Tellus (1995), 47A, 80–102
Printed in Belgium – all rights reserved

TELLUS
ISSN 0280–6495

On the sensitivity of a least-squares fit of discretized linear hyperbolic equations to data

By U. CALLIES* and D. P. EPPEL, *GKSS-Forschungszentrum, 21494 Geesthacht, Germany*

(Manuscript received 9 June 1993; in final form 27 January 1994)

ABSTRACT

Difficulties are investigated which occur when trying to specify a noise-free initial model state as the solution of a variational data assimilation problem. A linear shallow water model is used to investigate the existence and physical basis of the model fit to data. As in this context the shape of the cost function is of crucial importance, the interrelations between the cost function's Hessian and specific model-data configurations are investigated. Special emphasis is put on the influence of the temporal/spatial data distribution and the choice of the scheme used for numerical model integration. It is illustrated how such details may cause intolerable uncertainties for those aspects of the recovered solution that are related to very small eigenvalues of the curvature operator. Due to the shortcomings of descent algorithms, uncontrolled large-amplitude error modes may remain invisible if a limited number of minimization cycles is applied. However, to render the retrieved smooth fields stable with respect to further iterations, prior knowledege has to be taken into account in the cost function definition.

1. Introduction

The problem of merging measurements and numerical models in an appropriate way has been recognized to be a crucial point in order to exploit the skill of the advanced meteorological forecast models. Presently, the formulation of meteorological data assimilation as a variational problem is viewed as a method to optimally take advantage of the information content of data (cf., Ghil and Malanotte-Rizzoli, 1991). In four dimensional data assimilation this technique specifies the initial state which results in a model trajectory that best fits the available observations (e.g., Talagrand, 1981; Kapitza, 1991; Thépaut and Courtier, 1991). Long and Thacker (1989) applied this technique to an ocean model. However, it is not obvious that a specific variational problem of the type described has a reasonable solution (Bennet and Miller, 1991). Even if the number of available data is much larger than the number of unknown variables it may happen that the optimization problem turns out to be *ill*-posed. In this case, either no unique solution exists or an existing solution is extremely sensitive to the specific data set that constitutes the input for the variational problem (e.g., Tikhonov and Arsenin, 1977). Since real data always contain errors, such a solution would be useless.

The mathematical task in the context of variational data assimilation is to minimize some scalar cost function (or objective function) which quantifies the discrepancy between available data and a given model trajectory. The independent variables may be initial fields, boundary values or model parameters. In this paper we restrict ourselves to the consideration of an elementary example with periodic boundary conditions assuming perfect model equations. The model trajectory and the value of the cost function are assumed to depend only on the initial fields.

Courtier and Talagrand (1987) solved a more realistic problem of similar type when assimilating radiosonde observations into a vorticity equation model. One difficulty reported by Courtier and Talagrand was the appearance of unrealistic noise

* Corresponding author.

concentrated near the smallest scales resolved by the model. Bennet and Miller (1991) attributed this unphysical noise to ill-posedness and discussed the role of penalizing initial roughness for well conditioning the inverse problem. However, Bennet and Miller study the issue in a generalized and very formal way without going into details. The purpose of the present investigation is to perform simple numerical experiments in order to convey some idea of how techniqual aspects (like choice of the numerical scheme) may render an inverse assimilation problem ill-posed.

Mathematically, ill-posedness of the variational problem means that the curvature of the cost function J varies over many orders of magnitude for different directions in the space of control variables. An objective measure for ill- or well-posedness is provided by the condition number of the curvature matrix, the latter being closely related to the Hessian containing the second derivatives of J. The condition number is defined as the ratio of the largest to smallest eigenvalues of a matrix. After fitting some model to data, the condition number of the curvature matrix reflects the ratio of error variances of the worst-determined to the best-determined model aspects (Thacker, 1989).

