Tellus (1995), 47A, 974–997
Printed in Belgium – all rights reserved

TELLUS
ISSN 0280–6495

Four-dimensional data assimilation with a wide range of scales

By MONIQUE TANGUAY, PETER BARTELLO and PIERRE GAUTHIER*, *Meteorological Research Branch, Atmospheric Environment Service, Dorval (Québec) Canada H9P 1J3*

(Manuscript received 22 August 1994; in final form 12 December 1994)

ABSTRACT

A number of experiments investigating four-dimensional variational data assimilation using the adjoint method are presented. It has been proposed that the method will be able to produce improved initial conditions in data-sparse regions. In order to describe the flow in a region where there is little data, it is necessary for observational information either to be advected into the region, or to cascade downscale from larger scales characterizing the separation between observations. We focus on the latter and examine the method's ability to "fill in" small-scale detail determined dynamically from large-scale data. We choose to examine barotropic β-plane flow since it is one of the simplest geophysical settings involving a wide range of scales. In the limit of small error, predictability studies have shown that exponential error growth occurs along the gradients of two highly-correlated realizations. When the realizations have decorrelated, error statistics saturate at climatological levels. By appealing to the adjoint of the linearized equations, the adjoint method accounts for the former behaviour, but not the latter. When the assimilation period exceeds the validity timescale of the linearization, the assimilated fields show spectra which are spuriously shallow in the small scales, following the basic-state gradients. Moreover, it is essential to note that the validity timescale of the linearization is a function of lengthscale. Therefore, for a given assimilation period there is a scale below which useful initial conditions cannot be obtained. Equivalently, for a given model resolution, there is an assimilation period beyond which the exact initial conditions cannot be recovered. Some speculation on the optimal resolution at which to perform 4D data assimilation as a function of the assimilation period is offered.

1. Introduction

Recently, there has been a renewed interest in four-dimensional data assimilation as a potential remedy to some of the deficiencies of current statistical interpolation schemes. This has been spawned by the work of Lewis and Derber (1985), LeDimet and Talagrand (1986) (LDT86 hereafter) and Talagrand and Courtier (1987) (TC87 hereafter) who demonstrated that the use of adjoint methods made it feasible to perform a quadridimensional variational assimilation (4Dvar) with fairly realistic models (for a review, see Courtier et al. (1993)). The strong points of 4Dvar are several. With respect to the observations, it allows them to be assimilated at the time they are made and in their raw form. Moreover, 4Dvar is a constrained minimization in that initial conditions $\boldsymbol{X}_0$ at t_0 are sought to fit globally all observations contained within a time interval $(t_0, t_0 + T_a)$, where T_a is the assimilation period. These initial conditions act as control variables of the minimization and determine the model state $\boldsymbol{X}(t)$ at all times. Recently, Andersson et al. (1992) have shown how this introduces a dynamic coupling between the model's variables. For instance, their results showed that observations of humidity inferred information about the advecting wind field.

With respect to the fact that model resolutions are constantly increasing, observational networks will always remain in some sense coarse and the

* Corresponding author.

4D fitting of models to data will be expected to "fill in" the fine detail between data points in a manner which is dynamically consistent with the larger scales. To what extent is this possible? To answer, one must first consider the mechanism by which information is transferred from scale to scale. As a first step in this process, the most geophysically-relevant transfer mechanism was investigated in as simple a context as possible: the enstrophy cascade of barotropic β-plane homogeneous turbulence.

It would appear that fully-developed turbulent cascades present a challenge to the method. A functional measuring a spatial and temporal mean discrepancy between model and observation is minimized by appealing to the adjoint of the equations linearized about the previous guess. The next guess is chosen somewhere along a phase-space direction obtained from the functional's gradient. Exactly where, is determined by integrations of the direct (i.e., fully nonlinear) model. Clearly, the linearized equations can only be expected to approximate well the model equations for times less than the nonlinear timescale, τ_{NL}. In numerical weather prediction τ_{NL} characterizes the growth and decay of synoptic eddies and is of the order of 3 days. The problem is that τ_{NL} also governs the transfer of information from one scale to another. We are therefore faced with conflicting requirements. For significant scale redistribution of the observational information, T_a must be of order τ_{NL}, while for accuracy of the tangent linear model, $T_a \ll \tau_{NL}$. The situation is further complicated by the fact that, depending on the spectrum, the nonlinear timescale is a function of lengthscale.

The model details are presented in Section 2. In Section 3, we compare integrations of the linear model with differences obtained from two nonlinear-model integrations, to demonstrate that the limit of validity of the tangent linear approximation is a function of lengthscale, the smallest scales becoming nonlinear most quickly. The rest of the paper deals with data assimilation experiments where "observations" are saved from previous model integrations. The focus here is on the transfer of information in the enstrophy cascade. It follows that neither our observations nor our model have error: they are perfect to within the precision of our computations. In Section 4, a data assimilation experiment is presented in which observations are provided at all spatial and temporal grid points in order to examine the nature of the convergence of the method. It is found that convergence is rapid in the large scales throughout the assimilation period. In the small scales, convergence near the beginning of the assimilation period is slow. After $O(10^2)$ iterations, the bulk of the error is contained in the finest scales at the initial time. When the assimilation period is lengthened well beyond τ_{NL}, large-scale convergence is very slow, while the small scales diverge. Similar experiments are performed in Section 5 except that only observations of the largest model scales are provided. Here, we investigate the method's ability to supply small-scale detail dynamically consistent with the large-scale observations. We note that for small T_a, the method does not converge in the small scales early in the assimilation interval. However, during the course of the assimilation period, information from the large-scale observations is redistributed spectrally by the nonlinear transfer if T_a is large enough. At the end of the assimilation period the small-scale error is 2 orders of magnitude lower than at the beginning. When T_a exceeds the predictability timescale for all lengthscales, the smallest error in the initial conditions leads to a solution that has no resemblance to the truth. In this case the small-scale error approaches spuriously large values that can be interpreted in terms of the nonlinear transfer of predictability error in the enstrophy-cascade range. At the same time, the form of the functional becomes less and less quadratic as T_a is increased, eventually leading to multiple minima and a failure of the minimization. This is further illustrated in Section 6 by experiments involving a single observation (with complete spatial coverage) inserted at $t = T_a$. In Section 7, the Hessian matrix is computed at the absolute minimum and is seen to be a full matrix. Consequently, convergence is very slow, even with a quadratic functional, and it is far from obvious how to precondition the minimization. The conclusions with respect to the relation between optimal assimilation periods and resolution are discussed in Section 8.

2. The model and methodology

2.1. *The numerical model*

The barotropic vorticity equation on the β-plane is

$$\frac{\partial \zeta}{\partial t} + J(\Psi, \zeta) + \beta v = f - \mathscr{D}(\zeta), \tag{1}$$

where $\Psi = -U_o y + \psi$, $\zeta = \nabla^2 \psi$, $u = -\partial\psi/\partial y$, $v = \partial\psi/\partial x$, J is the Jacobian, f is a forcing term and $\mathscr{D}$ is a linear dissipation operator. Since geophysical fluids are at statistical stationarity, we apply an ad hoc forcing mechanism to the large scales. This affects our choice of $\mathscr{D}$ in that energy in barotropic flow is cascaded to larger scales, where it must be absorbed by the dissipation. Here, we apply a Rayleigh dissipation and an iterated Laplacian to dissipate the small scales selectively, $\mathscr{D} = \nu_o + \nu(\nabla^2)^8$. We have employed hyperviscosity in order to reduce the spectral width of the small-scale dissipation range to a minimum (see, e.g., Basdevant and Sadourny (1983)).

The numerical model is similar to that used by Bartello and Holloway (1991). It integrates (1) on a doubly-periodic domain of length 2π using pseudo-spectral methods (Orszag, 1971). The fields can be expressed in physical or Fourier space via discrete Fourier transforms. In the latter (1) can be written

$$\left[\frac{\partial}{\partial t} + i\omega_{\boldsymbol{k}} + \nu_{\boldsymbol{k}}\right] \hat{\zeta}_{\boldsymbol{k}} = \hat{f}_{\boldsymbol{k}} + \sum_{\boldsymbol{p}+\boldsymbol{q}=\boldsymbol{k}} A_{\boldsymbol{kp}} \hat{\zeta}_{\boldsymbol{p}} \hat{\zeta}_{\boldsymbol{q}}, \tag{2}$$

where a caret is used to denote a Fourier-space quantity, $\boldsymbol{k}$ is the wavevector, $k = |\boldsymbol{k}|$, $\omega_{\boldsymbol{k}} = k_x(U_o - \beta/k^2)$ is the linear Rossby-wave frequency, $A_{\boldsymbol{kp}} = \hat{\boldsymbol{z}} \cdot \boldsymbol{k} \times \boldsymbol{p}/p^2$, $\hat{f}_{\boldsymbol{k}}$ is the forcing and $\nu_{\boldsymbol{k}} = \nu_o + \nu k^{16}$ represents the dissipation operator $\mathscr{D}$. The forcing term was

$$\hat{f}_{\boldsymbol{k}} = \alpha \begin{cases} 1, & \text{if } \boldsymbol{k} = (0, \pm k_f) \text{ or } (\pm k_f, 0), \\ 0, & \text{otherwise,} \end{cases}$$

where $\alpha = 0.04 \in \Re$ and $k_f = 3$. A mean zonal wind $U_o = 0.3$ was also imposed. Although U_o does not interact with the other modes, it provides a net translation of the fluid over the stationary forcing field, effectively introducing a frequency $2\pi/k_f U_o \approx 6$. Unless otherwise specified, we used $\mathscr{N} = 64$ collocation points per dimension and applied circular Fourier-space truncation at $k_T = (\mathscr{N} - 3/2)/3$ in order to avoid aliasing errors. The extra term "3/2" conveniently assures that all of the wavebands used to define spectra below are complete. For the runs presented here we used a timestep, $\Delta t = 0.95/k_T$ with $\nu_o = 0.02$ and $\nu = 8.80/k_T^{16}$.

We note that the resolution is modest by today's standards. However, we are not attempting to simulate inertial ranges as such, but only require a range with significant nonlinearity. In addition, this resolution allows us to perform integrations, corresponding to several days in atmospheric terms, in just a few seconds of CPU time. This is a definite advantage since, in the following we present data-assimilation experiments with up to 500 iterations of the adjoint method. The calculations were performed using the NEC SX-3 at the Centre d'Informatique de Dorval.

2.2. *Model statistics and scaling*

Before introducing the adjoint method formalism, we examine the "climatology" of the model and specify the appropriate scaling that permits us to convert the results to dimensional atmospheric variables. In the absence of forcing and dissipation (1) conserves fluctuation energy, E_T and enstrophy, Z_T where

$$E_T = \frac{1}{2} \sum_{\boldsymbol{k}} \frac{|\hat{\zeta}_{\boldsymbol{k}}|^2}{k^2} \quad \text{and} \quad Z_T = \sum_{\boldsymbol{k}} |\hat{\zeta}_{\boldsymbol{k}}|^2.$$

The reality of ζ implies that $\hat{\zeta}_{\boldsymbol{k}}$ is the complex conjugate of $\hat{\zeta}_{-\boldsymbol{k}}$.

We form the instantaneous spectrum, $Q(k)$, of a quantity, Q_T, by summing its modal contributions, $\tilde{Q}_{\boldsymbol{k}}$, in k-space rings of unit thickness centred on wavenumber k, i.e.,

$$Q(k) = \frac{2\pi k}{M(k)} \sum_{k-(1/2)<|\boldsymbol{k}'|\leqslant k+(1/2)} \tilde{Q}_{\boldsymbol{k}'}, \tag{3}$$

where $M(k)$ is the number of discrete modes in the kth waveband and the factor $2\pi k/M(k)$ serves to smooth out the modal distribution in k-space. The energy spectrum, $E(k)$, and enstrophy spectrum, $Z(k)$ are obtained with $\tilde{Q}_{\boldsymbol{k}'} = |\hat{\zeta}_{\boldsymbol{k}'}|^2/2k'^2$ and $\tilde{Q}_{\boldsymbol{k}'} = |\hat{\zeta}_{\boldsymbol{k}'}|^2$, respectively.

