should the impact of a radar reflectivity observation from a particular beam angle at 0045 GMT be allowed to impact a model temperature variable located 150 km northeast of the radar at 300 hPa at 0100 GMT" (Snyder and Zhang 2003) need to be addressed in a systematic fashion.

Approaching the ensemble filtering problem as a Monte Carlo approximation to the Bayesian filtering problem (Jazwinski 1970, Tarantola 1987) leads to a generic solution for this problem. The need to limit observation impacts on state variables is directly related to sampling errors in the ensemble filter. Applying a hierarchical Monte Carlo method, in which a group of independent ensemble assimilations is done using the same model and observations, can quantify these sampling errors. In simple one-dimensional models, results are often similar to the physically motivated methods already in use. However, the new method can provide the appropriate width for Gaussian-like localizations. In large multivariate models, computed localization functions are often non-Gaussian and would be difficult to approximate with heuristic methods. The hierarchical Monte Carlo method may significantly enhance performance in realistic atmospheric and oceanic assimilation/prediction applications.

2. Sources of error in ensemble filters

Most of the ensemble (Kalman) filtering literature starts from the classical Kalman filter (Evensen 1994, Kalman 1960). Anderson (2003) developed an alternative framework starting from the probabilistic filtering equations (Jazwinski 1970). In this framework, one can describe the assimilation problem as the impact of a single scalar observation on a single scalar state variable without loss of generality.

Figure 1 depicts how an ensemble filter is implemented in this framework. First, a model is used to advance a Monte Carlo sample (ensemble) of state estimates from a previous time, t_k, to time, t_{k+1}, when the next observation becomes available (step 1); these integrations are represented by dashed lines. Next, a forward operator, H, is applied to each sample of the prior state estimate to obtain a prior sample estimate of the scalar observation variable, y (step 2). The observed value, y^o, and corresponding observational error distribution (gray density superposed on the y-axis) are

obtained from the instrument (step 3). Next, the prior sample and observation information are combined to give an updated sample estimate of y (thin tick marks on y-axis) and corresponding increments (vectors below the y-axis) for each of the prior samples (step 4). Differences in the details of step 4 distinguish most of the flavors of ensemble (Kalman) filters documented to date (Houtekamer and Mitchell 1998, Whitaker and Hamill 2002, Pham 2001). Finally, the joint prior sample of the observation variable, y, and a state variable, x_i, are used to compute the corresponding increment in each sample of the state variable (symbolized by the vectors at the end of the dashed model integrations at time t_{k+1}), (step 5). In general, this is done using a linear regression (this is implicit in the Kalman filter derivations of ensemble methods). When linear regression is used, each state variable can be sequentially updated independently (Anderson 2003). The result is a sample of the model state vectors conditional on the observation y^o.

Errors can be introduced at each step. Model error (Dee and Todling 2000, Hansen 2002), including the fact that model sub-grid scale parameterizations are often not stochastic (Buizza et al 1999), is introduced in step 1. Operational prediction centers and many researchers are dedicated to reducing prediction model errors, but there is no indication that model errors will become negligible in the foreseeable future. The forward operators, H, in step 2 are rife with error sources including time and space interpolation errors, representativeness errors, general errors in forward operator specification, etc. Step 3 introduces errors in retrieving and transmitting observations from instruments and the use of often poorly known observational error distributions. Algorithms for step 4 generally approximately model the prior distribution (a Gaussian assumption is most common) and make additional approximations when computing the updated conditional probability. Sampling error from small ensembles is also an issue here. Finally, in step 5 the model-generated relationship between observation and state variables certainly differs from the relation in the physical system. Errors are also introduced by assuming that a linear statistical relation accurately characterizes the relation between the observation and state variable increments. Sampling error in this linear regression is often the dominant source of error in the whole filtering procedure. This paper discusses ways to compensate for this regression sampling error. The sampling error in step 4 can be addressed in a similar fashion but will be discussed in a future report.

3. Dealing with error in Monte Carlo methods

Monte Carlo methods like the ensemble filter can converge to the answer for the underlying problem, even for small ensemble sizes. This can even happen in idealized situations for large non-linear models like atmospheric GCMs. However, in any realistic situation, errors introduced by sources outlined above are a serious obstacle to using sample statistics from a Monte Carlo method. There are a variety of ways to deal with sampling error in the regression step (or observation increment step) of an ensemble filter.

The first method for dealing with sampling errors is to ignore them (and treat results with less confidence). This often fails for atmospheric and oceanic applications because filters diverge from the true state of the system without some correction. Simple cases in low order models (section 5) can work when sampling error is ignored.

A second method is to make heuristic assumptions that reduce the confidence given to sample statistics during filter execution. Covariance inflation (Anderson 2001) in atmospheric and oceanic ensemble filters is an example. Covariance inflation can alleviate impacts of error from all the sources (section 2) and is predicated on the idea that serious errors in ensemble filters are those that lead to overconfidence in the ensemble estimates. In this case, the 'truth' is expected to be farther from the ensemble mean than is suggested by the prior sample. An overconfident prior can reduce the weight given to subsequent observations leading to further separation of the ensemble from the truth. This can eventually lead to filter divergence where the ensemble estimate is oblivious to the observations. Covariance inflation avoids this by increasing the prior variance. While many variants of this method have been applied, all are similar to the following. After the model is advanced in time (step 1 in Fig. 1), the prior sample variance of each model state variable, x_i, is increased by linearly 'inflating' the ensemble around its mean,

$$x_{i,j} = \sqrt{\gamma}(x_{i,j} - \bar{x}_i) + \bar{x}_i \qquad (1)$$

where the first subscript indexes the state variable element, the second indexes the ensemble member, an overbar is an ensemble mean, and γ is the covariance inflation factor. This method is surprisingly effective in a variety of applications. One can argue that much of the important information in the prior is retained by leaving the prior sample mean and correlations between state variables unchanged (Anderson and Anderson 1999). A covariance inflation of this form is required in many of the results in later sections.

A third method applies physically-based assumptions about the underlying probability distribution. Distance dependent covariance factors, also called covariance localization or just localization (Mitchell and Houtekamer 2000, Hamill et al 2001), are an example. Localization reduces the impact of an observation on a state variable in the regression (step 5 in Fig. 1) by a factor that is a function of some 'physical distance' between the two. The most common form for localization is the compactly supported Gaussian-like fifth order polynomial of Gaspari and Cohn (1999); called a GC envelope here. Another localization used in the literature is a boxcar function: all state variables within some distance of an observation are fully impacted by the observation and all other variables are not. This methodology was used by Anderson and Anderson (1999) and found to be inferior to the more smoothly varying GC method. A similar method has been used by Ott et al. (2004) and appears to produce good results. A fundamental problem with distance dependent localization is that a distance must be defined between an observation and each state variable. This study replaces heuristic localization with a more robust and theoretically motivated method for reducing errors in the regression step of filters.

A fourth method for dealing with sampling errors makes *a priori* statistical estimates of the error. This entails making a statistically-based estimate of the expected error in the regression coefficient determined from an ensemble given the numerical model and the set of observations. This appears to be extremely difficult in problems with large non-linear models and complicated forward observation operators.

A fifth method uses *a posteriori* statistical information from a filter to estimate corrections needed for a subsequent run of the same (or a similar) filtering problem. Assume that sample regression coefficients between an observation taken periodically at a fixed station and all state variables are available from a long successful ensemble filter assimilation. One can compute estimates of the sampling error under a variety of different statistical assumptions about the underlying 'true' distributions of the coefficients. The sampling error can then be corrected during a subsequent assimilation. The intrinsic difficulty with such methods is the required sample of the statistical relation between an observation and a state variable from a successful filter assimilation. In many cases, this begs the question since a good filter run cannot be made without knowing how to correct for the sampling errors. The approach may also require iterative development. Use of statistics from an initial poor assimilation lead to an improved assimilation; statistics from the improved assimilation might lead to a further improvement. The cost of such iterative methods may be prohibitive. A related method that is the closest published result to that

described here has been used by Houtekamer and Mitchell (2001) who split their ensemble into two parts and use statistics from one half to update the other half.

A sixth method uses a Monte Carlo technique to evaluate sampling errors in an ensemble filter. Since the ensemble filter is itself a Monte Carlo technique, this is referred to as a hierarchical Monte Carlo technique and is the subject of this report. Here, 'groups' of ensembles are used to understand regression sampling errors in the ensembles.

4. A hierarchical ensemble filter

Assume that m groups of n-member ensembles (m x n total members) are available. When using linear regression to compute the increment in a state variable, x, given increments for an observation variable, y^o, m sample values of the regression coefficient, β, are available. Neglecting other error sources, assume that the correct, but unknown, value of β is a random draw from the same distribution from which the m samples of β were drawn. The uncertainty associated with the sample value of β for a given ensemble implies that increments computed for state variable are also uncertain. One can minimize the impact of this uncertainty by introducing a confidence (weighting) factor, α, for the regression coefficients computed from the individual ensembles. α is chosen to minimize the expected RMS difference between the increment in a state variable and the increment that would be used if the 'correct' regression factor were used. The m-member sample of the regression coefficient is used to compute the α that minimizes

$$\sqrt{\sum_{j=1}^{m}\sum_{i=1,i\neq j}^{m}\left(\alpha\beta_i-\beta_j\right)^2} \tag{2}$$

where β_i is the regression factor from the ith group. This is equivalent to finding the α that minimizes

$$\sum_{j=1}^{m}\sum_{i=1,i\neq j}^{m}(\alpha^2\beta_i^2-2\alpha\beta_i\beta_j+\beta_j^2) \tag{3}$$

Taking a derivative with respect to α and setting to 0 gives

$$2\alpha\sum_{j=1}^{m}\sum_{i=1,i\neq j}^{m}\beta_i^2-2\sum_{j=1}^{m}\sum_{i=1,i\neq j}^{m}\beta_i\beta_j=0\,. \tag{4}$$

Noting that the first sum in (4) can be rewritten as

$$(m-1)\sum_{i=1}^{m}\beta_i^2 \tag{5}$$

and the second sum as

$$\left(\sum_{i=1}^{m} \beta_i\right)^2 - \sum_{i=1}^{m} \left(\beta_i^2\right) \tag{6}$$

gives

$$\alpha_{\min} = \left\{\left[\left(\sum_{i=1}^{m} \beta_i\right)^2 / \sum_{i=1}^{m} \beta_i^2\right] - 1\right\} / (m-1) \tag{7}$$

α_{min} can be computed directly from (7) as a function of the sample values β. It can also be written as a function of the ratio, Q, of the sample standard deviation to the absolute value of the sample mean

$$\alpha_{\min} = \max\left[\frac{m - Q^2}{(m-1)Q^2 + m}, 0\right]. \tag{8}$$

Figure 2 plots the regression confidence factor, α, as a function of the ratio Q for group sizes 2, 4, 8 and 16; if α is less than zero it is set to zero. Smaller groups have smaller values of the regression confidence factor (RCF), especially on the tail of the distribution. When the uncertainty is large (larger Q), small groups are unable to distinguish signal from noise and no information is contained in the regression coefficient sample.

The hierarchical Monte Carlo algorithm described here proceeds as follows. Each n-member ensemble is treated exactly as described in section 2 except for step 5, the regression computation. A regression coefficient, β_i, i = 1, …, m is computed for each of the m ensembles and the sample mean and standard deviation are computed, along with the ratio Q and the RCF is computed from (8). The regression is completed for each ensemble using its sample regression coefficient multiplied by α. The set of RCFs for a given observation and the set of model state variables is called a ‘regression confidence envelope’. As discussed below, the envelope can be viewed as a more rigorously derived replacement for heuristic localization approximations used in previous ensemble filter applications. It may also replace part of the covariance inflation required in these algorithms since some part of the regression error is often corrected by covariance inflation.

The hierarchical Monte Carlo approach applied here is philosophically similar to hierarchical closure schemes that have been applied to problems in geosciences, for instance turbulence closure schemes (Mellor and Yamada 1982). Like these, the hierarchical Monte Carlo technique must be 'closed' at some level. Here, a second level scheme in which 'groups' of ensembles are used is applied. A 'closure' is obtained by dealing with sampling error at level two (sampling error in the groups) using some other method. This makes sense only if the sampling error at level two is less severe than that from just using one of the methods above at level one (a single ensemble). The fourth method in section 3 can be used at the second level when the first, naïve method, is insufficient.

5. Regression confidence envelopes in the L96 model

A. Experimental design

The 40-variable model of Lorenz (Appendix 1), sometimes called the Lorenz-96 (L96) model, has been widely used in ensemble data assimilation research (Lorenz 1996). This model as configured here has 40 state variables equally spaced on a periodic one-dimensional domain and an attractor dimension of 13 (Lorenz and Emanuel 1998).

Results here are from 'perfect model' experiments in which a free integration of the model and a prescribed observational error distribution are used to generate synthetic observations that are then assimilated by the same model. A set of 40 randomly located 'observing stations' is used in most experiments. Observations from the 40 stations are available at every model time step. The 40 station locations are marked by the asterisks at the top of Fig. 4a. Forward observation operators, H, are linear interpolation between the two nearest model state variables while observational error distributions are prescribed as Gaussian with mean 0 for all stations. The observational error variance is varied between 10^{-7} and 10^{7} for different experiments.

As noted in section 2, most variants of previously documented ensemble filters are distinguished by the algorithm used to compute the observation variable increments in step 4. Here, the deterministic square root filter (Tippett et al., 2003) referred to as an Ensemble Adjustment Kalman Filter (EAKF) in Anderson (2001) is used. Most results do not change qualitatively when using other observation space update methods such as the classical ensemble Kalman filter

(Evensen 1994) or some of the more exotic observation update techniques discussed in Anderson (2003).

For efficient application of hierarchical filters small group sizes must produce accurate results. Group sizes of 2, 4, 8 and 16 have been evaluated and comparisons for different group sizes are examined in selected cases.

All assimilations are initiated with ensemble members selected from a climatological distribution of the L96 model generated by integrating slightly perturbed initial states for 100,000 time steps. This selection of the initial ensemble is the only 'random' part of this Monte Carlo algorithm. A 4000-step assimilation is performed; the first 2000 steps are discarded and results are shown from the second 2000 steps. A covariance inflation factor (selected from the set 1.0, 1.0025, 1.005, 1.0075, 1.01, 1.015, 1.02, 1.025, 1.03, 1.04, 1.06, 1.07, 1.08, 1.09, and 1.10 to 1.40 by intervals of 0.02) is tuned by experimentation to give the smallest time mean RMS error for the ensemble mean prior estimate over the final 2000 steps. Initial conditions for the second 2000 steps from the first group of the hierarchical ensemble filter are used as initial conditions for additional single filter assimilations discussed below (see section 5c for why).

RCF values are kept for each observation/state variable pair at each assimilation time and the time mean and median of these values are computed from the last 2000 steps. Additional assimilation experiments are performed using a traditional ensemble filter with the same ensemble size, but no traditional localization. Instead, the time mean (median) values of α from the hierarchical filter multiply the regression factors for each observation/state variable pair from the single ensemble. The covariance inflation factor for the time mean and median cases is selected so that it minimizes the time mean RMS error of the ensemble mean over the second 2000 steps of the single ensemble assimilation.

