The Geometry of Model Error

Kevin Judd, Carolyn A. Reynolds, Leonard A. Smith, and Thomas E. Rosmond

August 6, 2007

Abstract

This paper investigates the nature of model error in complex deterministic nonlinear systems, such as weather forecasting models. Forecasting systems incorporate two components, a forecast model and a data assimilation method. The latter projects a collection of observations of reality into a model state. Key features of model error can be understood in terms of geometric properties of the data projection and a model's *attracting manifold*. Model error can be resolved into two components, a *projection error*, which can be understood as the model's attractor being in the wrong location given the data projection, and *direction error*, which can be understood as the trajectories of the model moving in the wrong direction compared to the projection of reality into model space. Our investigation introduces some new tools and concepts, including, the *shadowing filter*, *causal* and *non-causal shadow analyses*, and various geometric diagnostics. Various properties of forecast errors and model errors are described with reference to low dimensional systems, like Lorenz's equations, then an operational weather forecasting system is shown to have the same predicted behaviour. The concepts and tools introduced show

promise for the diagnosis of model error and the improvement of ensemble forecasting systems.

1. Introduction

In operational numerical weather prediction (NWP), data assimilation is a process whereby a series of observations is transformed into a single best-guess model state or an ensemble of model states from which forecasts are to be launched. In the perfect model scenario, an ensemble would consist of a set of model states, each the end point of a model trajectory consistent with observations of the system. If the system evolves on an attractor, then the ensemble members should lie on the attractor. If they do not, then it can be easily shown for nonlinear systems that having states not on the attractor can significantly degrade forecasts (Judd 2003). Even when the model is imperfect, sampling the full state space of the model is less efficient than sampling the manifold of states consistent with the model dynamics; in high-dimensional models the difference in efficiency can be vast. And even a single "best guess" forecast can benefit from being consistent with longer term dynamics of the model. In this paper we present a series of arguments and numerical experiments to support three conjectures. First, attracting

manifolds exist in operational weather models. Secondly, model states can be found which lie much closer to the relevant manifold than the output of current data assimilation algorithms. Thirdly, these states are beneficial to model development and forecasting. This extends earlier work on low dimensional chaotic maps (Judd et al. 2004b) to the Navy Operational Global Atmospheric Prediction System (NOGAPS) (Baker et al. 1998).

In section 2 we first provide a brief overview of some important properties of nonlinear dynamical systems from a geometric point of view. It is noted that model states can be expected to evolve toward an *attracting manifold*, which is of lower dimensional than the entire state space. We argue that forecast errors can be resolved into two distinct components: one due to initial conditions not being on the attracting manifold, the other due to model error. We argue that using initial conditions on the attracting manifold (a *shadow analysis*) would be expected to provide benefits in a wide class of dynamical models, including improved forecasts and detection of model error. To establish our claims we discuss the signatures and geometry of various error growths. This enables us to highlight potential shortcomings of forecast systems.

It is not easy to visualise the decomposition of uncertainty in high-dimensional spaces. In section 3 we introduce methods to extract information about the dynam-

ics of uncertainty and error growth from different initial conditions in operational weather models. First a method for locating analyses on (or near) the attracting manifold is introduced. We then show how low order polygons (triangles and tetrahedra) can be employed to effectively extract information from a handful of trajectories in a high-dimensional space and displaying it in an intuitively accessible form. We derive the behaviour we expect to see in these graphs if indeed attracting manifolds are relevant to dynamics.

Section 4 presents the results from exploring these ideas in high-dimensional, operational weather model. We show that the dynamics indicate that an attracting manifold plays a significant role in the dynamics of NOGAPS, suggesting that operational forecasting systems might usefully take this into account. We also note how this style of analysis can provide strategic insight into the details of model inadequacy, in addition to improving tactical skill by sampling near the model attracting manifold, and only near that manifold. Conclusions are presented briefly in section 5.

2. Geometry, Statistics, and Model Error

Dissipative nonlinear dynamical systems can have a variety of geometric structures that can be used to help understand the system. These structures include invariant sets, slow manifolds, inertia manifolds, and attractors. Here our interest is *attracting manifolds*, which we will define as forwardly invariant manifolds that are attracting, in the sense that there is a neighbourhood of the manifold that trajectories enter and do not leave. In this section we begin by discussing the role attracting manifolds play in understanding model error, introducing the idea of *shadow analyses*, and exploring some of the dynamical behaviours and geometrical relationships associated with model error.

The geometry of linear systems is straight forward, and consequently most features of linear systems are revealed through appropriate choice of metric and basis vectors, which define a global coordinate system. The geometry of nonlinear systems is complex with a rich variety of structures. Features of nonlinear systems can sometimes be revealed by employing nonlinear coordinate systems. These useful coordinate systems are often defined by local properties of the system. For example, a coordinate system that moves with state or a coordinate system defined by local singular vectors. In the following discussion, particularly sections 2c and 2d, we employ different nonlinear coordinates where necessary, often implicitly.

When and how this is done is discussed in an appendix to avoid unnecessarily obscuring the main points with technical details.

2a. *Lessons from Lorenz*

Much is known about the properties of the Lorenz equations (Lorenz 1963; Sparrow 1982; Guckenheimer and Holmes 1983) and new discoveries continue to be made (Tucker 2000). The *attractor* of the Lorenz equations is a complex object and what is often referred to as its "butterfly-shaped attractor" might be better thought of as an attracting manifold. It is not necessary for us to be precise here; thinking of the "butterfly wings" as being a branched two-dimensional attracting manifold is sufficient to visualize the following.

One of the important properties of the Lorenz system is sensitivity to initial conditions: two states close together on an attracting manifold will move apart over time, until eventually they will be far apart. This implies that even with a perfect model any uncertainty in the initial state leads to growing forecast errors and eventual failure of the forecast. It is often stated that forecast errors will grow exponentially, generally this is not what happens (Smith et al. 1999), it is only what happens on average. Even on the attractor states can move closer together before moving apart.

Another important property of the Lorenz system is that almost all states quickly evolve to states close to an attracting manifold, and remain close. The attracting manifold represents *physically realisable states* of the system, that is, one always expects to find the Lorenz system in a state close to an attracting manifold.

Sensitivity to initial conditions and the existence of an attracting manifold are properties of many nonlinear dissipative systems, although the latter property may not be obvious. The Lorenz equations cannot be solved exactly, but can be numerically integrated to reveal the three dimensional structure of the attracting manifold. For NWP models it is impossible to visualise an attracting manifold, not only is the state space dimension enormous (on the order of 10^6–10^7 for operational models), but also the recurrence time (time to go "once around" the attractor) is even more enormous (on the order of 10^{30} years (van der Dool 1994)). Nonetheless the existence of an attracting manifold can be deduced. Dissipative systems must have an attractor, and sensitivity to initial conditions and the existence of attracting manifolds can be surmised from the spectrum of local Lyapunov exponents and singular values. Later we will provide evidence using new geometric methods.

2b. *Data assimilation, models, and attracting manifolds*

Data assimilation uses observations of reality to obtain an *analysis*, a state of the model. What the analysis represents is open to interpretation, especially for imperfect models. Here we consider the relationship between an analysis and the a model's attracting manifold.

Data assimilation should really be thought of as an aspect of the modelling process. The process of assimilating data implements a mapping from observations of reality into model states, and by doing so provides concrete meaning to the model variables. This mapping may involve some statistics to account for observational errors. With a perfect model there is an *isomorphism* between reality and model states[1].

If a deterministic model were *perfect*, then there is a *true* state, under the isomorphism. This true state would give perfect forecasts for all time. Indeed the isomorphism and determinism imply that the property of giving perfect forecasts can be taken as a definition of what a true state of a model means. This true state

[1]Isomorphism here means that for every state of reality there is a corresponding unique state of the model. Consequently, there is mapping between observed quantities and model variables. There are many perfect models, with different isomorphisms, but each is isomorphic to the others. Observations may be inaccurate, or incomplete, or both. Inaccuracy alone is sufficient to prevent determination of the true state (Judd and Smith 2001). If the system is finite dimensional, then Takens Theorem implies that generically the isomorphism can be achieved by time-delay embedding (Takens 1981).

must be a state on, or very close to, the attracting manifold of the model. When data is assimilated into a perfect model to obtain an analysis one expects some random variation of the analysis due to observational errors. Therefore, the analysis may be thought of as a *random variable* distributed about the true model state. Data assimilation for a perfect model is a statistical process of estimating the true state or an ensemble representation of our knowledge (or uncertainty) about the true state. If the assimilation of data were ideal, then the *expected location* of the analysis is the true state[2]. Hence, for a perfect model and ideal data assimilation, the analysis is on, or close to, the attracting manifold. There is no sensible alternative; failure to put the analysis close to the attracting manifold is a failure of the data assimilation method. (It is known that some data assimilation methods (optimal interpolation, 3D variational assimilation, Kalman filter) do not obtain an analysis close to the attracting manifold even in a perfect model scenario (Judd 2003), although later we describe a *shadowing filter* which appears to.) It should also be noted that the *expected location* of an analysis traces a path over time. With a perfect deterministic model and unbiased data assimilation, this path is a trajectory of the model dynamics.

[2] Achieving this may be difficult even with a perfect model (Judd and Smith 2001). It is possible that it can only be achieved in retrospect using observations from both the past and future (Ridout and Judd 2001). Even then there are situations when the true state cannot be determined (Judd 2007a).

If a deterministic model is *imperfect*, then there is no true state for the model; with an imperfect model there cannot be an isomorphism between reality and model, and there can be no state of the model that provides perfect forecasts. Data assimilation provides an analysis, but it is not clear how to interpret what the analysis is. It is certainly not valid to interpret the analysis as an estimate of a true state of the model[3]. The analysis will still be a random variable, but there need not be any special relationship between the *expected location* of an analysis and an imperfect model's attracting manifold.

In both the perfect and imperfect model scenario, the mapping from observations to model states that data assimilation provides should be considered part of the model, and so, the term *model error* may refer to errors of the mapping or the dynamics, and the source of error is not necessarily separable.

2c. *Two types of model error, and shadow analyses*

Our goal is to obtain a model and data assimilation method that are *useful* for forecasting given available resources, where useful might mean close to being perfect by some measure. The previous section implies that a perfect model re-

[3]Although it may be useful to interpret a state obtained by data assimilation as an approximation of the "true state" of the atmosphere, the model and the atmosphere are different in terms of state space and dynamics. Consequently, no state of an imperfect model can ever represent a "true state" of the atmosphere.

quires having a data assimilation method that produces analyses that lie close to the attracting manifold of the model, and that the expected location of the analyses over time traces a trajectory of the model.

It should be clear that there are essentially two different types of error that require tuning: one tuning ensures that the analyses lie on an attracting manifold of the model, the other tuning ensures that the expected location of the analyses over time traces a trajectory of the model. Figure 1 may help to illuminate the following discussion.

[Figure 1 about here.]

If the location of an attracting manifold is known for a model, then a *projection* can be defined that maps each model state to a corresponding state on the attracting manifold. This is typically a nonlinear projection[4]. Then for any analysis not on the attracting manifold, there is a corresponding *shadow analysis*; it is the shadow of the analysis on the attracting manifold under the projection. The difference between an analysis and its shadow analysis will be referred to as *projection error*.

Given a sequence of shadow analyses that are on an attracting manifold, one can test whether these analyses are a trajectory by computing a forecast from

[4]The term *projection* is used here in the sense of a *topological retract*. A retract is a continuous mapping of the entire space into a subspace. In the present context the subspace is the attracting manifold. This mapping will typically be nonlinear. Ideally the projection also retracts trajectories; that is, trajectories in the entire state space are mapped to trajectories on the attracting manifold.

a shadow analysis and compare this with the next (verifying) shadow analysis. These one-step forecast errors will be referred to as *direction error*.

Shadow analyses, along with projection and direction errors are illustrated in figure 1(a). Projection and direction errors will depend on the projection used to obtain shadowing analyses; indeed a badly chosen projection can produce spurious direction errors. (Ideally the projection should map trajectories to trajectories on the attracting manifold, which will avoid spurious direction errors.) Furthermore, one needs to take care when computing direction errors to minimize the effects of analyses being random variables influenced by observational errors.

In practice, using several original analyses to obtain a given shadow analysis will assist at minimizing random effects. It is useful to consider two different kinds of shadow analysis. A *causal shadow analysis* uses only information up to the present time, that is, it is obtained using only original analyses (or observations) from the past (possibly distant past) and present; it uses no information from future analyses or observations. A *non-causal shadow analysis* uses information from the past, present and future, including the distant past and far future. Obviously, non-causal shadow analyses cannot be used for real-time forecasting, but they are arguably the highest quality and most appropriate verifications. Non-causal shadow analyses play an important role in investigating model error; espe-

cially in the computation of direction errors. By incorporating information from the past and future a non-causal shadow analysis will have smaller random variation than a causal shadow analysis[5].

2d. *Forecast errors*

Forecast error is usually defined as the difference between a forecast and truth, which is, strictly speaking, impossible to calculate, because the truth is unknown. At best, one must substitute some proxy for truth. Two common proxies are observations of reality, or some analysis state. Working with observations of reality (verification in observation space) and working with analyses (verification in model space), each have their own advantages and disadvantages. When a model is perfect, then these two alternatives are equivalent, because there is an isomorphism between reality and model. When the model is imperfect, then verification against observations requires introducing a mapping from model space to observation space. This mapping need not be unique, nor need it be possible to "invert" the mapping that data assimilation provides. This mapping (like the mapping that data assimilation provides) should be considered as part of the model. The map-

[5]In fact, in a perfect model of a hyperbolic system, when observation errors are sufficiently small, it can be shown that as information is gathered from further in to the past and future a non-causal shadow analysis will converge to the true state (Ridout and Judd 2001).

ping is another source of model error that needs to be accounted for.