The model was spun up from random initial conditions with $E_T = 10^{-4}$ and integrated for 500 time units. Time series of the total energy and enstrophy are shown in Fig. 1. After an initial overshoot the nonlinearity begins to act and the system settles down to oscillations about the values $\bar{E}_T \approx 0.04$ and $\bar{Z}_T \approx 0.5$, where the overbar is used to denote an ensemble average. (Note that if the

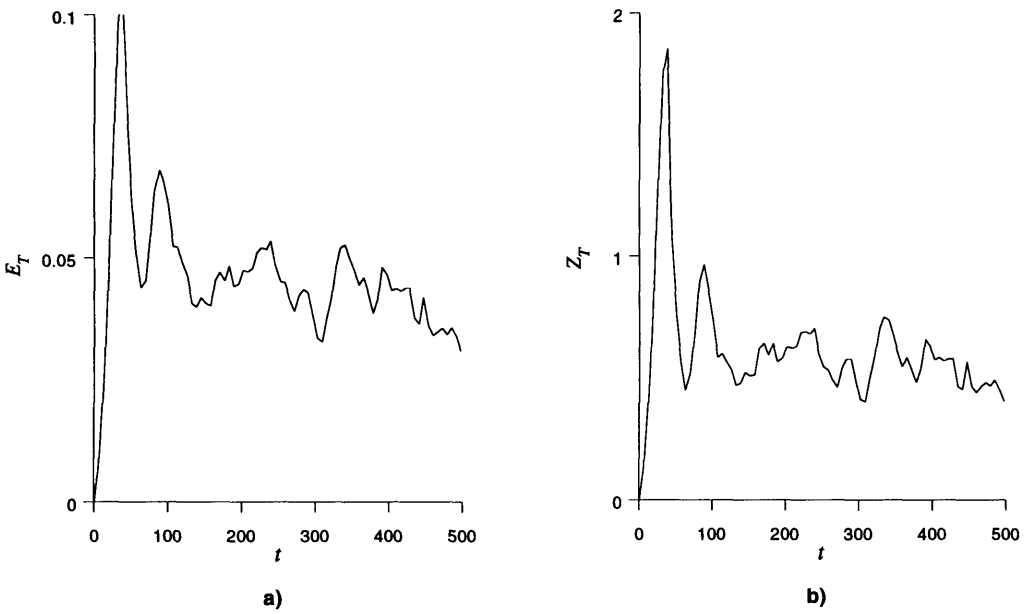

Fig. 1. Timeseries of (a) fluctuation energy and (b) enstrophy for the spin-up simulation with the direct model.

flow is statistically stationary and ergodic then time averages can be used as ensemble averages.) We use these values as the basis for our scaling arguments which are stated using one significant digit. The large-scale nonlinear turnover time can be identified with $\tau_{NL} \sim 2\pi/\zeta_{rms} = 2\pi/\bar{Z}_T^{1/2} \approx 9$. Using the synoptic-scale advective timescale, $\tau_{NL} \sim L/U \approx 3$ days, we can associate one model time unit with 0.3 days. The rms velocity of eddy motions was $U_{rms} = (2\bar{E}_T)^{1/2} \approx 0.3$, which is approximately the same value as the imposed mean flow U_o. The Rhines (1975) wavenumber, $k_\beta = (\beta/2U_{rms})^{1/2}$, can be used to define a lengthscale above which the flow is dominated by Rossby waves and below which turbulence prevails. We have $\beta = 0.5$ and therefore $k_\beta \approx 1$, while the mid-latitude atmosphere has $\beta \approx 1.6 \times 10^{-11}$ m^{-1} s^{-1} and $k_\beta \approx 2\pi/7000$ km if $U_{rms} \approx 10$ m s^{-1}. Therefore, our domain size corresponds roughly to 7000 km, and since $k_f = 3$, the forcing scale is 2000 km. As a first approximation, we view this as the injection of barotropic energy at the synoptic scale. The ratio of β to the eddy vorticity gradients can be expressed as in Bartello and Holloway (1991), $\hat{\beta} = (\beta U_{rms}/\zeta_{rms}) \approx 0.3$. The Rayleigh friction term results in a viscous dissipation of the most energetic scales with spin-down time, $\tau_R \sim \nu_o^{-1} \approx 50$ model time units or approximately 16 days. Since this is 6 τ_{NL}, it is safe to conclude that the Rayleigh dissipation is weak and that the model is dominated by nonlinearity at most scales.

In order to describe the model's statistically stationary state we performed an integration of $100\tau_{NL}$ duration starting from the end of the run presented in Fig. 1. In Fig. 2 we present the average energy spectrum $\bar{E}(k)$. Here, we note a relatively flat curve for $k < k_f$, above which the spectrum drops off roughly following $\bar{E}(k) \sim k^{-\alpha}$, with $\alpha \approx 4$. Although this is steeper than atmospheric observations, it is typical of that obtained with barotropic models (Basdevant et al., 1981). Above the dissipation wavenumber, $k_D \approx 14$, the spectrum falls off faster than algebraically in the dissipation range. These familiar results are typical of 2D turbulence simulations (Basdevant et al., 1981; Lesieur, 1990). A detailed examination of the nonlinear transfer term in the evolution equation for $\bar{E}(k)$ revealed that the model sustains an active enstrophy cascade, although its wavenumber range is small.

Tellus A/47/5:2-6

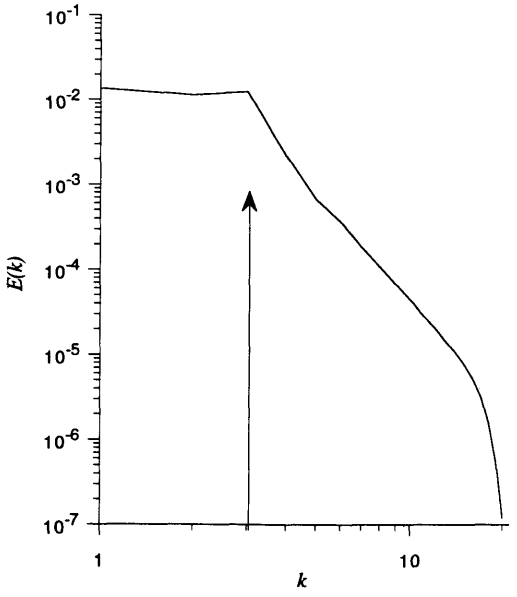

Fig. 2. Time-averaged energy spectrum of the direct model. Arrow indicates k_f.

2.3. *The adjoint method*

To apply the adjoint method, we introduce the tangent linear model by considering a reference solution to the nonlinear equations, ζ and a perturbation about that solution, $\delta\zeta$. We write (1) for the total field $\zeta + \delta\zeta$ and then subtract the equation for ζ to get

$$\left[\frac{\partial}{\partial t} + U_\circ \frac{\partial}{\partial x}\right] \delta\zeta + J(\psi, \delta\zeta) + J(\delta\psi, \zeta) + J(\delta\psi, \delta\zeta) + \beta\delta v = -\mathscr{D}(\delta\zeta). \tag{4}$$

The tangent linear model is obtained by neglecting the term $J(\delta\psi, \delta\zeta)$. It can be used to approximate the evolution of the error caused by a change $\delta\zeta_0$ in the initial conditions. If this change is small but finite, it will only be accurate for a given period of time, referred to as the limit of validity of the tangent linear model, and discussed in the next section.

If it is assumed that the error is governed by the linear model, then its evolution in Fourier space can be symbolically represented as

$$\delta\hat{\zeta}(t) = R(t, t_0)\, \delta\hat{\zeta}_0, \tag{5}$$

where $\delta\hat{\zeta}(t)$ represents the vector of Fourier coefficients whose components are $\{\delta\hat{\zeta}_k(t)\}$ and $R(t, t_0)$ is the resolvent, associated with the tangent linear model, that carries the solution from time t_0 to time t (see LDT86 and TC87). To assimilate data spread out over a finite time interval $(t_0, t_0 + T_a)$, one seeks the initial conditions $\hat{\zeta}_0$ leading to a solution $\hat{\zeta}(t)$ that will minimize the "distance" between this solution and the data as measured by the distance function (or functional)

$$J(\hat{\zeta}_0) = \tfrac{1}{2} \sum_{i=1}^{N} \langle \hat{\zeta}_i - O_i, \hat{\zeta}_i - O_i \rangle.$$

Here, $\hat{\zeta}_i = \hat{\zeta}(t_i) = \{\hat{\zeta}_k(t_i)\}$, O_i is an observation available at time t_i, while $\langle\ ,\ \rangle$ stands for the inner product chosen as a measure of the error. For example, inner products based on error energy or error enstrophy would be, respectively

$$\langle \hat{\zeta}_1, \hat{\zeta}_2 \rangle_E = \frac{1}{2} \sum_k \frac{\hat{\zeta}_{1k} \hat{\zeta}^*_{2k}}{k^2},$$

$$\langle \hat{\zeta}_1, \hat{\zeta}_2 \rangle_Z = \sum_k \hat{\zeta}_{1k} \hat{\zeta}^*_{2k}.$$

Here, the energy norm is used and, for the sake of simplicity, the observations O_i correspond to full spectral states unless otherwise specified.

The minimization of the distance function can be carried out iteratively using algorithms such as conjugate gradient or quasi-Newton methods, the latter being generally more efficient (Gill and Wright, 1981; Navon and Legler, 1987; Gilbert and Lemaréchal, 1989). All experiments in this paper were carried out with a variable storage quasi-Newton algorithm that uses information from past iterations to precondition the minimization by approximating the inverse of the Hessian matrix (see Section 7). At each iteration, the functional $J(\hat{\zeta}_0)$ and its gradient must be computed. After integrating the direct model to obtain the solution $\hat{\zeta}(t)$ the evaluation of J itself is straightforward.

In LDT86 and TC87, ideas from control theory were used to find an efficient way of computing ∇J. They noted that the variation δJ due to an *infinitesimal* change $\delta\hat{\zeta}_0$ of the initial conditions can be written as

$$\delta J = \sum_{i=1}^{N} \langle R(t_i, t_0)\, \delta\hat{\zeta}_0, \hat{\zeta}_i - O_i \rangle. \tag{6}$$

The fact that $R(t, t_0)$ is a linear operator, $\mathscr{L}$, makes it possible to use the property that for any $\mathscr{L}: H_1 \to H_2$, its adjoint $\mathscr{L}^*$ is defined by

$$\langle \alpha, \mathscr{L} b \rangle_2 = \langle \mathscr{L}^* \alpha, b \rangle_1$$

where H_1 and H_2 are vector spaces with inner products $\langle \alpha_1, \alpha_2 \rangle_1$ and $\langle b_1, b_2 \rangle_2$, respectively. Consequently, (6) can be rewritten as

$$\delta J = \left\langle \delta \hat{\zeta}_0, \sum_{i=1}^{N} R^*(t_i, t_0)(\hat{\zeta}_i - O_i) \right\rangle. \tag{7}$$

Therefore,

$$\nabla_E J(\hat{\zeta}_0) = \sum_{i=1}^{N} R^*(t_i, t_0)(\hat{\zeta}_i - O_i). \tag{8}$$

The subscript $_E$ indicates that the gradient is defined with respect to the energy norm. One is then left with the determination of the adjoint of $R(t_i, t_0)$. It has been shown in LDT86 and TC87 that $R^*(t_i, t_0)$ corresponds to the resolvent $S(t_0, t_i)$ of the adjoint model which is (see Appendix)

$$\left(\frac{\partial}{\partial t} + U_0 \frac{\partial}{\partial x} \right) \delta^* \zeta = - \mathscr{L}^*(\hat{\zeta})\, \delta^* \zeta,$$

with

$$\mathscr{L}^*(\hat{\zeta})\, \delta^* \zeta \equiv -\Delta J(\Delta^{-1}\, \delta^* \zeta, \Delta^{-1} \zeta) - J(\zeta + \beta y, \Delta^{-1}\, \delta^* \zeta) - \mathscr{D}(\delta^* \zeta). \tag{9}$$

Consequently, (8) implies that to obtain the gradient, the adjoint model is integrated *backward* in time from time t_i to t_0 using the difference, $\hat{\zeta}_i - O_i$ as an "initial" condition. Theoretically, the calculation of the gradient is exact since the tangent linear approximation is applied to only infinitesimal changes, $\delta \hat{\zeta}_0$ to the initial conditions. However, in practice, numerical errors are involved and it is necessary to use the adjoint of the numerical scheme instead of discretizing (9). This is what has been done here by employing a leap-frog scheme and its adjoint as described in TC87. Verification of the adjoint code was made by using the Taylor formula (Thépaut and Courtier, 1991).

The adjoint model involves twice the number of nonlinear terms so that its integration is approximately twice as costly. Given the fact that the minimization algorithm requires several iterations, each requiring one or more evaluations of the functional and its gradient, the cost of one iteration corresponds roughly to that required to integrate the model over a period of $3T_a$. As will be seen shortly, the number of iterations can be quite high, which makes this method an expensive proposition if the model used is complex.

3. Linear model integrations

To investigate the accuracy of the tangent linear model, we performed an integration of 50 model time units (16 days) with the nonlinear model. The result, $\zeta_1 = \zeta$, served as a basic state for the linearized model. A perturbation field, $\delta\zeta$, was then initialized randomly and the linear equations were integrated. For comparison we also performed a nonlinear-model integration starting from $\zeta_2 = \zeta + \delta\zeta$. The perturbation, $\delta\zeta(t=0)$, was confined to the single waveband $13.5 < k \leqslant 14.5 \approx k_D$, implying that it contained only the smallest active scales of the model. Its energy was arbitrarily set to approximately $3E(k=14)$. The energy spectra of both realizations (i.e., $E_1(k, t)$ corresponding to ζ_1 and $E_2(k, t)$ corresponding to ζ_2) as well as the energy spectrum of the difference field, $E_\delta(k, t)$, which is obtained from (3) using $\tilde{Q}_{k'} = |\delta\hat{\zeta}_{k'}|^2/2k'^2$, are displayed at two instants in Fig. 3. At $t = 28 \approx 3\tau_{NL}$ the linear model performed well. This may appear to be somewhat surprising since at wavenumber 14 the perturbation amplitude was initially of the order of the basic-state field, implying that the linearization should have been inaccurate from the start. However, near $k \approx k_D$ the small-scale dissipation acts as rapidly as the advective phase organization which precedes significant transfer (see, e.g., Lesieur, 1990). As a result, there was an initial phase in which $E_\delta(k)$ was considerably dissipated.