In addition, traditional ensemble filters with localization using a GC function are performed for each hierarchical filter case. The optimal value of the GC half-width is selected by searching from the set of values 0.025, 0.05, 0.075, 0.10, 0.125, 0.15, 0.20, 0.25, 0.30, 0.40, 0.50, 0.6, 0.75, 1.0 and 10^8 for that value producing the smallest time mean RMS error over the final 2000 steps of the experiment. For each GC half-width the optimal value of the covariance inflation is determined as for the other filters. Results from the combination of GC half-width and covariance inflation that minimizes the RMS are presented.

Time mean values of the RMS error of the prior ensemble mean state variables are presented as a rough measure of performance. The time mean of the RMS difference between ensemble members and the ensemble mean (a measure of the ensemble spread) is also computed. In ideal situations, error and spread values should be statistically indistinguishable. For most cases discussed here, the spread is slightly greater than the RMS error for the cases with the smallest RMS error (see Tables 1 to 6).

B. Small error amplitude results

Initially, tiny observational error variances of 10^{-7} are prescribed. Table 1 includes RMS error and spread values along with optimal values of the covariance inflation factor and GC localization half-width (for the standard filter cases) for a variety of ensemble and group sizes. Figure 3 compares time mean RMS errors for a variety of filters for ensemble sizes of 13, 8 and 5.

Results for any ensemble size n>13 combined with any number of groups m>1 are nearly quantitatively identical after long assimilations (for small groups, this may be much longer than the standard 2000 steps). Time median RCFs are 1.0 for nearly all observation/state variable pairs and time mean values are significantly greater than 0.99. Tiny amounts of variability in the RCFs as a function of time lead to slight variations in the error and spread results for the hierarchical filters (and the corresponding time mean and median cases) as a function of group size and ensemble size. This behavior indicates that for the hierarchical filter, the m different independent ensembles have all converged to nearly identical sample covariance estimates for the observations and the state variables. With such tiny error variances specified, the ensembles have nearly converged to the exact solution (the same solution would also be obtained with a classical Kalman filter in this case). All error sources outlined in section 2 are either eliminated (model error for instance) or reduced to levels so small that sampling error for both the regression and the observation increment becomes negligible. In other words, the assimilation is operating in the presence of trivial amounts of noise. Since all m samples of the regression factor are nearly the same, the true value is known nearly exactly. There is no need for the traditional distance dependent localization (optimal results for the standard filter in Table 1 are for no covariance localization).

While the sample covariances for ensembles larger than 13 converge to the same values, the sample covariance from any sub-sample of an ensemble does not converge. For instance, if one ran a case with a 100-member ensemble and computed the sample covariance (either between all the state variables or between the state variables and the observations) using just the first 20 ensemble members, this sub-sample covariance would generally be quite different from that for the whole ensemble sample. This is why the hierarchical ensemble filter technique in which independent groups of ensembles are used in the observation space increment is required, rather than partitioning a single large ensemble when computing RCFs.

As ensemble size is reduced below 14, the time means and medians of the RCFs demonstrate increasingly local structure. Figure 4a(b) shows the time mean (median) of the RCFs for an observation located at 0.6424, about 70 percent of the way between the 26th and 27th state variables, and all 40 state variables. For a 13-member ensemble, the median peaks at a value of 1.0 for state variables 26 through 28 and has a minimum of about 0.12 for state variables that are most remote from the observation. The mean peaks slightly below 1 and has a minimum of about 0.25. Reducing the ensemble size to 8 and 5 leads to progressively more localization of the impact of the observation. The median continues to have a maximum near 1 for state variable 27, but it is compact with non-zero values being confined to state variables progressively closer to the observation location. The maximum of the time mean is reduced and the RCF is increasingly sharply localized but does not go to 0 away from the observation location.

Figure 3 and Table 1 show that the time mean RMS error and spread increase as the ensemble size is reduced for all the ensemble methods (hierarchical, time mean, time median, standard). The optimal GC half-width for the traditional ensemble filter becomes smaller as the ensemble size is decreased, consistent with the behavior of the time mean and median RCFs.

For ensemble sizes less than 14, the sample covariance cannot accurately reflect the actual covariance because the L96 attractor is on a 13 dimensional manifold. Attempts to apply a traditional ensemble filter without localization with fewer than 14 ensemble members eventually leads to filter divergence. For the hierarchical ensemble filter, errors due to small ensemble size degeneracy can be characterized as noise for purposes of the regression. Smaller ensembles have larger errors in computing the regression coefficient and the corresponding RCFs are smaller (Fig. 4).

In the L96 model, there is remarkable similarity between the time mean/median RCF envelopes and the GC function. The GC localization with a half-width of 0.2 is displayed on Figures 4a/b for comparison; this is between the optimal half-widths of 0.3 for 8 ensemble members and 0.15 for 5 ensemble members. The central portion of the GC (which is nearly Gaussian) is very similar in shape to the time median regression confidence factors in Fig. 4b while the time mean shapes (Fig. 4a) are more sharply peaked. The time median envelopes fall off to zero as does the GC function while the flanks of the time mean envelopes are broader and do not go to zero. The RCF envelopes produced from 2000 step assimilations show evidence of sampling noise that makes them appear less smooth than the plots of the GC functions. This noise in the time mean and median RCFs is one factor that leads the time mean and median ensemble filters to produce time mean RMS errors that are in general slightly larger than those produced by the optimal GC localized ensemble filter (Figures 3, 5 and 7; Tables 1 to 6).

C. Varying observational error variance

Noise can also be introduced into the assimilations by increasing the uncertainty in the 'correct' solution distribution to the point where nonlinear effects become a significant factor. This can be done by reducing the information available to the assimilation by decreasing the number of observations, decreasing the frequency of observations, or increasing the observational error variance. Here, the observational error variance is increased and the resulting RCF envelopes examined.

Figures 6a(b) show the time mean (median) of the RCFs for observation location 0.6424 as the observational error variance is increased to 10^{-5}, 10^{-3}, 0.1, 1.0, 10.0 and 10^{7}. Table 2 shows the error, spread, and parameter settings for filters applied in these problems while Fig. 5 compares the time mean RMS errors. Results are for 4 groups and ensemble size 14. Qualitatively, as the error variance increases, the impact on the RCF envelopes is similar to that from reducing the ensemble size. Larger error variance leads to more compact time median RCFs and more strongly peaked time means. The case with error variance 10^{7} is associated with prior ensembles that have climatological distributions since the observations have negligible impact (so there is no point in comparing error results from different types of filters). The RCF envelopes in this case have an interesting double peaked structure. When beginning an assimilation from a climatological distribution (a safe and simple choice in many cases), this is the appropriate way to weight the regression factors for surrounding state variables.

Figure 6 also motivates the experimental design in section 5a. The RCFs for an ensemble with climatological error variance are very different from those for ensembles with smaller errors. If the standard, time mean, or time median filters are started with climatological initial conditions, they require narrow RCFs (or GC envelopes) in order to avoid assimilating noise initially. However, such narrow envelopes are highly sub-optimal for assimilating in the equilibrated cases with small error. The hierarchical ensemble filter is able to transition smoothly from large to small error situations during an assimilation. Initial conditions from the end of the first 2000 steps of the hierarchical ensemble filter were appropriate for starting high quality single filter assimilations.

Sampling error introduced into the regression by reducing the information available from the observations is qualitatively similar to that from degenerate ensembles (section 5b). In most previous work on ensemble filtering methods, the issues of degeneracy and noise introduced by lack of observational information were viewed as distinct problems possibly requiring independent analysis and solution. While gaining an a priori understanding of these two sources of error may require independent analysis, the hierarchical ensemble filter approach addresses both types of errors.

A significant advantage of the hierarchical ensemble filter over previous ensemble filters is that it does not require tuning of a localization function like the GC half-width. Heuristic tuning can require large numbers of iterations, even in one dimensional, univariate models like L96 with simple forward observation operators. Heuristic tuning of localization becomes much more difficult in multivariate three-dimensional models with complex forward observation operators.

D. Impact of group size on results

In order for hierarchical ensemble filters to be practical for many large applications, the number of groups required for good results must be small. Figure 7 shows RMS errors for different group sizes for the 40 random observation, 1.0 error variance case with 14 ensemble members while Table 3 shows more details on these assimilations. Figure 7 shows that for hierarchical group filters, increasing group size leads to a gradual reduction of error. The corresponding time mean and median filters also show this behavior, but the impact of group size appears to be less significant. Figures 8a (b) show the time mean (median) RCFs for group sizes 2, 4, 8 and 16.

Close to the observation, group size has almost no impact on the RCFs. Larger differences are seen in the tails where increasing group size leads to progressively smaller values of the RCF. This is sampling error in the groups (second level sampling error) in the hierarchical filter. Table 3 includes the time mean error and spread from the hierarchical filter and the corresponding time mean and median filters. Time mean error decreases with group size with groups of 8 and 16 showing smaller errors than does the best standard filter. For this particular assimilation, there is some suggestion that group sizes much above 8 will have relatively little additional impact. However, it is important to note that this behavior is clearly a function of the model and observations. With large enough models, hierarchical filters can have the same type of sampling problems as traditional filters, only at level two. While the hierarchical approach ameliorates the severity of this sampling error (in the same way that the use of GC localization does), there still may be a need for additional effort to eliminate the impact of second order sampling error. The most straightforward way to do this is to include a heuristic localization (like GC) in concert with the hierarchical group filter approach.

E. Time variation of regression confidence factors

There is considerable variability in time for the RCFs between a given observation and state variable. Figures 9a and 9b show the time evolution from assimilation steps 1000 to 1050 of the RCFs for the 1.0 observational error variance case with 14-member ensembles for groups of 16 and 2 respectively. The time mean (median) of the full 2000 steps can be seen in Figs. 6a (b). Close to the observation location, the median (Fig. 6b) indicates that the confidence factor is almost always close to 1. However, Figs. 9a and 9b show that there are occasions when the value is small, for instance near time 1033. Recall that the RCF is a function of the ratio of the sample standard deviation of the regression coefficients to the absolute value of the mean value (7). Areas of small RCF close to the observation location are associated with cases where the mean of the regression coefficient is small. This occurs when the dynamics of the model is dominated locally by a set of standing or propagating waves all of which happen to have a peak or trough near to the observation location.

The group 2 results (Fig. 9b) are noisier, with lots of significantly non-zero values for state variables remote from the observation. The relative lack of non-zero values for the group 16 results suggests that much of the non-zero part of the time mean RCF far from the observation is related to second order sampling error. There is also evidence of noise in the group 2 results close

to the observation where there are more localized areas of small values imbedded in values that are close to 1 than for the 16 group results.

F. What is noise?

It is important to remember that the definition of ‘noise’ for the assimilation algorithm depends on the details of the algorithm itself. Here, any aspect of the filtering system that is inconsistent with the choice to perform a linear regression to find the relation between the observation variable and state variable increments can be categorized as noise. One could find more complicated ways to relate the two variables, for instance local linear regressions (Anderson 2003) or nonlinear regression. It seems likely that very large ensembles would be required to perform regressions of this type. It is unclear how much error is introduced by the use of linear regression in atmospheric and oceanic problems of interest. More sophisticated error analysis should be applied to answer this question in the future.

6. Regression confidence factors for different observation types

The previous section showed that the shape of RCF envelopes depends on the observational error variance and the ensemble size for a given model. This section examines the impact of observation distribution and type. This is difficult to do in a meaningful way in a one-dimensional, univariate model like L96, but this section examines a few simple examples and relates them to observation types in more realistic models.

A. Spatially inhomogeneous observations

Although the advent of satellite remote sensing greatly reduced disparities in data density encountered in global atmospheric assimilation, there are still suggestions that regions like North America are significantly better observed than regions over the southern continents. In mesoscale prediction applications, radar reflectivity data is available in regions where hydrometeors are falling but unavailable in other regions. There is great disparity in the density of ocean observations as a function of location and depth although this, too, is slowly being rectified by new observing systems.

Figure 10 depicts the locations of 40 observations characterizing a well-observed and a poorly-observed region in the L96 system. 39 observations are equally spaced between 0.011 and 0.391 (they are offset slightly from the state variable locations to avoid the significant boost in performance that occurs when no interpolation is required for a forward observation operator) while the 40th observation is located at 0.701. All observations have an error variance of 1.0. The time mean RMS error of three different assimilations as a function of state variable is also depicted. In the data dense region, there is no visible difference between the error characteristics of a standard filter, a 4-group filter, and its corresponding time mean filter, all run with 14 ensemble members. However, in the data sparse region, the hierarchical filter and the time mean filter show reduced time mean error (also see Table 4).

Figures 11a (b) show the time mean (median) RCF envelopes for observations located at 0.011, 0.191, 0.391, and 0.701. RCF envelopes for the observation at 0.191, in the middle of the well-observed region where time mean error is small, are relatively wide, consistent with low error cases from the previous section. Observations located in higher time mean error areas away from the center of the densely observed region have progressively narrower RCF envelopes. The RCF for the observation at 0.701, in the middle of the poorly-observed region, is very narrow and displays the two-lobed structure that was found for very large errors in the previous section. The observation at 0.011, immediately downstream of the poorly observed region, has intermediate width but a triply-peaked structure not been seen in previous examples.

The optimized GC half-width for the traditional filter is 0.2, relatively broad compared to the RCFs in data sparse regions. The result is that the standard filter performs well in the data dense region but has significant sampling errors in the data sparse region. These sampling errors result in increased RMS error along with reduced spread in poorly sampled regions.

It seems likely that real atmospheric and oceanic prediction problems continue to present significant disparities in observation spatial density and expected assimilation time mean RMS error. Hierarchical filters can deal with these areas, but traditional filters would require spatially-varying localizations which would be very time-consuming to tune. The effects of temporal variations in observation density are similar and may also be significant for real assimilation problems. Traditional data is denser at 00Z and 12Z in the atmosphere while many remote sensing observations are only available during certain orbital periods or under certain atmospheric conditions. In these instances, a hierarchical filter might perform much better than a standard

filter or a time mean / time median filter. This is similar to arguments that are made for the use of standard filters; they owe some of their advantages to being able to resolve time-varying correlations between observations and state variables. One can similarly argue that a hierarchical filter may be able to resolve time-varying components of the sampling error associated with standard filters.

B. Spatially averaged observations

While most traditional observations of the ocean and atmosphere can be regarded as sampling at a discrete point, many remote sensing observations are more accurately viewed as a weighted sampling of a finite region of space (and possibly time). For instance, satellite observations measuring the total amount of water vapor in an atmospheric column are used in many operational assimilation systems. Many other satellite observations have fields of view that are not small compared to the spacing of model gridpoints (especially in the vertical). Forward operators for these observations must be viewed as a weighted average of a large number of model gridpoints.

Spatially averaged observations are simulated in the L96 model by defining a forward observation operator that averages a 0.375 wide domain of the state variables. Given an observation at x_o, the forward operator averages 15 standard forward observation operators located at $x_o + 0.025k$, where k = -7, -6, …, 6, 7. The 40 observing stations from section 5a are used with error variance of 4.0. Filters with 14-member ensembles are used to assimilate these observations.

Figure 12 shows the time mean and median RCF envelopes for the observation located at 0.6424. Both have a broad, relatively flat maximum with values around 0.6 centered on the observation location. The median has an abrupt drop to 0 near the edge of the averaging region while the mean decreases more gradually to a minimum of about 0.15.