For our present purposes, verification in model space has the convenient advantage that we do not need to contend with a third source of model error. Consequently, in the following we will use analysis states as proxies for truth when investigating forecast errors. A disadvantage of using analyses is that variances and co-variances of errors are not as readily available as they are for simple observations. Since we are more interested in dynamical features, we avoid this difficultly by working with an energy-like norm when comparing model states.

Define the *forecast error* of a given analysis to mean the difference $\|y(t) - x(t)\|$ between a forecast $x(t)$ and a suitably chosen verifying analysis $y(t)$ at given lead time t, where the forecast $x(t)$ is the trajectory starting at the analysis $x_0 = x(0)$ at $t = 0$. To understand the nature of these forecast errors we investigate the geometric relationship between analysis states, the attracting manifold, and shadow analyses. Figure 1(b) illustrates the essential geometrical relationship between a sequence of analyses, their corresponding shadow analyses, and forecast trajectories starting from the analysis and its shadowing analysis at $t = 0$. In this section we decompose the forecast errors into three different sources: sensitivity to initial conditions, entrainment with an attracting manifold, and accumulation of direction errors. These effects are illustrated in figure 2 as

we now describe.

[Figure 2 about here.]

If a sequence of analyses have fairly constant projection error[6], as illustrated in figure 1(a), then a forecast trajectory from an analysis at $t = 0$ will move towards the attracting manifold and away from the verifying analyses. In fact, we should predict that rather than seeing approximately exponential growth of errors, figure 2 curve (A), we will see errors that increase with an approximately inverted exponential decay as illustrated in figure 2 curve (B); see appendix for a more detailed discussion of this and the following arguments. This observation will be true regardless of whether the model includes direction error, the approximately inverted exponential decay only requires fairly consistent projection error, and sufficiently fast motion onto the attracting manifold.

Now consider forecasting future shadow analyses from the shadow analysis at $t = 0$, see figure 1(b). Since all these states are already on the attracting manifold, it follows that the principle source of forecast errors will be due to sensitivity to initial conditions, or accumulated direction error, or a combination of both. Sensitivity to initial conditions should result in more or less exponential increase in

[6]The projection error could result a model bias (attractor in wrong location) or from random errors in a high-dimensional space producing a chi-squared distribution with many degrees of freedom.

errors, but direction errors should accumulate to give a more or less linear increase in forecast errors, see figure 2 curves (A) and (C). Whether approximately exponential or linear increase of errors dominates depends on the relative magnitudes of the singular values of the tangent map and the magnitude of direction errors. Of course, the singular vectors can change direction and magnitude along the forecast trajectory, so other higher order (nonlinear) effects may be in evidence.

Finally, consider how the original analysis performs at forecasting the future shadow analyses. Because the forecast trajectory moves away from the future analyses toward the attracting manifold where their shadow analyses lie, one should anticipate that forecasts from the original analysis will be better at forecasting future shadow analyses than they are at forecasting future analyses. Initially we expect to see a decrease in distance between the analysis forecast and shadow analyses, but this distance should then increase as a result of sensitivity to initial conditions and accumulated direction error, see figure 2 curve (D). Of course, nonlinear effects could have an effect too.

Hence, figure 2 shows four different and distinctive error curves depending on where the forecast is started, and what is being used as the verification. The relative slopes of these graphs initially depend on the relative magnitudes of stable and unstable singular values, the magnitude of the direction errors, and the magnitude

of random variation of analyses due to observational error. Figure 2 corresponds to a situation where three conditions are met. Firstly, the largest magnitude of stable singular values are much larger than the largest magnitude of the unstable singular values, so that errors in stable directions dissipate faster than errors in unstable directions grow. Secondly, accumulation of direction errors dominates the unstable singular values. Thirdly, the magnitude of random variation of analyses due to observational error is small compared to the projection error, or just finite variance in a system with large dimension[7].

It might be noted that a similar behaviour to curves A and B in figure 2 has been noted in operational weather models by Lorenz (1982) and has been studied by many others (Dalcher and Kalnay 1987; Nicolis 2004a,b; Simmons et al. 1995; Simmons and Hollingworth 2002; Reynolds et al. 1994; Vannitsem and Toth 2002). Previous work has proposed algebraic models for the observed error growth, whereas we provide a geometric interpretation. Section 4 provides evidence in support of our interpretation for an operational weather model.

[7]If the observational errors have finite variance and the system has a large dimension, then projection error will have an approximately chi-squared distribution, with many degrees of freedom. Such a distribution is approximately Gaussian with small variance relative to the mean. This implies the projection error will be approximately constant.

3. Analytical tools and methods

In order to provide evidence that the phenomena described in the previous section are significant in an operational weather forecasting model, we have employed analytical tools and methods that are either new or not previously employed in a NWP context. These tools are used to obtain our shadow analyses and to investigate the geometric relationship of analyses, shadow analyses and forecasts to reveal the influence of an attracting manifold.

3a. *Shadowing filter*

A *shadowing filter* is a method of obtaining from an initial sequence of analyses another sequence of analyses that are closer to being a trajectory of a model. We will use the shadowing filter to obtain our shadow analyses. There is no guarantee that the shadowing analyses will be closer to the model's attracting manifold, but we will present evidence later that for the NOGAPS model this was indeed the case. The *shadowing filter* we use employs *gradient descent* of *indeterminism*.

Gradient descent of indeterminism is well established in filtering (Davies 1992, 1994; Grassberger et al. 1993; Grebogi et al. 1990; Hammel 1990). Originally introduced and demonstrated for simple chaotic systems, only recently has a good

theoretical understanding of its convergence been obtained (Ridout and Judd 2001; Judd 2007c). New theoretical and experimental results have shown that gradient descent of indeterminism could be practical for NWP (Judd et al. 2004b), in particular, experimental results have shown that in a perfect model scenario high quality shadowing pseudo-orbits could be obtained with a T21L3 quasi-geostrophic model. These results motivated the present investigation using NOGAPS, which is considerably more complex than any model previously analyzed by these methods.

Let f be a forecast model defined on a d-dimensional state space R^d, so that for $x \in \mathrm{R}^d$, $f(x)$ is the forecast for a fixed time period, which is typically 6 hours for operational data assimilation cycles of weather models. Let $\mathbf{x} = (x_0, \ldots, x_w)$, denote an arbitrary sequence of $w+1$ time ordered states $x_i \in \mathrm{R}^d$, running from the past to the present, with time separation being the forecast period of f. The quantity w is called the *window width*. The window width is an integer, but it is often more convenient to think of it in units of time, that is, multiply w by the forecast period of f, so that window width is the time period between the first and last states in the sequence.

Define the *mean squared mismatch*, or *indeterminism*, of $\mathbf{x}$ by

$$I(\mathbf{x}) = \frac{1}{w} \sum_{i=1}^{w} \|x_i - f(x_{i-1})\|^2. \tag{1}$$

Observe that $I(\mathbf{x}) = 0$ if, and only if, $\mathbf{x}$ is a trajectory. Furthermore, it can be shown that $I(\mathbf{x})$ has local minima only where $I(\mathbf{x}) = 0$ (Judd and Smith 2001; Ridout and Judd 2001).

Given an initial sequence of states $\mathbf{y} = (y_0, \ldots, y_w) \in \mathrm{R}^{(w+1)d}$ we can obtain a new sequence of states $\mathbf{x} = (x_0, \ldots, x_w)$ with smaller indeterminism by moving down the gradient of $I(x)$. For example, consider $\mathbf{x}$ to be a function of a scalar s and solve the differential equation,

$$\frac{d\mathbf{x}}{ds} = -\frac{\partial I}{\partial \mathbf{x}}, \quad \mathbf{x}(0) = \mathbf{y}, \tag{2}$$

where I is considered a function of $\mathbf{x}(s)$. That is, start at $\mathbf{y}$ and move continuously in the steepest descent direction of I. Solving the integration by a fixed-step Euler

method reduces equation (2) to the iteration

$$x_i \mapsto x_i - \frac{2\Delta}{w} \times \begin{cases} -\mathcal{A}(x_i)(x_{i+1} - f(x_i)), & i = 0, \\ (x_i - f(x_{i-1})) \\ -\mathcal{A}(x_i)(x_{i+1} - f(x_i)), & 0 < i < w, \\ (x_i - f(x_{i-1})), & i = w, \end{cases} \tag{3}$$

where Δ is the step size, and $\mathcal{A}(x_i)$ is a suitable approximation of the adjoint $df(x_i)^T$ (Judd et al. 2004b). That is, we have defined an iterative algorithm, where at each step every state x_i is moved slightly according to the mismatch of the forecast from the past $(x_i - f(x_{i-1}))$, and the mismatch of the forecast into the future $(x_{i+1} - f(x_i))$ pulled back through the adjoint $\mathcal{A}(x_i)$, see Judd et al. (2004b).

The effectiveness of a shadowing filter, or any other method for obtaining trajectories from a sequence of analyses, is limited by a number factors. There are the limitations of the algorithm itself, for example, convergence rates. Model error also plays a role. The spectrum of singular exponents of the tangent and adjoint model are important, because values close to zero will limit, and in practice halt, convergence (Ridout and Judd 2001; Judd 2007c). There are also pathological situations where the shadowing filter can give misleading indications. These pathological circumstances are atypical and in any case can be identified by in-

dependent tests. Misinterpretation can be avoided by suitable modification of the basic algorithm. All indications are that the experiments discussed here are in a typical situation far from pathology. A detailed discussion of these issues is beyond the scope of this paper. Technical investigations of the limitations and modification are being prepared (Judd 2007c,b), and some preliminary detail appears in Ridout and Judd (2001); Judd (2003); Judd et al. (2004a). Some readers may note an apparent similarity between gradient descent of indeterminism and weakly constrained 4D variational assimilation. The similarity is superficial and irrelevant to the current investigation of the nature of model error. A discussion of the differences of these methods is beyond the scope of the current investigation, and will appear elsewhere (Judd 2007b). Our purpose here is not to argue the merits of the shadowing filter, we merely report that this is method we used, and that it appears to achieve useful projections as desired.

3b. *Triangle and bi-triangle diagrams*

To assess the effectiveness of a shadowing filter we employ a number of geometrical constructs based on simple geometric figures such as triangles and tetrahedra. We first describe the construction and meaning of triangle and bi-triangle diagrams, then relate their properties to projection and direction errors.

Triangle diagrams: Suppose one has an analysis A_0, a forecast from this analysis fA_0, and the verifying analysis A_1. Ideally fA_0 should be identical to A_1, but in general the three states A_0, fA_0, and A_1 are the vertices of a triangle in state space, whose shape is completely defined by the distances between the states, as shown in figure 4 (upper left). If the goal is to obtain a sequence of analyses that are as close as possible to being a trajectory, then a measure of success is that the length of side b in figure 4 (upper left) is small relative to the length of side a for each pair of consecutive analyses. Comparing triangle diagrams of analyses and shadow analyses provides a clear visual indication of how close a sequence states is to being a trajectory.

Bi-triangle diagrams: Suppose S_0 is the shadow analysis of A_0, fS_0 is the forecast from S_0, A_1 is the verifying analysis of fA_0, and S_1 is the shadow analysis of A_1. Another useful comparison of analyses and shadow analyses is the location of fA_0 and A_1 relative to fS_0 and S_1. This relationship can be plotted as a *bi-triangle diagram* as shown in figure 6 (upper left). This diagram plots two triangles, with a common edge, obtained from computing the distances between the relevant states. This bi-triangle diagram reveals how close the analyses and shadow analyses are to being trajectories, and how close the shadow analyses are to the original analyses.

The properties of triangle and bi-triangle diagrams can be related to projection and direction errors as follows. If the shadow analyses lie on an attracting manifold of the model, then the distance between A_1 and S_1 is a measure of the projection error, and the distance between S_1 and fS_0 is a measure of direction error[8]. It should be stressed that these distances provide a measure of the projection error and direction error, but they may not be precise or free of artifacts. Circumstances can be contrived where simple application of a shadowing filter could give misleading indications; details of this will be discussed elsewhere (Judd 2007c,b). *On the other hand, the notion of "closeness" is dependent on the metric used to measure distance; the shadowing filter as described here does not guarantee that the shadow analysis is the state on the attracting manifold that is closest to the analysis or observations, although it may be modified to do so (Judd 2007b). In typical applications the shadowing filter described here has been found to be effective; in particular,* it appears to be in the case of the NWP application we discuss later. The only detail worth mentioning here is that the residual mismatch between

[8]The reason for this is not immediately obvious and requires detailed technical analysis to show. Theoretically it would appear that the shadowing filter, equation (2), should converge to a trajectory, in which case S_1 and fS_0 should be the same. More detailed theoretical analysis shows that in non-hyperbolic systems the rate of convergence becomes exceeding slow long before a trajectory is obtained (Ridout and Judd 2001; Judd 2007c). So some of the residual mismatch from applying the shadowing filter is the result of the algorithm not having fully converged for slow and nearly neutral modes. On the other hand, further theoretical analysis also shows that when model error is present the convergence rate is slowed to halt by direction errors.

