

Adaptive observations and assimilation in the unstable subspace by breeding on the data-assimilation system

By ALBERTO CARRASSI^{1,2*}, ANNA TREVISAN¹ and FRANCESCO UBOLDI³,
¹*ISAC-CNR, Bologna, Italy*; ²*Dept. of Physics, University of Ferrara, Italy*; ³*Novate Milanese, Italy*

(Manuscript received 11 January 2006; in final form 3 September 2006)

ABSTRACT

Results of targeting and assimilation experiments in a quasi-geostrophic atmospheric model are presented and discussed. The basic idea is to assimilate observations in the unstable subspace (AUS) of the data-assimilation system. The unstable subspace is estimated by breeding on the data-assimilation system (BDAS). The analysis update has the same local structure as the observationally forced bred modes.

Use of adaptive observations, taken at locations where bred vectors have maximum amplitude, enhances the efficiency of the procedure and allows the use of a very limited number of observations and modes.

The performance of the targeting and assimilation design is tested in an idealized context, under perfect model conditions. It is shown that the process of driving the control solution toward the true trajectory accomplished by the assimilation reduces the number and growth of unstable modes. By observing and assimilating the unstable structures it is then possible to stabilize the assimilation system so that few observations are sufficient to keep the analysis error within very low bounds, even in the presence of observational noise.

In an idealized limited area model configuration the number and frequency of observations necessary to control the system is shown to be related to the properties of its unstable subspace.

1. Introduction

Data assimilation and observation strategies, possibly encompassing adaptive (or targeted) observations, are connected by the aim of reducing the uncertainty in the estimate of the present and future state of the atmosphere (Daley, 1991; Ghil, 1997; Talagrand, 1997; Kalnay, 2002). In spite of the intense activity related to field experiments, the impact on forecast quality of adaptive observations is perhaps less satisfactory than expected and is neutral or negative in some cases (Szunyogh et al., 2002; Fourrié et al., 2006). Therefore, it seems worthwhile to perform further studies that may improve our understanding of the problem and possibly be useful for future applications.

From a theoretical point of view, it is well known that assimilation errors intrinsically project on the unstable manifold of the system (Pires et al., 1996) and that instabilities can be tracked by advanced dynamically based assimilation techniques (Todling and Ghil, 1994; Ghil and Todling, 1996). In a previous work, Trevisan and Uboldi (2004, hereafter TU) introduced the concept of assimilation in the unstable subspace (AUS) and the stability of the data-assimilation system. They estimated the instabilities

by assimilating observations in the perturbed trajectories, a procedure introduced by Lorenz and Emanuel (1998) and consistent with the perturbative equations of the data-assimilation system. The same method is used here to estimate the unstable directions of the data-assimilation system, and is referred to as breeding on the data-assimilation system (BDAS).

In the data-assimilation system, the analysis step is equivalent to a forcing added to the model equations and, as remarked by TU, this observational forcing in principle reduces the growth rate and the number of unstable directions with respect to those of the original system. The important result, confirmed by the present study, is that, given a system that possesses positive Lyapunov exponents, by means of a sufficiently large number and frequency of observations, the corresponding data-assimilation system can be stabilized, in the sense that the number of positive exponents is reduced, possibly to zero. The system becomes more controllable: its instabilities can be easily tracked and efficiently eliminated, particularly by means of properly located adaptive observations.

In the context of the 40-variable Lorenz (1996) model, TU showed that only a few observations and few unstable structures were sufficient to drastically reduce the error of the data-assimilation system solution. Further developments of the theory introduced by TU, are found in Uboldi et al. (2005, hereafter UTC) and Uboldi and Trevisan (2006, hereafter UT), who

*Corresponding author.
e-mail: carrassi@fe.infn.it
DOI: 10.1111/j.1600-0870.2006.00210.x

applied AUS and BDAS to the atmospheric quasi-geostrophic model of Rotunno and Bao (1996) and to a primitive equation ocean model (Bleck, 1978). Both studies were performed in the context of perfect model and perfect observations.

In the present follow-up study we extend the application in the QG-model to the noisy observations case. We then examine in more depth the question of stabilization induced by the assimilation of observations in a global model and by boundary forcing in an idealized limited area model.

The following simple arguments, at the basis of the present work, illustrate why it is conceivable that the same concepts can be applied in a more general context and that the results of this study may turn out to be relevant for atmospheric and oceanic prediction. Forecast models have a huge number of positive exponents. However, suppose we make an hypothetical experiment where all model variables at each analysis time are observed in all but a small portion of the domain. We expect the number of unstable modes that can grow in the data-assimilation system and their growth rate to be reduced since no instability can grow in the observed part of the domain. This is similar to what happens in a limited area model, where the influence of boundary conditions, that can be seen as a special case of observational forcing, inhibits error amplification in domains of small enough size (Vannitsem, 2003).

Even in global scale models we can expect that the assimilation of a large number of observations may have a significant stabilizing effect. Similarly, with reference to targeting, it is not surprising that the number of unstable modes that can grow in such adaptive data-assimilation systems can be greatly reduced if every time an unstable structure starts to grow we can detect it by special observations. This is exactly what happens both in the experiments of TU, UTC, UT with perfect observations and in the present ones that include observational noise. In particular, the present results give more insight in the stabilization process and quantify the dimension of the unstable subspace of the data-assimilation system.

The paper is organized as follows. In Section 2, we review the work of TU and discuss in detail the theory at the basis of assimilation in the unstable subspace (AUS); furthermore we discuss how the observability condition (Ghil, 1997) can be applied in the context of the tangent linear assumption. In Section 3, the stability of the data-assimilation system is discussed. Section 4 introduces breeding on the data-assimilation system (BDAS), a variant of the well established breeding technique (Toth and Kalnay, 1993, 1997). The validity of the BDAS technique goes beyond the present work context. In fact, BDAS constitutes a natural way to estimate how a particular observational network modulates the growth of free bred vectors and could replace the use of a masking function (Toth and Kalnay, 1997; Wang and Bishop, 2003). Section 5 illustrates technical aspects of the implementation with adaptive observations. In Section 6, the model and experimental setup are described. Section 7 illustrates the performance of AUS in comparison with 3DVar and presents

results relative to the stabilization obtained by observational forcing in the noisy observational case. In Section 8, we estimate the properties of the unstable subspace of an idealized limited area model in a perfect observational setting and investigate the observability condition. General conclusions are drawn in Section 9.

2. Assimilation in the unstable subspace

In this section, the basic elements of the method proposed by TU are reviewed and further evidence is provided to motivate its foundations. The perfect observations case is considered with particular attention to the observability of the system.

The system that we are considering is given by the model equations subject to observational forcing at each analysis step and its solution is the solution of the data-assimilation system, i.e. the complete analysis–forecast cycle. The model is assumed to be perfect.

In the present formulation (TU):

- the control trajectory forecast state constitutes the background state;
- the analysis solution is obtained by confining the analysis increment in the unstable subspace of the control trajectory;
- the unstable vectors are obtained as solutions of the perturbative equations of the observationally forced system (TU).

One advantage is that, when errors are sufficiently small, the analysis state is on the attractor of the system (Lorenz, 1984; Trevisan and Pancotti, 1998).