The condition number of the curvature matrix will be infinite, if some eigenvalues vanish. This indicates that the input data contain insufficient information to determine all aspects of the related initial model state (e.g., less data than model variables, or data not being independent). Extremely small eigenvalues of the curvature matrix may occur for physical reasons if the basic model equations are unable to transfer observational information backwards over the assimilation interval without excessive error growth. This would be the case, for instance, considering irreversible processes, if the time lag between model initialization and observations is much larger than the time scale of the system's approach to its equilibrium state for small scale disturbances. Considering the same configuration, the data may still contain valuable information about smooth initial field components since the strength of dissipation is scale dependent. When treating a perfectly reversible flow system by means of a computer, similar effects must be expected due to the truncated representation of the model state with respect to a discrete basis. It is a main goal of the present paper to illuminate several aspects of this problem using characteristic examples.

Of course, the curvature matrix and the specific value of its condition number depend on the actual formulation of the cost function. The ideal situation would be to find all eigenvalues of the curvature matrix equal to one, indicating perfect rotational symmetry of the cost function. Then, starting from an arbitrary first guess the minimum of a quadratic cost function can easily be found by a single step in the direction of its negative gradient. An important tool for improving the condition of a variational assimilation problem is appropriate scaling of the model variables. Scaling may be viewed as a mathematical transformation changing the metric over the space of model variables (cf., Tarantola, 1987). However, in Section 2 it is shown that the shape of the cost function generally depends on the sampling schedule if dynamical equations are used to transport information in time. Section 3 contains numerical investigations on several aspects of the model-data configuration which are influential on the structure of the cost function: Subsection 3.1 is about asynoptic data sets, each containing information on some component of the model state, Subsection 3.2 discusses problems that may arise from spatial interpolation of the model variables to data points, Subsection 3.3 demonstrates effects of the numerical integration scheme. The eigenvectors related to small eigenvalues of an ill-conditioned curvature matrix will contribute with large amplitudes to the overall error of the retrieved initial state when erroneous data are used. An example for this is presented in in Section 4. If the eigenvalues of the curvature matrix are widely spread, the convergence rate of most gradient based search algorithms will slow down considerably. The investigations of Section 5 show that this is mainly due to problems in the recovery of solution components related to small eigenvalues. Therefore one may take advantage of the deficiency of the algorithm in order to suppress the unrealistically large error components. However, in order to obtain a solution which is stable with respect to further steps of the minimization algorithm, some measure of the fit to prior knowledge has to be introduced into the definition of the cost function (Lorenc, 1986). In operational data assimilation systems, this is done by using background information obtained from numerical weather prediction models.

2. The cost function and its curvature

Let us assume that the state of some numerical model is represented by the vector w_0. Let us further assume that some data vector d is to be used to retrieve w_0 from observations. A scalar cost function J measuring the distance between the analyzed model state and the available information can be written in the following form,

$$J = \tfrac{1}{2}(m-d)^{\mathrm{T}} Q^{-1}(m-d), \tag{1}$$

where m represents the vector of calculated data counterparts. Q is an observation error covariance matrix encapsulating our knowledge about the relative precision of the data. The upper index T denotes matrix transpose.

For the sake of convenience let us assume that m depends linearly on the model state w_0 so that its evaluation can be represented by application of some matrix G,

$$m = Gw_0. \tag{2}$$

G describes an arbitrary mapping from the space of model states into the space of observations.

According to (1) and (2), $J(w_0)$ is a quadratic function. Therefore, disturbing some given initial model state by δw_0, the change of the cost function will be,

$$J(w_0+\delta w_0) = J(w_0) + \left(\frac{\mathrm{d}J}{\mathrm{d}w_0}\right)^{\mathrm{T}} \delta w_0 + \frac{1}{2}(H\,\delta w_0)^{\mathrm{T}}\,\delta w_0, \tag{3}$$

H denoting the Hessian matrix whose elements represent the second order derivatives of J with respect to the components of w_0.