The spreading of the error energy both towards larger and smaller scales is familiar from predictability studies by Lorenz (1969), Leith (1971), Leith and Kraichnan (1972), Litherland and Holloway (1984), Herring (1984) and Métais and Lesieur (1986). A word on these previous studies is in order here. Most have performed evaluations of statistical 2-point closure theories (such as the eddy-damped quasi-normal Markovian theory or EDQNM) for problems similar to the one

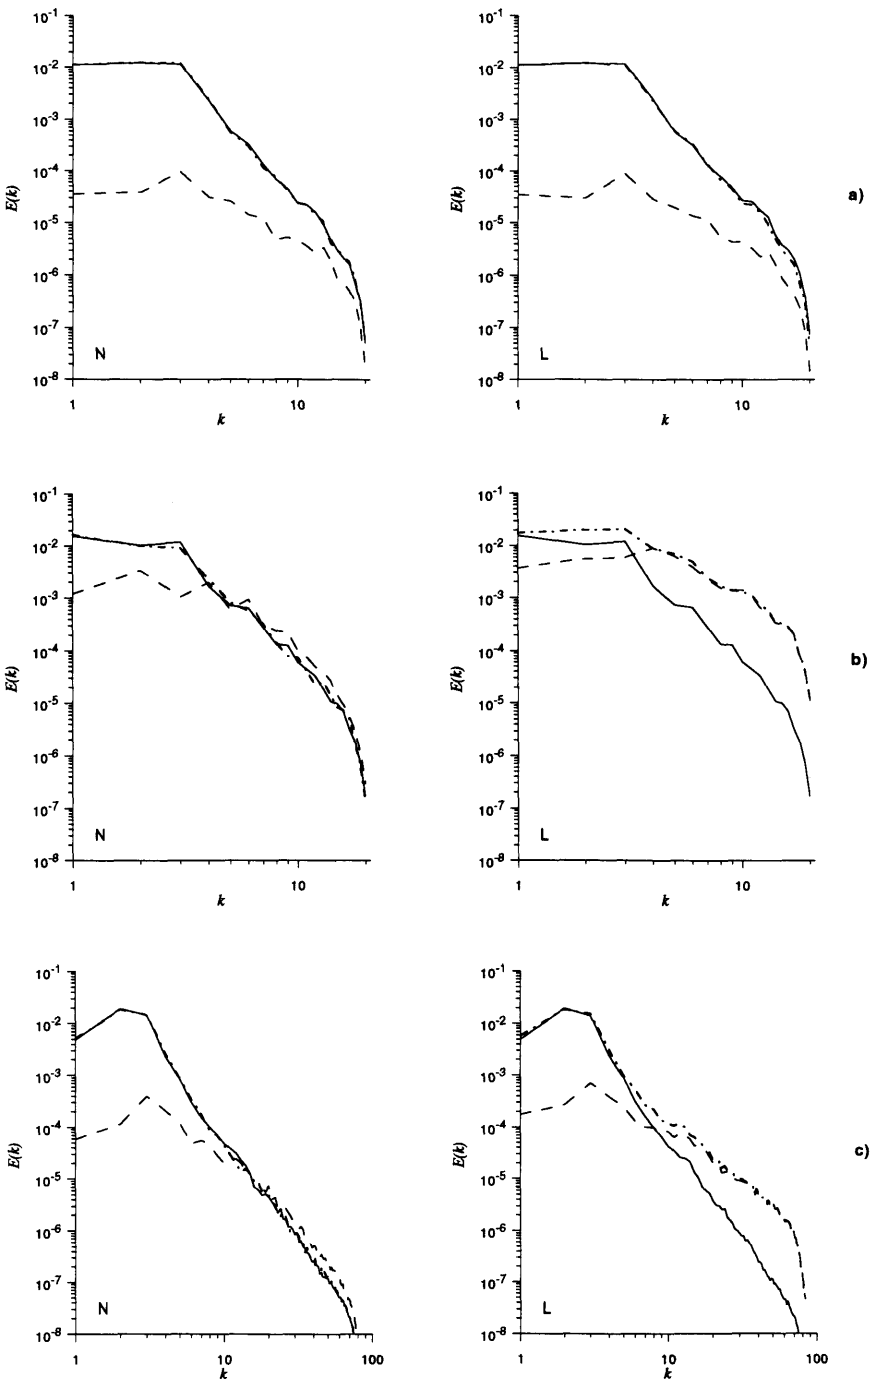

Fig. 3. Comparisons between predictability simulations made with two nonlinear integrations (left-hand panels) and with the tangent linear model (right-hand panels). (a) Resolution, 64^2 at $t = 28$; (b) 64^2 at $t = 57$ and (c) 256^2 at $t = 28$. The curves represent $E_1(k)$ (solid line), $E_2(k)$ (dashed-dotted line) and $E_\delta(k)$ (dashed line).

considered. These theories display $\bar{E}(k) \sim k^{-3}$ behaviour in the enstrophy range as opposed to the steeper spectra observed in simulations here and elsewhere. Nonetheless, many of their conclusions serve as a valuable guide in the interpretation of what follows. Particularly noteworthy is the fact that Herring (1984) and Métais and Lesieur (1986) found $\bar{E}_\delta(k) \sim k^{-1}$ in the small-error limit. In this limit Métais and Lesieur (1986) pointed out that the EDQNM spectral energy equation for the difference field in isotropic 2D turbulence can be approximated at large k by

$$\frac{\partial}{\partial t}\bar{E}_\delta(k) = 2\left[\int_0^\infty \theta_{kkq}\bar{E}_\delta(q)\,\mathrm{d}q\right] k^2\bar{E}(k) + \left[\int_0^\infty \theta_{kkq}q^2\bar{E}(q)\,\mathrm{d}q\right]\mathscr{A}[\bar{E}_\delta(k)] + \left[\int_0^\infty \theta_{kkq}q^2\bar{E}_\delta(q)\,\mathrm{d}q\right]\mathscr{A}[\bar{E}(k)],$$

where

$$\mathscr{A}[\ ] = \frac{1}{4}\frac{1}{k^2}\frac{\partial}{\partial k}\left(k^3\frac{\partial}{\partial k}(k[\ ])\right),$$

$$\theta_{kkq} = \frac{1-\exp[-(2\mu_k+\mu_q)\,t]}{2\mu_k+\mu_q},$$

$$\mu_k = \lambda\left[\int_0^k k'^2\bar{E}(k')\,\mathrm{d}k'\right]^{1/2} + \nu k^2.$$

The θ_{kkq} are model-dependent eddy damping rates which are only weak functions of k at large k if $\bar{E}(k)$ is sufficiently steep (Lesieur, 1990). Using the EDQNM prediction of $\bar{E}(k) \sim k^{-3}$, they noted that the first term generates $\bar{E}_\delta(k) \sim k^{-1}$ behaviour. This agreed with their solutions of the full theory since, in this case, $\mathscr{A}[\bar{E}(k)] = \mathscr{A}[\bar{E}_\delta(k)] = 0$. If however, we impose only the form $\bar{E}(k) \sim k^{-\alpha}$ and $\bar{E}_\delta(k) \sim k^{-\beta}$, where $\alpha, \beta > 0$, we obtain for the three terms

$$\frac{\partial}{\partial t}\bar{E}_\delta(k) \sim Ak^{2-\alpha} + Bk^{-\beta} + Ck^{-\alpha}.$$

Therefore, if $\bar{E}_\delta(k)$ is initially steeper than $k^{2-\alpha}$ at large k, we still obtain $\bar{E}_\delta(k) \sim k^2\bar{E}(k)$. This result is not very surprising, reflecting only the fact that the exponential growth of small differences is initially most rapid in regions of strong gradients. As the error grows, other terms enforcing simultaneous conservation in the two realizations become important. Ultimately, at complete decorrelation, we have $\overline{[\zeta_2-\zeta_1]^2} = \overline{\zeta_2^2} + \overline{\zeta_1^2} - 2\overline{\zeta_1\zeta_2} = 2\overline{\zeta^2}$ and mean error growth ceases at twice the climatological variance. We can conclude that there are two limiting behaviours: for very small difference fields the error growth is rapid and occurs along the gradients, yielding $\bar{E}_\delta(k) \sim k^2\bar{E}(k)$, while at complete decorrelation the error growth saturates at $\bar{E}_\delta(k) = 2\bar{E}(k)$. The important point here is that the linear equations can only display the former behaviour as they do not respect the necessary conservation properties to achieve error-growth saturation.

Referring back to Fig. 3b we note that at $t = 57$ the linear model's $\delta\zeta$ field shows anomalously large energies in the smallest scales, while the nonlinear-model error has saturated at twice the climatological spectrum. In Fig. 3c we present the results of similar parallel integrations of the linear and nonlinear equations, here using four times the resolution (256^2). These are displayed at the earlier time of Fig. 3a. It is noted that, although the low-resolution linear equations approximated well the nonlinear set, the high-resolution linearization performed poorly. Clearly, a limited number of conclusions can be drawn from such a comparison. In absolute terms, the validity timescale of the linearization is determined by the amplitude of the initial perturbation and its choice was completely arbitrary. However, in relative terms we can say that (i) for a given resolution there is a time after which the nonlinear equations saturate at twice the climatological spectrum $\bar{E}_\delta(k) = 2\bar{E}(k)$ and the linear set continues to display spurious growth with $\bar{E}_\delta(k) \sim k^2\bar{E}(k)$ and (ii) that this occurs first at large k. Not surprisingly, the linear model blows up in the small scales, where frequencies are highest and nonlinear saturation occurs first. As a result we conclude that the validity timescale of the tangent linear model is a function of resolution.

The limit of validity of the tangent linear approximation is more stringent when it is made to approximate the evolution of the forecast error, a finite quantity. For instance, an extended Kalman filter makes the approximation to describe the evolution of forecast-error covariances (Jazwinski, 1970; Cohn and Parrish, 1991; Gauthier et al., 1993; Bouttier 1993). In light of the results presented here, it is not surprising that in Gauthier

et al. (1993), the use of the tangent linear model was shown to lead to exponential growth of the forecast-error variance in regions of strong vorticity gradients. This was also acknowledged by Bouttier (1994), who modified the formulation to ensure that the variance did not exceed that of climatology. As explained in Miller et al. (1994), the nonlinearity also has some implications on the formulation of the extended Kalman filter, which is faced with a complicated closure problem similar to that of the EDQNM predictability studies cited above. The tangent linear model is also employed in the characterization of the singular vectors, which describe the initial conditions leading to maximum error growth over a given period of time (Lacarra and Talagrand, 1988; Molteni and Palmer, 1993).

4. The convergence of the method

This section describes the convergence of the adjoint method when complete coverage of perfect data is provided. The truth which we ask the method to seek is the flow evolution (over 10 time units ≈ 3 days) of the basic-state simulation described above. Only one element of difficulty is provided: the starting point of the minimization is a model field taken from a completely different, i.e., decorrelated, time. The error field is defined as $\varepsilon^{(n)} = \zeta^{(t)} - \zeta^{(n)}$, where $\zeta^{(t)}$ is the true field, $\zeta^{(0)}$ the starting point and $\zeta^{(n)}$, the nth iterate of the adjoint method. With respect to the linearization discussed above, it is clear that for the first guess $|\varepsilon^{(0)}|/|\zeta^{(t)}|$ is large. The tangent linear model cannot possibly represent accurately the evolution of the difference between the first guess and the truth, as the neglected terms are initially of the same order as the retained terms. As a result, the functional is not quadratic and multiple minima may exist. In this case, the end point of the minimization could depend on the starting point.

In Fig. 4a we display the energy spectrum of the error field, $E_\varepsilon^{(n)}(k, t)$, which is obtained from (3) using $\tilde{Q}_{k'} = |\hat{\varepsilon}_{k'}^{(n)}|^2/2k'^2$, for several iterations at the beginning of the assimilation period, $t = 0$. The curve labelled "0" shows essentially twice the ensemble-average spectrum since the correlation

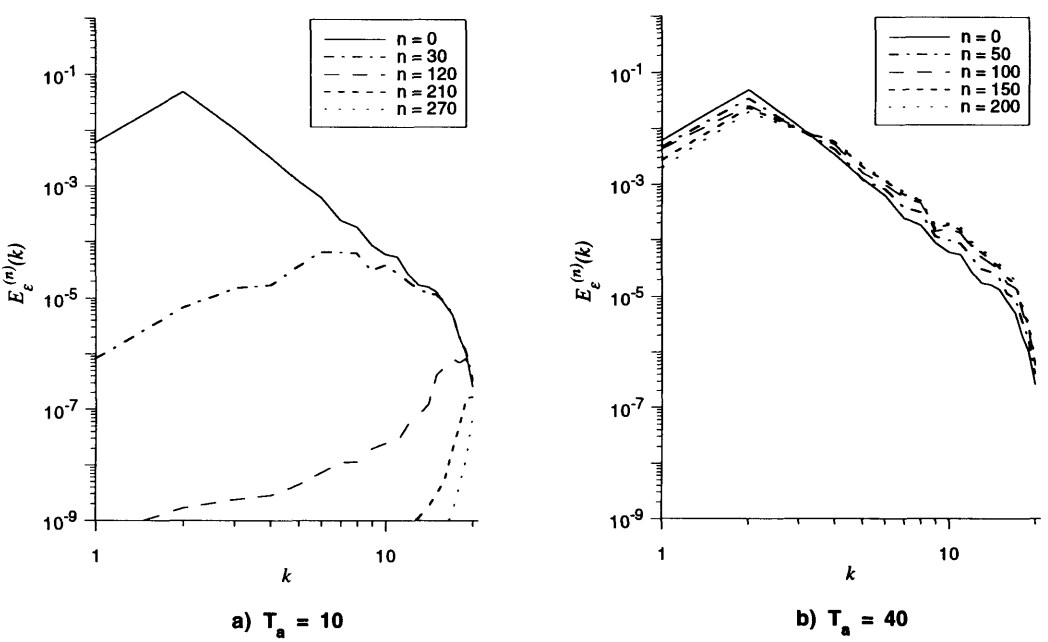

Fig. 4. Initial error spectra, $E_\varepsilon^{(n)}(k, t = 0)$, from the data assimilation experiment with observations at all scales at all times. The various curves represent different iteration numbers, n; (a) assimilation period, $T_a = 10$ model time units, (b) $T_a = 40$ units.

is negligeable. Right from the start, the method converges and the largest scales are adjusted to fit the observations. As the iterations proceed, the minimization focuses its attention on increasingly smaller scales until, after hundreds of iterations, virtually all of the error is in the dissipation range. In Fig. 5 a timeseries of the spatial-average error energy over the assimilation period is displayed for iteration number 270. It can be seen that almost all of the error disappears after a very short time. Since this error at $t = 0$ is confined to the smallest scales, we conclude that the initial small-scale field is only weakly dynamically coupled to the rest of the observational data (see also Thépaut and Courtier, 1991).