Following the conventions of Ide et al. (1997), \mathbf{y}^o is the M -dimensional observation vector, \mathbf{x}^f represents the background model forecast state, \mathbf{x}^a the analysis state. \mathbf{R} is the observational error covariance matrix, \mathbf{H} the linearized observation operator, \mathbf{P}^a and \mathbf{P}^f the analysis and background error covariance matrices, respectively.

Let $[\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_I]$ represent the, possibly non-orthogonal, Lyapunov vectors (Legras and Vautard, 1996; Trevisan and Pancotti, 1998) of the forced system, ordered by decreasing global exponent. If I is the number of degrees of freedom of the system, this set is a basis for the tangent space (Oseledec, 1968).

Let \mathbf{M}_k be the tangent linear propagator between time t_k and t_{k+1} . We can write:

$$\mathbf{M}_k = \mathbf{E}_{k+1} \mathbf{\Lambda}_k \mathbf{E}_k^{-1}, \quad (1)$$

where \mathbf{E} represents the matrix whose columns are the Lyapunov vectors, $\mathbf{\Lambda}_k$ is diagonal and its elements are the amplification factors, $\exp[\int_{t_k}^{t_{k+1}} \lambda_i(t) dt]$ where λ_i are the local Lyapunov exponents, in decreasing order, corresponding to the vectors \mathbf{e}_i .

Accordingly, during the forecast step:

$$\boldsymbol{\eta}_{k+1}^f = \mathbf{E}_{k+1} \mathbf{\Lambda}_k \mathbf{E}_k^{-1} \boldsymbol{\eta}_k^a, \quad (2)$$

where $\boldsymbol{\eta}^a$ and $\boldsymbol{\eta}^f$ represent the analysis and forecast error, respectively.

Following TU, we let the analysis increment be confined to the subspace spanned by the leading N Lyapunov vectors. Hereafter we denote by \mathbf{E} the $(I \times N)$ matrix formed by the leading N column vectors. The analysis solution obtained by minimizing the reduced-order cost function is:

$$\mathbf{x}^a = (\mathbf{I} - \mathbf{KH})\mathbf{x}^f + \mathbf{K}\mathbf{y}^a, \quad (3)$$

where, subject to the above constraint:

$$\mathbf{K} = \mathbf{E}\boldsymbol{\Gamma}(\mathbf{HE})^T [\mathbf{R} + (\mathbf{HE})\boldsymbol{\Gamma}(\mathbf{HE})^T]^{-1} \quad (4a)$$

or, equivalently:

$$\mathbf{K} = \mathbf{E}[\boldsymbol{\Gamma}^{-1} + (\mathbf{HE})^T \mathbf{R}^{-1} \mathbf{HE}]^{-1} (\mathbf{HE})^T \mathbf{R}^{-1} \quad (4b)$$

and $\boldsymbol{\Gamma}$ is a $N \times N$ symmetric positive definite matrix, representing the background error covariance matrix in the subspace spanned by the columns of \mathbf{E} . We observe that the matrices to be inverted in the expressions (4a) and (4b) have order M and N , respectively.

The distribution of data in space and time, in relation to the number, position and growth rate of the unstable structures, determines the possibility to control error growth. Full advantage of the present approach can be taken by means of an adaptive strategy, by which observations are located where the amplitude of the unstable structures is maximum.

The adaptive observation strategy, that is consistent with the assimilation in the unstable subspace and that is adopted in the present work, is to make measurements at locations where the unstable structures that appear in the gain matrix expression have maximum amplitude. Then, the same structures that are used to locate adaptive observations are also used to assimilate observations, according to (3) and (4).

As shown in previous works (TU, UTC), a progressive reduction of the error can be obtained even using a single vector \mathbf{e} and a single observation located in its maximum, at each analysis time.

2.1. Perfect observations and the observability condition

In this subsection, we investigate the consequences, on the analysis error and on its estimate, of confining the analysis increment in a given subspace. To this end we make the simplifying assumption of perfect observations.

The forecast error covariance matrix can be decomposed, in general, as:

$$\mathbf{P}^f = \mathbf{P}_E^f + \mathbf{P}_C^f, \quad (5)$$

where \mathbf{P}_E^f is estimated by means of N independent vectors, stored in the N columns, \mathbf{e}_n , of the matrix \mathbf{E} :

$$\mathbf{P}_E^f = \mathbf{E}\mathbf{E}^T. \quad (6)$$

Suppose that, neglecting the residual matrix \mathbf{P}_C^f , \mathbf{P}_E^f is used instead of \mathbf{P}^f to assimilate $M \geq N$ perfect observations, distributed

in such a way that the rank of the matrix \mathbf{HE} is maximum (N), so that each column vector \mathbf{e}_n is detected by at least one observation.

With perfect observations, the relation between the forecast and analysis error covariance matrices is:

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{KH})\mathbf{P}^f(\mathbf{I} - \mathbf{KH})^T. \quad (7)$$

In the hypothesis $N = M$, right multiplying (4a) by \mathbf{HE} after setting $\mathbf{R} = 0$, we have:

$$\mathbf{KHE} = \mathbf{E}. \quad (8)$$

If the error is actually confined in the subspace spanned by the column vectors of \mathbf{E} , then $\mathbf{P}_C^f = 0$. By setting $\mathbf{P}^f = \mathbf{P}_E^f$ in (7) and making use of (8) we obtain:

$$\mathbf{P}^a = 0. \quad (9)$$

This means that, in this idealized case, a number of observations equal to the dimension of the subspace, where the background error is confined, is sufficient to determine the state exactly. This is the observability condition under the above hypothesis.

If, instead, the number of unstable directions considered in the construction of the gain matrix is smaller than the dimension of the subspace where the actual background error is confined, then \mathbf{P}^a will not be zero, but it will be given by:

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{KH})\mathbf{P}_C^f(\mathbf{I} - \mathbf{KH})^T. \quad (10)$$

Only some components of the background error will be eliminated (in the perfect observation case) or at least reduced.

At the next analysis step, the error will strongly project on the subspace that we have ignored. This makes the recursive use of the same subspace undesirable and is the reason for introducing the *refresh* procedure described in Section 4. Further discussion on the need for a *refresh* procedure can be found in UT.

3. Stability of the data-assimilation system

The reason to consider the linearized equations of the data-assimilation system is twofold. On the one hand, we need to estimate the unstable vectors of the system because they appear in the analysis solution. On the other, the study of the stability of the analysis solution is necessary to ensure its independence from the initial guess and to establish the conditions of convergence to the truth (see TU and UT).

Given the model system equations:

$$\mathbf{x}_{k+1} = \mathcal{M}(\mathbf{x}_k) \quad (11)$$

the stability of the solution is studied by linearizing the equations about a nonlinear trajectory. The tangent linear propagator \mathbf{M} evolves perturbations from time t_k to time t_{k+1} :

$$\mathbf{x}'_{k+1} = \mathbf{M}(\mathbf{x}_k)\mathbf{x}'_k, \quad (12)$$

where the vector \mathbf{x}'_k represents a small perturbation of the nonlinear reference trajectory, solution of (11), at time step k . The

number of positive Lyapunov exponents of the model system may be quite large.