S_1 and $f S_0$ typically includes genuine direction errors and a component due to the shadowing filter not having fully converged. The later component will be composed of dynamically neutral modes, that is, perturbations that grow or decay only slowly.

3c. *Attracting manifolds and travelling tetrahedra*

A key element of our discussion of shadowing analyses has been the role of attracting manifolds of the model. The behaviour of forecast errors will have many influences other than that of an attracting manifold, and we wish to test the strength of the attracting manifold's role. In high-dimensional NWP models it is difficult to visualize attracting manifolds. Here we employ a geometrical investigation of the relationship between forecasts and verifying analyses to infer the influence of an attracting manifold, without explicitly finding it.

[Figure 3 about here.]

Reconsider the analyses and shadow analyses depicted in figure 1 in a new way, as depicted in figure 3. At $t = 0$ there is the original analysis A_0 and its (non-causal) shadow analysis S_0. Then at each lead time $t = 1, 2, 3$, there are four states of interest: a forecast from the original analysis $f^t A_0$, a forecast from

the shadow analysis $f^t S_0$, a verifying analysis A_t, and the shadow of the verifying analysis S_t, as illustrated in figure 3.

The goal is to show two things: (1) that A_t is not near an attracting manifold, the forecast $f^t A_0$ moves onto an attracting manifold; (2) that S_t is close to this attracting manifold, and the forecast $f^t S_0$ moves across this attracting manifold. We have to show this without explicitly finding the attracting manifold. The key to achieving our goal is to study the motion of $f^t A_0$ relative to motions of A_t, S_t, and $f^t S_0$. At each lead time these four states form the vertices of a tetrahedron, see figure 3. The relative motions of the four states can be inferred from the changing shape of the tetrahedra. (One can think of the tetrahedra as defining a local coordinate system.) The key observation is whether or not $f^t A_0$ moves away from A_t toward S_t and $f^t S_0$. To be precise, we look for three things. First, S_t and $f^t S_0$ are on, or close to, a (fairly flat) attracting manifold. Secondly, $f^t A_0$ moves away from A_t towards a hyperplane containing S_t and $f^t S_0$. Thirdly, the line (vector) between S_t and A_t is a (approximate) normal to the hyperplane. The presence of these three features is sufficient to demonstrate the two-part goal stated in the previous paragraph.

We use the tetrahedra (figure 3) to define a local (partial) rectilinear coordinate system as follows. The origin of the local coordinate system at time t will be

S_t. The first axis will be the line joining S_t and A_t. The second axis will be perpendicular to first axis and lie in the plane containing the first axis and $f^t S_0$. This has defined a plane at time t that contains the three points A_t, S_t and $f^t S_0$, and a rectilinear coordinate system on this plane, as described. We can imagine this plane (and coordinate system) moving through state space in time. In fact, the plane is just the extension of one face of the tetrahedron at t. To follow the motion of $f^t A_0$ we project $f^t A_0$ perpendicularly onto the our chosen coordinate plane, see figure 3 right panel. This technique is used to obtain figure 8 for experiments with an operational weather model discussed in the next section.

4. NOGAPS Experiments

NOGAPS is the Navy Operational Global Atmospheric Prediction System (Baker et al. 1998), which is used operationally by the United States Navy. Prior to October 2003 the operational system used optimal interpolation data assimilation, after which NAVDAS, a 3D variational assimilation method, has been used (Daley and Barker 2001). The experiments we describe have been performed with T47L24 and T79L30 NOGAPS models. Two types of analysis were used: interpolation to model resolution of one degree analysis fields obtained from the operational

T239 NOGAPS, and analyses produced from a NAVDAS 3DVAR data assimilation cycle at the model's resolution. In all computations an *analysis* refers to a state in spectral variables in the model's units (Baker et al. 1998). We note that the general characteristics of results were very similar regardless of the model resolution used, type of analysis, or analysis period (Judd et al. 2004a); all significant differences observed are reported. Generally speaking, direct assimilation into a T79L30 model using NAVDAS provided the analyses that were most consistent with the model's dynamics, as defined below.

Unless otherwise stated all displayed results are for calculations using the T79L30 model using NAVDAS data assimilation for the 7 day window 00:00 UTC, 1 October, 2003 to 00:00 UTC, 8 October, 2003 with 6 hour intervals, that is, a sequence of 29 states. For the purposes of displaying graphs the prognostic variables are sometimes scaled by a power of ten, as indicated.

The following calculations are for a particular distance metric. The results are not critically dependent on the metric used, because the geometric properties of invariant sets and model error are not metric dependent, although certain metrics may emphasize particular features. Ideally, one should either use non-dimensional coordinates, so that all variables are of order one, or a physically and dynamically relevant metric, such as energy. In the following we use the energy norm

for vorticity, divergence and temperature, plus the difference in specific humidity suitably scaled. We will refer to errors in this metric as *energy weighted errors*. *(The vorticity, divergence and potential temperature fields contribute respectively the rotational kinetic energy, divergent kinetic energy, and potential energy components of total energy. A* distance *in the this metric corresponds to the square root of the sum of these quantities, that is, the unit of distance is the square root of energy. Because a suitably scaled component of specific humidity is added, our energy weighted errors are effectively non-dimensional distances.)*

Shadow analyses were obtained from the original analyses by applying the shadowing filter, equation (3), with a window width of 7 days, and forecast step of 6 hours, so $w = 28$. The NOGAPS model has a dry adjoint, which we used to approximate the full adjoint. We chose $2\Delta/w = 0.1$ and iterated equation (3) for 30 to 100 steps. Most of the results shown in the following are for 30 iterations.

Causal shadow analyses are obtained from the last state (x_{28}) of the window when the gradient descent algorithm is stopped. Non-causal shadow analyses were obtained from the middle state (x_{14}) of the window.

4a. *Triangle and bi-triangle diagrams*

Triangle diagrams: Figure 4 shows triangle diagrams for a 7 day window of 6 hour forecasts. One triangle is plotted for each consecutive pair of analyses, giving 28 triangles. These triangles show that there is considerable forecast mismatch (b is not small relative to a) and that the mismatch is of a consistent magnitude. Figure 5 shows triangle diagrams for the non-causal shadow analyses. Comparing figures 4 and 5 it is seen that the shadow analyses are much closer to being a trajectory, having considerably smaller forecast mismatch.

[Figure 4 about here.]

[Figure 5 about here.]

Bi-triangle diagrams: Figure 6 shows bi-triangle diagrams of original analyses and non-causal shadow analyses. The shadow analyses may seem surprisingly far from the original analyses. The shadowing filter (3) does not constrain the shadow analysis to remain close to the observations, so the shadow analyses can wander away from the original analyses. Many may see this as a flaw of the shadowing filter, but it is not a flaw, it is a strength, because it allows the shadowing filter to reveal how far an analysis is from the attracting manifold. The success of the shadowing filter comes from the shadow analyses being close to an attracting

manifold, and hence the large distance between the original analyses and shadow analyses is an indication of the magnitude of the projection error.

[Figure 6 about here.]

4b. *Is this just a matter of balance?*

An important question is whether the differences between analyses and shadow analyses are just a matter of balance. Could it be that movement onto an attracting manifold merely represents geostrophic adjustment? To investigate this possibility, we examine the impact of nonlinear normal mode initialization as described by Errico et al. (1988). This procedure is designed to remove spurious gravity waves that may be present in model states; these gravity waves may have been introduced by data assimilation or interpolation from higher resolution states. As a further test we also examine surface pressure tendencies.

We assess the impact of nonlinear normal mode initialization by comparing the magnitude of the difference between uninitialized and initialized analyses, for both the original analyses and shadow analyses. The differences are summarized in table 1. Comparing the first column with columns two and three, it is seen that shadow analyses are more balanced than the original analyses; the effect of

nonlinear normal mode initialization of shadow analyses is less than half the effect on the original analyses. On the other hand, comparing the first three columns with the last two columns, it is seen that the effect of the shadowing filter is significantly larger than the effect of nonlinear normal mode initialization. We conclude that although shadowing analyses are more balanced, balancing can only account for a small fraction of the difference between analyses and shadow analyses.

To investigate the issue further we examined the global RMS of surface pressure tendencies over 4 day forecasts. Experience shows that global RMS values of around 0.5 hPa per time step are reasonable, whereas values in excess of 1.0 hPa per time step indicate significant spurious gravity wave activity. Neither original analyses, nor shadow analyses showed tendencies in excess of 0.5 hPa per time step.

[Table 1 about here.]

4c. *Forecast errors*

Figure 7 shows various forecast errors as discussed in section 2d and should be compared with figure 2 where the curves are labelled the same. Figure 7 shows error curves for the vorticity field. The errors for individual model layers, and for

other prognostic fields (divergence, temperature, specific humidity, surface pressure) are very similar. Curve A shows the approximately exponential increase of distance between two trajectories close to an attracting manifold, in this case trajectories from causal and non-causal shadow analysis, which are states likely to differ only in unstable directions. (Evidence that these states are close to an attracting manifold comes in the next section.) Curve B shows the original analysis forecasting future analyses, and shows an approximately inverted exponential decay as expected when the original analyser have significant projection error. Curve C shows the error of the non-causal shadow analysis forecasting future non-causal shadow analyses, which shows an initially fairly linear error growth that is significantly less than curve B; consistent with accumulated direction errors. Curve D shows the error of the original analysis forecasting future non-causal shadow analyses, which has a decrease then increase of error consistent with entrainment with an attracting manifold combined with effects of sensitivity to initial conditions or accumulated direction errors. The initial decrease of curve D reveals that the non-causal shadow analyses are not arbitrary states, they are states toward which forecasts from the analyses tend to move.

[Figure 7 about here.]

Curve E in figure 7 shows the error of the causal shadow analysis forecasting future non-causal shadow analyses, which is seen to be close to forecast errors of the non-causal shadow analysis shown in curve C. Note that curve A shows the divergence of these two forecasts. The fact that the distance between these forecast trajectories is far greater than the difference between curve C and E, implies that the accumulation of direction errors is more significant than sensitivity to initial conditions in this example.

4d. *Travelling tetrahedra and attracting manifolds*

Finally we demonstrate that A_t is not near an attracting manifold and the forecast $f^t A_0$ moves onto an attracting manifold, whereas S_t is close to this attracting manifold and the forecast $f^t S_0$ moves over this attracting manifold. In section 3c it was described how tetrahedra formed from forecast states and verifying analyses (A_t, fA_t, S_t and fS_t) provide local rectilinear coordinate system in which to track the relative motions of fA_t and fS_t.

[Figure 8 about here.]

Figure 8 shows the motion of S_t, A_t, $f^t S_0$, and $f^t A_0$ in the moving coordinate system. By construction, in our moving coordinate system, S_t is always fixed at

the origin, A_t moves only along the y-axis, but $f^t S_0$ and $f^t A_0$ could potentially move anywhere in the coordinate plane. In order to emphasize the relative motions of $f^t S_0$ and $f^t A_0$ we have connected points at consecutive times.

There are a number of things to observe in figure 8. First observe $f^t S_0$ moves more or less parallel to the horizontal axis. This implies that if S_t and $f^t S_0$ lie in the attracting manifold, then this manifold is (locally) fairly flat and the vertical axis (defined by S_t and A_t) is more or less perpendicular to the attracting manifold. We might also note that $f^t S_0$ moves away from S_t at a relatively constant rate, consistent with accumulation of direction errors. Secondly, observe that although A_t moves along the vertical axis, the motion its fairly restricted, so that the A_t remains at a fairly constant distance from the attracting manifold[9].

The most important observation to make about figure 8 is that when $f^t A_0$ is projected on the the coordinate plane it traces a path that moves away from A_t and down toward the path traced by $f^t S_0$, indeed once $f^t A_0$ gets close to the path traced by $f^t S_0$ it moves in similar ways. This is strong evidence that $f^t A_0$ moves

[9]It should be noted that restricted motion of A_t and the vertical axis being perpendicular to the attracting manifold is consistent with the differences $A_t - S_t$ (projection errors) having random mean zero errors that are largely independent for each component. This is consistent because these are the usual properties of such random vectors in high-dimensional spaces. To be more specific such vectors are always nearly perpendicular to low dimensional subspaces. The length of such vectors have a chi-squared distribution with many degrees of freedom, which are asymptotically Gaussian with mean $n\sigma^2$, where n is the dimension of the space and σ^2 the variance of each component. It is also certainly the case that the projection errors will have some random errors of this type, arising from observational errors, but we will see also that projection errors also have a large systematic component.

toward an attracting manifold that contains S_t and $f^t S_0$. We conclude that the traditional analysis A_0 initializes the model far from the attracting manifold, and the first part of the forecast $f^t A_0$ is dominated by motion toward the attracting manifold.

Readers may note that the characteristic shape of curve B has been observed previously in the context of forecast errors and has been attributed to what is termed nonlinear saturation of errors. It should be clear from the discussion of section 2 that movement onto an attractor is a distinct process from nonlinear saturation of errors; this can certainly be seen using the Lorenz equations as a example. In high-dimensional systems different processes can act simultaneously on different scales. Movement onto an attractor could be accompanied by nonlinear saturation of errors. In our NOGAPS experiments nonlinear saturation of errors is almost certainly occurring at smaller scales, but we have not tried to confirm its presence or investigated the magnitude of its effects. Since both movement onto an attractor and nonlinear saturation of errors have the same characteristic error growth, the curve B in figure 7 is not sufficient to identify either process, or determine which has the dominant effect. On the other hand, the authors are unable to see how nonlinear saturation of errors could account for the other curves in figure 7; in particular the non-monotonic curve D. Furthermore, the fact that

movement onto an attractor is seen so clearly in figure 8 leads us to conclude that movement onto an attractor is the dominant process leading to the graphs seen in figure 7. It would be an interesting experiment to attempt to determine the relative effect of nonlinear saturation of errors.