Table 5 shows the time mean error results for a 4-group hierarchical filter, its time mean and time median, and a standard filter. Relatively large values of the covariance inflation parameter were required for the hierarchical, time mean and median filters suggesting that the level of sampling error in this problem is larger than in previous examples. The time median filter had somewhat larger time mean RMS error. No pair of GC localization half-widths and covariance inflation

factors was found for which the standard filter did not diverge in this case. This does not prove that a 14-member standard filter that works for this problem cannot be found, but it does mean that it would work only for a very narrow range of parameters. On the other hand, the 4-group hierarchical filter did not diverge and produced roughly similar RMS error for a wide range of covariance inflation factors.

Causes of the standard filter's problem are apparent. The GC functions are defined so that state variables close to the observation location receive the full impact of the observation, while the hierarchical filter indicates that this is inappropriate for the averaged observations. The result is that the standard filter increments for a given state variable and a relatively small GC half-width are too heavily influenced by a group of close observations and insufficiently influenced by more distant observations. The local over-weighting can be corrected by large covariance inflation, but only by sacrificing even more of the information from more distant observations. Observations for which the forward operators involve averaging in time, or combinations of spatial and temporal averaging should also prove challenging to traditional filters.

C. Combinations of physically separated observations

Even in the spatially averaged case of the previous section, the RCF envelopes in the L96 model retain a Gaussian shape. This sometimes allows the relative performance of the traditional GC localizations to remain competitive. This section presents a somewhat unrealistic forward operator designed to generate highly non-Gaussian RCF envelopes. The forward operator for a given observation location is the standard interpolation plus the interpolation to the location on the opposite side of the cyclic domain (0.5 away); observational error variance is 4.0. It is conceivable that some remote sensing observations sample two physically separated locations like this. For instance, some limb sounding satellite instruments may be sensitive to a specific altitude or radar reflectivities may be subject to aliasing from range folding.

Figure 13 shows the time mean and median RCF envelopes for the observation at 0.6424 from a 4-group, 14-member hierarchical filter assimilating 40 observations with the standard locations. Both have two distinct peaks with maximum value about 0.95 for the median and 0.85 for the mean. It is impossible to closely approximate these shapes with a GC localization, although it is possible to approximate one of the peaks. Table 6 shows the time mean RMS error and spread for the various filter assimilations in this case. While the hierarchical filter and the time mean and

median are able to produce normalized errors of approximately 1.5, the best standard filter is for a very narrow GC localization of 0.10 with a comparatively large covariance inflation and still produces errors of 3.9. The standard filter is able to transfer information from an observation to half of the associated state variables in an approximately correct fashion. Because it misses the other half, it has a larger overall error, leading to a narrower optimized GC half-width and an error that is more than twice as large as that from the hierarchical filter.

7. Assimilation of observations from different times

There are many applications in which assimilation of observations from times other than the time associated with the state estimate available from model integrations is of interest. In some real-time forecast implementations, there are observations that may not be available at forecast centers until well after they are taken. It may be necessary to begin integrating the forecast model before all the observations at the time of the forecast initial conditions have been received. However, one could assimilate these observations to modify the state of an ongoing forecast at a later time (i.e. use observations from 12 GMT to modify the forecast state at 18 GMT).

As observations become increasingly ‘distant’ from the state estimate in time, one expects sampling noise to become an increasingly large problem (Majumdar et al 2002). This is examined using a 4-group, 14-member hierarchical filter. The 40 randomly located observations used in section 5a with observational error variance 1.0 are assimilated. In addition, a single extra observation, located at 0.6424, is available from a previous time for every assimilation time. The time lag between when this observation was taken and when it is available for assimilation is varied from 0 to 100 assimilation times in a series of 101 additional experiments. The forward observation operator is applied to the state at the time the delayed observation was taken and archived until the time at which it is to be assimilated. In real applications, the ensemble of state estimates from previous times could be archived if necessary and the forward operator applied to the old states when the observation was received.

Figure 14 shows the time mean RCF envelopes for the delayed observation as a function of the lag time by which receipt of the observation is delayed. For short lag times, a horizontal cross section through Fig. 14 looks very similar to the thick solid curve in Fig. 6a; that experiment differs only in not having an additional lagged observation available. As the lag time increases, the maximum of the RCF shifts downstream and is gradually reduced. This reflects the advection

of ‘information’ by the model from the delayed observation location as time advances. The amplitude decrease with lag time reflects increasing noise in the sample regressions between the lagged observation and the state as the expected relationship becomes increasingly weak.

This type of behavior would present significant problems for a traditional filter with a GC localization. While a GC function that closely approximates the appropriate RCFs for short lag times can be constructed, it could be difficult to determine how to modulate the amplitude and shift the GC localization downstream in the appropriate fashion.

Using appropriate RCFs for observations taken at different times is also relevant to other important ensemble applications. First, there is interest in the development of ensemble smoothers that use observations both from the past and the future to develop an accurate analysis of the state of the system (Li and Navon 2001). Since models of interest usually cannot be integrated backward in time, and since one of the advantages of ensemble methods is that there is no need for linearized or adjoint models, fixed lag smoothers appear to be the method of choice for ensemble application. In fixed lag smoothers, an initial ensemble filter is applied out to some time $t_0 + t_{lag}$. Observations from times past t_0 out to t_0+t_{lag} can then be assimilated into the state estimate at time t_0. In this case, a hierarchical filter could be used to compute appropriate RCFs for an observation in the future impacting state estimates in the past. A plot of the impact of a future observation at location 0.6424 as a function of lag would look very similar to Fig. 14 reflected around a vertical line at 0.6424.

Second, ensembles are a natural tool for targeted observation experiments (Bishop et al 2001, Khare and Anderson 2004). These experiments assume that there exist certain observations whose deployment can be controlled (Bergot et al 1998). For instance, a plane carrying dropsondes can be dispatched to a particular location of interest to obtain additional atmospheric soundings. Normally, there is some delay involved in deploying targeted observations. For instance, a plane must fly to the appropriate location from its base of operations. Hence, targeted observation experiments normally involve using forecasts initiated at time t_0 to determine the deployment of observations at time $t_0 + t_{tar}$ in order to improve some forecast element at time $t_0 + t_{ver}$, which is even further in the future (Berliner et al 1999, Reynolds et al 2000).

The reduction in expected spread for the function of the state variables at time t_0+t_{ver} can be computed by regressing the expected reduction in the spread of the ensemble estimate of an

observation at time $t_0 + t_{tar}$ onto the verification quantity using the ensemble sample joint distribution of the potential targeted observation and the function at the verification time. This regression will be subject to sampling noise (Hamill and Snyder 2001). RCFs, like those displayed in Fig. 14, are required to compute meaningful estimates of the impact of an observation at the targeting time on the function of the state at the verification time.

8. Estimating non-local model parameters

Although few practical applications have yet resulted, the use of assimilation to obtain values for model parameters has been discussed in the literature (Derber 1989). Anderson (2001) used assimilation in the L96 model to obtain the value of the parameter F found in eq. A1. In the perfect model control run, the value was constant at 8.0, but the assimilating model did not know this and each ensemble member had its own estimate of the parameter that was modified only by the assimilation.

Using traditional ensemble methods required addressing the question ‘what is the distance between the parameter and an observation at a given location’. Clearly, an RCF may be required to avoid incorporating unnecessary noise into the estimate of a parameter, especially in larger multi-dimensional models.

A modified version of the L96 model in which the forcing parameter, F, is treated as a 41^{st} model variable is used for a hierarchical filter assimilation. F is fixed at 8.0 in the perfect model integration that generates synthetic observations. Observations are the same as in section 5a with an observational error variance of 1.0. A 4-group 14-member ensemble filter produces a time mean RMS error of 1.714 with spread 1.868. The time mean RMS error in the estimate of the parameter F is 0.01212 (initial values for F are randomly selected uniformly from the range 7 to 9). Time mean values of the RCF between the individual observations and the ‘state variable’ F vary between 0.18 and 0.24 implying that F is weakly influenced by any single observation. This is not a surprise given that F impacts the state globally while the observations are correlated only with a local portion of the state. It is intriguing that this experiment gives the lowest time mean RMS error for the state variables of any case examined (even the 16-group case and the 56-member traditional ensemble), even though all other cases know the value of F exactly. Apparently the uncertainty introduced into the prior estimates by the varying values of F corrects for other sources of error. Extending this result to parameters in large realistic models may

demonstrate that uncertainty in parameters leads to a significant improvement in assimilation RMS error.

9. PE Dynamical Core on a sphere

As noted, most simple examples of assimilation in the L96 model result in RCF envelopes that are approximately Gaussian. It is difficult for hierarchical group filters to perform better than the best heuristically-tuned GC localized filters in these cases. When observations that lead to less Gaussian RCF envelopes are used, the relative performance of the hierarchical group filters is much better.

More complicated, multivariate models in higher dimensions might be expected to produce less Gaussian RCFs that might in turn provide additional motivation for using hierarchical ensemble filters for assimilation. Here, the dynamical core of the GFDL B-grid climate model is used to do a preliminary exploration of this issue.

The B-grid core (Wyman 1996, Anderson et al. 2004) is configured in one of the lowest resolutions that supports baroclinic instability in a Held-Suarez configuration (Held and Suarez 1994) with a 30 latitude by 60 longitude grid and 5 vertical levels. Assimilation with traditional ensemble filters has been explored in this model in Anderson et al. (2005).

1800 randomly located surface pressure stations are fixed on the surface of the sphere and provide observations every 12 hours with an observational error standard deviation of 1 hPa. This set of surface pressure observations is sufficient to constrain not only the surface pressure field, but also the rest of the model variables.

Assimilations are performed over a 100-day period, starting with ensemble members drawn from a climatological distribution. Results here are for a 4-group 20-ensemble hierarchical ensemble filter. No covariance inflation is applied in this model.

Figure 15 shows the RCF envelope for a surface pressure observation at 22.7 N 61.4E with surface pressure state variables while Fig.16 shows the RCFs for the same observation and the V field at model levels 2 to 5. The RCFs, especially for V, are not well described by Gaussians, with multi-modal structures being apparent. The V RCFs do not have maxima of 1.0 at any level. The

structure in the vertical is also non-Gaussian with a minimum at the middle levels. The vertical structure of the RCF in experiments with realistic GCMs (for instance NCAR's CAM2.0) is even more complex and suggests that vertical correlation errors may be a major source of error in traditional ensemble filters. Clearly, the use of hierarchical filters has some advantages in more complex hierarchical models like the B-grid core.

A complete discussion of hierarchical filters in the B-grid model has been completed and will appear in a later publication. However, it is important to note that the RCFs are even more complex than indicated in Fig. 15 and Fig. 16 because they vary considerably depending on the latitude of the observation.

10. Discussion and conclusions

Given ensembles of state variable assimilation results from a successful assimilation and a description of the observing system, it is possible to approximate the RCFs without using a group filter. Knowing the distribution of the correlation between an observation and a state variable allows the computation of the expected time mean RCF. However, in order to compute the RCF, one needs a successful assimilation that in turn requires a high-quality localization of observation impact. The hierarchical ensemble filter presented here provides a mechanism for producing high quality assimilations without any a priori notion of how the impact of observations on state variables should be localized. Instead, a Monte Carlo technique is applied to limit the impacts of ensemble sampling error. This technique can deal with situations in which the appropriate 'localization' of the impact of an observation on a state variable is a complicated function of both observation type, state variable type, spatial location of both the observation and the state variables, and time.

In the low order L96 results that comprise most of this report, the performance of hierarchical filters and traditional ensemble filters that use a prescribed localization is roughly equivalent. However, the traditional ensemble filters require many experiments to tune. In addition, when starting from climatological distributions (which is perhaps the safest way to design assimilation experiments of this sort) the traditional ensemble filter is unable to deal with the initial phases of the assimilation with large prior error spread while still providing high quality assimilations once the initial error is reduced. An hierarchical filter application is required to provide initial conditions for the traditional ensemble filter in such cases.

In complex multivariate models like atmospheric prediction models, a priori specification of localization functions becomes problematic. Even the appropriate distance between spatially and temporally collocated observations and state variables becomes unclear when the observation and state are of different variable types. When the added complexity of non-local forward operators, vertical and horizontal separation, and temporal separation are considered, the problem becomes very complex indeed. Matters are only made worse by the fact that ensemble size and error size also come into play. While naively localized ensemble filters have produced decent results in large multivariate models, it appears likely that performance can be enhanced by applying hierarchical filters.

Cost, of course, is an important consideration. While results from even small numbers of groups appear to lead to good estimates of sampling error in ensembles, this added expense may be the straw that breaks the camel's back for operational application. However, using short experiments to produce statistics for creating localization functions is probably affordable. The resulting ensemble filters using these statistics for localization cost no more than traditional ensemble filters.

Hopefully, addressing sampling error and other error sources in ensemble filters will continue to make them even more competitive with other existing assimilation methods and easier for non-experts to apply.

Appendix: The Lorenz-96 model

The L96 (Lorenz 1996) model has N state variables, X_1, X_2, …, X_N, and is governed by the equation

$$dX_i/dt = (X_{i+1} - X_{i-2})X_{i-1} - X_i + F, \qquad \text{(A1)}$$

where i = 1, …, N with cyclic indices. Here, N is 40, F = 8.0, and a fourth-order Runge-Kutta time step with dt = 0.05 is applied as in Lorenz and Emanuel (1998).

Figure Captions:

1. Schematic representation of the implementation of the ensemble filter used here with possible error sources marked by numbers 1 through 5.

2. Regression confidence factors as a function of the ratio Q of regression sample standard deviation to the absolute value of the sample mean for 2 (thin solid), 4 (thick dashed), 8 (thin dashed) and 16 (thick solid) groups.

3. 2000-step time mean RMS error (normalized by the observational error standard deviation) for observational error variance of 1e-7 with 13, 8 and 5 member ensembles for standard Gaspari Cohn localized filter, 4 group filter and corresponding time mean and time median filter, and 8 group filter and corresponding time mean and median filter.

4. 2000-step time mean (a) and time median (b) regression confidence factors for Lorenz-96 model assimilations for an observation located at 0.6424 and all 40 state variables. The asterisks at the top of (a) indicate the position of 40 randomly located observations with observational error variance of 10^{-7}. Results are from hierarchical ensemble filters with 4 groups and ensemble sizes of 5 (thick dashed), 8 (thick solid) and 13 (thin solid). Also shown is a Gaspari-Cohn localization function for a half-width of 0.2 (thin dash-dotted).

5. 2000-step time mean RMS error (normalized by the observational error standard deviation) for observational error variances of 10^{-5}, 10^{-3}, 0.01, 1.0 and 10.0 for standard Gaspari Cohn localized filter, 4-group filter and corresponding time mean and time median filters, all with ensemble size 14.

6. 2000-step time mean (a) and time median (b) regression confidence factors for Lorenz-96 model assimilations for an observation located at 0.6424 and all 40 state variables. Results are from hierarchical ensemble filters with 4 groups and 14 ensemble members. The observational error variance for 40 randomly located observations is 10^{-5} (thick dash-dotted), 10^{-3} (thick dashed), 10^{-1} (thick solid), 1.0 (thin dash-dotted), 10.0 (thin dashed) and 10^{7} (thin solid).