Turning now to the analysis cycle, its solution is subject to observational forcing at the analysis step. The evolutive equation of the system forced by the assimilation process at regular observation time intervals, can be written as:

$$\mathbf{x}_{k+1}^a = (\mathbf{I} - \mathbf{K}\mathcal{H})\mathcal{M}(\mathbf{x}_k^a) + \mathbf{K}\mathbf{y}_{k+1}^o, \quad (13)$$

where the new analysis state is obtained by applying the assimilation (3) to the forecast state obtained by the nonlinear model evolution (11) of the previous analysis state and \mathcal{H} represents the (possibly nonlinear) observation operator.

In the hypothesis that any perturbed trajectory undergoes the same forecast and assimilation steps, with the same (possibly noisy) observations as the control analysis solution, the linear equations describing the stability of the assimilation system at the analysis step read:

$$\mathbf{x}_{k+1}^{a'} = (\mathbf{I} - \mathbf{K}_{k+1}\mathbf{H})\mathbf{M}\mathbf{x}_k^{a'}, \quad (14)$$

where $\mathbf{M} = \mathbf{M}(\mathbf{x}_k^a)$ is the Jacobian of \mathcal{M} , $\mathbf{H} = \mathbf{H}(\mathbf{x}_{k+1}^f)$ is the Jacobian of \mathcal{H} and, as in eq. (12), ‘primed’ variables represent perturbations. These equations describe the linear perturbative dynamics of the complete observation–analysis–forecast cycle.

Comparing eq. (14) with (12), a stabilization with respect to the original model system is expected.

Furthermore we notice that, under perfect model conditions and perfect observations, if the control analysis solution is a sufficiently accurate approximation of the true trajectory, the analysis error will obey the same eqs. (14) as the perturbations to the control. If the maximum Lyapunov exponent of the data-assimilation system is negative, the analysis solution will be independent of the initial condition, a necessary condition for convergence to the truth.

The full spectrum of Lyapunov exponents associated to the assimilation cycle can be obtained with standard techniques (Benettin et al., 1980): a set of small orthogonal perturbations, superimposed to a given reference nonlinear trajectory, is evolved by the forced tangent linear eqs. (14) and orthonormalized at regular time intervals.

4. Breeding on the data-assimilation system (BDAS)

In applications, some of the simplifying hypothesis of Section 3 need to be re-examined, to confront with real situations.

When dealing with a system such as the atmosphere or the ocean, the chain from theory to application is not necessarily straightforward and it becomes necessary to find compromises between basic concepts and real complexities. A notable example is the delicate passage from Lyapunov vectors to bred vectors. The breeding method for generating ensemble perturbations devised by Toth and Kalnay (1993, 1997) contains the

basic elements for computing Lyapunov vectors, except that the rescaling amplitude of bred vectors is not infinitesimal. In fact it is necessary to calibrate the rescaling amplitude in atmospheric NWP models in order to prevent the fast growth of the small scale instabilities to become dominant.

In the QG model where the only instabilities are synoptic scale, a distinction between Lyapunov vectors and bred vectors becomes irrelevant. However, since the final goal is to develop assimilation methods that may work in an operational environment, we must use tools that are adequate in that context. Therefore, regardless of the rescaling amplitude that, in the QG model, can be chosen to be infinitesimal, breeding is the viable choice that we turn to for future applications.

As in TU, the bred vectors are meant to represent the instabilities of the data-assimilation system. Therefore, an important difference from standard breeding is that here the bred modes are considered to be solution of the perturbative equations of the data-assimilation system (14), rather than those of the original model (12). As discussed in Section 2, this means that the perturbed trajectories undergo the same forecast and analysis steps as the control trajectory.

As in the normal breeding procedure, a small random isotropic perturbation is superimposed to the control variables. The perturbed trajectories are then integrated with the non-linear model. At the analysis time both the control and perturbed trajectories assimilate the observations. The perturbations are then re-normalized to the initial amplitude and the cycle is repeated.

As usual, if the perturbations were orthonormalized and their amplitude was infinitesimally small, in the limit of infinite breeding time the bred vectors of the data-assimilation system would converge to the Lyapunov vectors of the same system (Benettin et al., 1980).

A second point regards the breeding time. It cannot be infinite, but it has to be made long enough to capture the structure of the most unstable modes. These are the instabilities we wish to control to reduce the forecast error and we also expect the background error to have a large component in the unstable subspace. In fact, along the analysis cycle there may be errors that have had a virtually infinite time to grow. If additional errors are not introduced through the observations (or in the analysis step), after a certain number of steps the only errors will be those that projected on some unstable direction at the first step of the cycle and which could not be eliminated.

In our application, we have optimized the breeding time to obtain the largest projection of the forecast error on the bred vectors; the amplitude of the perturbations has been chosen to be small enough that their evolution can be considered linear.

In the following, we describe how the BDAS procedure is implemented in the case of a single bred vector. At each analysis time a new perturbation is introduced. After the breeding time—typically much longer than the assimilation interval—has elapsed, the bred vector is used in the assimilation procedure as a proxy of an unstable mode. At the next analysis step the

correction is computed along a different vector, previously introduced and bred to represent a different unstable direction. According to their definition, all bred perturbations are subject to the number of assimilations that take place during their breeding time. As discussed in Section 2, at a particular analysis step the assimilation virtually eliminates the error along the direction of the currently used bred vector. Therefore, after being used, the bred vector is discarded and is replaced by a new random perturbation, that starts a new breeding cycle. This *refresh* operation aims at a systematic spanning of the unstable subspace of the system. As a consequence of the *refresh* operation, the observation-forced perturbations are naturally kept independent and do not need to be orthogonalized. The described procedure has been generalized to consider a set of bred vectors rather than a single one (UT).

5. Use of a single bred vector for targeting and assimilation

In all results presented here we apply AUS by making use of a single bred vector at each analysis step: $N = 1$ and the matrix \mathbf{E} reduces to the single column vector \mathbf{e} .

In what follows, we discuss in detail the case of a single *adaptive* observation, ($M = 1$), while in Appendix A we describe the case of more than one observation ($M > 1$).

As discussed in Section 2 and in TU, the adaptive observation should measure the largest component of the bred vector \mathbf{e} . In practice, if we measure a particular variable, the adaptive observation will be placed at the location where the bred vector attains the maximum value of that variable.

If the background error along the single unstable direction \mathbf{e} has an estimated amplitude γ , the background error covariance matrix appearing in the gain matrix will be:

$$\mathbf{P}^f = \gamma^2 \mathbf{e} \mathbf{e}^T. \quad (15)$$

Then, relation (3) reduces to:

$$\mathbf{x}^a = \mathbf{x}^f + \frac{\gamma^2 (\mathbf{H} \mathbf{e})^2}{\gamma^2 (\mathbf{H} \mathbf{e})^2 + \sigma^2} \frac{(y^o - \mathbf{H} \mathbf{x}^f)}{\mathbf{H} \mathbf{e}} \mathbf{e}, \quad (16)$$

where σ^2 is the observation error variance of the single observation. In the present application (see Section 6), σ depends both upon the type of variable being observed and the observation level (Morss, 1999) and the observation operator is inherently linear. We notice that in the particular case discussed here, ($M = 1$, $N = 1$), the vectors \mathbf{y}^o , $\mathbf{H} \mathbf{x}^f$ and $\mathbf{H} \mathbf{e}$ reduce to scalars.