4e. *Projection and direction errors*

In the authors' opinion the results discussed thus far have provided strong evidence for the existence of an attracting manifold in the NOGAPS model, and that this attracting manifold influences forecast errors as described in section 2. Thus we interpret the difference between an analysis and its shadow analysis as *projection error*, and the mismatch of a one-step shadow analysis forecast and the following shadow analysis as *direction error*. We investigate how resolving errors in this way has potential utility in diagnosing model error. As an illustration, we compare projection and direction errors of NOGAPS models of different spatial resolution that use different data assimilation schemes. Specifically, a T47L24 model using optimal interpolation to assimilate data, and a T79L30 model using NAVDAS 3D variational assimilation.

[Figure 9 about here.]

[Figure 10 about here.]

Figures 9 and 10 show zonal averages of projection and direction errors for the specific humidity field of T47L24 and T79L30 models. (Other fields are discussed in an appendix.) In these plots the errors have been averaged over a 7 day window. For T47L24 the first week of March 2003, and for T79L30 the first week of October 2003.

The projection errors (figure 9) of the two situations are quite different. Note in particular the different sign of the projection errors near the surface. Some of the difference in variance might be attributed to seasonal differences, but most of the significant differences can be attributed to the data assimilation methods projecting observations into model space differently.

The direction errors (figure 10) of the two situations are much smaller than the corresponding projection errors, and much more similar. Note for example that the distribution of sign is now very similar, although the direction errors in T79L30 NAVDAS are more negative near the surface. Some of what we call direction error may be residual mismatch from movement onto the attracting manifold, resulting from incomplete convergence of the shadow filter, because the mismatches were still decreasing when the algorithm was terminated. Also, some of the residue may be due to slow convergence of the shadowing filter for nearly neutral modes.

This is quite possibly the case for the largest residual mismatches in the vorticity field around the 200mb height in the extra-tropics (see appendix). Such an interpretation is not obviously related to the residual mismatches of tropical specific humidity seen in figure 10.

Whatever the interpretation of the direction errors (residual mismatches) shown in figure 10, it is clear that the sign and magnitude of direction errors are very similar at all but the lowest levels. That the direction errors are similar in the two situations is perhaps surprising, and certainly interesting. Firstly, it implies that the shadowing filter's determination of direction error is fairly immune to the much larger projection errors the different data assimilation methods introduce. Secondly, it implies the direction errors are really a property of the model, which appears to be present at both model resolutions and in different seasons. This observation may assist model development.

5. Conclusions

We have presented evidence that operational weather models evolve onto attracting manifolds of lower dimension than the entire state space, and argued the implications this holds for data assimilation and the diagnosis of model error. The

shadowing filter has been introduced and shown to locate states (shadow analyses) that are more consistent with the model dynamics than traditional analyses defined by optimal interpolation and three-dimensional variational assimilation. Also non-causal shadow analyses provide a new option for verifications that aim for simultaneous consistency with the dynamics of the model and the observations both past and future. It is not clear that avoiding the initial collapse onto the attracting manifold will offer tactical forecast improvement. In ensemble prediction systems, however, advantages could be massively increased, as ensembles on the attracting manifold would sample a lower dimensional space than those distributed in the full state space. Inasmuch as shadowing analyses exploit a long window of observations (long relative to a three or four dimensional variational assimilation window), they provide information on model mis-behaviours which cannot be gleaned either from one step tendencies or from free-running model integrations.

Using novel geometric methods to investigate and visualise the dynamics of several trajectories in high-dimensional spaces, we have verified dynamics in an operational weather model which resemble what would be expected in systems with attracting manifolds. Trajectories starting from states not on the manifold are seen to approach trajectories started from shadow analyses located on or closer to

the attracting manifold. This is illustrated in figures 1 and 2, while the changing shape of triangles and tetrahedra in figures 4, 5 and 6, are consistent with this kind of dynamics being realised in NOGAPS. Further evidence for this argument is obtained by examining "error" growth: the divergence of trajectories started from various initial states shown in figure 7 is consistent with the expectations of our geometric interpretation.

Acknowledging the difference between having an initial condition not on the attracting manifold (projection error) and the systematic inability of the model dynamics to shadow the observations (direction error) allows new insight of value to model improvement. The projection errors shown in figure 9 reveal systematic bias in the combination of model and data assimilation. On the other hand, the similar direction errors in figure 10, despite different model resolution and seasons, could indicate aspects of the model physics that require attention. When such shortcomings are known, theoretical and numerical resources can be deployed to reduce them.

These results are now being extended to other modelling scenarios, with the aim of comparing the skill of ensembles of initial conditions on, or near, the attracting manifold with traditional methods of ensemble formation. Further work requires a careful reconsideration of the preferred method for evaluating forecasts:

it would seem non-causal shadow analyses provide the most relevant targets for an assimilation scheme — yet these will differ in each model. The presence of projection errors and model imperfections appear to pose a fundamental limitation to verification using model states as targets. Verification against observations introduces additional complications and potential errors in translating probabilistic model (or ensemble) output back into observational space. Arguably, nonlinearities imply that the entire prediction systems, from assimilation of observations, to probabilistic prediction of future observations, can only be meaningfully evaluated as a whole. Respecting the geometric constraints due to the model dynamics may move us closer to more internally consistent and operationally valuable systems.

Acknowledgements

CAR and TER gratefully acknowledge the support of the Office of Naval Research (ONR) through Program Element 0602435N and 0601153N. KJ acknowledges the essential support of ONRIFO with a NICOP award N000140510668, ARC grant DP0662841, the US Naval Research Laboratory, Monterey, California, and the London School of Economics. The authors would like to thank the anonymous

reviewers for their comments and suggestions.

Appendix 1: Technical details

This appendix provides some of the mathematical details under-pinning the arguments of Section 2. For brevity we assume some familiarity with standard mathematical techniques found in any good text on modern dynamical systems theory (Guckenheimer and Holmes 1983; Katok and Hasselblatt 1995).

For the purposes of NWP the atmosphere is generally modelled as a partial differential equation describing multi-phase flow with non-local coupling; for example, the Navier-Stokes equations and physical state equations, with radiative coupling. The state space of these partial differential equations is an infinite dimensional Banach space. To make these equations manageable, NWP applies spatial discretization or basis truncation. This reduces the model to a finite dimensional ordinary differential equation:

$$\dot{x} = f(x), \quad x_t \in \mathcal{B} \subseteq \mathrm{R}^d, \tag{A.1}$$

where $\mathcal{B}$ is a subspace of Euclidean space. Discretization of time reduces this

model to a finite dimensional nonlinear difference equation:

$$x_{t+1} = f(x_t), \quad x_t \in \mathcal{B} \subseteq \mathrm{R}^d, \tag{A.2}$$

where the time variable t is typically chosen to count units of the integration time step (typically around 15 minutes), or the time interval between data assimilation cycles (typically 6 hours). The two models (A.1 and (A.2) are not equivalent (Teixeira et al. 2007). In the following we generally refer to the difference equation model (A.2). In figure 2 we plot filled circles to represent analyses and open circles to represent forecasts of the model; the forecasts come from a difference equation (A.2). On these plots there is also a background of arrow-headed lines that represent solutions of the ordinary differential equation model (A.1).

Since the work of Poincaré it has been known that many nonlinear dynamical equations are analytically unsolvable. To deal with this Poincaré introduced qualitative analysis, which applies principles and techniques of topology and geometry to provide qualitative and semi-quantitative descriptions of a system's behaviour, rather than a full quantitative solution. The arguments of Section 2 are of this type. The principle tools used to reveal features of interest are local linearization about a trajectory, and various nonlinear changes of coordinates to transform

chosen trajectories into straight lines.

Partially hyperbolicity: To allow a discussion of invariant structures and properties of trajectories (like movement onto attracting manifolds) some restrictions on the properties of f are required. A restriction of convenience is to assume f is a diffeomophism (differentiable invertible map whose inverse is differentiable). For such maps the Jacobian derivative $df(x)$ is defined and continuous for all $x \in \mathcal{B}$. Given the Jacobian derivative one can consider linearization about a trajectory, as in Floucet theory and differential geometry. These techniques can be applied as in Ridout and Judd (2001), but here the sightly more general formulation of Judd (2007c) is used.

The atmosphere appears to display sensitivity to initial conditions, so we can restrict attention to models with this property. We will assume the map f is *partially hyperbolic* as defined below. Partial hyperbolicity allows discussion of concepts like stable and unstable growth of perturbations, and local attracting manifolds. The definition is very broad and applies to a wide class of models. NWP models are likely to fall into this class, with at most minor modification.

We will first give the mathematical definition of partial hyperbolicity, then describe in general terms what the conditions of the definition mean. It is not necessary to master the definition to understand what follows. A diffeomorphism f

on $\mathcal{B}$ is *partially hyperbolic* if there exists an interval $(\lambda_0, \lambda_1) \subset (0, 1]$, such that for all $\lambda \in (\lambda_0, \lambda_1)$:

1. For each $x \in \mathcal{B}$ there is a splitting $T_x\mathcal{B} = E_\lambda^{(-1)}(x) \oplus E_\lambda^{(0)}(x) \oplus E_\lambda^{(+1)}(x)$;

2. The splitting is continuous with $x \in \mathcal{B}$;

3. The splitting is invariant, $df(x)(E_\lambda^\kappa(x)) = E_\lambda^\kappa(f(x))$ for $\kappa \in \{-1, 0, +1\}$;

4. $\forall v \in E_\lambda^\kappa(x), \quad \|df(x)^\kappa v\| \leq \lambda \|v\|$ for $\kappa \in \{-1, +1\}$;

5. $E_\lambda^\kappa(x) \neq 0$ for $\kappa \in \{-1, +1\}$.

The definition of partial hyperbolicity can be understood as follows. Let $x \in \mathcal{B}$ be a state of the model f, and let $x + v$ be a perturbation of this state. Property 1 says the perturbation v can be decomposed into a sum three components; one in each of the three subspaces labelled $\kappa = -1,\ 0,\ +1$. Property 2 says that this decomposition of components varies continuously as x is varied. Properties 3 and 4 say that the three components of v correspond to growing, decaying, and neutral modes. That is, if one investigates the "forecast error" at lead time t $\|f^t(x+v) - f^t(x)\|$, then for sufficiently small $\|v\|$ and t, (relative to the attractor diameter and recurrence time),

- if $v \in E_\lambda^{(-1)}(x)$, then $\|f^t(x + v) - f^t(x)\| \leq \lambda^t \|v\|$,

- if $v \in E_\lambda^{(+1)}(x)$, then $\|f^t(x+v) - f^t(x)\| \geq \lambda^{-t}\|v\|$,
- if $v \in E_\lambda^{(0)}(x)$, then $\lambda^t\|v\| \leq \|f^t(x+v) - f^t(x)\| \leq \lambda^{-t}\|v\|$.

Property 5 says that for any state there are perturbations that grow and others that decay. The *neutral modes* in $E_\lambda^{(0)}(x)$ tend to grow or decay only slowly, or fluctuate about a fairly constant value.

Curve A: The beauty of partial hyperbolicity is that complex nonlinear systems are seen to have local properties similar to linear systems, at least for perturbations that are not too large. In particular, error growth and decay can be *bounded* by an exponential growth or decay, even though the actual behaviour of errors is nonlinear. This justifies the discussion of section 2 that equates sensitivity to initial conditions with an exponential growth of errors (curve A). The main text indicated that in nonlinear systems the growth is not strictly exponential, but partial hyperbolicity implies it can be initially bounded below by an exponential growth.

Curve B: An attracting manifold M can be thought of as a set of states such that if a state $x \in M$ is perturbed off of M, then the trajectory of the perturbed state moves back toward M. If $T_xM \subseteq T_x\mathcal{B}$ represents the tangent space to M at x, then in a partially hyperbolic system a necessary condition for being an attracting manifold is that for some λ, $E_\lambda^{(+1)}(x) \subseteq T_xM$ for all $x \in M$, and a sufficient

condition is $E_\lambda^{(+1)}(x) \oplus E_\lambda^{(0)}(x) \subseteq T_x M$. The necessary condition says that no perturbation off of M grows faster than $1/\lambda^t$, and the sufficient condition says no perturbation off of M grows at all; in fact, they decay at least as fast as λ^t. If the sufficient condition were to apply for an attracting manifold M, then this justifies the discussion of section 2 that equates movement toward an attracting manifold with an exponential decay, which leads to an inverted exponential decay when considering the forecast errors (curve B). The sufficient condition may be too strong for some nonlinear systems, because there may be slowly growing modes. Even if just the necessary condition were to apply to an attracting manifold M, then a perturbation off of M would have components that decay and others that grow only slowly.

Curve C: According to the fundamental theorem of flows (Guckenheimer and Holmes 1983) in an open simply-connected region of state space containing no fixed points, there is differentiable change of coordinates so that the vector field $f(x)$ in (A.1) is mapped to a constant vector field. (As depicted by the background arrowed lines in the lower left panel of figure 2.) From the fundamental theorem of flows it follows that on local regions of state space without fixed points there always exists a nonlinear projection such that the model is locally perfect. Of course, this "perfect" projection may be unnatural and impossible to determine.