7. 2000-step time mean RMS error (normalized by the observational error standard deviation) for observational error variance of 1.0 and hierarchical filters with group sizes of 2, 4, 8 and 16 along with the corresponding time mean and time median filters.

8. 2000-step time mean (a) and time median (b) regression confidence factors for Lorenz-96 model assimilations for an observation located at 0.6424 and all 40 state variables. The observational error variance for 40 randomly located observations was 1.0. Results are for 2 groups (thick solid), 4 groups (thin dash-dotted), 8 groups (thin solid) and 16 groups (thin dashed) of 14 member ensembles.

9. Regression confidence factor for Lorenz-96 model assimilations for an observation located at 0.6424 and all 40 state variables as a function of time between assimilation steps 1000 and 1050 of a sixteen group (a) and two group (b) hierarchical filter with ensemble size of 14. The contour interval is 0.2 with values greater than 0.4 shaded.

10. 2000-step time mean RMS error as a function of model state variable for 14 member ensemble assimilations of 40 observations with observational error variance of 1.0 whose location is indicated by the asterisks at the top of the plot. Results are plotted for a 4-group hierarchical filter (thin dashed), a filter using the time mean regression confidence factors from the 4-group filter (thin solid), and a traditional ensemble filter with a Gaspari-Cohn localization with half-width 0.2.

11. 2000-step time mean (a) and time median (b) regression confidence factors for Lorenz-96 model 4-group hierarchical filter with 14 member ensembles. The observations are as for Fig. 10 and the locations are marked with asterisks at the top of the plot. Regression confidence factors are plotted for the observation at location 0.011 (thick solid), 0.191 (thin dashed), 0.391 (thin solid) and 0.701 (thick dashed).

12. 2000-step time mean (thick) and time median (thin) regression confidence factors for Lorenz-96 model assimilations with a 4-group 14-member ensemble filter for an observation located at 0.6424 and all 40 state variables. The forward observation operators are the average of 15-point observations surrounding the central location (for instance at locations $0.6424 + 0.025k$, where $k = -7, -6, \ldots 6, 7$). The observational error

variance is 4.0 and the mid-points of the 40 observation locations are the same as marked in Fig. 4a.

13. 2000-step time mean (thick) and time median (thin) regression confidence factors for Lorenz-96 model assimilations with a 4-group 14-member ensemble filter for an observation located at 0.6424 and all 40 state variables. The forward observation operators are the average of the observation location and the location 0.5 removed from the location (two points on opposite sides of the cyclic domain). The observational error variance is 4.0 and the mid-points of the 40 observation locations are the same as marked in Fig. 4a.

14. 1000-step time mean regression confidence factors for simulated time-lagged observation at location 0.6424 for assimilations with a 4-group 14-member hierarchical ensemble filter. The base observation set is the 40 random observations as marked in Fig. 4a with error variance 1.0. The plot shows the regression confidence factors for an observation that was taken n assimilation steps prior to the time at which it was assimilated. The contour interval is 0.1.

15. Time mean regression confidence factor for a pressure observation at 22.7 N 61.4 E and surface pressure state variables in the GFDL B-grid AGCM. Contour interval is 0.1.

16. Time mean regression confidence factors for the same surface pressure observation as in Fig. 15 but with v at each of the model levels 2 through 5. Contour interval is 0.1.

Table Captions:

1. Comparative RMS error and spread for assimilations with 40 randomly located observations with 10^{-7} error variance for ensemble sizes >13, 13, 8 and 5 and for 4-group and 8-group filters with corresponding time mean and time median filters and a traditional filter with Gaspari Cohn localization (base). Error and spread values are normalized by the observational error standard deviation.

2. Comparative RMS error and spread for assimilations with 40 randomly located observations with various error variances. 14-member ensembles are used for 4-group

and 8-group filters with corresponding time mean and time median filters. Traditional filters (base) for ensemble sizes of 14 and 56 are also included. Error and spread values are normalized by the observational error standard deviation.

3. Comparative RMS error and spread for assimilations with 40 randomly located observations with 1.0 error variance for ensemble size 14 for 2, 4, 8 and 16 groups with corresponding time mean and time median filters and a traditional filter with Gaspari Cohn localization (base).

4. Comparative RMS error and spread for assimilations with 40 observations located as in Fig. 4 to form a data dense and data void region with 1.0 error variance for 14-member ensembles.

5. Comparative RMS error and spread for assimilations with 40 randomly located observations with forward observation operators being the average of 15-point observations surrounding the central location (for instance at locations 0.6424 + 0.025k, where k = -7, -6, …6, 7). The observational error variance is 4.0 and the mid-points of the 40 observation locations are the same as marked in Fig. 4a.

6. Comparative RMS error and spread for assimilations with 40 randomly located observations with forward observation operators being the average of the observation location and the location 0.5 removed from the location (two points on opposite sides of the cyclic domain). The observational error variance is 4.0 and the observation locations are the same as marked in Fig. 4a.

Bibliography

Anderson, J. L. and S. L. Anderson, 1999: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Wea. Rev*., **127**, 2741-2758.

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. and 33 contributing authors, 2004: The new GFDL global atmosphere and land model AM2/LM2: Evaluation with prescribed SST simulations. Accepted in *J. Climate*.

Anderson, J. L., B. Wyman, S. Zhang and T. Hoar, 2005: Assimilation of Surface Pressure Observations using an Ensemble Filter in an Idealized Global Atmospheric Prediction System. Submitted to *J. Atmos. Sci*., May, 2004.

Bergot, T., G. Hello, A. Joly and S. Malardel, 1998: Adaptive observations: A feasibility study. *Mon. Wea. Rev*., **127**, 743-765.

Berliner, L. M., Z.-Q. Lu, and C. Snyder, 1999: Statistical design for adaptive weather observations. *J. Atmos. Sci*., **56**, 2536-2552.

Bishop, C. H., B. J. Etherton and S. J. Majumdar, 2001: Adaptive sampling with the ensemble transform Kalman filter. Part 1: Theoretical aspects. *Mon. Wea. Rev*., **129**, 420-436.

Buizza, R., M. Miller and T. N. Palmer, 1999: Stochastic representation of model uncertainties in the ECMWF ensemble prediction system. *Quart. J. Roy. Meteor. Soc*., **125**, 2887-2908.

Dee, D. P. and R. Todling, 2000: Data assimilation in the presence of forecast bias: the GEOS moisture analysis. *Mon. Wea. Rev*., **128**, 3268-3282.

Derber, J. C., 1989: A variational continuous assimilation technique. *Mon. Wea. Rev*., **117**, 2437-2446.

Evensen, G., 1994: Sequential data assimilation with a nonlinear quasigeostrophic model using Monte Carlo methods to do forecast error statistics. *J. Geophys. Res*.**, 99(C5),** 10143-10162.

Gaspari, G. and S. E. Cohn, 1999: Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteor. Soc*., **125**, 723-757.

Hamill, T. M. and C. Snyder, 2001: Using improved background error covariances from an ensemble Kalman filter for adaptive observations. *Mon. Wea. Rev*., **130**, 1552-1572.

Hamill, T. M., J. S. Whitaker and C. Snyder, 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.

Held, I. M. and M. J. Suarez, 1994: A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models. *Bull. Amer. Meteor. Soc*., **75**, 1825-1830.

Houtekamer, P. L. and H. L. Mitchell, 1998: Data assimilation using an ensemble Kalman filter technique. *Mon. Wea. Rev*., **126**, 796-811.

Houtekamer, P. L., and H. L. Mitchell, 2001: A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Wea. Rev*., **129**, 123-137.

Houtekamer, P. L., H. L. Mitchell, G. Pellerin, M. Buehner, M. Charron, L. Spacek, and B. Hansen, 2004: Atmospheric data assimilation with the ensemble Kalman filter: results with real observations. *Mon. Wea. Rev*., in review.

Jazwinski, A. H., 1970: Stochastic processes and filtering theory. Academic Press, 376 pp.

Kalman, R. E., 1960: A new approach to linear filtering and prediction problems. *Transactions of the AMSE Journal of Basic Engineering*, **82D**, 35-45.

Keppenne, C. L. and M. M. Rienecker, 2002: Initial testing of a massively parallel ensemble Kalman filter with the Poseidon isopycnal ocean general circulation model. *Mon. Wea. Rev*., **130**, 2951-2965.

Khare, S. P. and J. L. Anderson, 2004: A Bayesian framework for adaptive observing network design: An analysis of deterministic ensemble square root filter implementations. In press with *Mon. Wea. Rev*.

Li, Z. and I. M. Navon, 2001: Optimality of variational data assimilation and its relationship with the Kalman filter and smoother. *Quart. J. Roy. Meteor. Soc*., **127**, 661-683.

Lorenc, A. C., 2003: The potential of the ensemble Kalman filter for NWP. *Quart. J. Roy. Meteor. Soc.*, **129**, 3183-3204.

Lorenz, E. N., 1996: Predictability: A problem partly solved. Proc. ECMWF Seminar on Predictability, Vol. I, Reading, United Kingdom, ECMWF, 1-18.

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

Majumdar, S. J., C. H. Bishop, B. J. Etherton, I. Szunyogh and Z. Toth, 2002: Can an ensemble transform Kalman filter predict the reduction in forecast-error variance produced by targeted observations? *Quart. J. Roy. Meteor. Soc*., **127**, 2803-2820.

Mellor, G. L. and T. Yamada, 1982: Development of a turbulent closure model for geophysical fluid problems. *Rev. Geophys. Space Phys*., **20**, 851-875.

Mitchell, H. L. and P. L. Houtekamer, 2000: An adaptive ensemble Kalman filter. *Mon. Wea. Rev*., **128**, 416-433.

Mitchell, H. L., P. L. Houtekamer and G. Pellerin, 2002: Ensemble size, balance, and model-error representation in an ensemble Kalman filter. *Mon. Wea. Rev*., **130**, 2791-2808.

Ott, E., B. Hunt, I. Szunyogh, A. Zimin, E. Kostelich, M. Corazza, E. Kalnay, D. Patil, and J. Yorke, 2004: A local ensemble Kalman filter for atmospheric data assimilation. To appear in *Tellus* (refs should be available by completion of reviews).

Pham, D. T., 2001: Stochastic methods for sequential data assimilation in strongly non-linear systems. *Mon. Wea. Rev*., **129**, 1194-1207.

Reynolds, C. A., R. Gelaro and T. N. Palmer, 2000: Examination of targeting methods in a simplified setting. *Tellus*, **52A**, 391-411.

Snyder, C. and F. Zhang, 2003: Assimilation of simulated Doppler radar observations with an ensemble Kalman filter. *Mon. Wea. Rev*., **131**, 1663-1677.

Tarantola, A., 1987: Inverse Problem Theory. Elsevier Science, 613 pp.

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.

Wyman, B. L., 1996: A step-mountain coordinate general circulation model: Description and validation of medium range forecasts. *Mon. Wea. Rev*., **124**, 102-121.

Ensemble size	Group Size & Type	GC Half-width	Covariance Inflation	Time Mean RMS Error	Time Mean Spread
>13	4 Groups	None	None	.2335	.2481
13	8 Groups	None	1.04	.2412	.2903
	Mean	None	1.03	.2487	.2837
	Median	None	1.03	.2456	.2748
	4 Groups	None	1.04	.2544	.2950
	Mean	None	1.03	.2562	.2877
	Median	None	1.03	.2596	.2838
	Base	0.5	1.04	.2527	.2851
8	8 Groups	None	1.04	.2651	.3297
	Mean	None	1.05	.3017	.3145
	Median	None	1.05	.2845	.3171
	4 Groups	None	1.05	.2830	.3274
	Mean	None	1.07	.3003	.3523
	Median	None	1.05	.2996	.3308
	Base	0.3	1.06	.2890	.3113
5	8 Groups	None	1.10	.3084	.4114
	Mean	None	1.16	.3895	.4637
	Median	None	1.14	.3830	.4872
	4 Groups	None	1.12	.3497	.4392
	Mean	None	1.16	.3877	.4723
	Median	None	1.07	.3698	.4007
	Base	0.15	1.12	.3792	.4332
3	8 Groups	None	1.24	.6395	.8450
	Mean	None	1.28	.5965	.7755
	Median	None	Any	Diverges	
	4 Groups	None	1.30	.9075	1.208
	Mean	None	1.28	.7235	.9566
	Median	None	Any	Diverges	
	Base	0.1	1.36	.7595	.8155

1. Comparative RMS error and spread for assimilations with 40 randomly located observations with 10^{-7} error variance for ensemble sizes >13, 13, 8 and 5 and for 4-group and 8-group filters with corresponding time mean and time median filters and a traditional filter with Gaspari Cohn localization (base). Error and spread values are normalized by the observational error standard deviation.

Observation Error Variance	Ensemble Size	Group Size & Type	GC Half-width	Covariance Inflation	Time Mean RMS Error	Time Mean Spread
1e-5	14	4 Groups	None	1.02	.2258	.2564
	14	Mean	None	1.03	.2380	.2693
	14	Median	None	1.05	.2472	.2928
	14	Base	None	1.03	.2478	.2713
	56	Base	None	1.005	.2153	.2348
1e-3	14	4 Groups	None	1.02	.2314	.2584
	14	Mean	None	1.03	.2352	.2721
	14	Median	None	1.02	.2252	.2528
	14	Base	0.4	1.03	.2486	.2800
	56	Base	None	1.0075	.2177	.2372
0.1	14	4 Groups	None	1.02	.2488	.2776
	14	Mean	None	1.04	.2692	.3113
	14	Median	None	1.03	.2630	.2931
	14	Base	0.3	1.03	.2657	.2952
	56	Base	None	1.01	.2396	.2575
1.0	14	4 Groups	None	1.03	.2901	.3230
	14	Mean	None	1.03	.3121	.3377
	14	Median	None	1.05	.3146	.3652
	14	Base	0.3	1.05	.3080	.3426
	56	Base	0.5	1.01	.2816	.2885
10.0	14	4 Groups	None	1.05	.3782	.4294
	14	Mean	None	1.05	.4052	.4482
	14	Median	None	1.04	.4075	.4502
	14	Base	0.2	1.06	.4285	.4245
	56	Base	0.25	1.02	.3560	.3712
1e7	14	4 Groups	None	None	23.21*	23.02*

2. Comparative RMS error and spread for assimilations with 40 randomly located observations with various error variances. 14-member ensembles are used for 4-group and 8-group filters with corresponding time mean and time median filters. Traditional filters (base) for ensemble sizes of 14 and 56 are also included. Error and spread values are normalized by the observational error standard deviation.

Group Size & Type	GC Half-width	Covariance Inflation	Time Mean RMS Error	Time Mean Spread
2 Groups	None	1.05	.3066	.3505
Mean	None	1.06	.3243	.3980
Median	None	1.05	.3297	.3787
4 Groups	None	1.03	.2901	.3230
Mean	None	1.03	.3122	.3377
Median	None	1.05	.3146	.3652
8 Groups	None	1.03	.2854	.3346
Mean	None	1.04	.3150	.3550
Median	None	1.05	.3096	.3640
16 Groups	None	1.03	.2795	.3210
Mean	None	1.04	.3080	.3523
Median	None	1.03	.3042	.3414
Base	0.3	1.05	.3080	.3426

3. Comparative RMS error and spread for assimilations with 40 randomly located observations with 1.0 error variance for ensemble size 14 for 2, 4, 8 and 16 groups with corresponding time mean and time median filters and a traditional filter with Gaspari Cohn localization (base).