The gradient of J and the Hessian H can easily be calculated from our definition of the cost function. One obtains

$$\frac{\mathrm{d}J}{\mathrm{d}w_0} = G^{\mathrm{T}}Q^{-1}(m-d), \tag{4}$$

$$H = G^{\mathrm{T}}Q^{-1}G. \tag{5}$$

Both, the gradient and the Hessian matrix, still contain physical dimensions, generally different for different components. Therefore, in order to use the gradient of the cost function in minimization algorithms, a metric in the space of model states has to be defined, which provides the link between the gradient of J and search directions in the space of the control variables (Tarantola, 1987). Similarly, to obtain an estimate of the reliability of the results based on information contained in the Hessian matrix (Thacker, 1989), a metric tensor has to be introduced in order to transform the Hessian into a nondimensional curvature matrix.

A similar argument can be made from a more mathematical point of view. The pseudo-matricial notation of transposed vectors in (3) represents linear forms from the model space into the real line. The derivative of J in (3) with respect to contravariant vector components gives covariant components. As a consequence the Hessian matrix is twice covariant, while error covariance matrices are twice contravariant. An appropriate tool for the assessment of a matrix of second derivatives would be the consideration of its eigenvalue spectrum in connection with the structure of the related eigenvectors. However, direct application of this method to the Hessian matrix is impossible since eigenvalues exist only for tensors of mixed type (once covariant/once contravariant). Tarantola (1987) illustrates this fact.

Introducing a metric tensor C_{W} gives rise to the definition of a scalar product $(\ldots, \ldots)_{\mathrm{W}}$ for arbitrary elements v, w of the model space W:

$$(v, w)_{\mathrm{W}} = v^{\mathrm{T}} C_{\mathrm{W}}^{-1} w. \tag{6}$$

Then, the norm $\|w\|$ of some model state w is obtained from

$$\|w\|^2 = (w, w)_{\mathrm{W}}. \tag{7}$$

Having defined the scalar product for elements of W we can specify the curvature tensor $\hat{H}$ which relates to the Hessian matrix H in the following way,

$$(w, \hat{H}w)_{\mathrm{W}} = w^{\mathrm{T}} H w. \tag{8}$$

Considering definition (6), one arrives at

$$\hat{H} = C_{\mathrm{W}} H = C_{\mathrm{W}} G^{\mathrm{T}} Q^{-1} G. \tag{9}$$

$\hat{H}$ is not symmetric but it is selfadjoint with respect

to the above scalar product and has real eigenvalues.

For the remainder of this section we will confine ourselves to problems with data and model space being identical. For the operator G, we choose the time evolution operator of a dynamical model w, and for w_0 the initial state to be recovered. Observations of the entire model state shall be given at one specific time. Let w obey a linear dynamic equation,

$$\dot{w} = Aw, \tag{10}$$

where A is a $n \times n$ matrix and $\dot{w}$ denotes the time derivative of w. After scaling of the model state,

$$w = Sw', \tag{11}$$

the governing equation for the nondimensional variables reads

$$\dot{w}' = Bw' \qquad \text{with} \qquad B = S^{-1}AS. \tag{12}$$

For $t_0 = 0$ the model state m is related to the initial state w_0 by

$$m = w(t) = G(t)\, w_0$$

with

$$G(t) = e^{At} = Se^{Bt}S^{-1}. \tag{13}$$

Scaling of the dynamical equations is a priori independent of the optimization problem. To define S, one may primarily regard the specific nature of the governing model equations. Here we are interested in hyperbolic systems. Such systems are perfectly reversible, i.e., information of the model state can be propagated forward and backward in time. However, Browning et al. (1990) point out that it is difficult to solve a hyperbolic system numerically if it contains spatial derivative terms with different orders of magnitude. Truncation errors may easily destroy the numerical solution when integration is performed using limited computational accuracy. To avoid this, the system variables should be scaled in a way that the dynamical matrix shows perfect skew-symmetry. Let us assume that S has been chosen in order to meet this constraint, i.e.,

$$B^{\mathrm{T}} = -B \quad \Leftrightarrow \quad S^{\mathrm{T}}A^{\mathrm{T}}(S^{\mathrm{T}})^{-1} = -S^{-1}AS. \tag{14}$$