What is more important, is such "perfect" projections are not generic; that is, arbitrarily small changes to the model or projection destroy the perfection. From the transversality theorem (Guckenheimer and Holmes 1983) it follows that generic projections result in transverse intersections of model trajectories and the projection of true trajectories; that is, direction errors. Since direction errors are zeroth order, their error growth is initially linear in time, until variation in the direction error and other nonlinearities take effect.

Appendix 2 : NOGAPS experiments

This appendix provides some additional plots of projection and direction errors in order to better appreciate the nature of these errors. The plots are for the T79L30 model using NAVDAS 3D variational assimilation for a 7 day window being the first week of October 2003.

Zonal averages of the projection and direction errors, averaged over the 7 day window, for the vorticity, divergence and temperature fields are shown in figures 11, 12 and 13, respectively. The corresponding specific humidity plots were shown in figure 9(b) and figure 10(b). Some features of note are the following. The vorticity projection and direction errors, shown in figure 11, appear to

be closely correlated. There are reasons to believe that for the vorticity field the residual mismatch after applying the shadowing filter may not be largely direction error, but rather, may reflect more the effects of the shadowing filter having not fully converged for some of this field. It is known that convergence of the shadowing filter is slowest for nearly neutral modes (Ridout and Judd 2001; Judd 2007c). The residual mismatch of the vorticity is largest in the jet stream, which is likely to have strong wave motions that may be nearly neutral. Further investigation is needed to establish whether the residual mismatch of the vorticity field are associated with nearly neutral modes. On the other hand, the projection and direction error of the divergence field seen in figure 12 also show a close correlation, but the largest residual mismatch (in the tropics around 200mb) is not associated with any obvious neutral modes. Furthermore, the residual mismatch of the divergence field correlates with that of the temperature field in this region, see figure 13(b). The residual mismatch of temperature in this region is clearly not the result of incomplete convergence, because it is not associated with a large projection error in this region, see figure 13(a). We conclude that the residual mismatch of the divergence and temperature fields are, like the specific humidity field, an indication of direction error, but those of part of the vorticity field may not be.

[Figure 11 about here.]

[Figure 12 about here.]

[Figure 13 about here.]

Some understanding of the spatial distribution of projection errors can be obtained from averages over a 7 day window for a fixed model level. Here we consider model level 24, which corresponds to the nominal 850mb level. Figure 14 shows the average for the vorticity and divergence fields, while figure 15 shows the average for the temperature and specific humidity fields.

Figure 14 reveals that the projection error in the vorticity and divergence fields do not, on a 7 day average, show any large scale features. The plots do show a lot of small scale features, and these features are closely associated for the two fields. The projection error appears to be of two forms. One form are large localized errors that appears to result from topographic influence, or intense weather systems near the tropics. The second form are wide-spread small-amplitude errors, which overall make a significant contribution to the total projection error. Detailed study of the evolution of the projection error fields show that the small amplitude projection errors (especially over oceans) is mainly the result of the shadow analyses having much greater spatial-temporal consistency than the original analyses. Figure 15 shows the projection error of the temperature and specific humidity fields, which, unlike vorticity and divergence, show large scale features on a 7 day aver-

age. Some of the stronger small scale features are associated with projection error features of the vorticity and divergence fields. Detailed study of the evolution of the temperature field shows that the some of the average projection error results from the shadow analyses having a larger diurnal range.

[Figure 14 about here.]

[Figure 15 about here.]

References

Baker, N. E., R. Daley, R. Gelaro, J. Goerss, T. Hogan, R. Langland, R. Pauley, M. Rennick, C. Reynolds, G. Rohaly, T. Rosmond, and S. Swadley, 1998: The navy operational global atmospheric prediction system: A brief history of past, present, and future developments. Technical report, Naval Research Laboratory, Monterey, California.

Dalcher, A. and E. Kalnay, 1987: Error growth and predictability in operational ECMWF forecasts. *Tellus*, **39**, 474–491.

Daley, R. and E. Barker, 2001: NAVDAS: Formulation and diagnostics. *Monthly Weather Review*, **129**, 869–883.

Davies, M., 1992: Noise reduction by gradient descent. *International Journal of Bifurcation and Chaos*, **3**, 113–118.

— 1994: Noise reduction schemes for chaotic time series. *Physica D*, **79**, 174–192.

Errico, R., E. Barker, and R. Gelaro, 1988: A determination of balanced normal modes for two models. *Monthly Weather Review*, **116**, 2717–2724.

Grassberger, P., R. Hegger, H. Kantz, C. Schaffrath, and T. Schreiber, 1993: On noise reduction methods for chaotic data. *Chaos*, **3**, 127.

Grebogi, C., S. M. Hammel, J. A. Yorke, and T. Sauer, 1990: Shadowing of physical trajectories in chaotic dynamics: Containment and refinement. *Physical Review Letters*, **65**, 1527–1530.

Guckenheimer, J. and P. Holmes, 1983: *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, volume 42. Springer-Verlag, New York.

Hammel, S. M., 1990: A noise-reduction method for chaotic systems. *Physics Letters A*, **148**, 421–428.

Judd, K., 2003: Nonlinear state estimation, indistinguishable states and the extended Kalman filter. *Physica D*, **183**, 273–281.

— 2007a: Failure of maximum likelihood methods for chaotic dynamical systems. *Physical Review E*, **75**, 036210.

— 2007b: Forecasting with imperfect models, dynamically constrained inverse problems, and geometrically modified gradient descent. *Physica D*, **to appear**.

— 2007c: Shadowing pseudo-orbits and gradient descent noise reduction. *Journal of Nonlinear Science*, ***to appear***.

Judd, K., C. Reynolds, and T. Rosmond, 2004a: Toward shadowing in operational weather prediction. Technical Report NRL/MR/7530-04-18, Naval Research Laboratory, Monterey, CA, 93943-5502 USA.

Judd, K. and L. Smith, 2001: Indistinguishable states I : perfect model scenario. *Physica D*, **151**, 125–141.

Judd, K., L. Smith, and A. Weisheimer, 2004b: Gradient free descent: shadowing and state estimation with limited derivative information. *Physica D*, **190**, 153–166.

Katok, A. and B. Hasselblatt, 1995: *Introduction to the modern theory of dynamical systems*, volume 54 of *Encylopedia of Mathemtics and its Applications*. Cambridge Univeristy Press.

Lorenz, E., 1982: Atmospheric predictability experiments with a large numerical model. *Tellus*, **34**, 505–513.

Lorenz, E. N., 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.*, **20**, 130–141.

Nicolis, C., 2004a: Dynamics of model error: some generic features. *Journal Atmospheric Science*, **60**, 2208–2218.

— 2004b: Dynamics of model error: the role of the unresolved scales revisited. *Journal Atmospheric Science*, **61**, 1740–1753.

Reynolds, C. A., P. Webster, and E. Kalnay, 1994: Random error growth in NMC's global forecasts. *Monthly Weather Review*, **122**, 121–150.

Ridout, D. and K. Judd, 2001: Convergence properties of gradient descent noise reduction. *Physica D*, **165**, 27–48.

Simmons, A. and A. Hollingworth, 2002: Some aspects of the improvement in skill of numerical weather prediction. *Quarterly Journal of the Royal Meteorological Society*, **128**, 647–677.

Simmons, A., R. Mureau, and T. Retroliages, 1995: Error growth and estimates of predictions from the ECMWF forecating system. *Quarterly Journal of the Royal Meteorological Society*, **121**, 1739–1771.

Smith, L., C. Ziehmann-Schlumbohm, and K. Fraedrich, 1999: Uncertainty dynamics and predictability in chaotic systems. *Q. J. R. Meteorol. Soc.*, **125**, 2855–2886.

Sparrow, C. T., 1982: *The Lorenz Equations: Bifurcations, Chaos and Strange Attractors*, volume 41. Springer, New York.

Takens, F., 1981: Detecting strange attractors in turbulence. *Dynamical Systems and Turbulence*, D. A. Rand and L. S. Young, eds., Springer, Berlin, volume 898, 365–381.

Teixeira, J., C. Reynolds, and K. Judd, 2007: Time-step sensitivity of nonlinear atmospheric models: numerical convergence, truncation error growth and ensemble design. *Journal of Atmospheric Science*, **64**, 175–189.

Tucker, W., 2000: The lorenz attractor exists. *Nature*, **406**, 948–949.

van der Dool, H. M., 1994: Searching for analogues, how long must we wait? *Tellus*, **46**, 314–324.

Vannitsem, S. and Z. Toth, 2002: Short-term dynamics of model errors. *Journal Atmospheric Science*, **59**, 2594–2604.

List of Figures

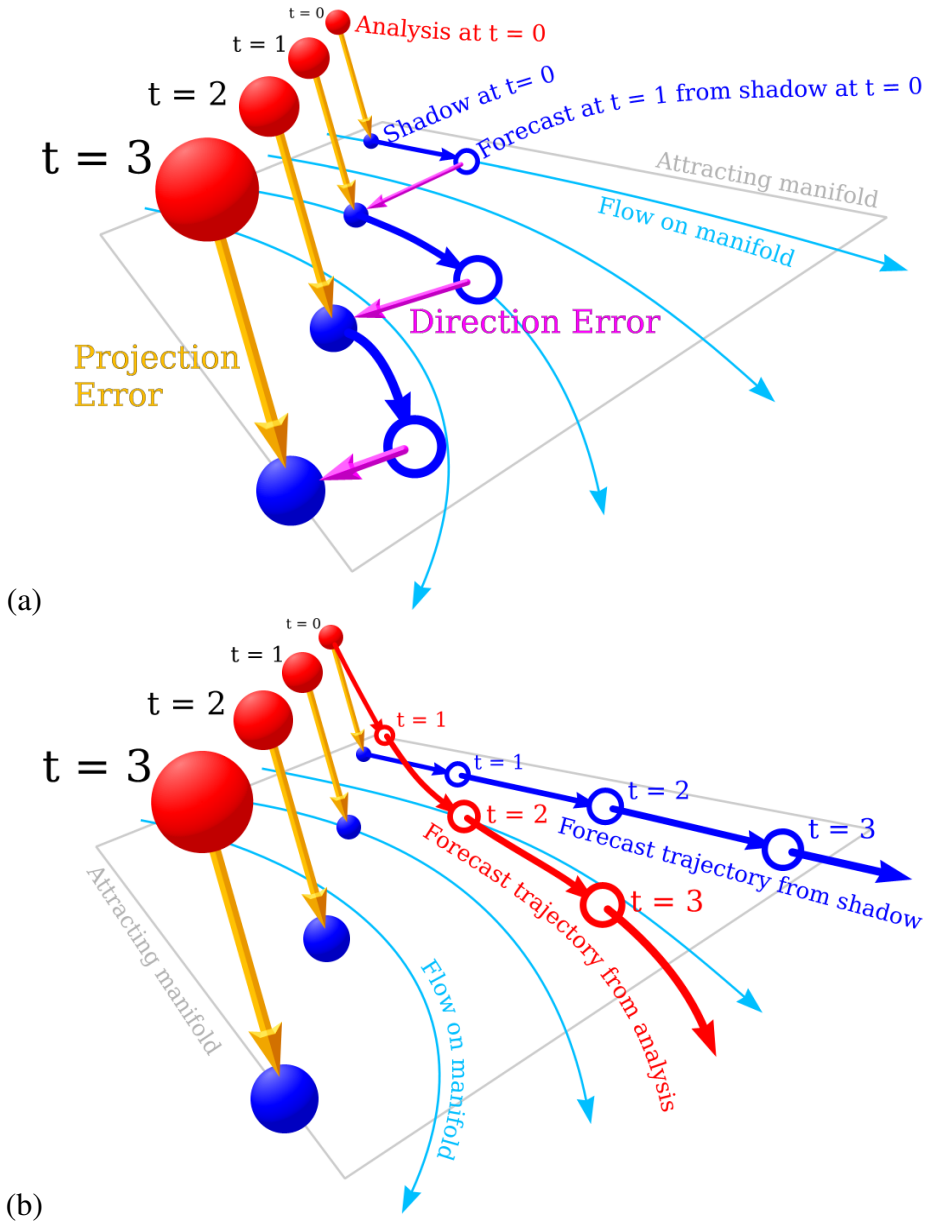

Figure 1: *(a) Schematic representation of an attracting manifold, a sequence of four analyses, their shadow analyses, forecasts of three shadow analyses, the projection errors and direction errors. (b) Analyses, shadow analyses, and forecast trajectories from analysis and shadow analysis at $t = 0$, showing how a forecast trajectory of analysis moves down onto the attracting manifold.*