The first ratio on the right hand side of (16) represents the relative weight to be given to the observation with respect to the background, when the error along the direction \mathbf{e} is projected on the observation subspace. It only modifies the amplitude of the correction, not its direction, which is given by the vector \mathbf{e} .

When observations are perfect, $\sigma = 0$ and (16) becomes:

$$\mathbf{x}^a = \mathbf{x}^f + \frac{y^o - \mathbf{H} \mathbf{x}^f}{\mathbf{H} \mathbf{e}} \mathbf{e} \quad (17)$$

and the correction in the direction \mathbf{e} fits the observation exactly. In the experiments with perfect observations direct use of (17) is made to assimilate the adaptive observations (see also UTC), while in the experiments with noisy observations we turn to (16). In the latter case, besides estimating the direction \mathbf{e} , we need to evaluate γ .

The amplification of the error in the forecast step and its reduction at the analysis step, needed to evaluate the forecast error projection along \mathbf{e} , could be recursively evaluated if the same direction \mathbf{e} were maintained during the analysis cycle. The estimate of the forecast error covariance would then follow from (15). Instead, for the reasons explained above, we adopted a *refresh* operation. Therefore we chose to make a statistical evaluation of the background error along the current (and renewed) direction used for the assimilation.

More details on the estimate of γ can be found in Appendix B.

6. Experiments with the QG Model

6.1. The model

The test ground of the method is a quasi-geostrophic model (Rotunno and Bao, 1996) already used by several authors in data-assimilation and adaptive observations studies (Morss et al., 2001; Hamill and Snyder, 2000, 2002; Hamill et al., 2002; Corazza et al., 2003). The model exhibits dynamical behaviour similar to real atmospheric flow, but is simple enough to make long runs feasible.

It is a β -plane periodic channel model with rigid lids, with 64 longitude (16 000 km), 33 latitude (8000 km) gridpoints and a depth of 9 km. Potential vorticity is defined on five inner vertical levels and potential temperature at top and bottom boundaries, so that the number of degrees of freedom of the model is ($64 \times 33 \times 7 = 14784$). It is forced by relaxation to a zonal mean state with constant stratification and damped by a ∇^4 horizontal diffusion and by an Ekman pumping at the surface.

The model, with the present resolution, has 24 positive Lyapunov exponents; the leading exponent is approximately $.31 \text{ d}^{-1}$, corresponding to a doubling time of about 2.2 d. The Kaplan–Yorke attractor dimension is approximately 65.2.

6.2. Implementation of BDAS and AUS

In all the observation system simulation experiments described below, the model is considered to be perfect and a long reference model trajectory represents the true state evolution from which observations are taken.

In order to simulate an irregular observation distribution, with densely and poorly observed areas, the domain has been divided into two sectors. A first portion, referred to as ‘land’, is

completely covered by observations while the remaining part represents an ‘ocean’ void of observations except for the possible addition of observations located at a particular gridpoint adaptively chosen at each assimilation time (Lorenz and Emanuel, 1998, TU, UTC)

As discussed in Sections 2 and 4, the bred vectors are obtained by BDAS which means that at each assimilation time all, perturbed and unperturbed, trajectories are subject to the same assimilation process. In the targeting-assimilation application that follows, BDAS is used for two purposes: to identify the location for the adaptive observations and to estimate the unstable vectors that enter the AUS analysis solution. The adaptive observation location is chosen as the horizontal gridpoint where the current bred vector attains its maximum value (in absolute value, for the variable being observed).

The breeding time has been optimized to obtain the largest projection of the forecast error on the bred vectors. By varying the breeding time from 1 to 10 d, the forecast error projection on the bred modes increases significantly up to 6 d, while gradually smaller improvements are obtained by using longer breeding times. In view of these results, the breeding time has been set to 10 d.

We recall that the bred vectors are discarded after being used and new random perturbations are introduced. Due to this refresh operation, the number of bred vectors that are simultaneously computed is equal to the ratio between the breeding time and the assimilation interval: for example, for a breeding time of 10 d, an assimilation interval of 12 hr implies that 20 bred vectors are simultaneously computed.

TU already noted, in the context of the Lorenz–Emanuel model, the necessity to regionalize the correction made using the unstable structure in order to avoid the following unwanted effect. If the background error and the bred mode were correlated over the whole domain there would be no need of a regionalization. Instead, the background error is, in general, the superposition of different modes and the bred vector itself (that only ideally represents a single Lyapunov vector) present several spatially separated structures. Consequently, a perfect correlation cannot be expected throughout the domain and the assimilation only leads to a reduction of the error in the area surrounding the observations. The presence of secondary structures in the bred mode, in regions distant from the observations, with opposite correlation with the background with respect to the observed structure, can lead to a correction with the wrong sign. We found the same problem during the experimentation with the QG model and also in the present work, we make use of a modulating function, choosing a simple Gaussian with a decay scale of 2500 Km (10 gridpoints). This scale is large enough to preserve the structure of the bred mode and consequently of the analysis increment in a wide area around the observation, nevertheless this modulation is sufficient to overcome the problems encountered.

7. Results

This section is aimed on one hand at estimating the impact of the BDAS targeting strategy, on the other at comparing the performance of AUS with that of a standard 3DVar assimilation scheme, when only poor and irregular information about the state of the system is available. To this end, as explained in Section 6, we intentionally divided the model domain into a land area, where errors are kept small by a complete observation coverage, and an ocean area where no observations are present and errors are free to grow, unless adaptive observations are introduced. Land observations consist of soundings (vertical profiles of temperature and horizontal velocity) and are assimilated by means of 3DVar (Parrish and Derber, 1992; Morss, 1999). The land covers one third of the domain (gridpoints 1–20). A single adaptive observation over the ocean, located by means of BDAS (see Section 6) is introduced at each assimilation step and consists either of a sounding or of the measurement of a single scalar quantity, such as temperature or a velocity component, at the targeted location and level. Using a similar observational network setting and perfect observations UTC already obtained a dramatic improvement when AUS is used instead of 3DVar to assimilate the adaptive observations.

All experiments described below consist of 2-yr assimilation cycles and are initialized with the solution of a long analysis cycle that assimilates only land observations by means of 3DVar. The random observation errors that are added to the true values are produced using a Gaussian distribution and the eigenvectors and eigenvalues of the observation error covariance matrix. \mathbf{R} is the same as in Morss (1999) who adapted it from the operational error covariance described in Parrish and Derber (1992). The assimilation interval is set to 6 hr.

7.1. 3DVar—AUS comparison

The following two experiments (3DVar–BDAS and AUS–BDAS) are aimed at comparing the relative performance of AUS and 3DVar when the same observational network and adaptive strategy is used. In the first, 3DVar is used to assimilate an adaptive vertical sounding while in the second the adaptive observation, consisting of a single temperature measurement, is assimilated with AUS via (16). In both experiments, fixed land observations are assimilated by means of 3DVar and the adaptive observation location is identified using BDAS. Thus the main difference between the two experiments is in the way the adaptive observations are assimilated and in the dimension of the *adaptive* observation vector which, in the first experiment, is 21 (three scalar quantities at seven levels) while it is only one in the second.