The skew-symmetry of B relates to time invariance of the integrated squared state w',

$$\frac{\mathrm{d}}{\mathrm{d}t}\frac{w'^{\mathrm{T}}w'}{2} = w'^{\mathrm{T}}\frac{\mathrm{d}w'}{\mathrm{d}t} = w'^{\mathrm{T}}Bw' = 0, \tag{15}$$

a property which can always be obtained with hyperbolic systems. Since

$$w'^{\mathrm{T}}w' = w'^{\mathrm{T}}I^{-1}w' = w^{\mathrm{T}}(SS^{\mathrm{T}})^{-1}\, w,$$

$$I = \text{identity}, \tag{16}$$

the purely dynamical point of view suggests to choose the metric tensor in definition (6) as

$$C_{\mathrm{W}} = SS^{\mathrm{T}} \qquad \text{or} \qquad C'_{\mathrm{W}} = S^{-1}C_{\mathrm{W}}S = I, \tag{17}$$

depending on the choice of the reference system. Then, according to (15) any evolution of the model trajectory is represented in phase space by rotations of the state vector. Referring to the new basis, the scalar product (6) becomes Euclidean, and the representation $\hat{H}'$ of the curvature matrix becomes the Hessian matrix itself (cf., Tarantola, 1987):

$$\hat{H}' = S^{-1}\hat{H}S = S^{\mathrm{T}}HS = H'. \tag{18}$$

Note that the symmetric matrix $\hat{H}'$ has the same eigenvalues as $\hat{H}$. In the remainder of this paper we will always consider the cost function to depend on nondimensional variables and thus be able to abandon the inconvenient discrimination between the curvature operator and the Hessian of the cost function.

In order to accelerate minimization algorithms, we would like to choose the metric tensor C_{W} such that the curvature for different directions is the same (then a descent algorithm would find the minimum in one step). To do this, let us first consider the example, where a complete data set is available at $t_0 = 0$. Then the data counterparts are calculated with $G(0) = I$, so that according to (5) and (18) the nondimensional Hessian takes the form

$$H' = S^{\mathrm{T}}Q^{-1}S. \tag{19}$$

H' will equal the identity if

$$Q = SS^{\mathrm{T}}. \tag{20}$$

However, for all practical applications the

Hessian will depend on the time lag between model initialization and observations. We obtain the following general expression for H',

$$H' = e^{B^T t} S^T Q^{-1} S e^{Bt}. \tag{21}$$

It is obvious from (21) that a time independent Hessian can only be achieved when S obeys (14) and (20) simultaneously. Then, using the fact that a skew-symmetric matrix is normal, i.e., it commutes with its transpose, we obtain

$$H' = e^{B^T t} e^{Bt} = e^{(B^T + B)t} = e^0 = I. \tag{22}$$

Generally, a matrix S satisfying both (14) and (20) will not exist. The conclusion is that for realistic assimilation problems the reliability of the retrieved results in their different components will depend on the temporal spacing of the available measurements.

3. Numerical studies of the cost function's shape

Our numerical studies are based on a simple 2-dimensional linear shallow water system, whose evolution in time t is governed by the following set of coupled equations for two velocity components $u_i = u, v$ aligned with coordinates $x_i = x, y$ and a height field h,

$$\frac{\partial u_i}{\partial t} + g \frac{\partial h}{\partial x_i} = 0, \tag{23}$$

$$\frac{\partial h}{\partial t} + h_r \sum_i \frac{\partial u_i}{\partial x_i} = 0, \tag{24}$$

g and h_r denoting the gravity acceleration and the constant mean water depth, respectively. We have disregarded the Coriolis term in (23) since its incorporation has little influence on the phenomena discussed in the remainder of this paper. If scaling parameters for the model variables are chosen according to

$$u_i = \sqrt{gh_r}\, u_i', \qquad h = h_r h', \tag{25}$$

one ends up with the skew-symmetric set of equations:

$$\frac{\partial u_i'}{\partial t} + \sqrt{gh_r} \frac{\partial h'}{\partial x_i} = 0, \tag{26}$$

$$\frac{\partial h'}{\partial t} + \sqrt{gh_r} \sum_i \frac{\partial u_i'}{\partial x_i} = 0. \tag{27}$$