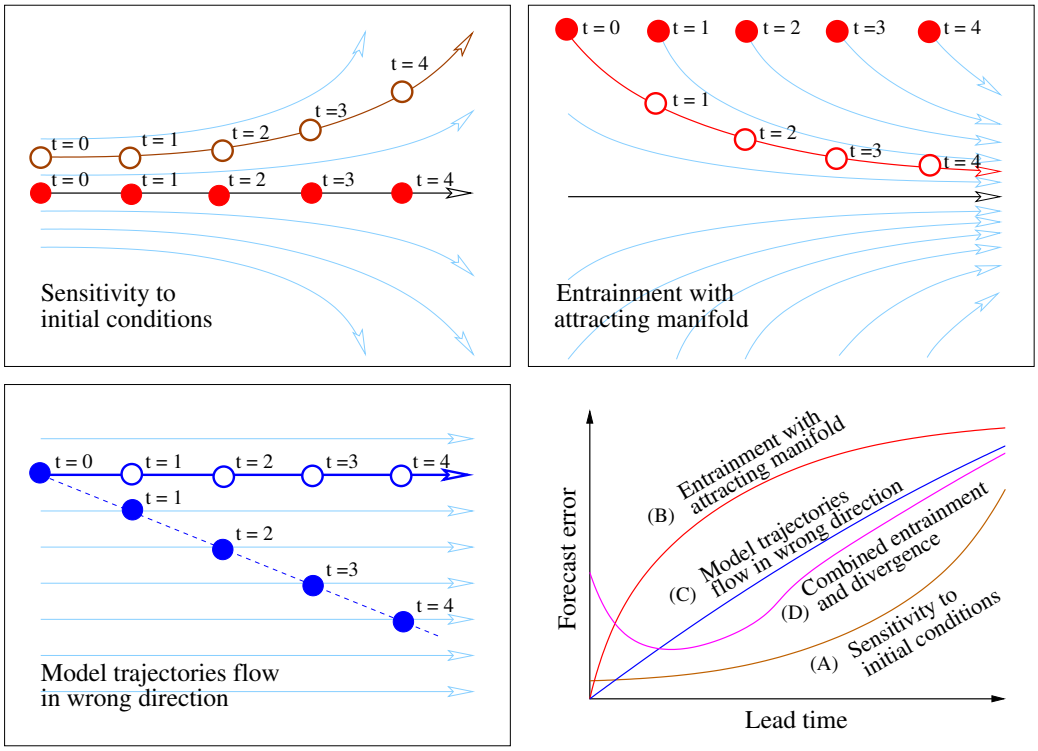

Figure 2: *Forecasting errors. Boxed panels represent states and trajectories in model space with analyses as filled circles and forecast states as open circles. The graph on the lower right represents error between forecast and verification (distance between filled and open circle at lead time t) under different circumstances. (A) Sensitivity to initial conditions, (B) Entrainment with an attracting manifold. (C) Model trajectory flow in the wrong direction, (D) Analysis forecasting a verifying shadow analysis showing the combined effects of entrainment with an attracting manifold and divergence due to sensitivity to initial conditions, or model flow in the wrong direction, or both.*

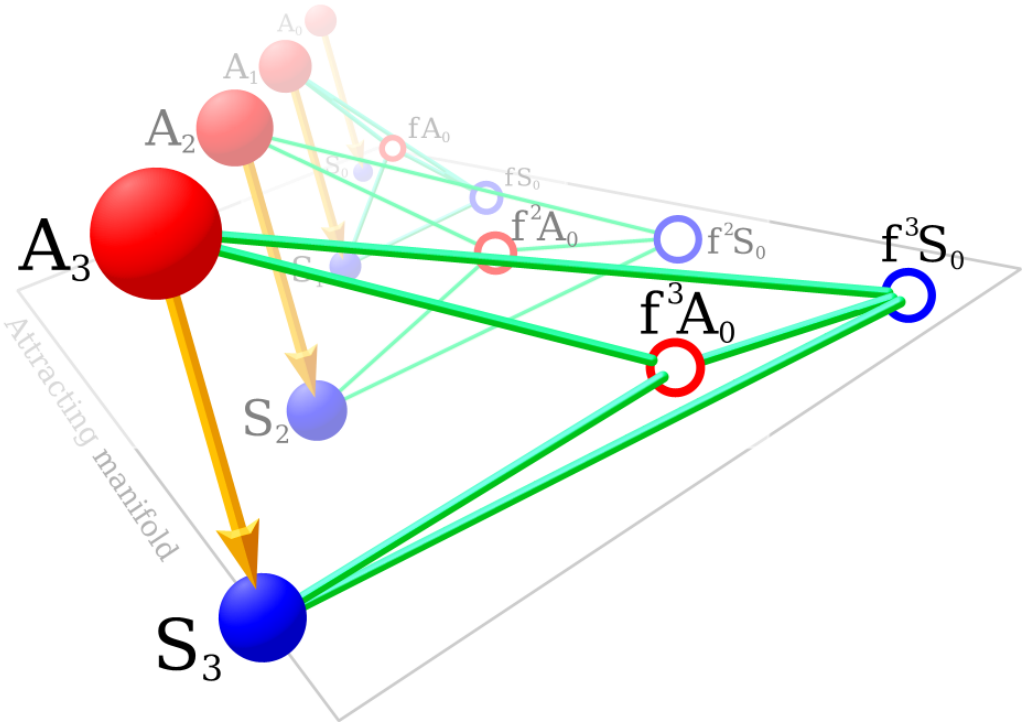

Figure 3: *Travelling tetrahedra: Initially at $t = 0$ there is an analysis A_0 and its shadow analysis S_0. At any time $t > 0$ there are four points A_t, $f^t A_0$, S_t, and $f^t S_0$. These points are the vertices of a tetrahedron, which can be used to define a local coordinate system in which to view the relative motions of $f^t A_0$ and $f^t S_0$.*

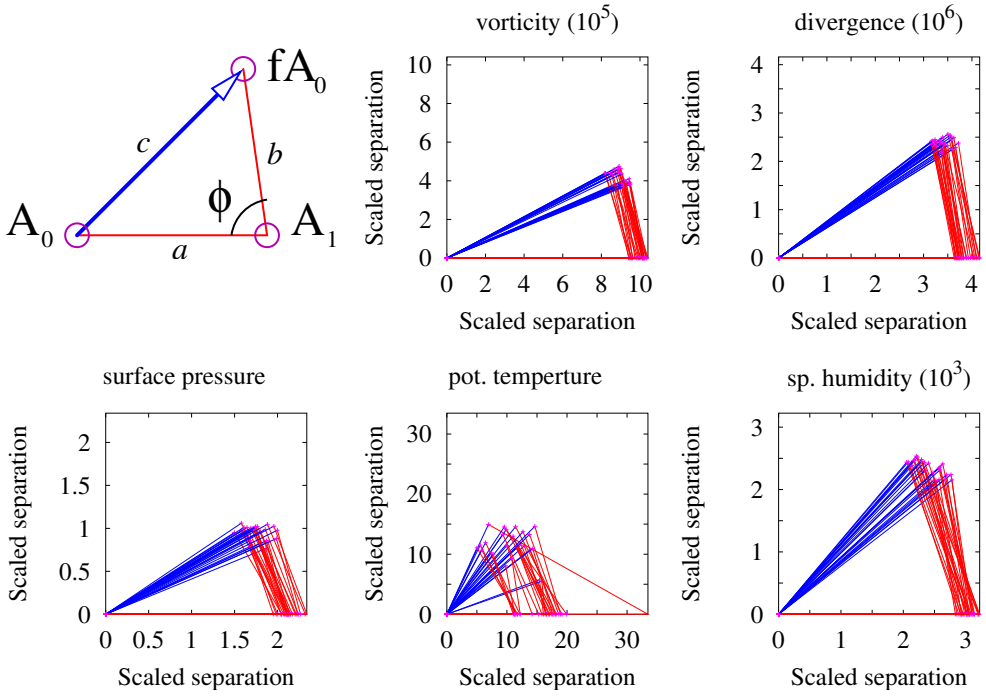

Figure 4: *Triangle diagrams that reveal the relationship between analyses, their forecasts, and the verifying analysis. Triangle diagram for original analyses for a 7 day window of 6 hour forecasts. One triangle is plotted for each consecutive pair of analyses, giving 28 triangles. (Scaled separation means the separation in scaled by the amount stated in the title.)*

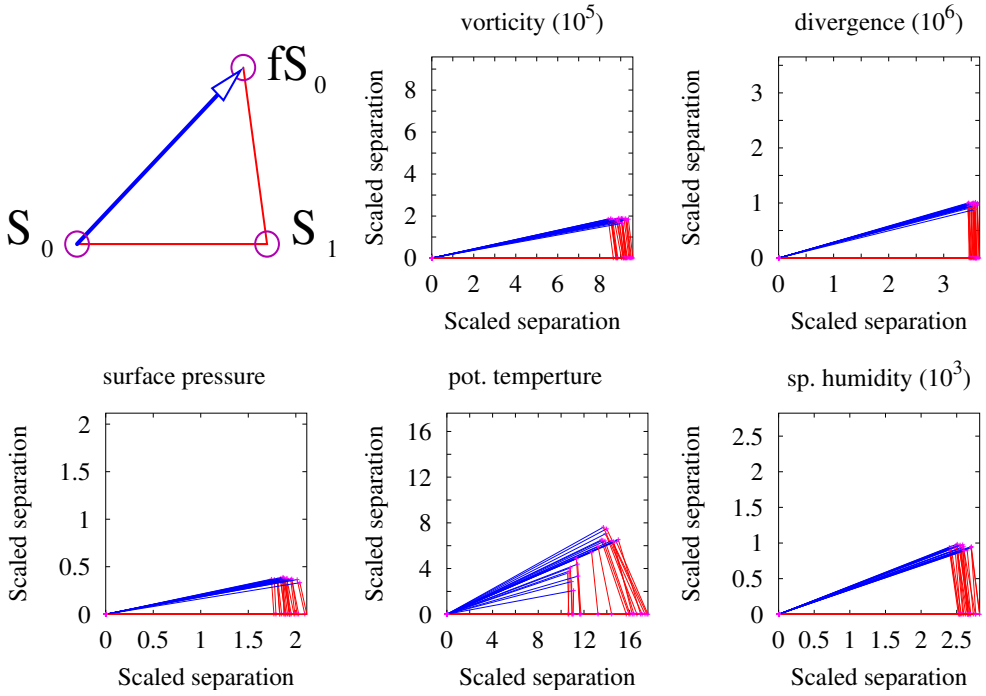

Figure 5: *Triangle diagrams that reveal the relationship between shadow analyses, their forecasts, and the verifying analysis. Triangle diagram for the shadow analyses for a 7 day window of 6 hour forecasts. One triangle is plotted for each consecutive pair of analyses, giving 28 triangles. (Scaled separation means the separation in scaled by the amount stated in the title.)*

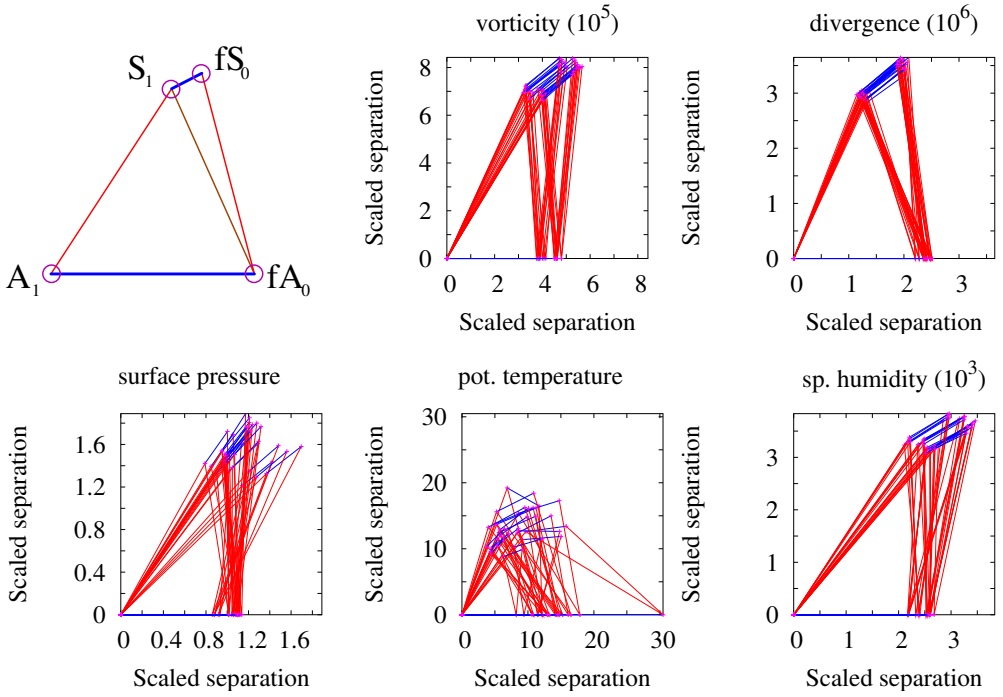

Figure 6: *Bi-triangle diagrams reveal the relationship between the mismatch of an analysis and the mismatch of the corresponding shadow analysis. Bi-triangle diagrams show the separation between analyses A_i, shadowing states S_i and forecasts of these, fA_i and fS_i. This figure shows that the shadowing states are not much further from original analyses than the forecast from the original analyses is from the verifying analysis. (Scaled separation means the separation in scaled by the amount stated in the title.)*