Group Size & Type	GC Half-width	Covariance Inflation	Time Mean RMS Error	Time Mean Spread
4 Groups	None	1.015	12.75	13.33
Mean	None	1.03	13.33	13.55
Median	None	1.01	13.59	14.32
Base	0.15	1.02	13.77	13.86

4. Comparative RMS error and spread for assimilations with 40 observations located as in Fig. 4 to form a data dense and data void region with 1.0 error variance for 14 member ensembles.

Group Size & Type	GC Half-width	Covariance Inflation	Time Mean RMS Error	Time Mean Spread
4 Groups	None	1.12	2.030	2.797
Mean	None	1.14	2.182	3.116
Median	None	1.12	2.508	3.159
Base	Any	Any	Diverged	

5. Comparative RMS error and spread for assimilations with 40 randomly located observations with forward observation operators being the average of 15-point observations surrounding the central location (for instance at locations 0.6424 + 0.025k, where k = -7, -6, …6, 7). The observational error variance is 4.0 and the mid-points of the 40 observation locations are the same as marked in Fig. 4a.

Group Size & Type	GC Half-width	Covariance Inflation	Time Mean RMS Error	Time Mean Spread
4 Groups	None	1.04	1.337	1.567
Mean	None	1.06	1.517	1.856
Median	None	1.05	1.479	1.678
Base	0.10	1.10	3.985	4.723

6. Comparative RMS error and spread for assimilations with 40 randomly located observations with forward observation operators being the average of the observation location and the location 0.5 removed from the location (two points on opposite sides of the cyclic domain). The observational error variance is 4.0 and the observation locations are the same as marked in Fig. 4a.

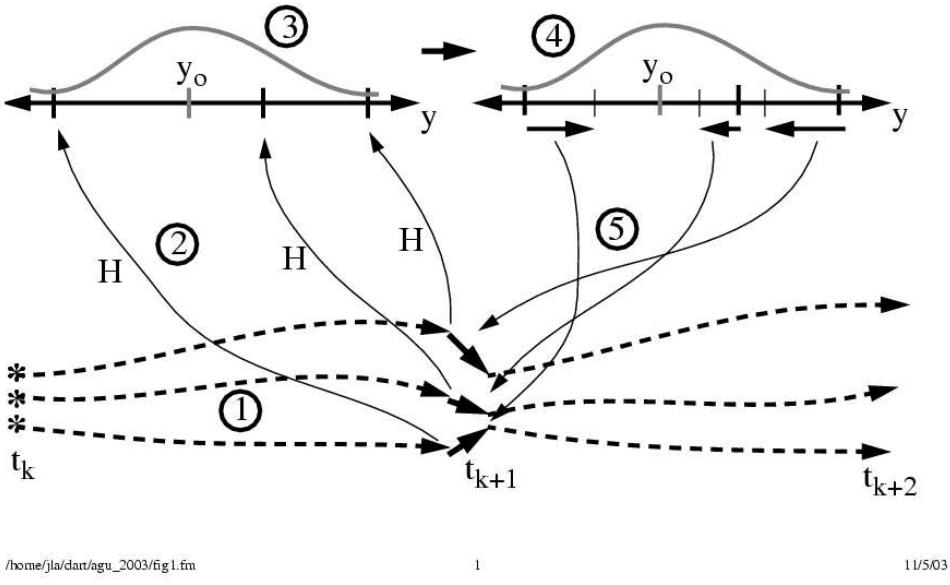

1. Schematic representation of the implementation of the ensemble filter used here with possible error sources marked by numbers 1 through 5.

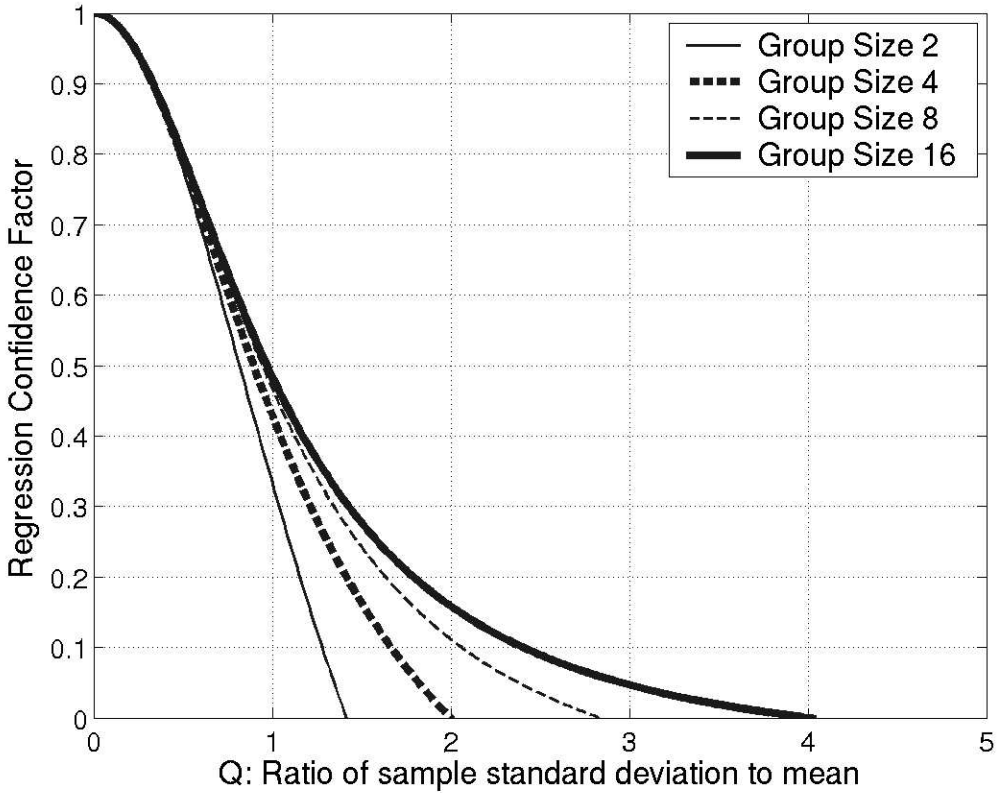

2. Regression confidence factors as a function of the ratio Q of regression sample standard deviation to the absolute value of the sample mean for 2 (thin solid), 4 (thick dashed), 8 (thin dashed) and 16 (thick solid) groups.

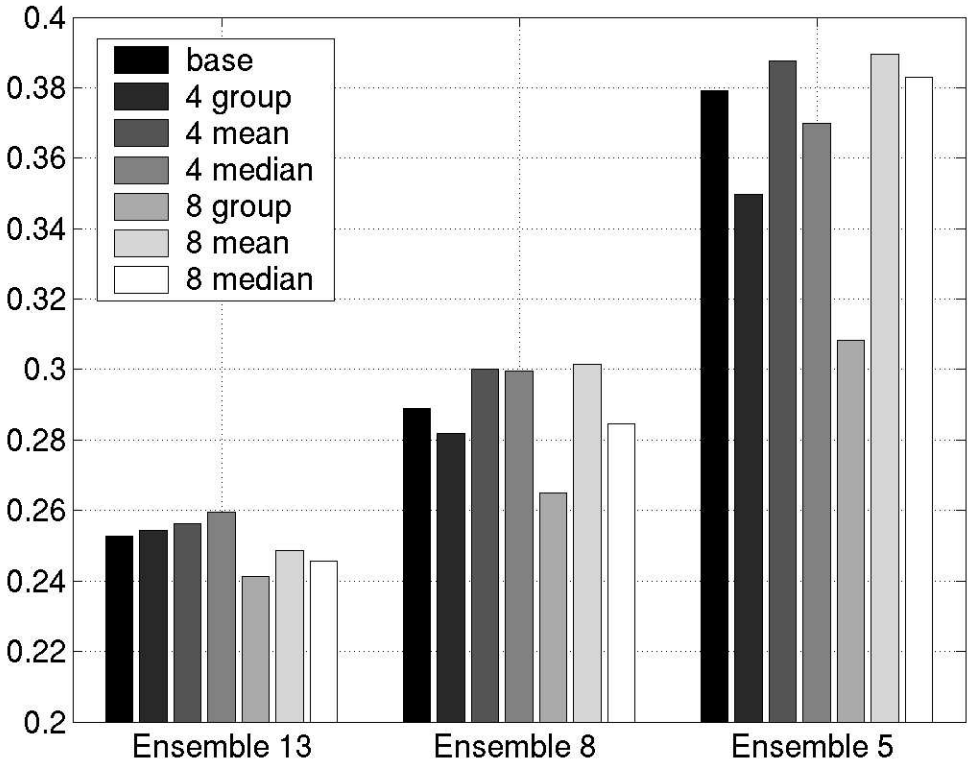

3. 2000-step time mean RMS error (normalized by the observational error standard deviation) for observational error variance of 1e-7 with 13, 8 and 5 member ensembles for standard Gaspari Cohn localized filter, 4 group filter and corresponding time mean and time median filter, and 8 group filter and corresponding time mean and median filter.

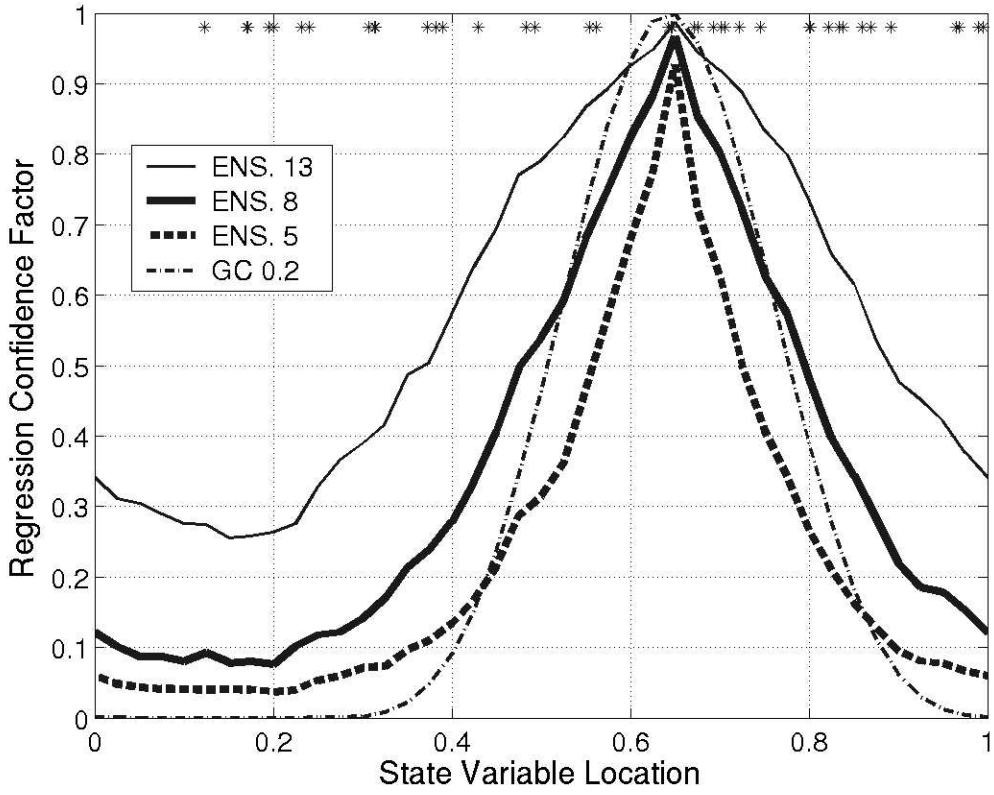
ENS. 13
ENS. 8
ENS. 5
GC 0.2
Regression Confidence Factor
State Variable Location
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0.2
0.4
0.6
0.8

Fig 4a

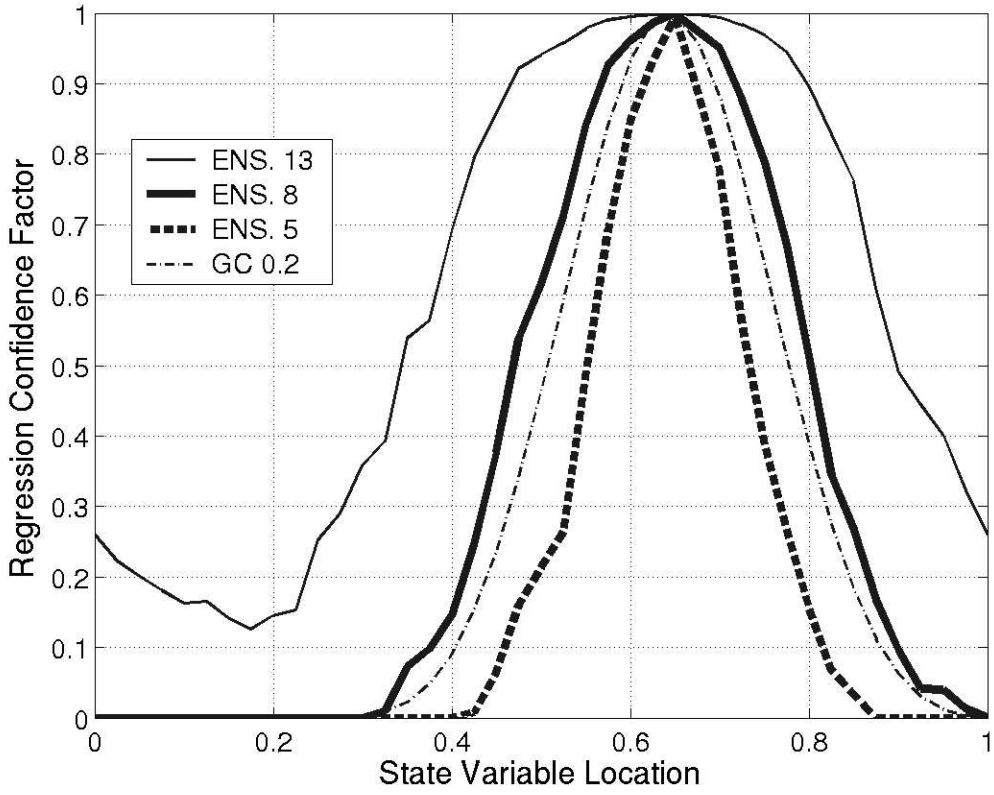

4. 2000-step time mean (a) and time median (b) regression confidence factors for Lorenz-96 model assimilations for an observation located at 0.6424 and all 40 state variables. The asterisks at the top of (a) indicate the position of 40 randomly located observations with observational error variance of 10^{-7}. Results are from hierarchical ensemble filters with 4 groups and ensemble sizes of 5 (thick dashed), 8 (thick solid) and 13 (thin solid). Also shown is a Gaspari-Cohn localization function for a half-width of 0.2 (thin dash-dotted).