For this particular realization, we have tried several similar experiments starting from random initial fields, $\zeta^{(0)}$, and have always obtained convergence for $T_a = 10$. It would seem that, although the adjoint of the linearized equations returns a phase-space direction of the functional's gradient which is unlikely to point to the true minimum the method's use of the fully nonlinear model to explore the functional's variability along this direction is enough to reduce significantly the functional's value. The fact that convergence is obtained implies that the functional is rather smooth in phase space. This conclusion is at variance with the Lorenz-model study of Gauthier (1992), implying that the conclusions reached can depend on several factors. These two models differ in their dynamics, and also by the fact that there are only three real degrees of freedom in the Lorenz model and over a thousand here. It is therefore difficult to compare the length of the assimilation periods in the two studies.

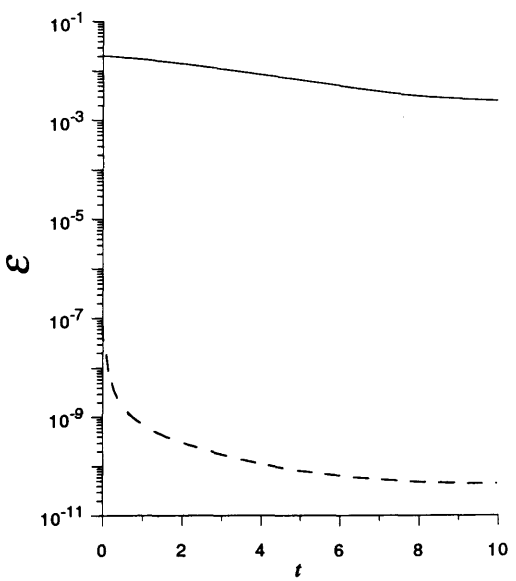

Fig. 5. The evolution of the total error $\mathscr{E} = \int_0^\infty E_\varepsilon^{(n)}(k, t)\, dk$ with time at the last iteration ($n = 270$) for the experiment of Fig. 4; $T_a = 40$ (solid line) and $T_a = 10$ (dashed line).

In order to investigate convergence when the assimilation period is long compared to the nonlinear timescale of the model, we performed an assimilation experiment, as above with full data coverage, over a period of 40 model time units (≈ 13 days). As shown in Fig. 4b, the convergence in the large scales is very slow and, more importantly, the method seems to be diverging in the small scales. The consequence of having an assimilation period that exceeds the decorrelation time at all scales is that the solution at the end of the interval has essentially lost its memory of the initial conditions and this constitutes the ultimate limit of the method. As will be seen below, the nonlinearities lead to secondary minima to which the minimization may converge. Here, it is observed that the minimum attained is consistent with a systematic trend towards a spectrum more shallow than the statistical mean of the model, as in the linear-model simulations discussed in Section 3. In the following, we investigate the more relevant case in NWP where there is no observational information at small scales.

5. Observations with restricted spatial coverage

In this section, we present experiments performed in the same manner as those of Section 4 with the exception of the spectral coverage of the observations. We now supply observations of the modes with $k \leqslant k_c = 4.5$ only. A more realistic experiment would have been to provide accurate observations on a coarse grid. This would truncate the observational information at k_c. However, the small-scale signal would have been aliased on to the larger scales. Our approach was to truncate the true spectral coefficients, thereby removing observational error associated with aliasing, in order to focus on the nonlinearity. The method is asked to supply the small-scale detail in accord with the

observations of the large-scale flow. We consider assimilation periods, T_a, ranging from 5 to 40 model time units (from 2 to 13 days). These are rather long periods compared to experiments such as those presented in Thépaut and Courtier (1991) and Thépaut et al. (1993) who used an assimilation window of 24 h, but their aim was to test the method within an operational setting to demonstrate the feasibility of its implementation. This point is discussed further below.

In Fig. 6 we plot the error energy spectra $E_\varepsilon^{(n)}(k, t)$ at $t = 0$ for several different iterations,

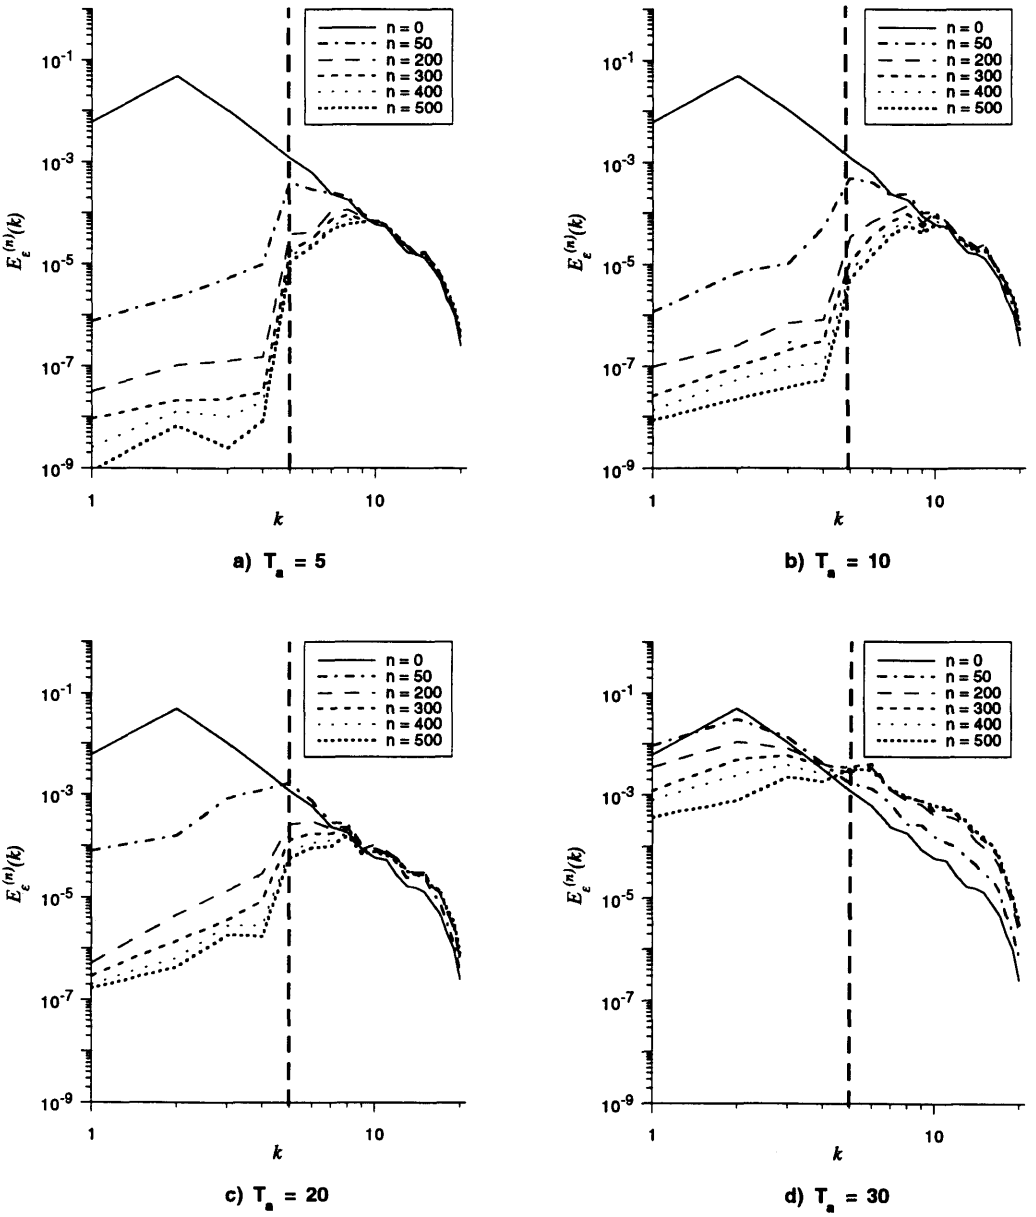

Fig. 6. As in Fig. 4 for experiments where observations are supplied at large scales only: (a) $T_a = 5$, (b) $T_a = 10$, (c) $T_a = 20$ and (d) $T_a = 30$.

n, for some of the experiments. For the shorter assimilation periods the method is essentially unable to reduce the error in the unobserved small scales at $t = 0$. In this range, the first-guess field remains almost unaltered, implying that the gradient of the functional is weak in the phase-space directions corresponding to the small scales. For longer assimilation periods convergence is slow even in the larger scales and a spurious approach to $E_\varepsilon(k) \sim k^2 E(k)$ is noted in the unobserved small scales. In Fig. 7 are displayed the error spectra $E_\varepsilon^{(n)}(k, t)$ at $n = 500$ iterations for the

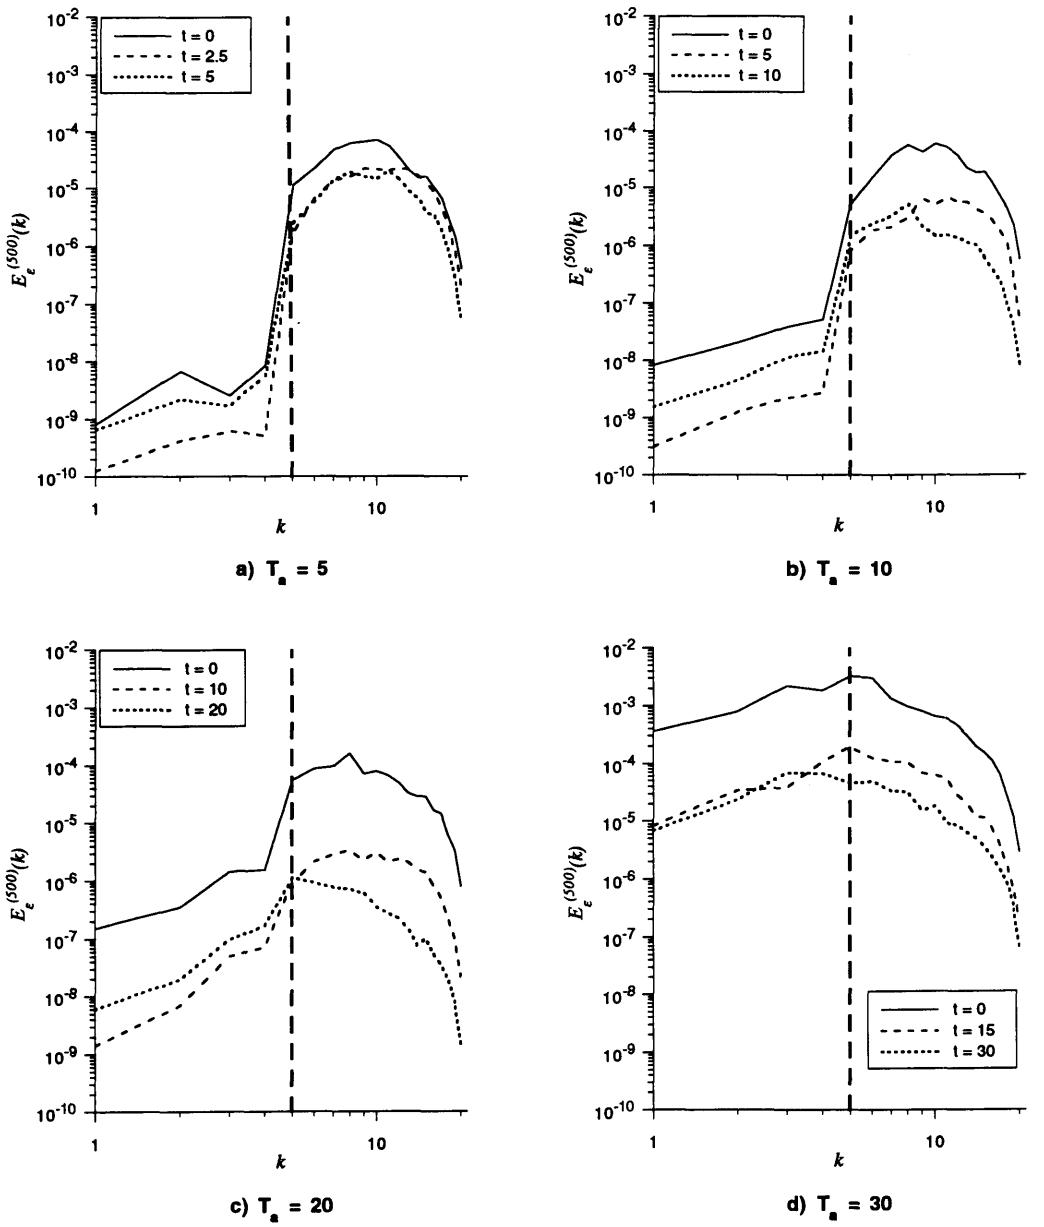

Fig. 7. Error spectra at the beginning, middle and end of the assimilation period for the experiments of Fig. 6 at $n = 500$: (a) $T_a = 5$, (b) $T_a = 10$, (c) $T_a = 20$ and (d) $T_a = 30$.