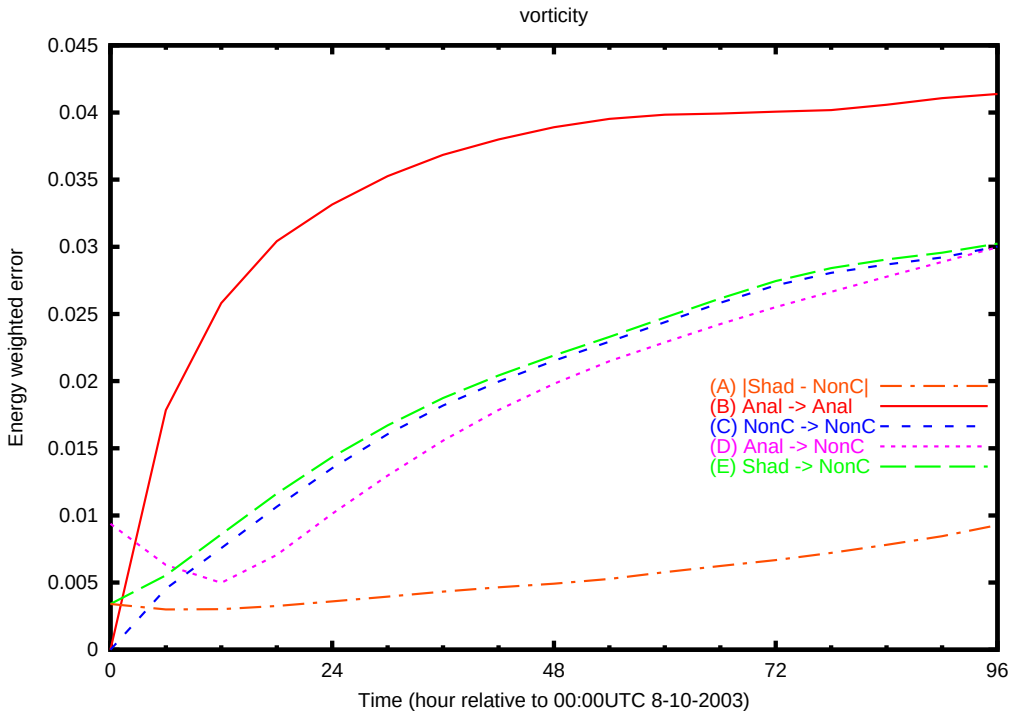

Figure 7: *Energy weighted error for NOGAPS that should be compared to the equivalent labelled curve of figure 2. Plotted are the results for vorticity, similar results were obtained for other prognostic fields.*

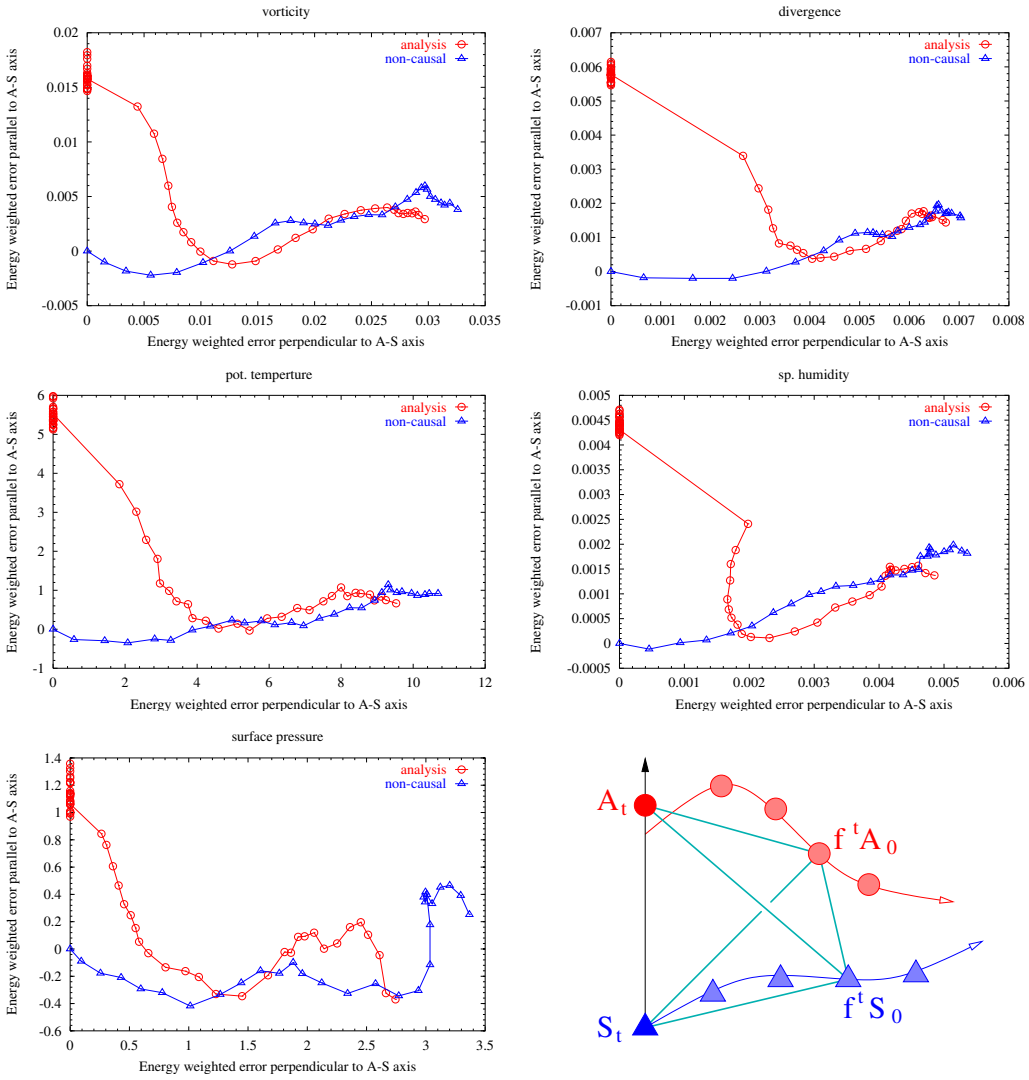

Figure 8: *The relative motion of S_t, A_t, $f^t S_0$, and $f^t A_0$ to the moving coordinate system defined by these points. Where S_t is at the origin, A_t on the y-axis, and $f^t A_0$ projected perpendicularly onto the plane through S_t, A_t, and $f^t S_0$. By connecting consecutive $f^t S_0$, and $f^t A_0$ the relative motions of these is revealed.*

(a)

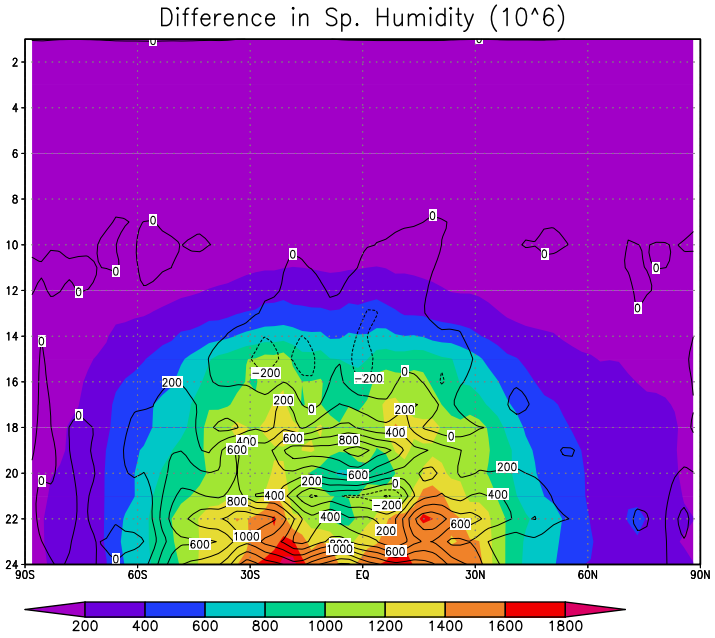

(b)

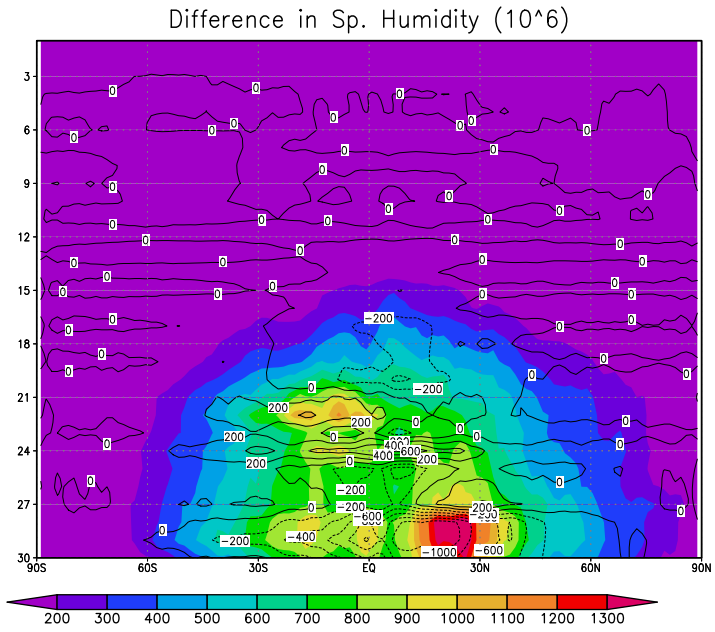

Figure 9: *Projection error (expressed here as shadow minus analysis) for (a) a T47L24 model using optimal interpolation to assimilate data, and (b) a T79L30 model using NAVDAS 3D variational assimilation. Plots show zonal average averaged over a 7 day window. Contour lines show mean error and shading shows standard deviation.*

(a)

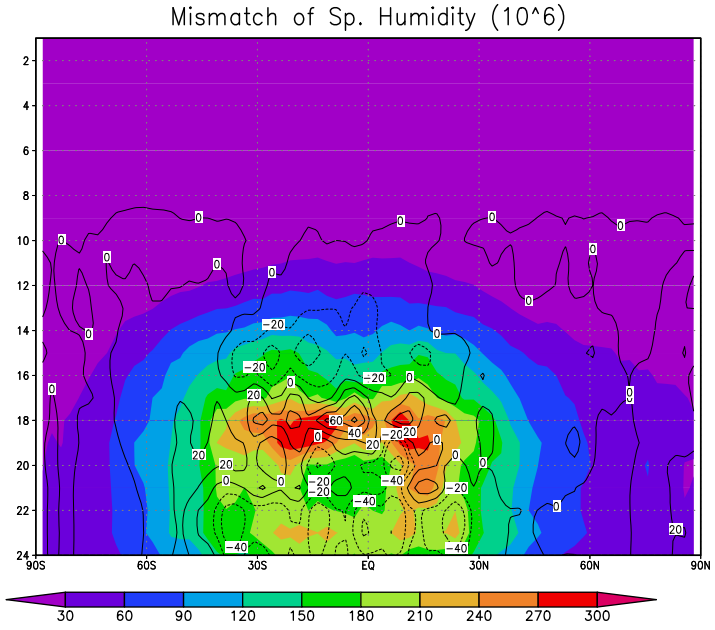

(b)

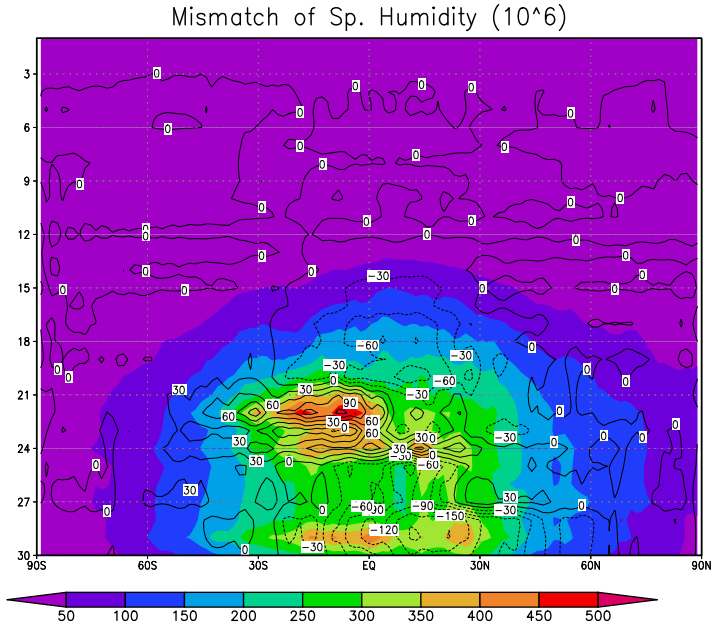

Figure 10: *Direction error for T47L24 and T79L30 models. Contour lines show mean error and shading shows standard deviation. Details as in figure9*

(a)

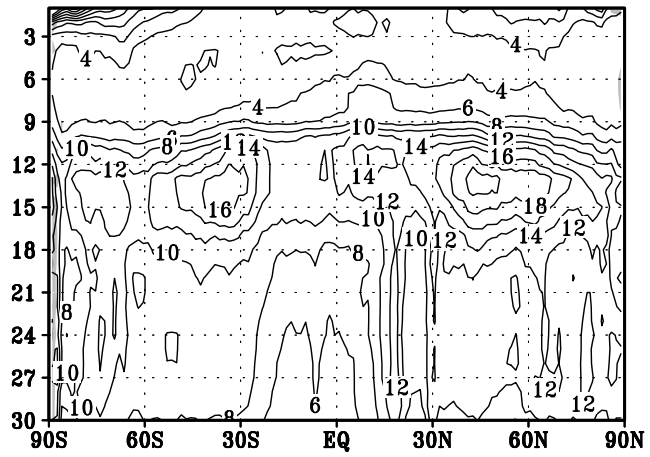

(b)

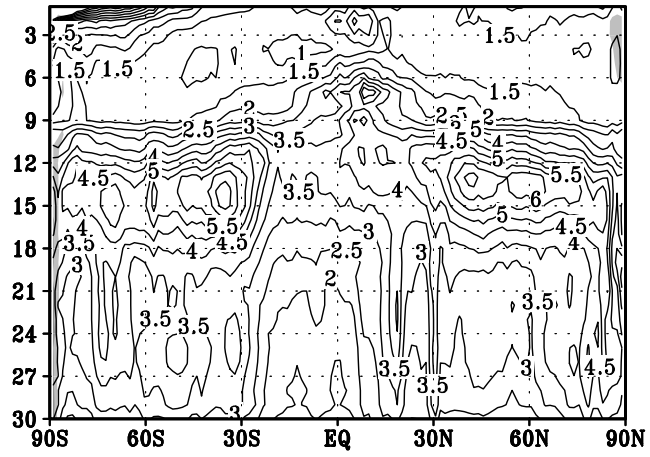

Figure 11: T79L30 vorticity field: (a) projection error, (b) direction error. Lines show standard deviation, shading shows mean error in half standard deviation increments.