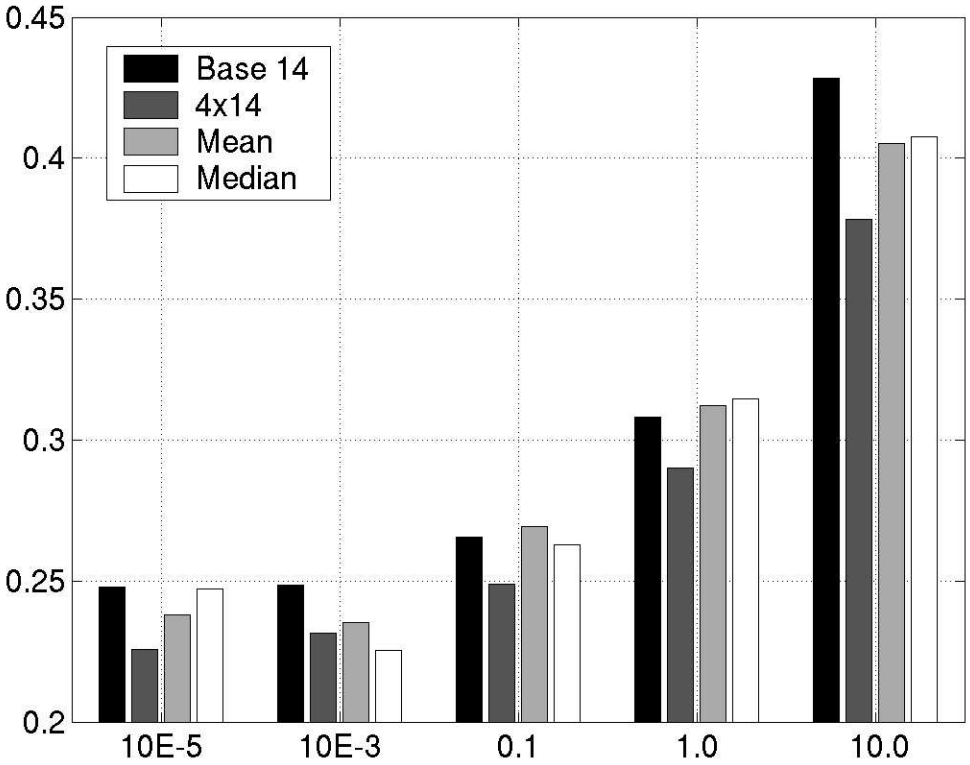

5. 2000-step time mean RMS error (normalized by the observational error standard deviation) for observational error variances of 10^{-5}, 10^{-3}, 0.01, 1.0 and 10.0 for standard Gaspari Cohn localized filter, 4-group filter and corresponding time mean and time median filters, all with ensemble size 14.

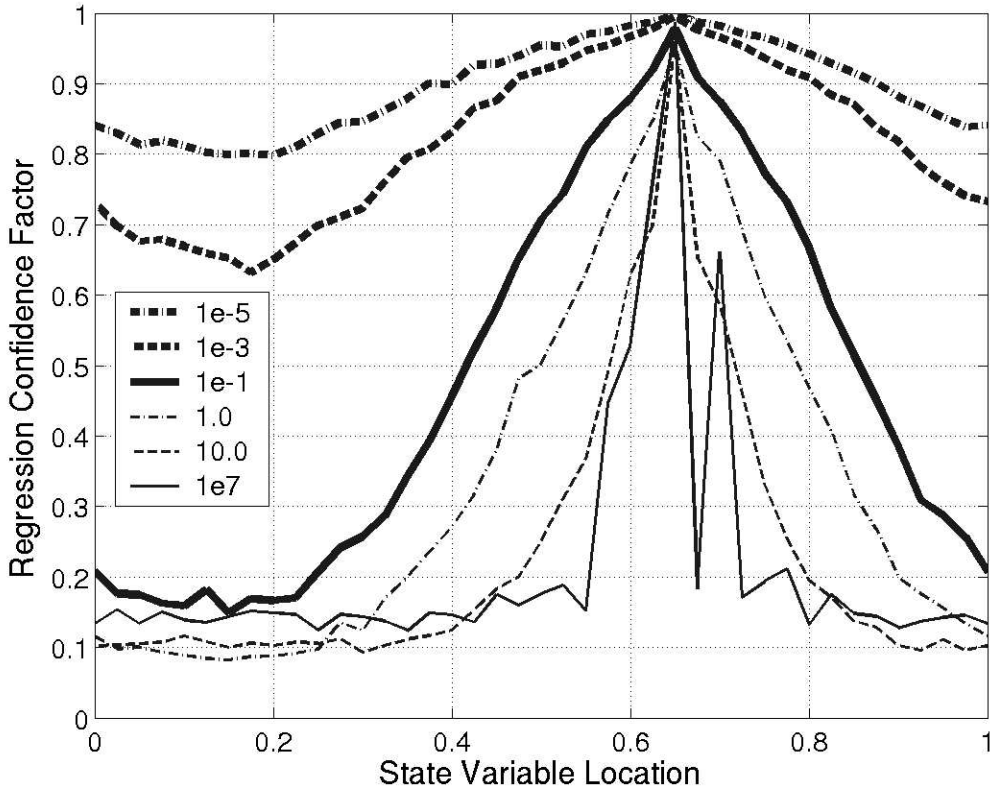
Regression Confidence Factor
State Variable Location
1e-5
1e-3
1e-1
1.0
10.0
1e7
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
0.2
0.4
0.6
0.8
1

Fig6a

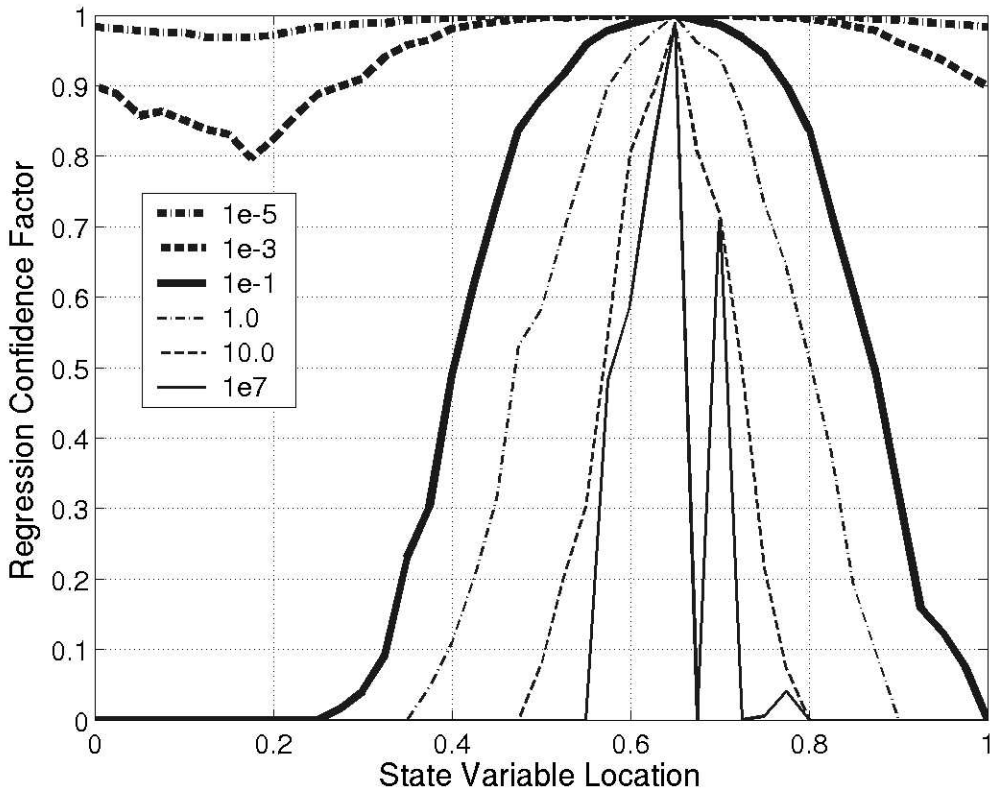

6. 2000-step time mean (a) and time median (b) regression confidence factors for Lorenz-96 model assimilations for an observation located at 0.6424 and all 40 state variables. Results are from hierarchical ensemble filters with 4 groups and 14 ensemble members. The observational error variance for 40 randomly located observations is 10^{-5} (thick dash-dotted), 10^{-3} (thick dashed), 10^{-1} (thick solid), 1.0 (thin dash-dotted), 10.0 (thin dashed) and 10^{7} (thin solid).

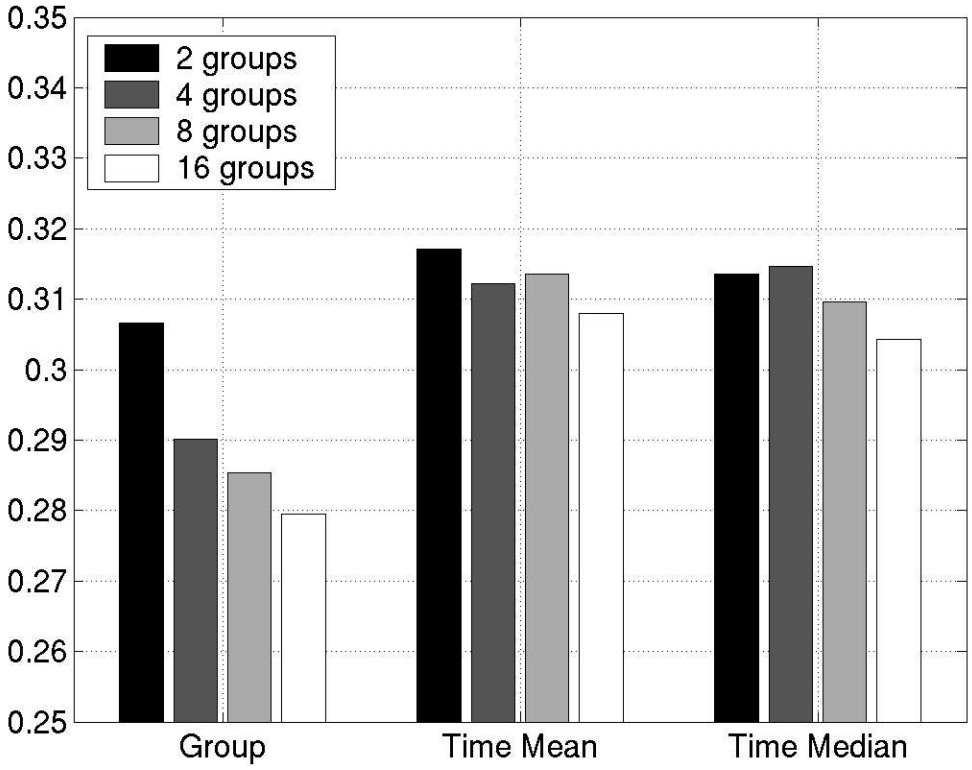

7. 2000-step time mean RMS error (normalized by the observational error standard deviation) for observational error variance of 1.0 and hierarchical filters with group sizes of 2, 4, 8 and 16 along with the corresponding time mean and time median filters.

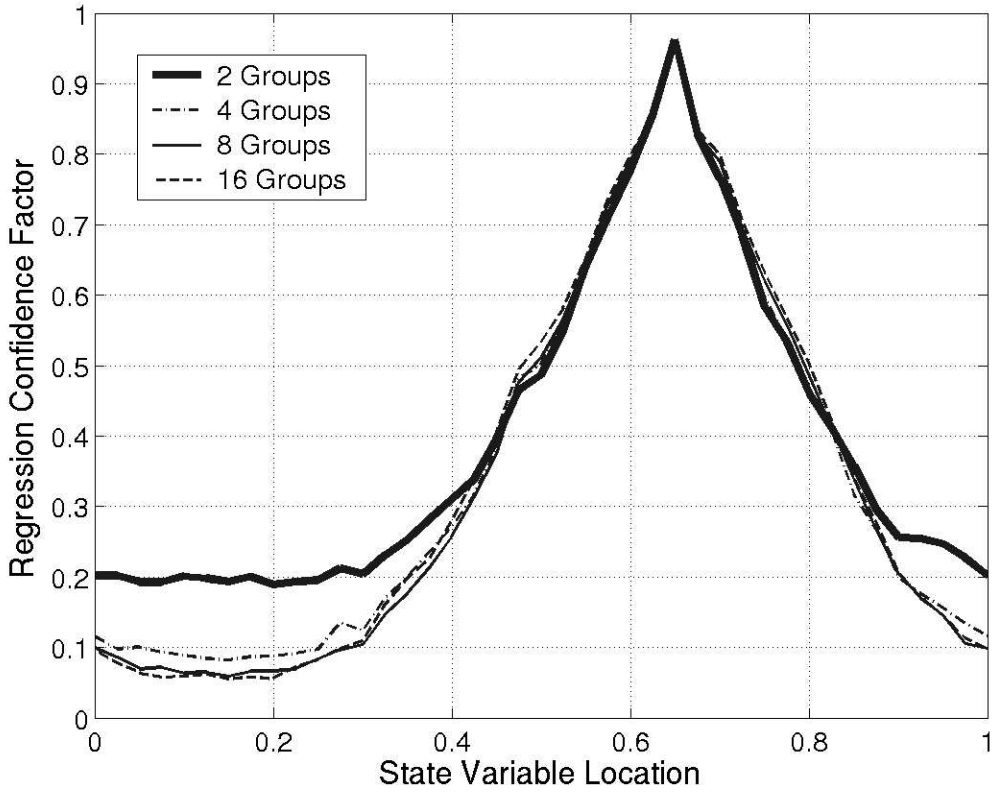
2 Groups
4 Groups
8 Groups
16 Groups
Regression Confidence Factor
State Variable Location
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
0.2
0.4
0.6
0.8
1

Fig8a

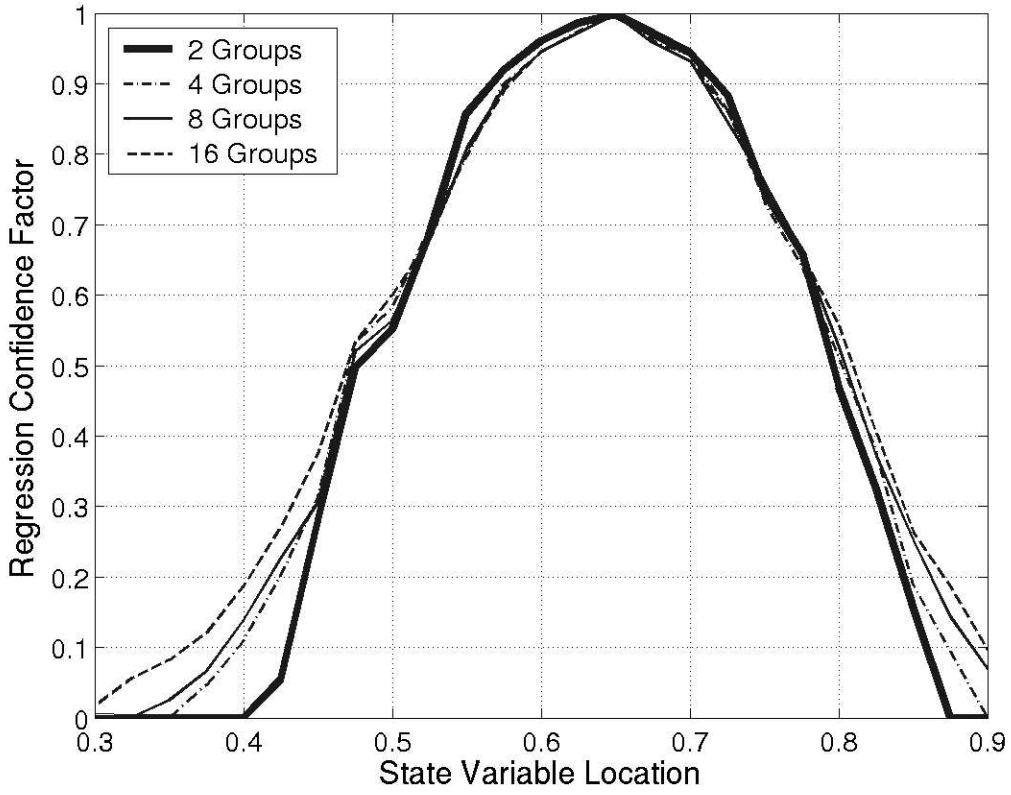

8. 2000-step time mean (a) and time median (b) regression confidence factors for Lorenz-96 model assimilations for an observation located at 0.6424 and all 40 state variables. The observational error variance for 40 randomly located observations was 1.0. Results are for 2 groups (thick solid), 4 groups (thin dash-dotted), 8 groups (thin solid) and 16 groups (thin dashed) of 14 member ensembles.