Approximating first order derivatives of the non-dimensional variables by central finite differences preserves skew symmetry of the dynamic matrix which is related to conservation of total energy.

Our data assimilation experiments are based on a simple definition of the cost function being sufficient for our purposes:

$$J(w_0') = \frac{1}{2} \sum_v \frac{[m_v(w_0') - d_v]^2}{\sigma_v^2}. \tag{28}$$

The short notation w_0' stands for the nondimensional initial fields u_0', v_0', h_0', and m_v are the modelled counterparts of the data d_v. The index v is understood to count single measurements.

In the more general considerations of the previous chapter we presumed a single synoptic data set specifying all model variables at one common time. Under these very special circumstances it was shown that the curvature of the cost function (i.e., the reliability of the retrieved initial state) may be independent of the observation time if the discretized dynamic model equations and the observation error covariance matrix are closely related. Though this situation has to be judged as completely unrealistic, we choose it as kind of reference situation for the further investigations. In what follows we are interested in studying the effect of asynoptic and spacially scattered data on the variational problem. For such an investigation it is convenient to have a homogeneous eigenvalue spectrum equal one for a single complete synoptic data set. This is achieved by scaling according to (25) corresponding to a diagonal matrix S in notation (11). Relation (20) will be satisfied by assuming a diagonal matrix Q with elements Q_{ii},

$$Q_{ii} = \sigma_v^2 = gh_r \qquad \text{or} \qquad Q_{ii} = \sigma_v^2 = h_r^2, \tag{29}$$

depending on whether the corresponding datum represents a velocity or a height observation, respectively. A common factor for all elements of Q changing the absolute magnitude of all error variances simultaneously is irrelevant for the nature of the minimization problem.

All numerical experiments are performed on a periodic grid with 8×8 active points. To start the model, three initial fields must be specified resulting in 192 unknowns to be optimized and thus in a Hessian containing 192×192 second order derivatives.

The integration scheme is chosen in order to avoid numerical diffusion which would spoil the hyperbolic character of the model equations. We apply the time-centered implicit method (e.g., Anderson et al., 1984) which consists of second-order central differences in space and trapezoidal or Crank-Nicolson differencing in time. The constant depth h_r and the grid increments $\Delta x = \Delta y$ are specified as 8 km and 3000 km, respectively. A time increment of 1000 s is chosen resulting in a Courant number $v_p \, \Delta t/\Delta x$ of less than 0.1, with phase velocity $v_p = \sqrt{gh_r} = 280$ m/s.

The minimum of J is found by solving the equation $dJ/dw'_0 = 0$. Because the dependence $m_v(w'_0)$ is linear comprising the effects of temporal integration, spatial/temporal interpolation and descaling (note that the measured values d_v remain unmodified, i.e., they keep their physical dimension), the cost function J will be quadratic and a Taylor series establishes a linear relationship between the gradient of J and the deviation of the actual initial fields from their optimum values $w'_{0,\mathrm{opt}}$,