beginning, middle and end of the assimilation periods for the same T_a's as in Fig. 6. It can be seen that the small-scale error is largest at $t = 0$ and that, during the course of the assimilation period, it decreases. The rapidity with which it does so depends on the length of the assimilation period. The experiment with $T_a = 20$ time units shows a two-order-of-magnitude decrease in small-scale error at the end of the period. At the same time, the large-scale error actually increases, reflecting the two-way cascade of predictability error. In Fig. 8, the $n = 500$ vorticity field for this experiment is plotted in physical space along with the true field, at the three different times referred to in Fig. 7c. The field at $t = 0$ (Fig. 8a) does not at all look like a model vorticity field, since at this time, the large

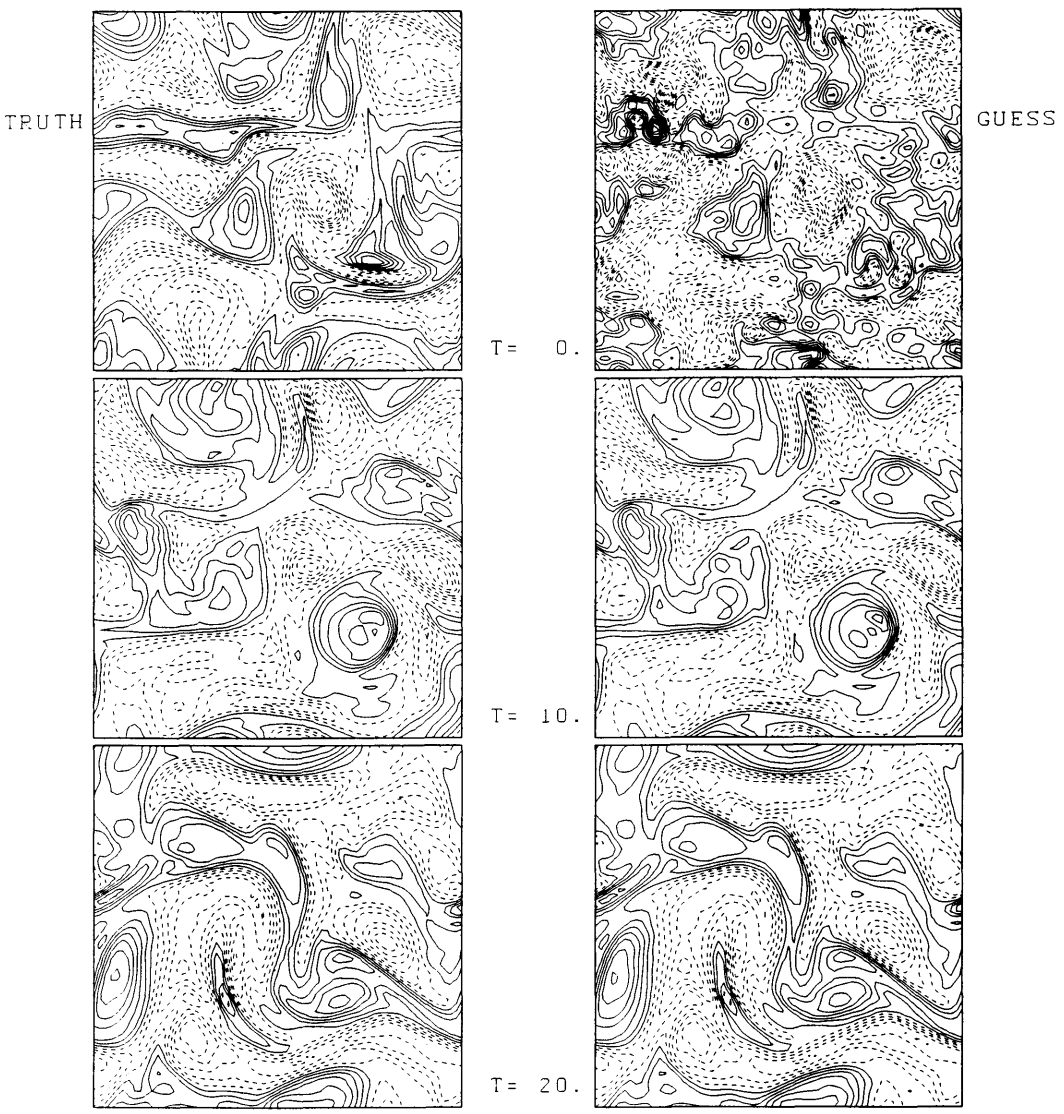

Fig. 8. Vorticity fields for the experiment of Fig. 7 with $T_a = 20$ at the beginning, middle and end of the assimilation period; left-hand panels are the truth, right-hand panels were obtained from the adjoint method using large-scale observations only; (a) $t = 0$, (b) $t = 10$, (c) $t = 20$.

scales reflect the observations, while the small scales are taken from the completely decorrelated initial guess. Not surprisingly, the mixture displays a lack of coherence. Although discrepancies are noted at $t = 10$, the error has been considerably reduced and the field is much more coherent. This trend continues and at $t = 20$ (Fig. 8c) the assimilated field is quite accurate.

It is important to ask the extent to which it suffices to have the initial large-scale flow treated accurately in order to obtain reasonable fields at $t = 20$. To investigate, we performed a set of control runs which resemble 3D data assimilation in that observational information is injected at one instant only. In Fig. 9 we plot the covariance between the forecast and the truth (normalized by the truth)

$$\mathscr{C}(t) = \frac{\langle \hat{\phi}, \hat{\zeta}^{(t)} \rangle_Z}{\langle \hat{\zeta}^{(t)}, \hat{\zeta}^{(t)} \rangle_Z},$$

as a function of time, where ϕ is, in turn, $\zeta^{(500)}$ and the various control runs. The first of these was initialized with Fourier coefficients set to those of the true field at the observed scales ($k < k_c$) and set to zero for $k \geq k_c$. The second has perfect data for $k < k_c$ and the true Fourier amplitudes at $k \geq k_c$, but with a phase scrambling yielding zero correlation with the truth at these scales. Lastly, we supplied perfect data for $k < 8.5$ such that $\mathscr{C}(0)$ was approximately equal to that obtained with $\zeta^{(500)}$,

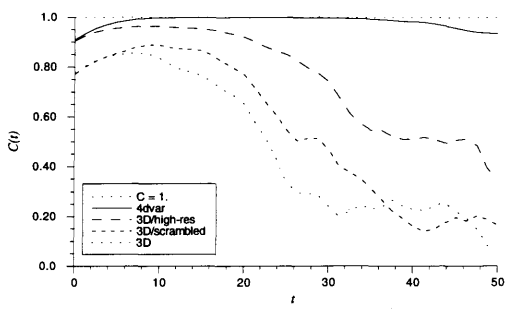

Fig. 9. Covariance between the forecast and the truth (normalized by the truth) as a function of time; 4Dvar: experiment of Fig. 8 with $T_a = 20$ at $n = 500$, 3D: control run with Fourier coefficients initialized perfectly for $k < k_c$ and set to zero for $k \geq k_c$, 3D/scrambled: control run with Fourier coefficients initialized perfectly for $k < k_c$ and with scrambled phases for $k \geq k_c$, 3D/high-resolutions: control run with Fourier coefficients initialized perfectly for $k < 8.5$ and set to zero for $k \geq 8.5$.

corresponding to a 3D assimilation benefitting from improved spatial coverage of observational data as compared with the 4D assimilation. It can be seen in Fig. 9 that the adjoint method significantly outperforms all of the control runs, indicating that important dynamical information has been transmitted to wavenumbers $k \geq k_c$ in the assimilation experiment.

In Fig. 10, we display the total error energy, $\mathscr{E}_T = \int_0^\infty E_\varepsilon(k)\,dk$, the large-scale error energy, $\mathscr{E}_L = \int_0^{k_c} E_\varepsilon(k)\,dk$ and the small-scale error energy, $\mathscr{E}_S = \int_{k_c}^\infty E_\varepsilon(k)\,dk$ at the end of the assimilation period (iteration = 500) as a function of T_a. The terms large- and small-scale are here being used merely to indicate the presence or absence of data. Our observational information was truncated at $k_c = 4.5$, which corresponds to 1600 km. We propose to use $\mathscr{E}_S$ as a means to assess the method's ability to supply small-scale detail in accord with large-scale observations. Recall that it is from the end of the assimilation period that one normally commences a forecast-model integration. In Fig. 10 it can be seen that $\mathscr{E}_S$ displays a minimum at $T_a = 20 \approx 2.2\tau_{NL}$, which corresponds to 6 days. Significantly poorer performance is noted at both

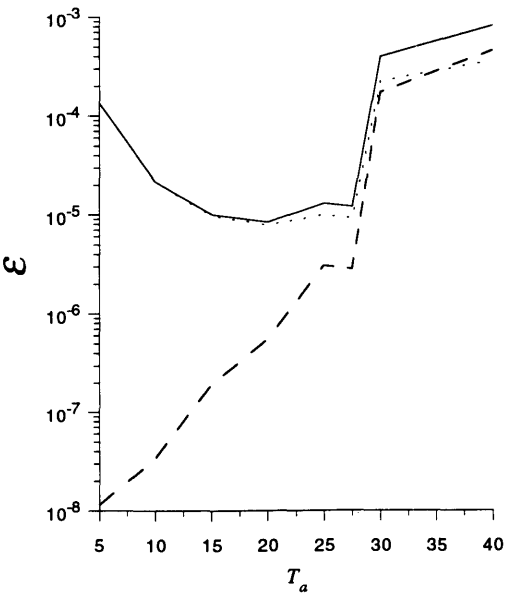

Fig. 10. Total error, $\mathscr{E}_T$ (solid line), large-scale error, $\mathscr{E}_L$ (dashed line) and small-scale error, $\mathscr{E}_S$ (dotted line) at the end of the assimilation period as a function of the assimilation period.

shorter and longer assimilation periods. For T_a small, the transfer has not had time for significant spectral redistribution of observational information, while for T_a large, the spurious small-scale behaviour of the linear model has caused a relaxation to a very shallow spectrum.

6. Observations with restricted temporal coverage

In this section, we consider 4Dvar experiments where initial conditions are adjusted to fit an observation with full spatial coverage, but inserted at just one instant, $t = T_a$. As discussed in Section 3, infinitesimal differences in initial conditions cause two realizations eventually to become completely decorrelated as neighbouring trajectories in phase space diverge. The success of the present experiments would seem to depend on the ratio of T_a to the decorrelation, or predictability, timescale τ_{NL}. If $T_a/\tau_{NL} \ll 1$ at all scales, then the quality of the fit to the data at the end of the assimilation period is a relatively smooth function of the control variables (i.e., the initial conditions). On the other hand if T_a is much larger than τ_{NL}, the situation appears hopeless, since the system at $t = T_a$ will have essentially lost its memory of the control variables. For intermediate values, smaller scales will have become decorrelated, while progress may still be possible at eliminating large-scale error. Experiments of this type are now presented and are similar to those of Thépaut and Courtier (1991) (referred to as the inversion problem). The aim is to test the ability of the method to recover the initial conditions from the full model state at the final time. Since only one observation is involved, the functional reduces to a measure of the error at $t = T_a$:

$$J(\zeta_0) = \langle \zeta(T_a) - \zeta^{(t)}(T_a), \zeta(T_a) - \zeta^{(t)}(T_a) \rangle, \quad (10)$$

while the initial error is $\langle \zeta_0 - \zeta^{(t)}(0), \zeta_0 - \zeta^{(t)}(0) \rangle$. It is convenient to rewrite (10) in terms of $\Delta\zeta(t) = \zeta(t) - \zeta^{(t)}(t)$ to obtain

$$J_2(\Delta\zeta_0) = \langle R(T_a, t_0)\,\Delta\zeta_0, R(T_a, t_0)\,\Delta\zeta_0 \rangle + \text{higher-order terms in } \Delta\zeta_0. \quad (11)$$

If the starting point of the minimization is such that $\Delta\zeta_0$ is sufficiently small to be within the limit of validity of the tangent linear approximation, then (11) implies that the functional is approximately quadratic and consequently that the minimum is locally unique. On the other hand, as the validity timescale of the tangent linear model is exceeded, J can no longer be considered quadratic and multiple minima become a distinct possibility.