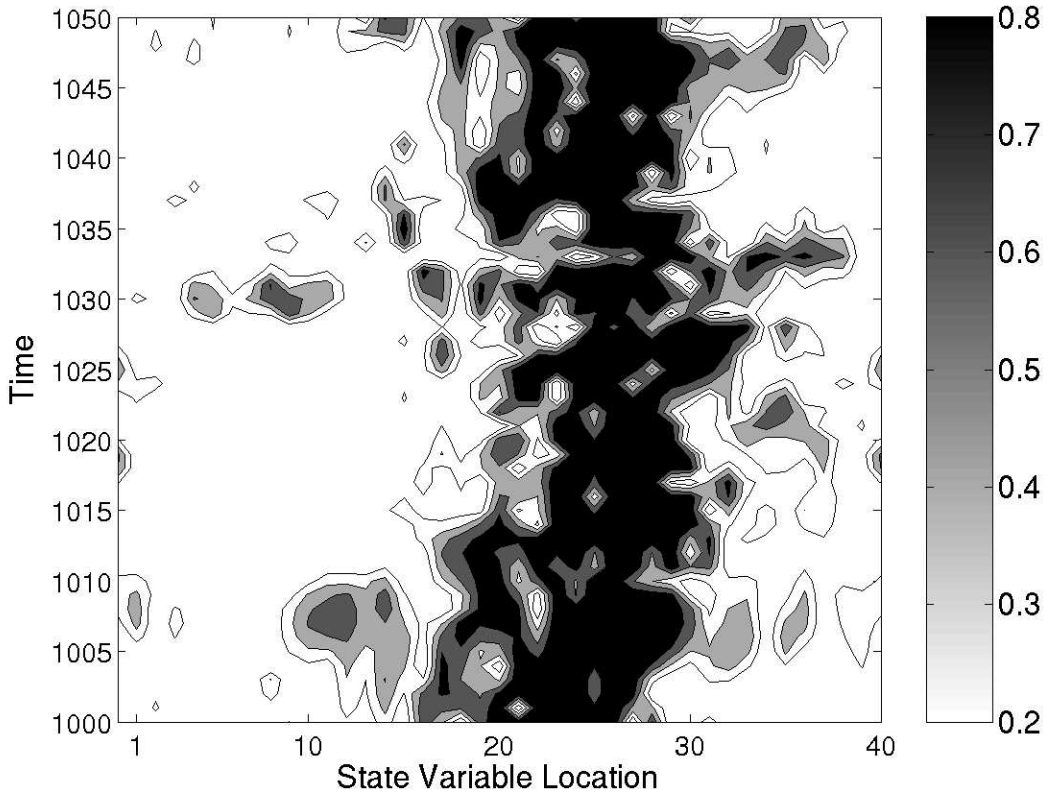
Time
1050
1045
1040
1035
1030
1025
1020
1015
1010
1005
1000
1
10
20
30
40
State Variable Location
0.8
0.7
0.6
0.5
0.4
0.3
0.2

Fig9a

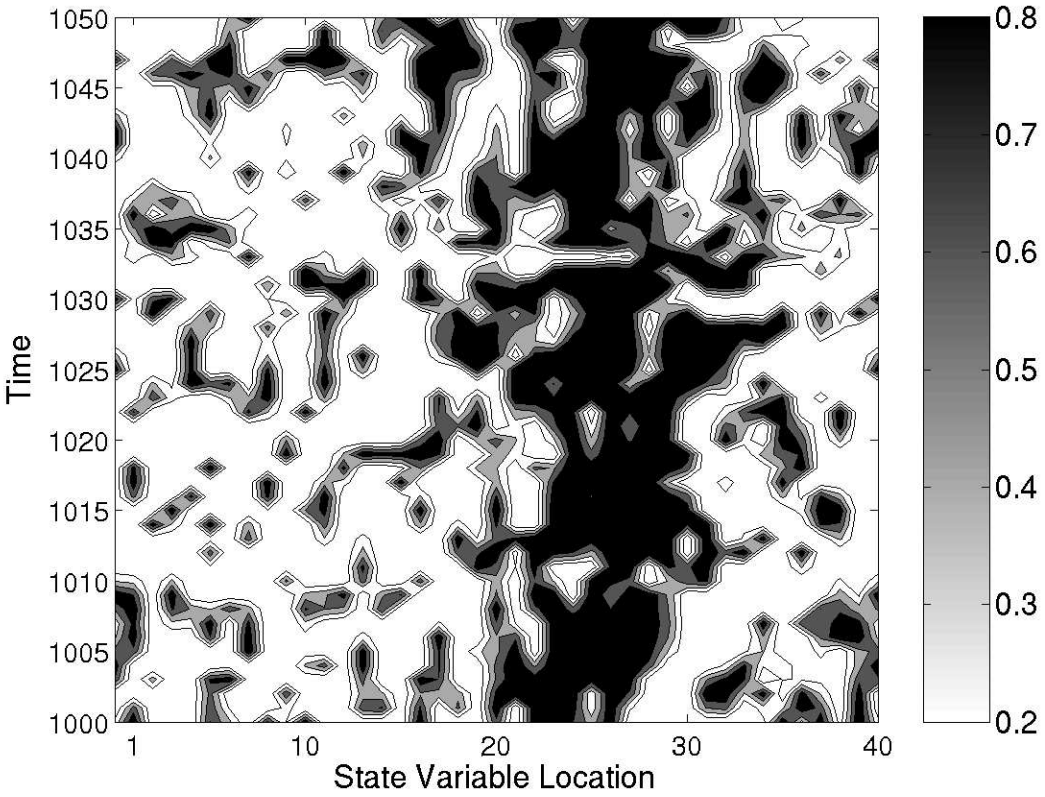

9. Regression confidence factor for Lorenz-96 model assimilations for an observation located at 0.6424 and all 40 state variables as a function of time between assimilation steps 1000 and 1050 of a sixteen group (a) and two group (b) hierarchical filter with ensemble size of 14. The contour interval is 0.2 with values greater than 0.4 shaded.

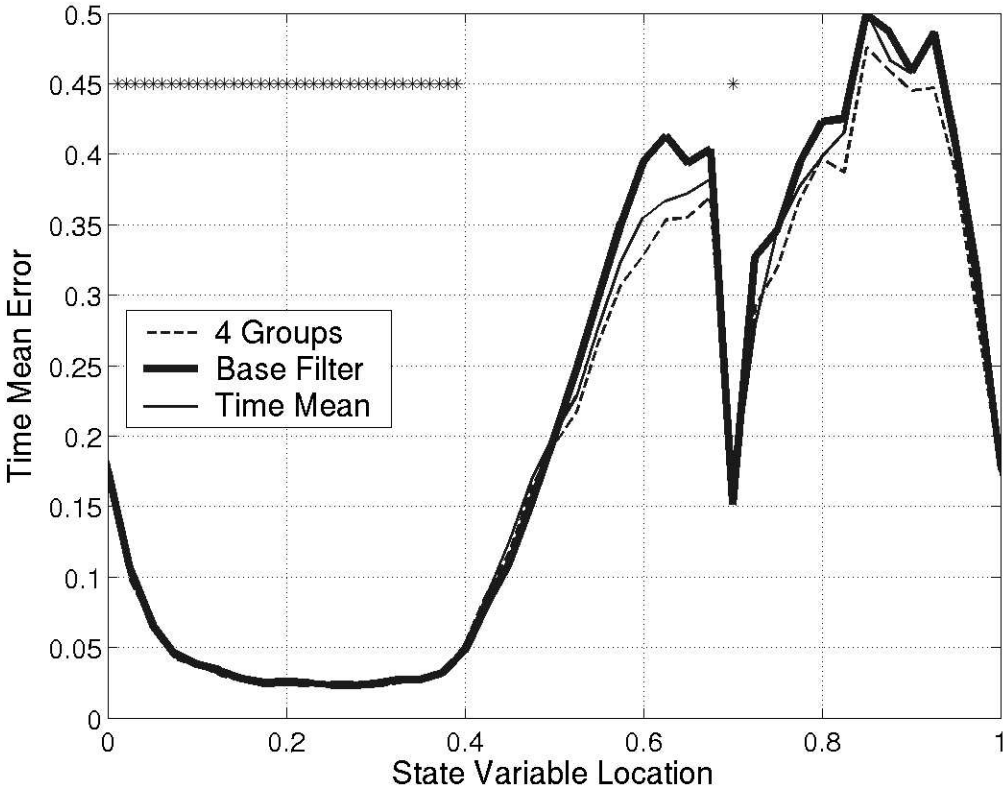

10. 2000-step time mean RMS error as a function of model state variable for 14 member ensemble assimilations of 40 observations with observational error variance of 1.0 whose location is indicated by the asterisks at the top of the plot. Results are plotted for a 4-group hierarchical filter (thin dashed), a filter using the time mean regression confidence factors from the 4-group filter (thin solid), and a traditional ensemble filter with a Gaspari-Cohn localization with half-width 0.2.

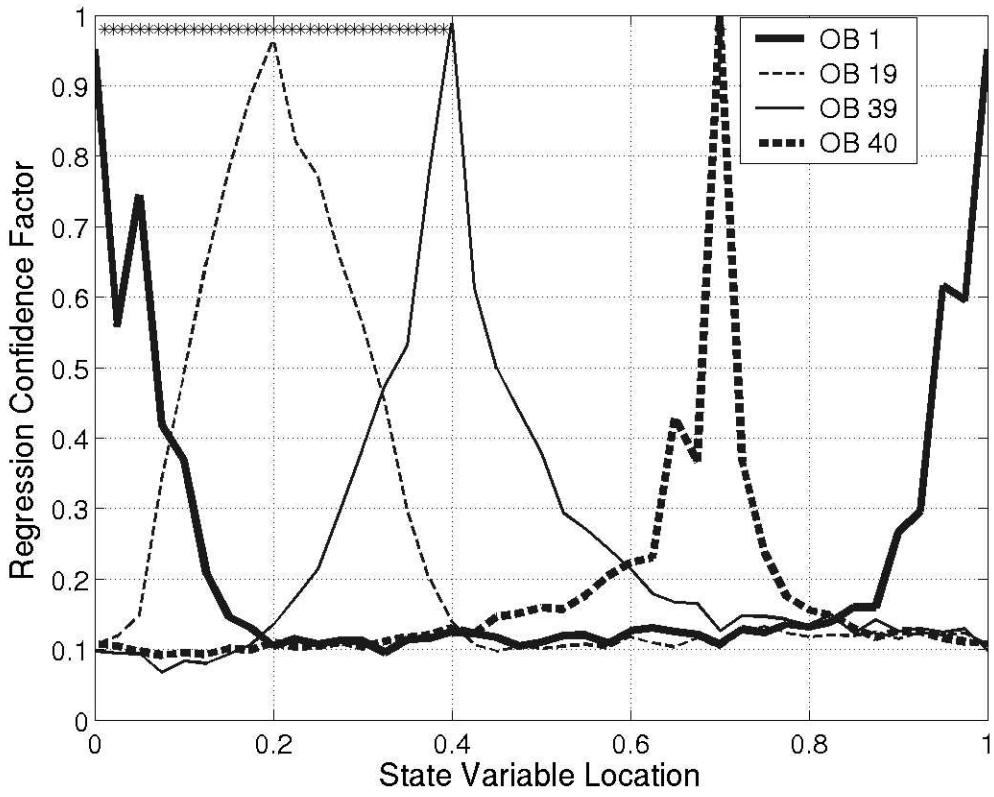
OB 1
OB 19
OB 39
OB 40
Regression Confidence Factor
State Variable Location
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
0.2
0.4
0.6
0.8
1

Fig11a

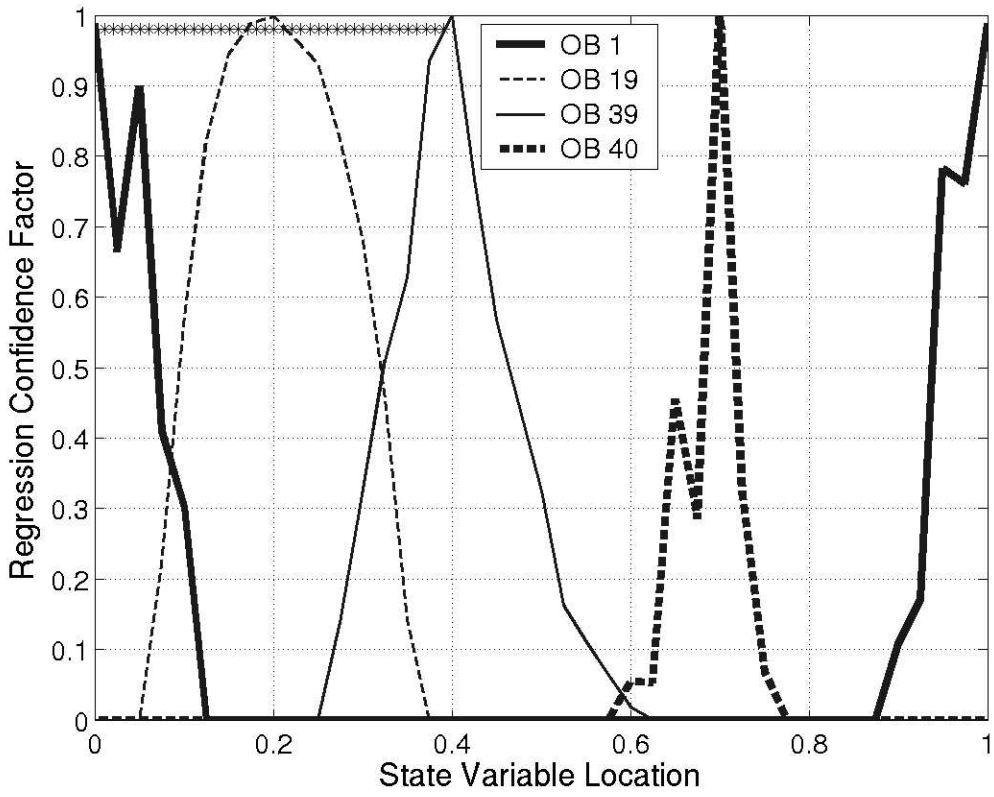

11. 2000-step time mean (a) and time median (b) regression confidence factors for Lorenz-96 model 4-group hierarchical filter with 14 member ensembles. The observations are as for Fig. 10 and the locations are marked with asterisks at the top of the plot. Regression confidence factors are plotted for the observation at location 0.011 (thick solid), 0.191 (thin dashed), 0.391 (thin solid) and 0.701 (thick dashed).