Figure 1 shows the normalized (with respect to natural variability) root mean square analysis error (potential enstrophy norm) as a function of time. The dashed line is relative to the experiment with 3DVar (adaptive vertical profile assimilated




Fig. 1. Normalized RMS analysis error as a function of time. The error is expressed in the potential enstrophy norm and it is normalized by the natural variability (RMS difference between decorrelated states). Dashed line: experiment with 3DVar–BDAS. Continuous line: experiment with AUS–BDAS.

using 3DVar), while the continuous line refers to the experiment with AUS (one temperature observation assimilated using eq. 16). After about two weeks the error in experiment with AUS drops drastically and thereafter is confined within a range of values below 10% of the natural variability. The improvement with respect to 3DVar appears even more impressive considering that AUS assimilates only a temperature observation rather than a sounding. When a single temperature observation, rather than a sounding, is assimilated with 3DVar, the average error is about one order of magnitude larger than that obtained with AUS. On the other hand, only minor improvements are obtained when assimilating a complete sounding with AUS.

The panels of Fig. 2 show the horizontal structure of the variability-normalized, vertically and time averaged (over the last 580 d) analysis error. In both panels the error spatial structure clearly reflects the specific observational configuration, with the land covered by observations. It is possible to note a strong error maximum over the eastern part of the ocean, and a partial downstream advection over the western portion of land in correspondence to the jet stream. But apart from these similarities, the error in the two panels differs enormously. In panel (b), that refers to AUS–BDAS experiment, the maximum error over the ocean is one third of the corresponding maximum in experiment with 3DVar.

Figure 3 shows the longitudinal dependence of the error in the two experiments obtained from the fields of Fig. 2, after averaging also in latitude. It appears that by combining an adaptive observation strategy with the dynamic assimilation (AUS), the unstable components of the forecast error can be efficiently eliminated to the extent that the error over the ocean becomes comparable on average to the error over land of experiment with 3DVar.

These results clearly reveal that bred modes constructed by BDAS share a strong similarity with the actual flow dependent forecast error. Such a good estimate of the forecast error is in turn at the basis of AUS ability to exploit a smaller amount of information by spreading the analysis correction in a dynamically consistent manner. Moreover, once the errors have been reduced in the course of the assimilation, the unstable vectors become more representative of the background errors which can then be efficiently controlled. We have argued that the analysis error reduction is accompanied by the reduction of the dimension of the unstable subspace; as a consequence, the forecast error is expected to be described by few bred modes.

The stability analysis discussed in Section 3 is now applied and the leading exponent of the forced assimilation system is computed. Figure 4 shows instantaneous and time-average values of the growth rate for experiments 3DVar–BDAS and AUS–BDAS. As can be seen, while the instantaneous values of the growth rate for the 3DVar experiment shows large fluctuations, the dynamical consistency of AUS is reflected in the more regular behaviour of the corresponding exponent. Moreover the leading exponent for the AUS experiment is negative. A negative growth rate formally means that the unstable subspace of the system is zero-dimensional, a necessary condition for the independence of the analysis solution from the initial condition. However, stability guarantees uniqueness of the analysis solution but not convergence to the truth.

As a final point, we would like to comment on the use of results from stability analysis in relation to the quality of an assimilation scheme. As discussed in Section 3, a connection between the level of error and the stability of the solution can be expected only if the control analysis solution is a sufficiently accurate estimate of the true trajectory. Only in this hypothesis,




Fig. 2. Time and vertically averaged RMS analysis error. Upper panel: experiment with 3DVar-BDAS. Lower panel: experiment with AUS-BDAS.

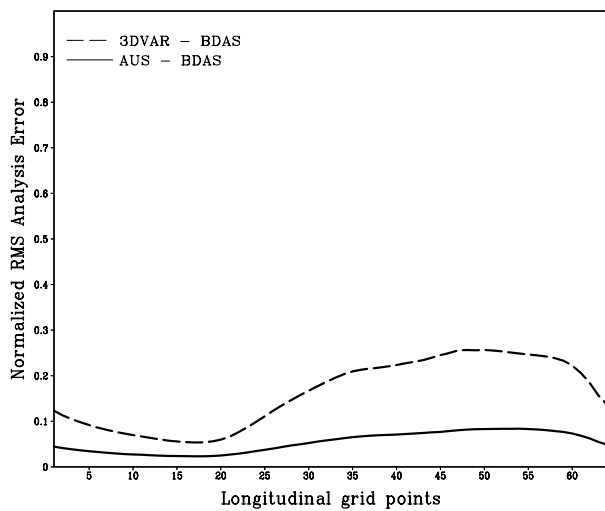


Fig. 3. Longitudinal dependence of normalized RMS analysis error. Dashed line: experiment with 3DVar-BDAS. Continuous line: experiment with AUS-BDAS.




Fig. 4. Stability analysis. Upper panel: mean growth rate as a function of time for 3DVar-BDAS (dashed line; leading exponent = 0.006 d^{-1}) and for AUS-BDAS (continuous line; leading exponent = -0.052 d^{-1}). Middle panel: instantaneous growth rate as a function of time for experiment 3DVar-BDAS. Lower panel: instantaneous growth rate as a function of time for experiment AUS-BDAS.




Fig. 5. Time and domain average RMS forecast error as a function of forecast range. Dashed line: experiment with 3DVar-BDAS. Continuous line: experiment with AUS-BDAS.

perturbations of the reference trajectory (the ‘truth’) and those of the control analysis solution will obey approximately the same tangent linear equations.

To complete the comparison of AUS with 3DVar, time and domain averaged forecast error are shown in Fig. 5. It should be stressed that the error doubling time in the first 12 hr of forecast in experiments with AUS is about 6 d and becomes comparable to the doubling time of 2.2 d only after 60 hr of integration. This means that the most unstable modes were effectively eliminated by the AUS scheme. In experiments with 3DVar the analysis error is already too large for any comparison with doubling times of infinitesimal errors to be significant.

7.2. Performance of BDAS

In order to provide an estimate of the relative benefit of the adaptive observations strategy based on BDAS we perform the following three experiments where observations are assimilated with 3DVar:

- LO: Observations only over land;
- RO: Land observations plus a randomly located ocean observation (vertical sounding);
- FO: Land observations plus a fixed observation (vertical sounding) located at a point where a forecast error maximum is expected.

The ocean observation location of experiment FO, has been chosen as the point where a maximum forecast error is found in experiment LO: $x = 56$, $y = 16$. In Table 1 the RMS analysis error and the leading exponent for LO, RO, FO, 3DVar-BDAS and AUS-BDAS, are shown.

The insertion of a supplemental fixed sounding over the ocean (FO) reduces the RMS analysis error by about 25% but the sys-

tem is still unstable. Experiment RO gives similar results although a small improvement with respect to FO is achieved. Although, when placed where maximum error is statistically expected, a fixed observation should give the maximum benefit, it is however less beneficial than a randomly located observation because it provides redundant information in the specific area that is repeatedly observed. When BDAS is used to identify the location of the adaptive observations, its impact on the analysis is evident. The RMS error is reduced by about 60% with respect to LO and the leading exponent of the analysis-forecast system solution is also significantly reduced. Finally, when AUS is used in combination with BDAS the improvement is remarkable: the error is reduced to 6% of the natural variability and the forced system is stabilized.