Fig. 11 summarizes the results of experiments in which the assimilation interval T_a is gradually increased, by displaying the spectrum of the error as a function of iteration. It is interesting to observe that the lack of convergence occurs more rapidly than in Section 5. At $T_a = 10$, the error has been slightly reduced in the large scales but is increased in the smaller scales. Since the method seems to have converged, but not to the absolute minimum $\zeta^{(t)}(0)$, the functional was evaluated along the phase-space direction joining $\zeta^{(t)}(0)$ to the point of convergence ζ_0^*. This reduces J to a function of a single variable α defined as

$$f(\alpha) = J(\zeta^{(t)}(0) + \alpha(\zeta_0^* - \zeta^{(t)}(0))).$$

By definition, the spectrum of the change in the initial conditions has the form of the error spectrum at the final iteration as shown in Fig. 11. Fig. 12 shows $f(\alpha)$ for $-2 \leqslant \alpha \leqslant 2$ as a function of the distance with respect to $\zeta^{(t)}(0)$. At $T_a = 2.5$, the functional shows a single minimum, indicating that convergence has not yet saturated. But for $T_a = 5$, 6.25 and 10, the minimization converged to a secondary minimum of $f(\alpha)$ and further iterations may not have improved the quality of the analysis. It is important to stress that this is only a strong indication of a true secondary minimum in the complete phase space, since the gradient, although very weak at convergence, may not be exactly zero. Fig. 12 shows also that the distance between the secondary and absolute minima increases with T_a, and the fact that the value of $J(\zeta_0^*)$ is also increasing with T_a, indicates that the fit at the end of the assimilation period is degrading as well. This is corroborated by Fig. 13, which shows the spectrum of error $\zeta^*(t)$ (with $\zeta^*(0) = \zeta_0^*$) at several times. Fig. 11 at $T_a = 5$ and 6.25 shows that the change in the initial conditions required to find the secondary minimum has a larger component in the small scales, but when $T_a = 10$, the initial error is spread over a wider

range of scales. This is to be expected given the results presented above, demonstrating that the tangent linear approximation ceases to be valid sooner in the small scales. This translates to a more nonlinear behavior of J in directions with large small-scale components. At $T_a = 10$, the large-scale component is equally important, as the limit of validity of the tangent linear approximation is being approached at all scales.

In the previous experiments, assimilations were

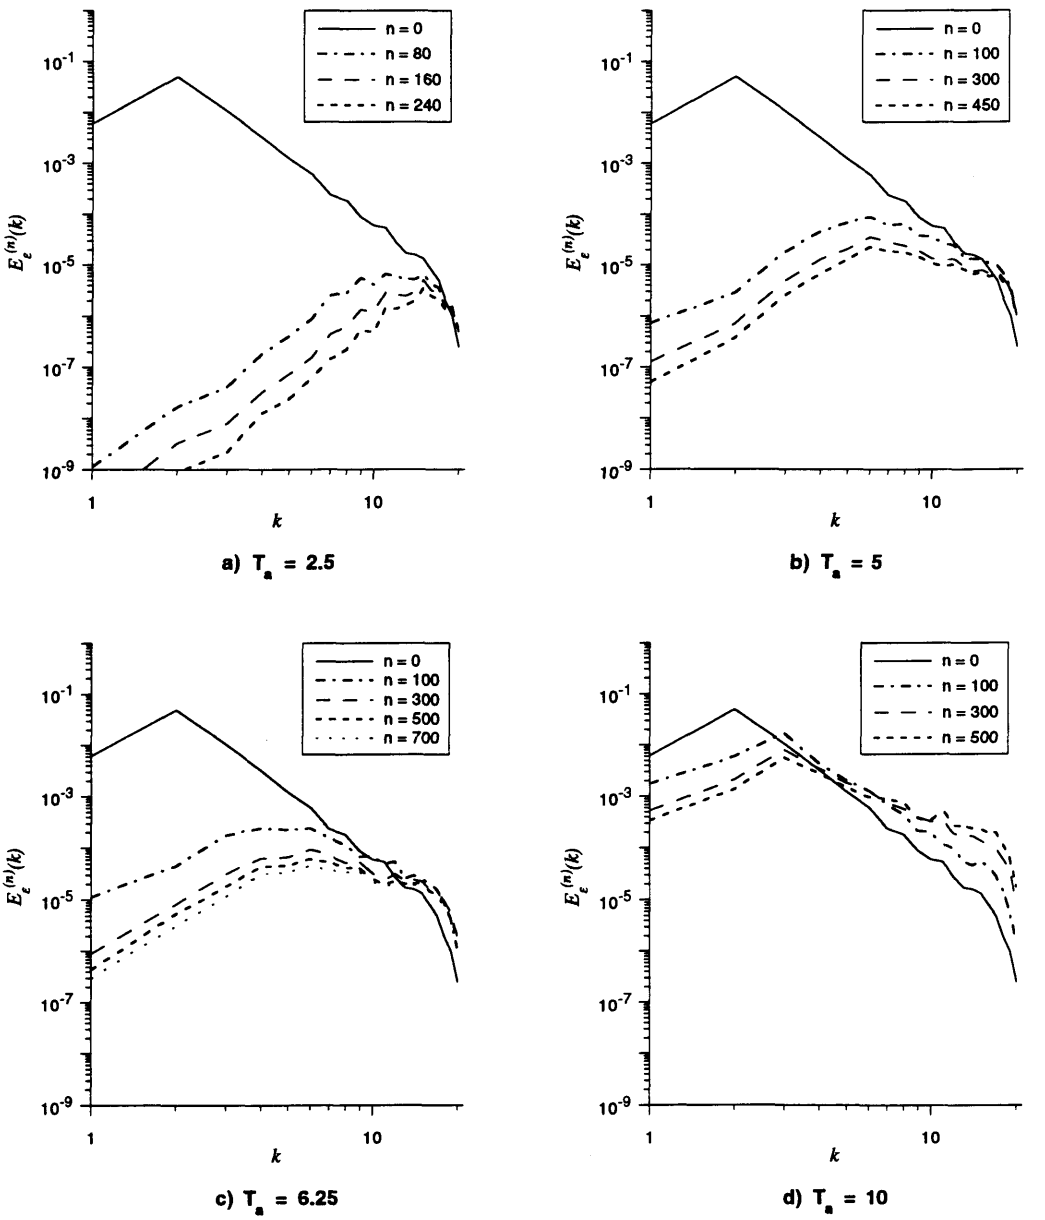

Fig. 11. As in Fig. 4 for experiments where observations at all scales are supplied only at $t = T_a$; (a) $T_a = 2.5$, (b) $T_a = 5$ (c) $T_a = 6.25$ (d) $T_a = 10$.

carried out for longer periods of time before the decorrelation time was reached. There, observations were provided at every timestep, so that even though the assimilation intervals were long, observations close to the initial time were available. This is not the case here, resulting in an earlier convergence failure. The impact of the spacing in time between observations would appear to be as important as the spatial resolution of the observational network. This makes it difficult to make

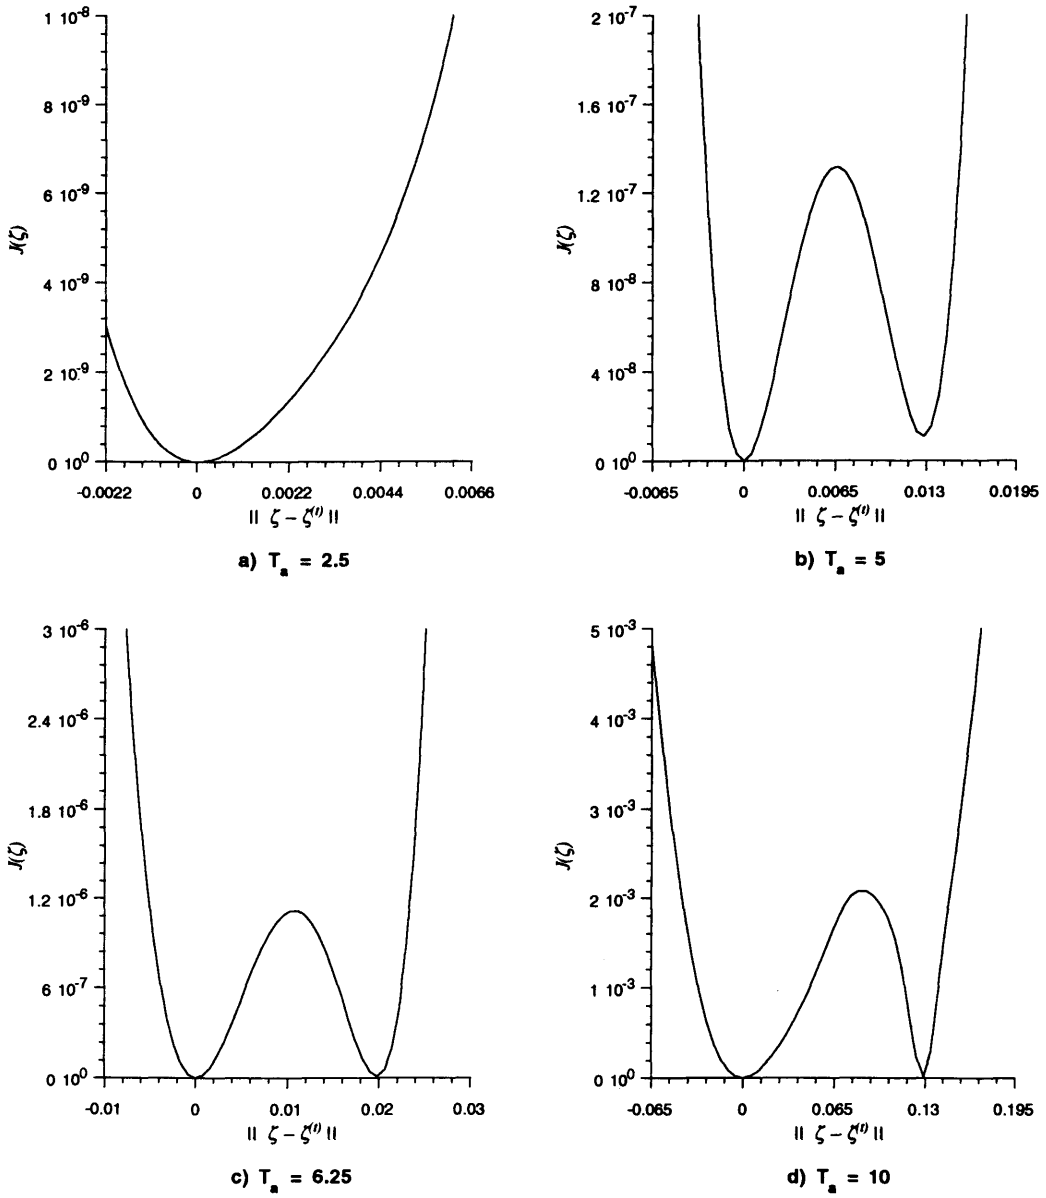

Fig. 12. Representation of the functional along the segment joining $\zeta^{(t)}(0)$, the true minimum and the point to which the minimization converged for assimilation with a single observation at $t = T_a$ for (a) $T_a = 2.5$, (b) $T_a = 5$, (c) $T_a = 6.25$ and (d) $T_a = 10$.

general statements about the convergence of the method, which also depends on the local dynamics of the flow (e.g., presence of local instabilites in phase space). However, there is a common factor to all experiments in that a large number of iterations were necessary to ensure convergence of the minimization. It is therefore relevant to examine common preconditioning techniques to see whether they can speed up the minimization. This is the object of Section 7.

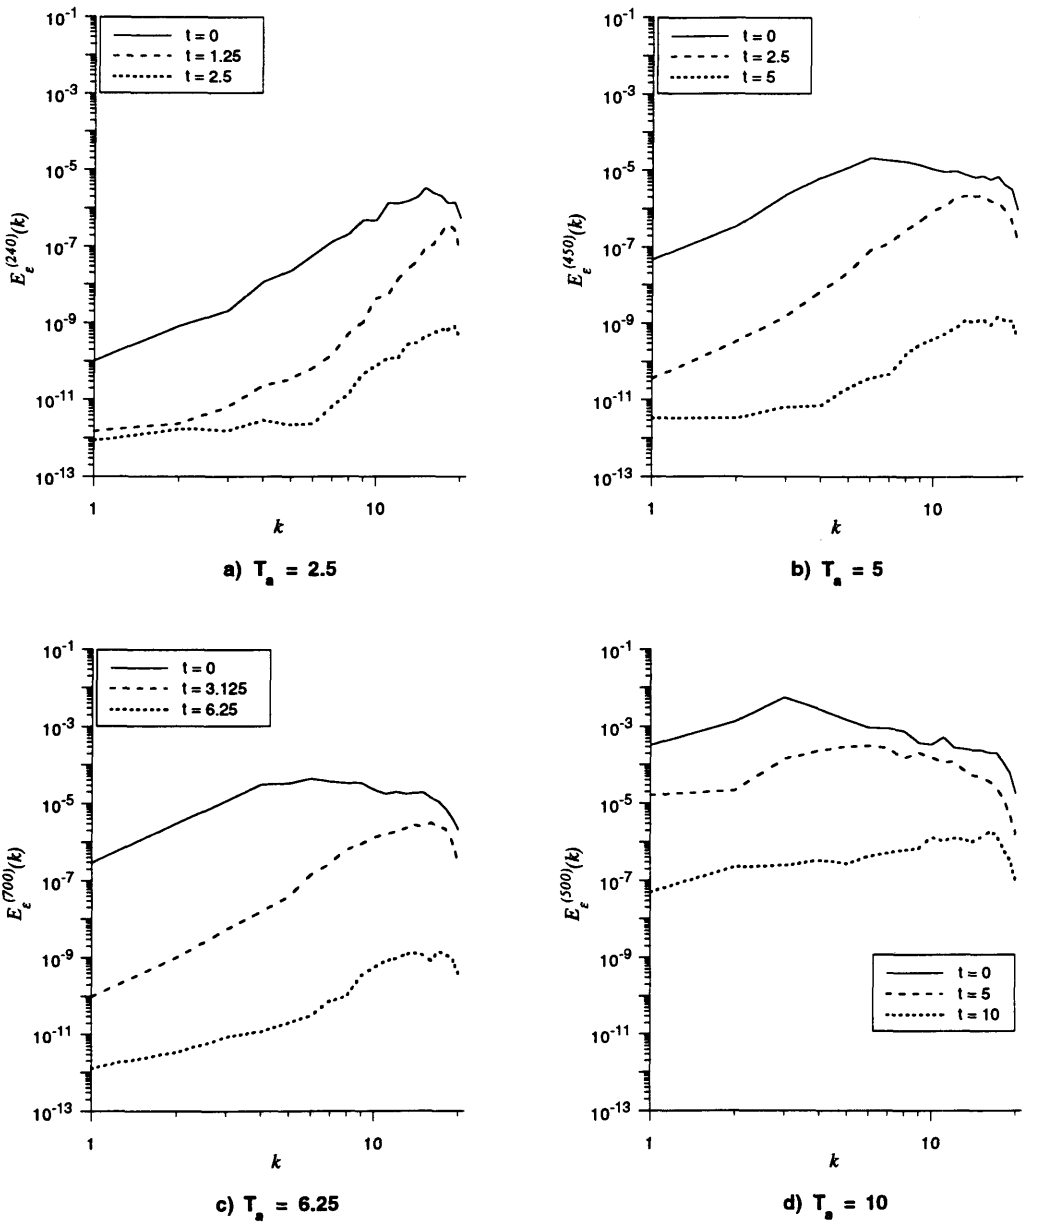

Fig. 13. Error spectra at the beginning, middle and end of the assimilation period for the experiments of Fig. 11. (a) $T_a = 2.5$, (b) $T_a = 5$ (c) $T_a = 6.25$ (d) $T_a = 10$.