(a)

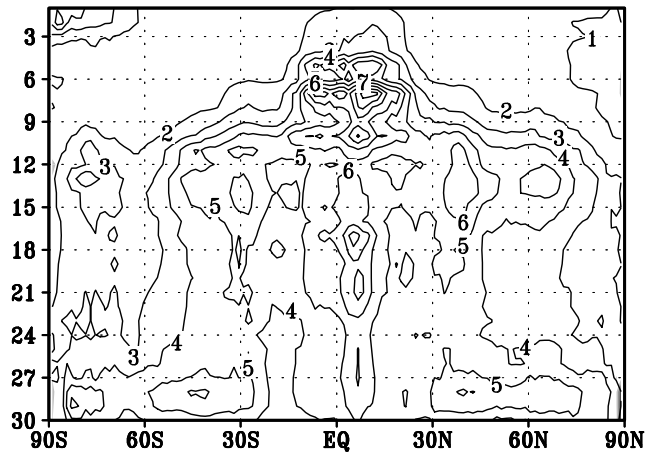

(b)

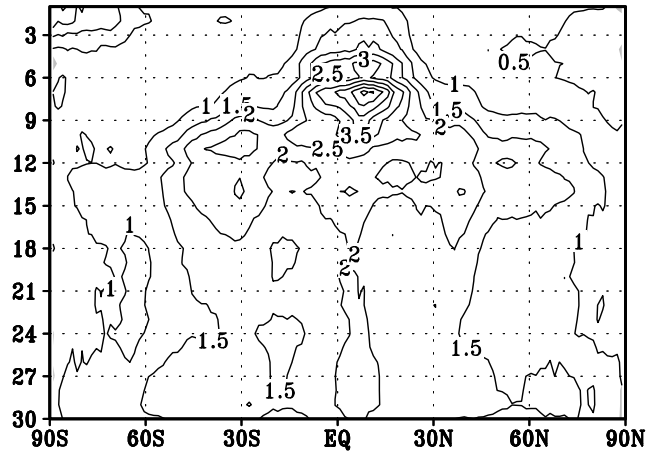

Figure 12: T79L30 divergence field: (a) projection error, (b) direction error. Lines show standard deviation, shading shows mean error in half standard deviation increments.

(a)

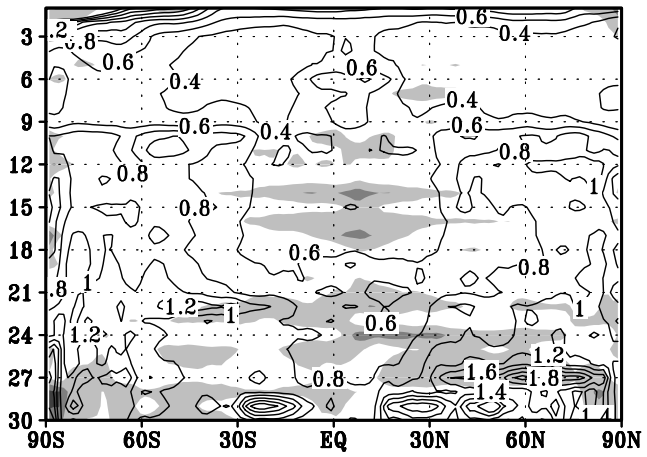

(b)

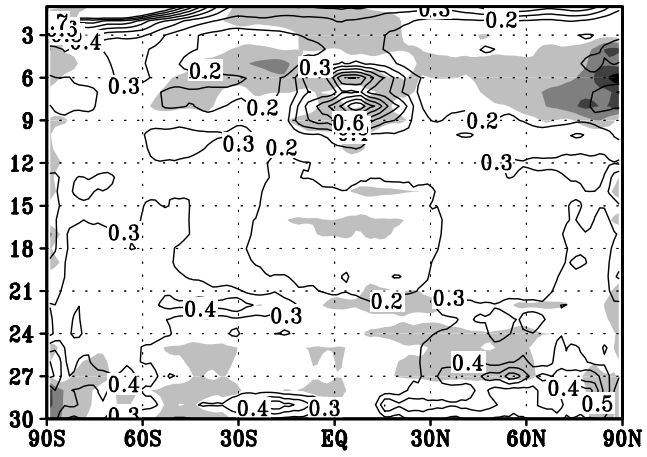

Figure 13: T79L30 temperature field: (a) projection error, (b) direction error. Lines show standard deviation, shading shows mean error in half standard deviation increments.

(a)

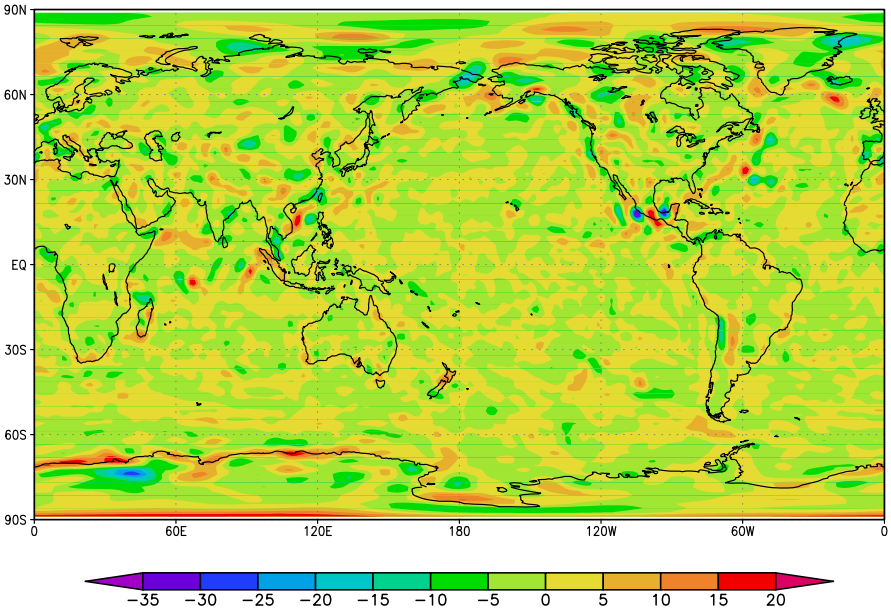

(b)

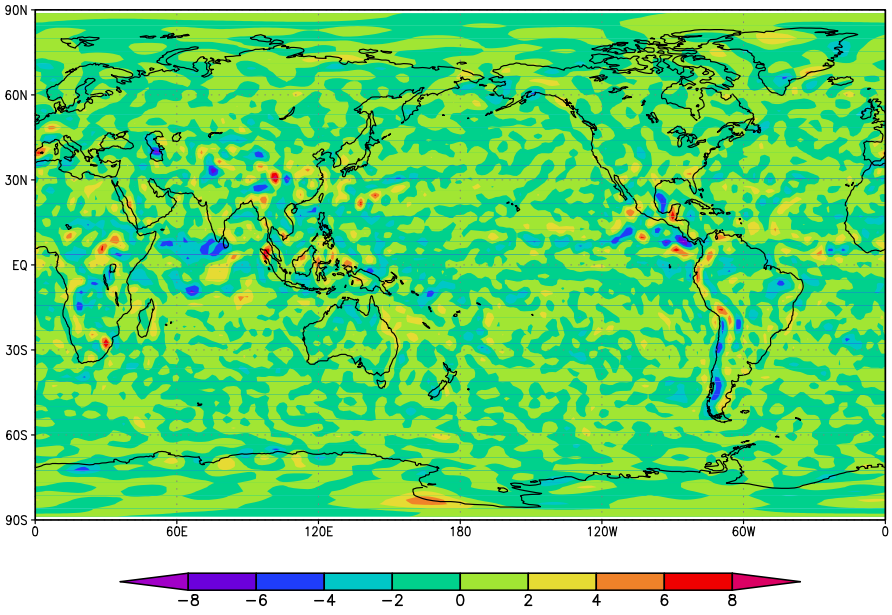

Figure 14: *The temporal average of the projection error (expressed here as shadow minus analysis) for (a) the vorticity field ($\times 10^6$) and (b) the divergence field ($\times 10^6$), at level 24, nominal 850mb level.*

(a)

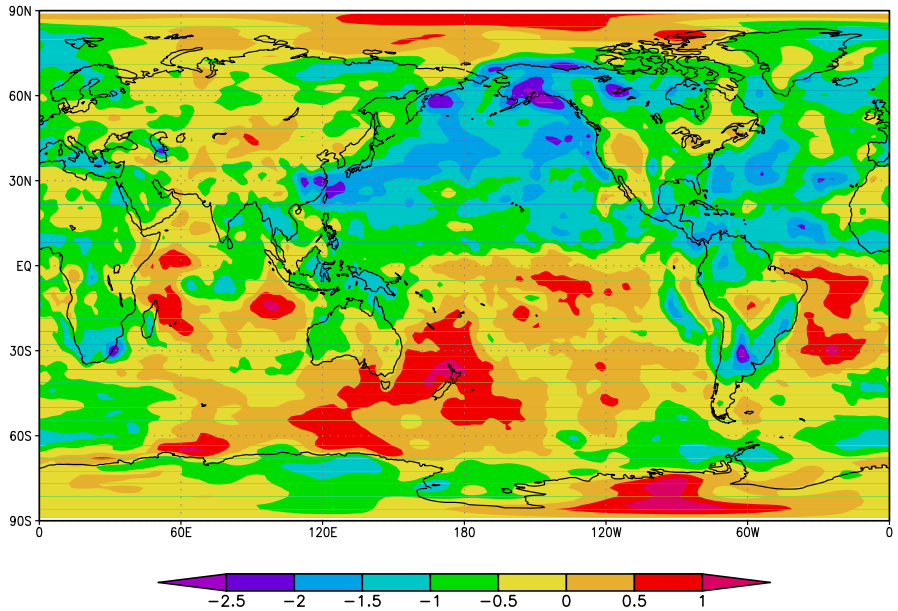

(b)

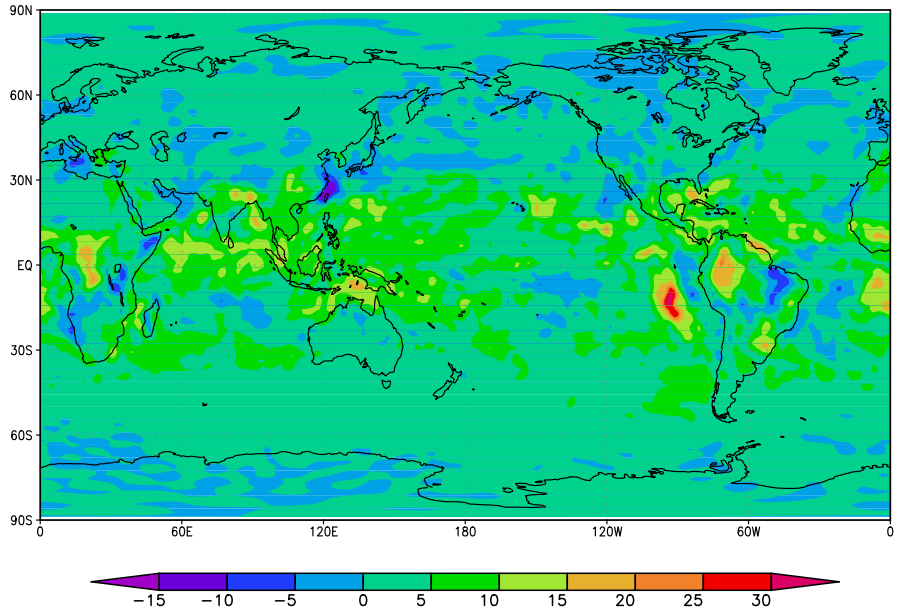

Figure 15: *The temporal average of the projection error (expressed here as shadow minus analysis) for (a) the temperature field and (b) the specific humidity field ($\times 10^8$), at level 24, nominal 850mb level.*

List of Tables

	$iA - A$	$iS - S$	$iN - N$	$A - S$	$A - N$
Vorticity	0.000436	0.000187	0.000225	0.00952	0.00938
Divergence	0.00146	0.000261	0.000390	0.00490	0.00591
Pot. Temperature	0.822	0.337	0.423	3.74	3.83
Sp. Humidity	0	0	0	0.00366	0.00367
Surface Pressure	0.271	0.131	0.139	0.465	0.441

Table 1: The impact of nonlinear normal mode initialization on analysis and shadow analysis states (first three columns) and the impact of shadowing filter (last two columns). The numbers record the magnitude of the difference of the specified field for the specified states. Symbols: A, original analysis; S, shadow analysis; N, non-causal shadow analysis; prefix i indicates state after nonlinear normal mode initialization.