8. Stabilization of the assimilation system: experiments in a limited area model configuration

The following experiments provide evidence in support of the argument that the assimilation cycle can be considered as a forced system and of the conclusion that the observational forcing, acting on the system at the analysis step, results in a reduction of the dimension of the unstable subspace.

In particular, the more efficient is the observational network and the analysis scheme in driving the control trajectory towards the true one, the more drastic is the stabilization of the assimilation system.

The first example is given by simulations with a hypothetical limited area model of variable size where the time dependent boundary conditions constitute the observational forcing. In the second, the same model is endowed with additional observations and with targeting and assimilation algorithms (AUS-BDAS). We restrict our attention to the perfect observation case that is relevant for establishing the observability condition (see Section 2.1).

Starting from the periodic channel QG model, a limited area model of variable east-west extension is constructed by reducing the number of internal longitudinal gridpoints. Perfect lateral boundary conditions for the limited area model are prescribed in the following way. A trajectory obtained by integrating the channel model represents the truth. The lateral boundary conditions for the limited area model are prescribed by setting the values of the model variables equal to the true values at all gridpoints in a longitudinal portion of the channel. By varying the number of gridpoints, a limited area model of variable longitudinal extension is obtained. A cubic interpolation relaxation band is introduced at the east-west boundaries. Without any additional observations, the instabilities are free to grow within the limited area domain and we expect the number of unstable modes that can grow to be dependent on the domain size. Therefore when, in addition to the boundary forcing, we introduce observations within the limited area domain, we

Table 1. Error statistics and stability analysis for experiment LO, FO, RO, 3DVar–BDAS and AUS–BDAS (see text for details). Errors are expressed in potential enstrophy norm and are normalized with natural variability. Leading exponents are expressed in d^{-1} . The values are relative to the last 580 d of 2-yr-long runs with noisy observations

Experiment	Ocean observation type	Ocean observation location	Ocean observation assimilation	RMS error	Leading exponent
LO	–	–	–	0.462	0.092
FO	Sounding	Fixed ($x = 42, y = 16$)	3DVar	0.338	0.076
RO	Sounding	Random	3DVar	0.311	0.063
3DVar–BDAS	Sounding	BDAS	3DVar	0.184	0.006
AUS–BDAS	Temperature	BDAS	AUS	0.060	–0.052

Table 2. Error statistics for experiments NOA (no observations) and AUS–BDAS implemented with different assimilation interval (τ), as a function of limited area model longitudinal extension (DLE) expressed in number of gridpoints. Errors are expressed in potential enstrophy norm and are normalized with natural variability. The values are relative to the last 580 d of 2-yr-long runs

DLE domain longitudinal extension	NOA no observation	AUS ($\tau = 24$ hr)	AUS ($\tau = 12$ hr)	AUS ($\tau = 6$ hr)	AUS ($\tau = 3$ hr)
60	.756	.742	.729	0.153	.024
55	.653	.623	.611	0.026	.020
50	.616	.542	.036	0.020	.019
45	.523	.336	.019	0.016	.016
40	.452	.026	.016	0.014	.013
35	.220	.016	.014	0.012	.012
30	.092	.010	.010	0.010	.010

Table 3. Leading exponent for experiments NOA (no observation) and AUS–BDAS implemented with different assimilation interval (τ), as a function of limited area model longitudinal extension (DLE) expressed in number of gridpoints. Leading exponents are expressed in d^{-1} . The first number in parenthesis indicates the unstable subspace dimension, while the second is the Kaplan–Yorke dimension. The values are relative to the last 580 d of 2-yr long runs

DLE domain longitudinal extension	NOA no observation	AUS ($\tau = 24$ hr)	AUS ($\tau = 12$ hr)	AUS ($\tau = 6$ hr)	AUS ($\tau = 3$ hr)
60	.160 (7; 15.5)	.066 (5; 9.1)	.021 (3; 5.7)	–.024	–.107
55	.143 (6; 12.5)	.064 (4; 7.7)	.019 (1; 3.4)	–.057	–.131
50	.120 (5; 9.9)	.024 (2; 3.3)	–.037	–.103	–.166
45	.076 (3; 6.2)	–.019	–.070	–.137	–.196
40	.062 (3; 5.1)	–.055	–.125	–.189	–.253
35	.051 (1; 2.4)	–.100	–.172	–.242	–.331
30	–.007	–.192	–.256	–.342	–.384

expect a smaller number and frequency of observations to be needed to control the system when the domain is of smaller size.

In a first set of experiments, that are referred to as NOA (no assimilation of observations), the limited area model equations are integrated without additional observations. The dimension of the unstable subspace and the leading exponent as a function of the domain size are computed. In a second set of experiments, perfect adaptive observations are assimilated with a different assimilation interval: 3, 6, 12 or 24 hr. A single adaptive observation is assimilated at each analysis time. The location of the adaptive observations is determined by BDAS and the observa-

tions are assimilated with AUS. The analysis error is estimated in all cases as an average over the last 580 d of 2 yr analysis cycles.

Tables 2 and 3 summarise, respectively the error statistics and stability analysis of the experiments. The longitudinal extension of the limited area domain varies from 7500 km (domain length, $DLE = 30$ gridpoints) to 15 000 km ($DLE = 60$). In the experiments indicated as NOA, the assimilation system is unstable for all limited-area extensions considered, except for the smallest one: when the system becomes stable (leading exponent, $\lambda_{\max} = -0.007$), the analysis error drops from 22% to 9% of the natural variability.

We notice that, even if the model and the boundary conditions are perfect the error is not zero. This means that errors are introduced that prevent the solution to converge to the truth, even if the model and boundary conditions are perfect: stability is a necessary but not sufficient condition for convergence to the truth. The dimension of the unstable subspace decreases gradually from 7 to 0, by decreasing the domain size from its maximum to its minimum extension. The leading Lyapunov exponent also gradually decreases from a value of .160 (DLE = 60) to a negative value (DLE = 30). The Kaplan–Yorke attractor dimension in turn decreases from 15.548 (DLE = 60) to 2.448 (DLE = 35).

Tables 2 and 3 summarise the results of the experiments with adaptive observations taken in the limited area domain assimilated by AUS. We see that the degree of stabilization and reduction of the error with respect to the reference experiments NOA, depend upon the frequency of observations. In general the magnitude of the leading exponent as well as the dimension of the unstable subspace and Kaplan–Yorke attractor dimension are further and further reduced by decreasing the assimilation interval. Accordingly, the error exhibits a progressive reduction and drops to remarkably low levels in correspondence to the system stabilization. Therefore by decreasing the extension of the domain and decreasing the assimilation interval, the system becomes more and more stable and the analysis error reaches values as low as 1% of natural variability. In particular, in the experiment with adaptive observations assimilated every 6 or 3 hrs, the leading exponent is negative and the error is quite small even in the case of the largest domain (DLE = 60).