7. Preconditioning of the minimization

For a quadratic functional of the form $J(x) = x^T \boldsymbol{Q} x + c^T x$, a slow convergence is related to the conditioning of the Hessian matrix $\boldsymbol{Q}$. If $\boldsymbol{Q}$ were known and invertible, it would be possible to speed up the convergence by redefining the inner product as $\langle f, g \rangle = f^T \boldsymbol{Q} g$ in which case the minimization can be made to converge in a single iteration. This assumes that the full Hessian matrix is known, a situation that is rarely encountered in large-scale minimization problems. In practice, information from previous iterations is used to approximate the Hessian matrix for use in the preconditioning, by modifying the inner product or alternatively, through a change of variables. The former approach is taken here, since the algorithm we used allows the user to define the inner product.

For a nonlinear problem, the preconditioning is based upon a diagonal scaling that uses only an approximation to the diagonal of the Hessian matrix and this is usually all that can be done (Gill and Wright, 1981). For our problem, it is possible, although costly, to compute explicitly the Hessian and from it, to extract the diagonal in order to define the inner product. Minimization experiments with three different inner products were conducted to evaluate their efficiency in improving the convergence. In all cases however, the functional is defined as the energy of the error. If $\langle f, g \rangle = f^T \boldsymbol{D} g$ stands for the inner product with $\boldsymbol{D}$ a diagonal matrix, then two of them were defined with respect to the energy and enstrophy norms, so that $\boldsymbol{D}_k = \boldsymbol{D}_{1k} = 1/k^2$ and $\boldsymbol{D}_k = \boldsymbol{D}_{2k} = 1$, respectively. The third was defined as $\boldsymbol{D} = \boldsymbol{D}_3 = \mathrm{diag}(\boldsymbol{H})$, $\boldsymbol{H}$ being the Hessian, computed for the case where only one observation was provided at the end of the assimilation as in Section 6. In that case, in the vicinity of the true minimum ζ_0, $\nabla J(\zeta_0) = 0$, and

$$J(\zeta_0 + \Delta\zeta_0) - J(\zeta_0) \equiv J(\Delta\zeta_0) = \langle \Delta\zeta(T_a), \Delta\zeta(T_a) \rangle$$
$$\approx \langle R^*(T_a, t_0)\, R(T_a, t_0)\, \Delta\zeta_0, \Delta\zeta_0 \rangle$$
$$+ \text{higher-order terms.} \qquad (12)$$

The Hessian matrix $\boldsymbol{H}$ can then be obtained by integrating the tangent and adjoint models a number of times equal to the number of real degrees of freedom in the model. In view of the results presented earlier, T_a was chosen to correspond to 0.5 time units. Due to the nature of the dynamics, the Hessian matrix is not diagonal but we nevertheless have a quadratic functional by neglecting the higher-order terms in (12). This quadratic problem involving a non-diagonal matrix was used to test the preconditioning. Fig. 14 presents the variation of the cost function during the minimization for the three cases. Inner products based on the energy and the Hessian clearly have a positive impact: a reduction of three orders of magnitude in just 2 iterations, while 10 were needed when the enstrophy was used to define the inner product. As the minimization continued, the improvement was best with the inner product based on the enstrophy, which put more emphasis on the larger scales of the flow.

After 30 iterations the minimization brought us to a point ζ_0 and if the preconditioning were perfect, the gradient at that point should be parallel to the vector $\Delta\zeta_0 = \zeta_0^{(i)} - \zeta_0$ in phase space. Fig. 15 compares the spectrum of this vector (solid line) to that of the gradient (thick dotted line). Even though an inner product based on enstrophy ends up giving the best results in terms of the minimization, the gradient is not very well correlated with the correct direction. Choosing the energy leads

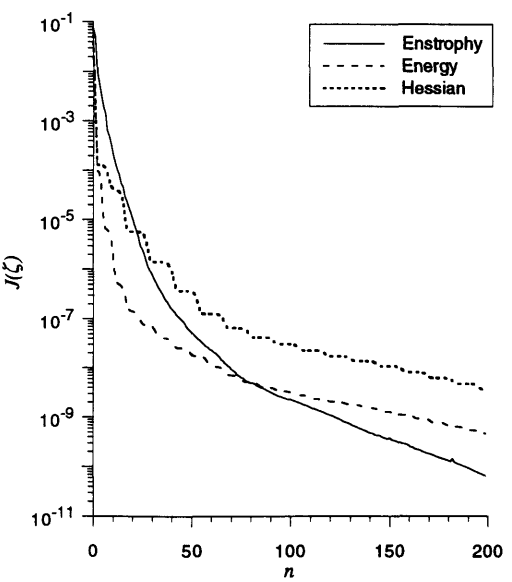

Fig. 14. Variation of the value of the functional as a function of iteration for experiments using an inner product defined with respect to enstrophy (solid line), energy (dashed line) and the diagonal of the Hessian matrix (dotted line).

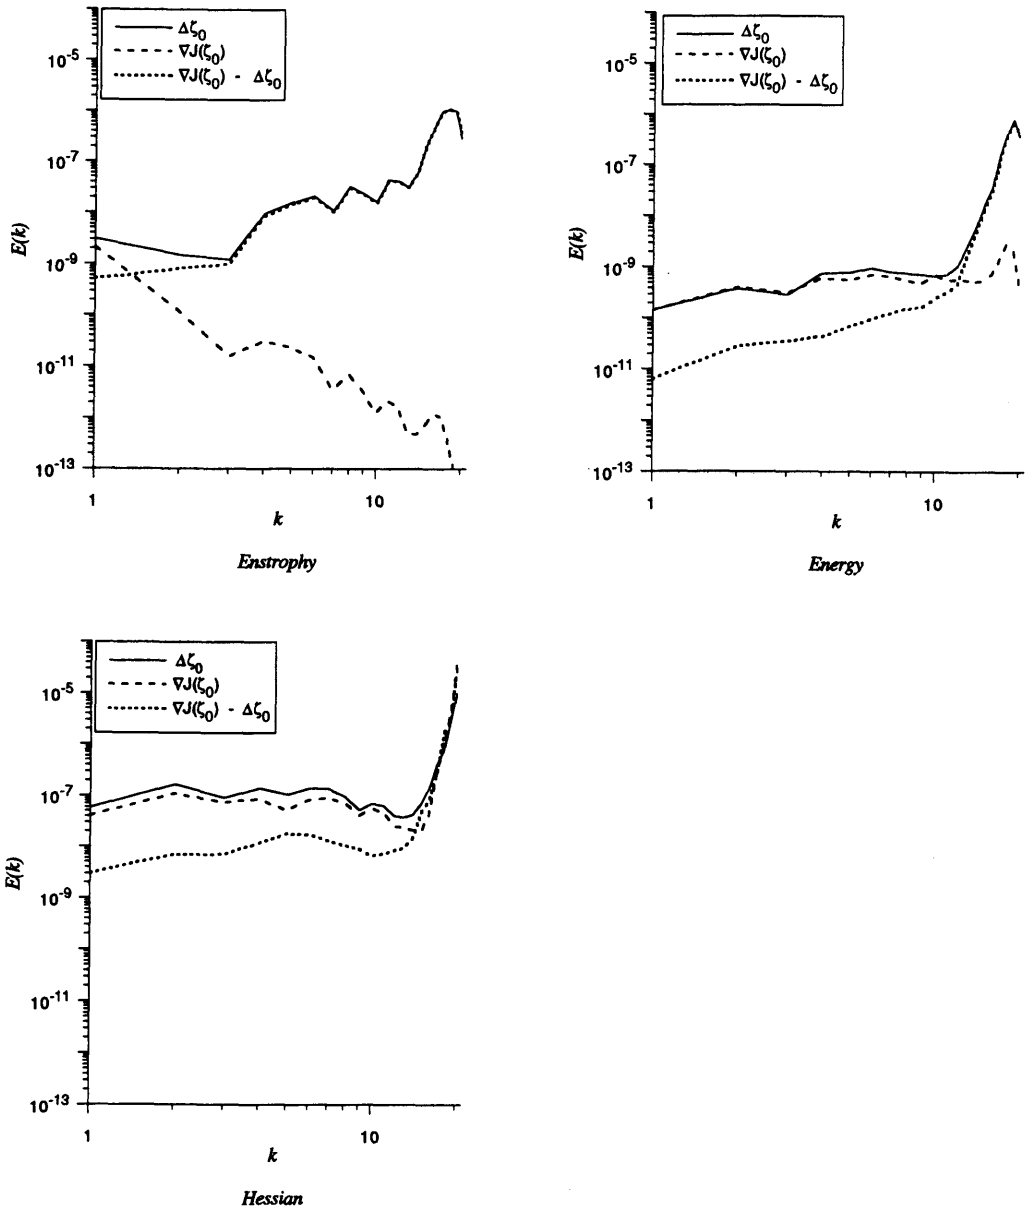

Fig. 15. Spectra of the direction between the point in phase space reached by the minimization after 30 iterations ($\Delta\zeta$, solid line), the gradient at that point (∇J, dashed line) and of the difference between the two ($\nabla J - \Delta\zeta$, dotted line).

to better agreement, except in the small scales. Finally, choosing the diagonal of the Hessian improves the agreement in the small scales. However, a comparison of spectra considers only amplitudes and disregards differences in phase. It is possible to obtain some information about these differences by considering the spectrum of the difference between the gradient and the true direction (thin dotted line in Fig. 15). It is clear that what appeared to be a good match in the

Tellus A/47/5:2-7

small scales, when diag($\boldsymbol{H}$) was used to define the inner product, has some important differences in phase that decorrelate the gradient with the true direction.

What is to be concluded from this? First, if the Hessian had been diagonal, then the preconditioning would have been a lot better than observed. Given that the functional is defined as the energy of the error, the Hessian is $J'' = R^* \boldsymbol{D}_1 R$ and, when $T_a \ll \tau_{NL}$, R is close to the identity matrix. Then the Hessian and energy inner products become identical, while using the enstrophy yields a gradient that goes as $1/k^2$, which explains the slope observed in the spectrum of ∇J for the enstrophy case in Fig. 15. The consequence is that using the enstrophy produces a gradient that puts less emphasis on the small scales than the other two, resulting in improved convergence when the functional measures the total energy. This is not so surprising given that the bulk of the energy is contained in the large scales.

8. Conclusions

Our goal was to test the ability of the adjoint method to fill in small-scale detail in accord with large-scale observational information in a context of fully-developed turbulent dynamics. In order to test the utility of the linearized equations, we considered only the simplest possible data assimilation experiments. We have introduced neither observational nor model error, focussing only on the critical limitations imposed by the method's dependence on linearizations, rather than consider realistic assimilations in an operational setting.

The synoptic-scale eddy turnover, or advective, timescale $\tau_{NL} \approx 3$ days is of the order of the limit of validity of the tangent linear model. Although this statement is undeniably true, it provides only an upper limit. In turbulent flows such as the atmosphere, the validity timescale of the linearization is a function of lengthscale. We have demonstrated that for a given resolution, the linearization becomes inaccurate first in the smallest model scales. This is manifested by a spurious approach to a shallow spectrum, reflecting the linearized predictability-error growth's adherence to basic-state gradients, even when the true error growth has saturated in the fully nonlinear equations. Consequently, for a given assimilation period, T_a, there is a scale below which the linearization fails. This failure occurs earlier when the approximation is applied to a finite quantity such as the evolution of the forecast error in the extended Kalman filter, or to the characterization of the singular modes, which describe the initial conditions leading to maximum error growth over a finite time interval (Lacarra and Talagrand, 1988). These modes have been used in predictability studies to improve Monte-Carlo simulations (Molteni and Palmer, 1993).