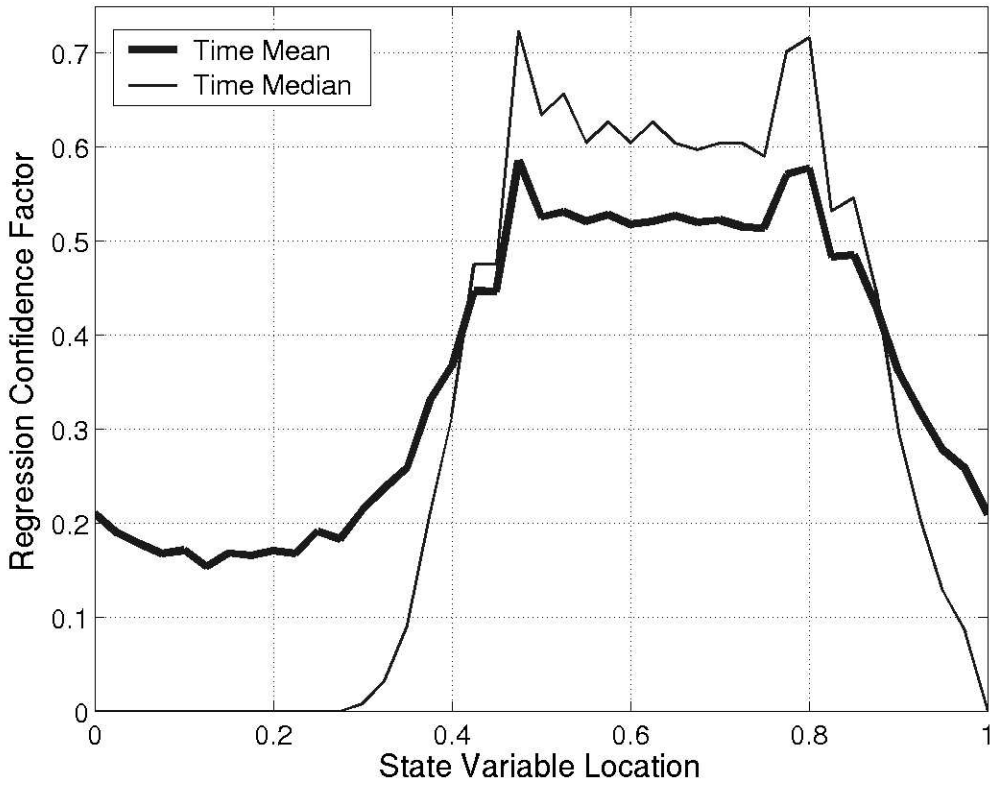

12. 2000-step time mean (thick) and time median (thin) regression confidence factors for Lorenz-96 model assimilations with a 4-group 14-member ensemble filter for an observation located at 0.6424 and all 40 state variables. The forward observation operators are the average of 15-point observations surrounding the central location (for instance at locations 0.6424 + 0.025k, where k = -7, -6, …6, 7). The observational error variance is 4.0 and the mid-points of the 40 observation locations are the same as marked in Fig. 4a.

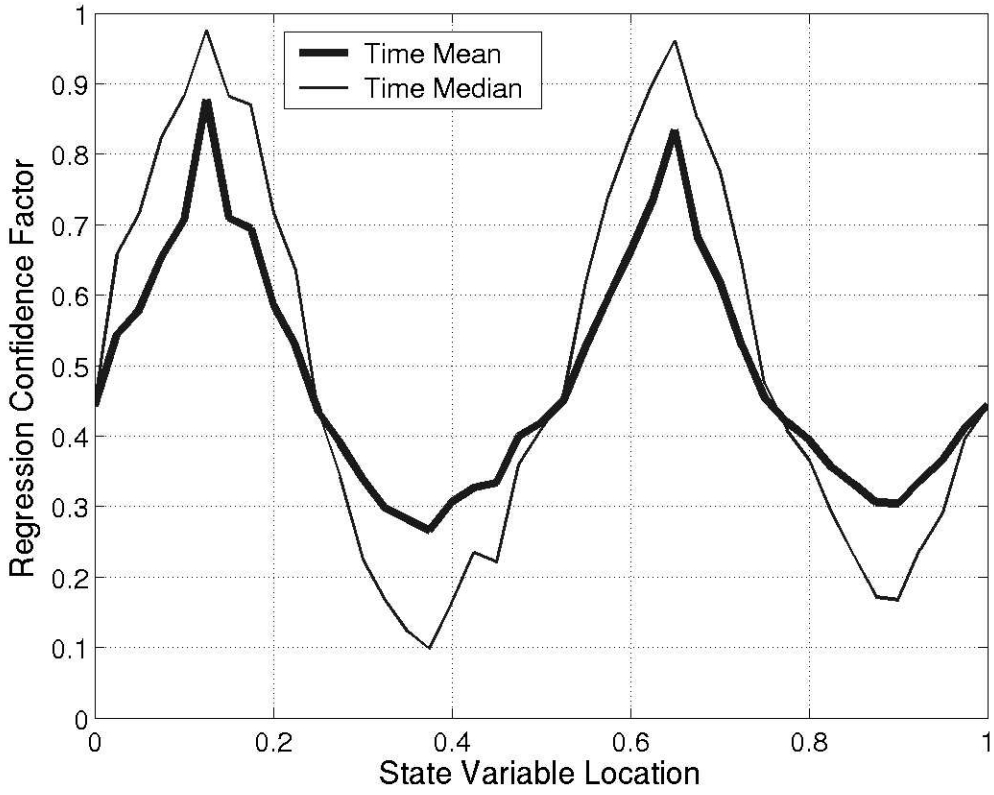

13. 2000-step time mean (thick) and time median (thin) regression confidence factors for Lorenz-96 model assimilations with a 4-group 14-member ensemble filter for an observation located at 0.6424 and all 40 state variables. The forward observation operators are the average of the observation location and the location 0.5 removed from the location (two points on opposite sides of the cyclic domain). The observational error variance is 4.0 and the mid-points of the 40 observation locations are the same as marked in Fig. 4a.

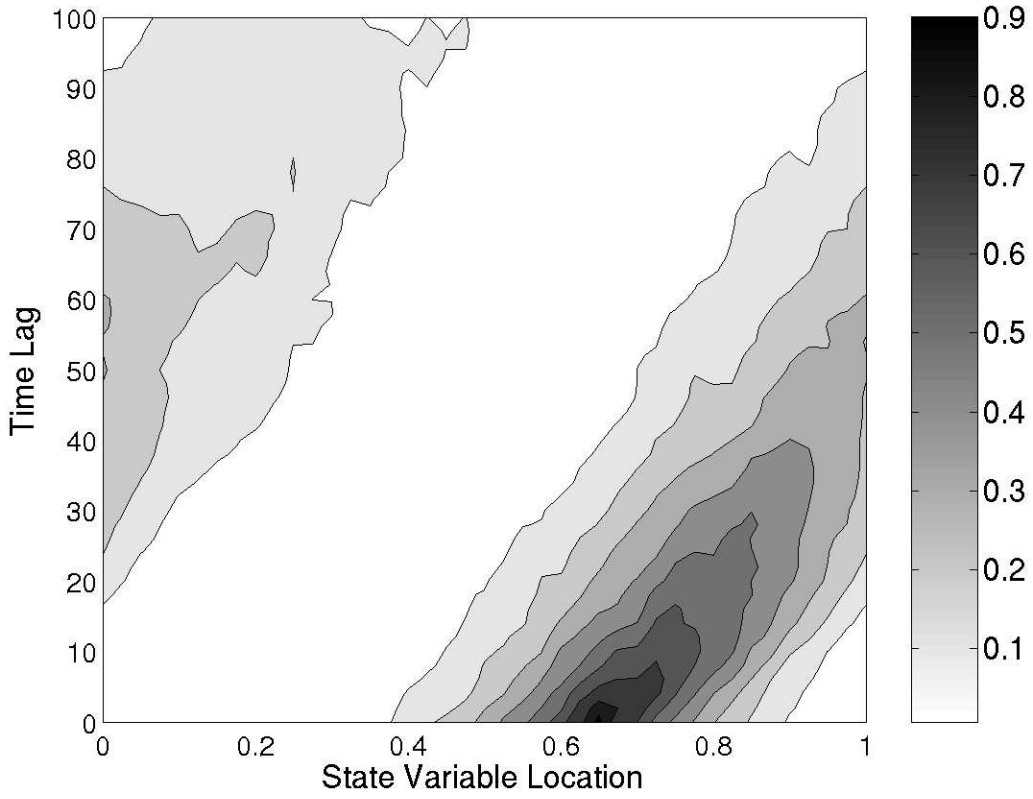

14. 1000-step time mean regression confidence factors for simulated time-lagged observation at location 0.6424 for assimilations with a 4-group 14-member hierarchical ensemble filter. The base observation set is the 40 random observations as marked in Fig. 4a with error variance 1.0. The plot shows the regression confidence factors for an observation that was taken n assimilation steps prior to the time at which it was assimilated. The contour interval is 0.1.

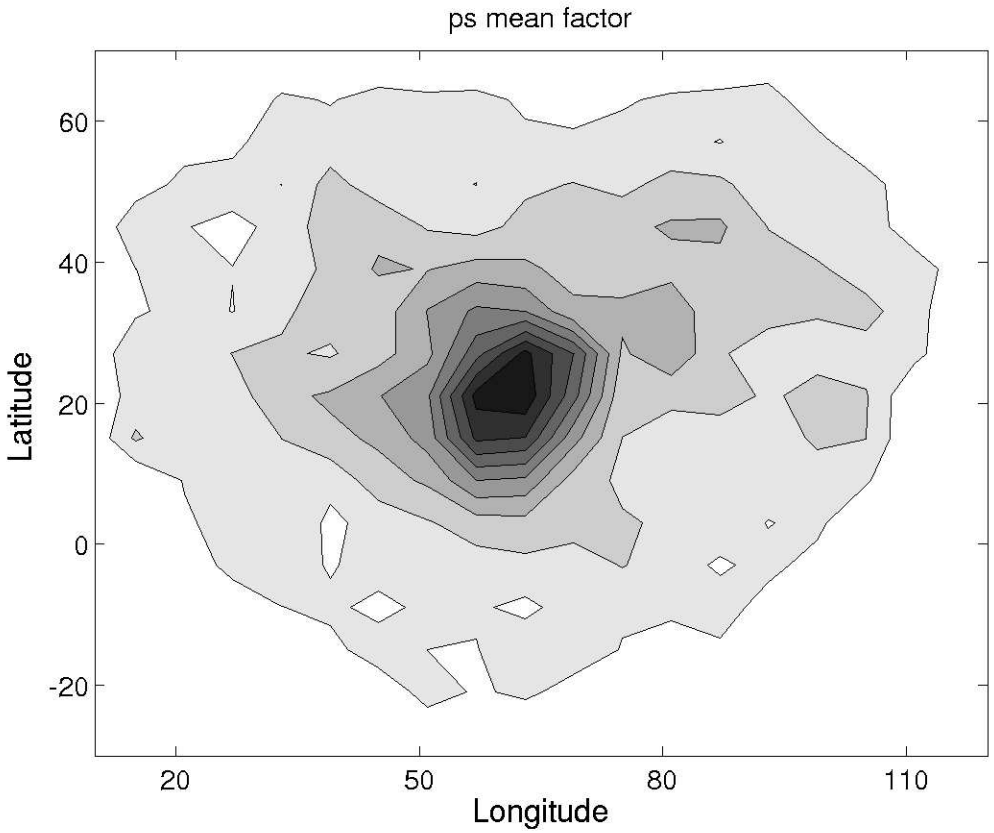

15. Time mean regression confidence factor for a pressure observation at 22.7 N 61.4 E and surface pressure state variables in the GFDL B-grid AGCM. Contour interval is 0.1.

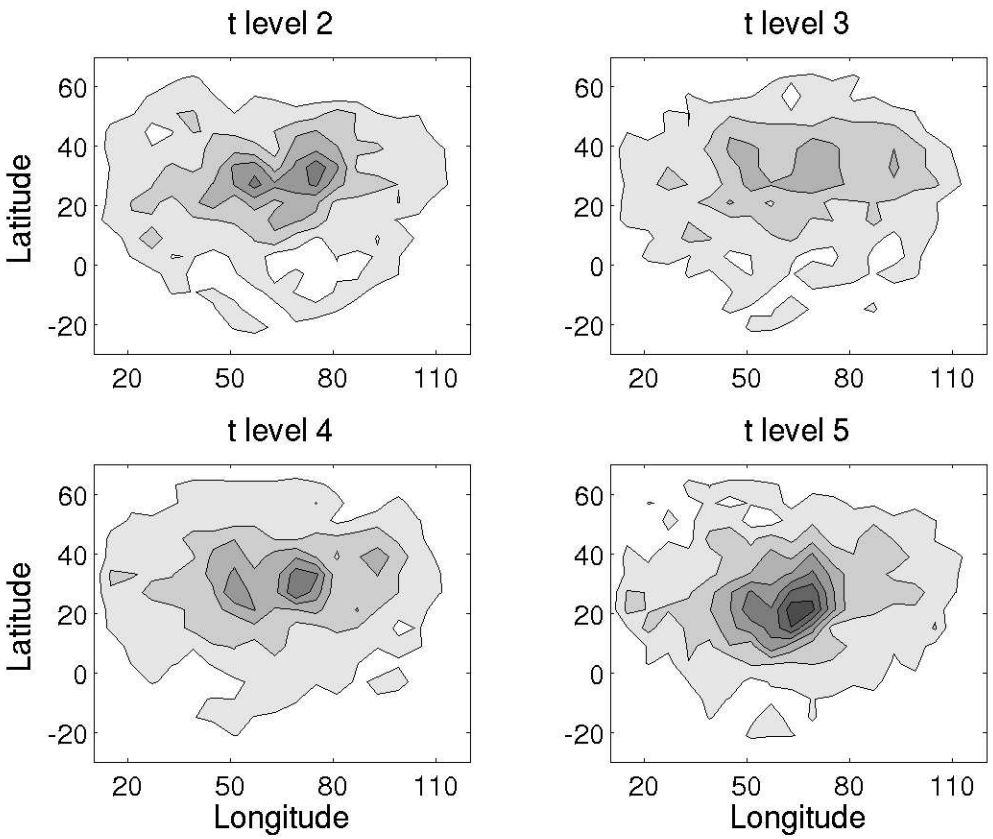

16. Time mean regression confidence factors for the same surface pressure observation as in Fig. 15 but with v at each of the model levels 2 through 5. Contour interval is 0.1.