In summary, Tables 2 and 3 reveal the following picture. When, relative to the number and growth of the unstable vectors that characterise a given limited area domain, the frequency of observations is too low to capture the growing modes, the targeting assimilation process is inefficient and the level of error high. Let us confine our attention to the successful experiments, defined as those where a system stabilization is obtained by AUS–BDAS and error is reduced to less than 1% of natural variability. A non-dimensional parameter can be constructed by multiplying the number of positive exponents (N^+) by the leading one (λ_{\max}) and by the assimilation interval (τ). Figure 6 shows the average error for the successful experiments as a function of this parameter. The error level at which the system stabilizes appears to be a monotonic function of this parameter with an approximate linear dependence. This result supports our claim that the number (frequency) of observations needed to stabilize the system and efficiently reduce the analysis error depends upon the number and growth rate of the unstable directions.

9. Conclusions

The assimilation in the unstable subspace developed by TU, applied to a QG model in an adaptive observation configuration,




Fig. 6. Time and domain average RMS error as a function of the non-dimensional parameter $N^+ \lambda_{\max} \tau$ (see text for details).

with observational noise, proved capable of reducing the analysis errors to a remarkably low level.

The key to this success is in the mutual enhancement of two beneficial effects obtained by:

- the strategy of observing the most unstable components of the data-assimilation system (BDAS);
- the effective reduction of the error in the unstable subspace achieved by the assimilation (AUS).

This is because, when errors in the unstable directions are efficiently reduced, the error becomes smaller and behaves more linearly. Therefore, the unstable directions themselves become more representative of the actual error, the unstable structures take a longer time to build up and a smaller number and lower frequency of observations becomes sufficient to control their growth.

The estimate of the unstable directions, consistent with the stability analysis of the data-assimilation system, is obtained by a modified breeding technique, BDAS, that naturally incorporates the information on the observational network, the assimilation system and its dynamical instabilities.

In the QG-model experiments, the benefit obtained by the targeting strategy based on BDAS and the assimilation in the unstable subspace (AUS) is estimated.

The comparison with a 3DVar algorithm, extended here to the case of noisy observations, confirms the ability of the BDAS technique to estimate the actual unstable modes of the assimilation system and the efficiency of the dynamically based assimilation (AUS).

Furthermore results show that, as predicted by the theory, the observational forcing reduces the dimension of the unstable subspace and stabilizes the assimilation system; in particular, the number of observations needed to stabilize the system is related to the instability properties of the free system.

It remains to be seen how this method works in an operational environment. Experience with this model and with a primitive equation ocean model (UTC, UT) suggests that the method has a good performance in different contexts, with minor details of implementation depending on the model and the specific observational configuration. In fact, the introduction of adaptive observations, based on the same principles of the present work, in the complex, primitive equation, large dimension ocean system, leads to a remarkable reductions of errors. These results suggest some optimism with regard to the applicability of this approach to more complex atmospheric and oceanic prediction models and real observational networks.

11. Acknowledgments

The authors are thankful to two anonymous reviewers for their comments and suggestions.

12. Appendix A: Analysis solution for $M > N = 1$

From (3) and the appropriate expression for \mathbf{K} when $M > N$, (4b), the analysis solution takes the form:

$$\mathbf{x}^a = \mathbf{x}^f + \mathbf{E} [\Gamma^{-1} + (\mathbf{H}\mathbf{E})^T \mathbf{R}^{-1} (\mathbf{H}\mathbf{E})]^{-1} (\mathbf{H}\mathbf{E})^T \mathbf{R}^{-1} (\mathbf{y}^o - \mathbf{H}\mathbf{x}^f). \quad (\text{A1})$$

Consider the case when many observations and a single unstable direction are used, $M > N = 1$.

The matrix $\mathbf{E} = \mathbf{e}$ has just one column and the matrix $\Gamma = \gamma^2$ reduces to a scalar. The matrix $\mathbf{H}\mathbf{E}$ is a column vector of length M : $\mathbf{y}^e = \mathbf{H}\mathbf{e}$. The analysis state is then:

$$\mathbf{x}^a = \mathbf{x}^f + \frac{(\mathbf{y}^e)^T \mathbf{R}^{-1} (\mathbf{y}^o - \mathbf{H}\mathbf{x}^f)}{\gamma^{-2} + (\mathbf{y}^e)^T \mathbf{R}^{-1} \mathbf{y}^e} \mathbf{e}. \quad (\text{A2})$$

Appendix B: Estimate of γ for $M \geq N = 1$

The innovation, $\mathbf{d} = \mathbf{y}^o - \mathbf{H}\mathbf{x}^f$, and the vector $\mathbf{y}^e = \mathbf{H}\mathbf{e}$ are vectors of length M . In the case $M > N = 1$, with $\mathbf{R} = \sigma^2 \mathbf{I}$, the analysis (A2) is:

$$\mathbf{x}^a = \mathbf{x}^f + \frac{\sigma^{-2} (\mathbf{y}^e)^T (\mathbf{y}^o - \mathbf{H}\mathbf{x}^f)}{\gamma^{-2} + \sigma^{-2} (\mathbf{y}^e)^T \mathbf{y}^e} \mathbf{e} \quad (\text{B1})$$

or, equivalently:

$$\mathbf{x}^a = \mathbf{x}^f + \frac{\gamma^2 (\mathbf{y}^e)^T \mathbf{d}}{\sigma^2 + \gamma^2 (\mathbf{y}^e)^T \mathbf{y}^e} \mathbf{e}. \quad (\text{B2})$$

After multiplying and dividing by $(\mathbf{y}^e)^T \mathbf{y}^e$ and rearranging we obtain:

$$\mathbf{x}^a = \mathbf{x}^f + \frac{\gamma^2 (\mathbf{y}^e)^T \mathbf{y}^e}{\sigma^2 + \gamma^2 (\mathbf{y}^e)^T \mathbf{y}^e} \cdot \frac{(\mathbf{y}^e)^T \mathbf{d}}{(\mathbf{y}^e)^T \mathbf{y}^e} \mathbf{e}. \quad (\text{B3})$$

In (B3), the background error variance appears in a form that is convenient for a direct evaluation because the first ratio on the

right hand side is a term in the form of a weight that we want to estimate.

If, as usual, observations and forecast error are assumed to be uncorrelated we have:

$$\langle \mathbf{d}^T \mathbf{d} \rangle \approx \langle \gamma^2 (\mathbf{y}^e)^T \mathbf{y}^e \rangle + \sigma^2, \quad (\text{B4})$$

where $\langle \rangle$ represents the expectation operator.

Clearly, at the beginning of the experiment we do not have this statistics available. Furthermore, during the experiment, large fluctuations of the background error occur. In addition to the familiar flow dependence of the forecast error, the analysis error is subject to transient behaviour before it stabilizes (see Fig. 1). Therefore, in the estimate of γ , it is convenient to consider a time mean over the recent past:

$$\langle \gamma^2 (\mathbf{y}^e)^T \mathbf{y}^e \rangle_\tau \approx \langle \mathbf{d}^T \mathbf{d} \rangle_\tau - \sigma^2, \quad (\text{B5})$$

where $\langle \rangle_\tau$ represents the average over an appropriate time interval τ .