It was also demonstrated that if only the large scales are observed, T_a must be large enough for significant information to be transferred downscale in order to fill in the fine-scale detail. The question we set out to answer concerned the feasability of the adjoint method given these two constraints. Interestingly, our best assimilation experiment with only large-scale data had a duration of over $2\tau_{NL} \approx 6$ days. The linearization was clearly not accurate over this timescale and alarming spurious excitation of the small scales was noted initially in the assimilation period (Fig. 6). This error was almost completely absent at the end of the assimilation period and the subsequent forecast of approximately ten days duration was of excellent quality when compared with various control runs (Fig. 9). Somewhat surprisingly, even though the initial fields did not respect the model's climatology, the subsequent evolution collapsed very near to the correct trajectory on the model's attractor by the end of the period. Therefore, useful information must have been transferred to the small scales by the adjoint method, even though the initial small-scale analysis was too energetic. In order to translate these successful results to the atmosphere, we must use an assimilation period measured in days and not hours. When the assimilation period was many times larger than τ_{NL} at even the largest scales, the results were not as good, suggesting that there is an optimal resolution at which the downscale transfer of large-scale observational information reaches a minimum scale, before the exponential blow-up of the linear model inhibits convergence. Below that length scale the method is not effective.

The relationship between turbulent dynamics and multiple minima of the penalty functional was clarified in Section 6. Here, we examined the inversion problem, i.e., one complete set of spatial observations was provided only at the end of the

assimilation period. The ergodic nature of turbulence implies that the system progressively loses its memory of the initial conditions, beginning in the small scales. As T_a is increased, the functional therefore becomes an increasingly unsmooth function of the initial conditions, implying multiple minima and likely affecting the convergence of the adjoint method. Here, we provided substantial evidence for the existence of multiple minima at large T_a (Fig. 12). As T_a is increased, increasingly large scales become uninvertable. As a result it can be concluded that the temporal resolution of the data network is as important an issue as the spatial resolution, in determining the minimum scale that can be accurately initialized by the adjoint method. The determination of this scale therefore depends sensitively on the details of the observational network.

It should be kept in mind that we examined only barotropic quasi-geostrophic dynamics characteristic of the largest atmospheric scales. Extrapolation to smaller scales, where ageostrophic effects such as vortex tube stretching may be important, is not justified. At these scales spectra are less steep and timescales decrease more rapidly at smaller lengthscales. For this reason, successful assimilations at $T_a > \tau_{NL}$ may be significantly more difficult than in the strictly two-dimensional turbulence considered here. This fact, combined with less complete data coverage, would appear to make meso- and small-scale 4Dvar rather challenging.

In all of the experiments described we employed hundreds of iterations of the adjoint method. Since this is far too costly for NWP, we also examined the performance, in a turbulent setting, of different preconditioning strategies. It was concluded that convergence was improved if more emphasis was put on the large-scale component of the gradient. This is an argument in favour of using the incremental approach of Courtier et al. (1994), who suggested reducing the resolution of the adjoint model in the computation of the gradient.

9. Acknowledgements

The authors would like to thank C. Chouinard, P. Courtier, O. Métais and an anonymous reviewer for their comments on the manuscript. They are also grateful to Jean-Charles Gilbert of the Institut National de Recherche en Informatique et en Automatique (INRIA), France who made his minimization code (module M1QN3) available to us.

10. Appendix. Derivation of the adjoint model equations

For the sake of completeness, the arguments and derivations presented in LDT86 and TC87 are briefly recalled in the context of our β-plane model. Setting the problem in a moving reference frame $x' = x - U_0 t$, $t' = t$, the tangent linear model is written as

$$\frac{\partial}{\partial t'}\delta\zeta = -J(\Delta^{-1}\delta\zeta, \zeta + \beta y) - J(\Delta^{-1}\zeta, \delta\zeta) - \mathscr{D}(\delta\zeta) \equiv \mathscr{L}(\zeta)\,\delta\zeta \tag{A.1}$$

with $\Delta \equiv \nabla^2$. The linearity of this model implies that $\delta\zeta(t) = R(t, t_0)\,\delta\zeta_0$. Defining the adjoint model as

$$\frac{\partial}{\partial t'}\delta^*\zeta = -\mathscr{L}^*(\zeta)\,\delta^*\zeta \tag{A.2}$$

with

$$\langle \delta^*\zeta, \mathscr{L}(\zeta)\,\delta\zeta\rangle = \langle \mathscr{L}^*(\zeta)\,\delta^*\zeta, \delta\zeta\rangle$$

for a given inner product, it then follows that

$$\frac{\partial}{\partial t'}\langle \delta\zeta, \delta^*\zeta\rangle = \left\langle \frac{\partial \delta\zeta}{\partial t'}, \delta^*\zeta \right\rangle + \left\langle \delta\zeta, \frac{\partial \delta^*\zeta}{\partial t'} \right\rangle = 0$$

implying that

$$\langle \delta\zeta_k, \delta^*\zeta_k\rangle = \langle \delta\zeta_0, \delta^*\zeta_0\rangle. \tag{A.3}$$

The adjoint model is linear and therefore, its solutions can also be written as $S(t, t_0)\,\delta^*\zeta_0$ and (A.3) becomes

$$\langle R(t_k, t_0)\,\delta\zeta_0, \delta^*\zeta_k\rangle = \langle \delta\zeta_0, R^*(t_k, t_0)\,\delta^*\zeta_k\rangle \equiv \langle \delta\zeta_0, S(t_0, t_k)\,\delta^*\zeta_k\rangle.$$

Consequently, $R^*(t, t_0) = S(t_0, t)$ and, to apply the adjoint of the resolvent, it suffices to integrate (A.2) backward in time from the final time t to the initial time.

As mentioned in Section 2, the energy norm is used to define the inner product as

$$\langle \zeta_1, \zeta_2 \rangle = \int_S \nabla \Delta^{-1} \zeta_1 \cdot \nabla \Delta^{-1} \zeta_2 \, dS$$

$$\equiv - \int_S \zeta_2 \, \Delta^{-1} \zeta_1 \, dS \equiv - \int_S \zeta_1 \, \Delta^{-1} \zeta_2 \, dS,$$

the integral being over the model's domain. With this definition, it is straightforward to show that

(i) $\langle \zeta_1, \Delta \zeta_2 \rangle = \langle \Delta \zeta_1, \zeta_2 \rangle$,

(ii) $\langle \zeta_1, \Delta^{-1} \zeta_2 \rangle = \langle \Delta^{-1} \zeta_1, \zeta_2 \rangle$,

and these operators are self-adjoint. On the other hand, the Jacobian operator is such that

$$\int_S J(f, g) \, h \, dS = \int_S f J(g, h) \, dS.$$

These results are used to show that

$$\langle \delta^* \zeta, \mathscr{L}(\zeta) \, \delta \zeta \rangle \equiv - \int_S (- J(\Delta^{-1} \delta \zeta, \zeta + \beta y) - J(\Delta^{-1} \zeta, \delta \zeta) - \mathscr{D}(\delta \zeta)) \, \Delta^{-1} \, \delta^* \zeta \, dS$$

$$= - \int_S (- \mathscr{D}(\delta^* \zeta) - \Delta J(\Delta^{-1} \, \delta^* \zeta, \Delta^{-1} \zeta) - J(\zeta + \beta y, \Delta^{-1} \, \delta^* \zeta)) \, \Delta^{-1} \, \delta \zeta \, dS$$

$$= \langle \mathscr{L}^*(\zeta) \, \delta^* \zeta, \delta \zeta \rangle$$

and therefore,

$$\mathscr{L}^*(\zeta) \, \delta^* \zeta = - \mathscr{D}(\delta^* \zeta) - \Delta J(\Delta^{-1} \, \delta^* \zeta, \Delta^{-1} \zeta) - J(\zeta + \beta y, \Delta^{-1} \, \delta^* \zeta).$$

Reverting to the original frame of reference yields

$$\left(\frac{\partial}{\partial t} + U_0 \frac{\partial}{\partial x} \right) \delta^* \zeta = - \mathscr{L}^*(\zeta) \, \delta^* \zeta.$$

REFERENCES

Andersson, E., Pailleux, J., Thépaut, J.-N., Eyre, J. R., McNally, A. P., Kelly, G. A. and Courtier, P. 1992. Use of cloud-cleared radiances in three/four-dimensional variational data assimilation. *Q. J. R. Meteorol. Soc.* **120**, 627–653.

Bartello, P. and Holloway, G. 1991. Passive scalar transport in β-plane turbulence. *J. Fluid Mech.* **223**, 521–536.

Basdevant, C., Legras, B., Sadourny, R. and Béland, M. 1981. A study of barotropic model flows: Intermittency, waves and predictability. *J. Atmos. Sci.* **38**, 2305–2326.

Basdevant, C. and Sadourny, R. 1983. Modélisation des échelles virtuelles dans la simulation numérique des écoulements turbulents bidimensionels. *J. Mec. Theor. et Appl., Numéro Spécial*, 243–269.

Bouttier, F. 1993. The dynamics of error covariances in a barotropic model. *Tellus* **45A**, 408–423.

Bouttier, F. 1994. A dynamical estimation of error covariances in an assimilation system. *Mon. Wea. Rev.* **122**, 2376–2390.

Cohn, S. E. and Parrish, D. F. 1991. The behavior of forecast error covariances for a Kalman filter in two dimensions. *Mon. Wea. Rev.* **119**, 1757–1785.

Courtier, P., Derber, J. C., Errico, R. M., Louis, J. F. and Vukićević, T. 1993. Important literature on the use of adjoint, variational methods and the Kalman filter in meteorology. *Tellus* **45A**, 342–357.

Courtier, P., Thépaut, J.-N. and Hollingsworth, A. 1994. A strategy for operational implementation of 4D-var using an incremental approach. *Q. J. R. Meteorol. Soc.* **120**, 1367–1387.

Gauthier, P. 1992. Chaos and quadri-dimensional data assimilation: a study based on the Lorenz model. *Tellus* **44A**, 2–17.

Gauthier, P., Courtier, P. and Moll, P. 1993. Assimilation of simulated wind Lidar data with a Kalman filter. *Mon. Wea. Rev.* **121**, 1803–1820.

Gilbert, J. C. and Lemaréchal, C. 1989. Some numerical experiments with variable-storage quasi-Newton algorithms. *Mathematical Programming* **45**, 407–435.

Gill, P. E. and Wright, M. H. 1981. *Practical optimization*. Academic Press, 401 pp.

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

Herring, J. R. 1984. The predictability of quasigeostrophic flows. In: *Predictability of Fluid Motions, AIP Conference Proceedings*, no. 106, Ed.: G. Holloway and B. J. West. American Institute of Physics.

Lacarra, J. F. and Talagrand, O. 1988. Short-range evolution of small perturbations in a barotropic model. *Tellus* **40A**, 81–95.

LeDimet, F. X. and Talagrand, O. 1986. Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus* **38A**, 97–110.

Leith, C. E. 1971. Atmospheric predictability and two-dimensional turbulence. *J. Atmos. Sci.* **28**, 145–161.

Leith, C. E. and Kraichnan, R. H. 1972. Predictability of turbulent flow. *J. Atmos. Sci.* **29**, 1041–1058.

Lesieur, M. 1990. *Turbulence in fluids*, 2nd edition. Kluwer, Dordrecht, 412 pp.

Lewis, J. M. and Derber, J. C. 1985. The use of adjoint equations to solve a variational adjustment problem with advective constraints. *Tellus* **37A**, 309–322.

Litherland, J. and Holloway, G. 1984. Statistical theory of the predictability of equivalent barotropic motion on a β-plane. In: *Predictability of fluid motions, AIP Conference Proceedings*, no. 106, ed.: G. Holloway and B. J. West. American Institute of Physics.

Lorenz, E. N. 1969. The predictability of a flow which possesses many scales of motion. *Tellus* **21**, 289–307.

Métais, O. and Lesieur, M. 1986. Statistical predictability of decaying turbulence. *J. Atmos. Sci.* **43**, 857–870.

Miller, R. N., Ghil, M. and Gauthiez, F. 1994. Advanced data assimilation in strongly nonlinear dynamical systems. *J. Atmos. Sci.* **51**, 1037–1056.

Molteni, F. and Palmer, T. N. 1993. Predictability and finite-time instability of the Northern winter circulation. *Q. J. R. Meteorol. Soc.* **119**, 269–298.

Navon, I. M. and Legler, D. M. 1987. Conjugate-gradient methods for large scale minimization in meteorology. *Mon. Wea. Rev.* **115**, 1479–1502.

Orszag, S. A. 1971. Numerical simulation of incompressible flows within simple boundaries (I). Galerkin (spectral) representations. *Stud. in Appl. Math.* **50**, 293–327.

Rhines, P. B. 1975. Waves and turbulence on a beta-plane. *J. Fluid Mech.* **69**, 417–433.

Talagrand, O. and Courtier, P. 1987. Variational assimilation of meteorological observations with the adjoint vorticity equation (I). Theory. *Q. J. R. Meteorol. Soc.* **113**, 1311–1328.

Thépaut, J.-N. and Courtier, P. 1991. Four-dimensional data assimilation using the adjoint of a multilevel primitive equation model. *Q. J. R. Meteorol. Soc.* **117**, 1225–1254.

Thépaut, J.-N., Vasiljevic, D., Courtier, P. and Pailleux, J. 1993. Variational assimilation of conventional observations with a multi-level primitive-equation model. *Q. J. R. Meteorol. Soc.* **119**, 153–186.