For finite τ there is no guarantee that the term on the right hand side of (B5) is positive. Therefore we decided to neglect σ^2 in (B5) and let our estimate be obtained from:

$$\langle \gamma^2 (\mathbf{y}^e)^T \mathbf{y}^e \rangle_\tau \approx \langle \mathbf{d}^T \mathbf{d} \rangle_\tau. \quad (\text{B6})$$

In practice, in experiments of type II we use:

$$\mathbf{x}^a = \mathbf{x}^f + \frac{\langle \mathbf{d}^T \mathbf{d} \rangle_\tau}{\sigma^2 + \langle \mathbf{d}^T \mathbf{d} \rangle_\tau} \cdot \frac{(\mathbf{y}^e)^T \mathbf{d}}{(\mathbf{y}^e)^T \mathbf{y}^e} \mathbf{e}. \quad (\text{B7})$$

The time average is relative to the last time interval τ that, after tuning, has been set equal to 8 d.

We point out that this is just one, and possibly not the best, among all viable choices.

References

- Benettin, G., Galgani, L., Giorgilli, A. and Strelcyn, J. M. 1980. Lyapunov characteristic exponents for smooth dynamical system and for Hamiltonian system; a method for computing them. *Meccanica* **15**, Part I: theory 9-20–Part II: application 21–30.
- Bleck, R., 1978. Simulation of coastal upwelling frontogenesis with an isopycnic coordinate model. *J. Geophys. Res.* **83C**, 6163–6172.
- Corazza, M., Kalnay, E., Patil, D. J., Yang, S. C., Morss, R. and co-authors. 2003. Use of the breeding technique to estimate the structure of the analysis ‘error of the day’. *Nonlin. Processes Geophys.* **10**, 233–243.
- Daley, R. 1991. *Atmospheric Data Analysis*, Cambridge University Press, Cambridge, 457 pp.
- Fourrié, N., Marchal, D., Rabier, F., Chapnik and Desroziers, G. 2006. Impact study of the 2003 north Atlantic THORPEX regional campaign. *Quart. J. Roy. Meteorol. Soc.* **132**, 275–295.
- Ghil, M. and Todling, R. 1996. Tracking atmospheric instabilities with the Kalman filter. Part II: Two-layer results. *Mon. Wea. Rev.* **124**, 2340–2352.
- Ghil, M. 1997. Advances in sequential estimation for atmospheric and oceanic flows. *J. Meteorol. Soc. Japan* **75**, 289–304.

- Hamill, T. M. and Snyder, C. 2000. A hybrid ensemble Kalman filter-3D-variational analysis scheme. *Mon. Wea. Rev.* **128**, 2905–2919.
- Hamill, T. M. and Snyder, C. 2002. Using improved background-error covariances from an ensemble Kalman filter for adaptive observations. *Mon. Wea. Rev.* **130**, 1552–1572.
- Hamill, T. M., Snyder, C. and Morss, R. E. 2002. Analysis-error statistics of a quasigeostrophic model using three-dimensional variational assimilation. *Mon. Wea. Rev.* **130**, 2777–2790.
- Ide, K., Courtier, P., Ghil, M. and Lorenc, A. C. 1997. Unified notation for data-assimilation: Operational, sequential and variational. *J. Meteorol. Soc. Japan* **75**, 181–189.
- Kalnay, E. 2003. *Atmospheric Modelling, Data Assimilation and Predictability*, Cambridge University Press, Cambridge, 341 pp.
- Legras, B. and Vautard, R. 1996. A guide to Lyapunov vectors. In: *Predictability Volume 1, Seminar Proceedings*, ECMWF, Reading, UK, pp. 143–156.
- Lorenz, E. N. 1984. The local structure of a chaotic attractor in four dimensions. *Physica* **13D**, 90–104.
- Lorenz, E. N. 1996. Predictability: a problem partly solved. In: *Predictability Volume 1, Seminar Proceedings*, ECMWF, Reading, UK, pp. 1–18.
- Lorenz, E. N. and Emanuel, K. A. 1998. Optimal sites for supplementary weather observations: Simulation with a small model. *J. Atmos. Sci.* **55**, 399–414.
- Morss, R. E. 1999. *Idealized sampling strategies for improving numerical weather prediction*. PhD Thesis, Massachusetts Institute of Technology, Cambridge, MA. Available from UMI Dissertation Services, 300 N.Zeeb Rd., Ann Arbor, MI 48106, 225 pp.
- Morss, R. E., Emanuel, K. A. and Snyder, C. 2001. Idealized adaptive observation strategies for improving numerical weather prediction. *J. Atmos. Sci.* **58**, 210–232.
- Oseledec, V. I. 1968. A multiplicative ergodic theorem: Lyapunov characteristic numbers for dynamical systems. *Trans. Moscow Math. Soc.* **19**, 197–231.
- Parrish, D. F. and Derber, J. D. 1992. The National Meteorological Center spectral statistical interpolation analysis system. *Mon. Wea. Rev.* **120**, 1747–1763.
- Pires, C., Vautard, R. and Talagrand, O. 1996. On extending the limits of variational assimilation in chaotic systems. *Tellus* **48A**, 96–121.
- Rotunno, R. and Bao, J. W. 1996. A case study of cyclogenesis using a model hierarchy. *Mon. Wea. Rev.* **124**, 1051–1066.
- Snyder, C. and Hamill, T. M. 2003. Leading Lyapunov Vectors of a Turbulent Baroclinic Jet in a Quasigeostrophic Model. *J. Atmos. Sci.* **60**, 683–688.
- Szunyogh, I., Toth, Z., Zimin, A. V., Majumdar, S. J. and Persson, A. 2002. Propagation of the effect of targeted observations: the 2000 Winter Storm Reconnaissance Program. *Mon. Wea. Rev.* **130**, 1144–1165.
- Talagrand, O. 1997. Assimilation of observations, an introduction. *J. Meteorol. Soc. Japan* **75**, 191–209.
- Todling, R. and Ghil, M. 1994. Tracking atmospheric instabilities with the Kalman filter. Part I: Methodology and one-layer results. *Mon. Wea. Rev.* **122**, 183–204.
- Toth, Z. and Kalnay, E. 1993. Ensemble forecasting at NMC: The generation of perturbations. *Bull. Amer. Meteorol. Soc.* **74**, 2317–2330.
- Toth, Z. and Kalnay, E. 1997. Ensemble forecasting at NCEP: the breeding method. *Mon. Wea. Rev.* **125**, 3297–3318.
- Trevisan, A. and Pancotti, F. 1998. Periodic orbits, Lyapunov vectors and singular vectors in the Lorenz system. *J. Atmos. Sci.* **55**, 390–398.
- Trevisan, A. and Uboldi, F. 2004. Assimilation of Standard and Targeted Observations within the Unstable Subspace of the Observation-Analysis-Forecast Cycle System. *J. Atmos. Sci.* **61**, 103–113.
- Uboldi, F., Trevisan, A. and Carrassi, A. 2005. Developing a dynamically based assimilation method for targeted and standard observations. *Nonlin. Processes Geophys.* **12**, 149–156.
- Uboldi, F. and Trevisan, A. 2006. Detecting unstable structures and controlling error growth by assimilation of standard and adaptive observations in a primitive equation ocean model. *Nonlin. Processes Geophys.* **13**, 67–81.
- Vannitsem, S. 2003. Intrinsic error growth in a large-domain eta regional model. *Mon. Wea. Rev.* **131**, 2697–2704.
- Wang, X. and Bishop, C. H. 2003. A comparison of breeding and ensemble transform Kalman filter ensemble forecast scheme. *J. Atmos. Sci.* **60**, 1140–1158.