$$\frac{dJ}{dw'_0} = H'[w'_0 - w'_{0,\mathrm{opt}}]. \tag{30}$$

The 192 eigenvalues of the nondimensional Hessian H' describe the variability of the objective function's curvature for different directions in the space of scaled initial values. Thus, the eigenvalue spectrum is the suitable tool to assess the reliability of different components of the solution of the optimization problem (Thacker, 1989). According to (30), even small gradients of the cost function can correspond to large deviations of the initial model state from its minimum if some eigenvalues of the Hessian become extremely small compared to others. Then, small model-data differences will have large effects on the recovered initial fields and the problem will be ill-posed. The magnitudes of the eigenvalues may correlate with a characteristic spatial scale of the corresponding eigenvectors of the Hessian matrix. This is to be expected, for instance, when considering diffusive phenomena like heat conduction. During temporal evolution of the system, small scale temperature fluctuations will be smoothed out relative to the large scale distributions prohibiting a reliable recovery of non-smooth initial fields from noisy data by means of inverse modelling over too long an assimilation interval. In Subsection 3.3 we shall see that similar difficulties may arise in the numerical treatment of reversible hyperbolic systems. However, before investigating this point in more detail, the following two subsections will address the question of how the shape of the cost function is modified if the assimilated observations do not contain enough information in order to specify each model component in every grid point at the respective observation time.

3.1. Asynoptic data

One restricting assumption among others underlying the general considerations in Section 2 was that the data are synoptic and describe all aspects of the model state. In this subsection we shall demonstrate how the eigenvalue spectrum of the nondimensional Hessian matrix will be affected if different model components are measured at different times.

In order to provide a link to Section 2, we begin by investigating the situation where a complete data set is available, i.e., it contains all three model fields. The 192 measurements (3 at each of 64 grid-points) shall be taken simultaneously after 32,000 s so that the relation $m_v(w'_0)$ comprises descaling and integration over 32 time steps. During these 32 time steps with Courant number 0.093, propagating waves will move approximately 3 grid increments. In agreement with our former general result, a degenerate spectrum of eigenvalues is found with all eigenvalues being equal to one.

Let us now modify this reference run by assuming that the measurements of height h and velocity (u, v), respectively, have been made at different times. In Fig. 1 the spectrum of 192 discrete values has been depicted quasi-continuously. The solid line shows the resulting eigenvalue spectrum if u- and v-data are available after 16 time steps while the h-data are given after 32 time steps as before. The effect of the asynoptic data distribution is that 120 of the 192 eigenvalues become different from 1. Though the mean of the eigenvalue spectrum remains unchanged, the condition number of the Hessian matrix is now 660. To explain this feature,

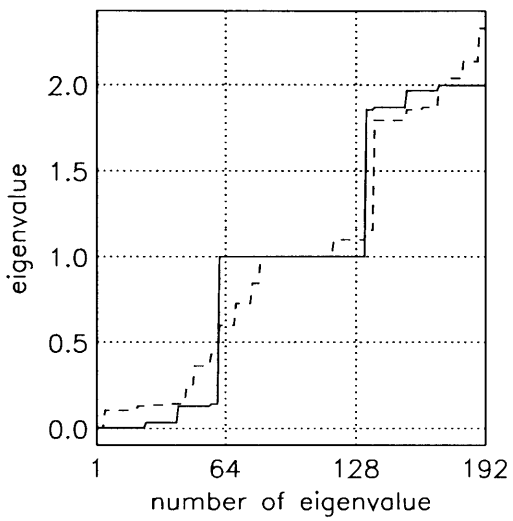

Fig. 1. The eigenvalues of the nondimensional Hessian matrix for one complete asynoptic system observation being available at each grid point. Results are shown for two different sampling schedules: (a) (u, v)-data after 16 time steps, h-data after 32 time steps (solid line), (b) u-data at initial time, v- and h-data given as before (dashed). In both cases the model has been integrated using a time-centered implicit scheme with Courant number 0.093.

one has to regard that in general, following a model trajectory, the conserved total energy will manifest itself alternating between the potential and the kinetic energy form, respectively. Therefore the sum of both components will be constant only for simultaneous measurements of the partial energies. Nevertheless, there are still specific initial states, in our case 72, giving rise to a model trajectory along which potential and kinetic energy are conserved individually. These flow patterns remain related to eigenvalues equal 1 of the Hessian even for asynoptic measurements of the two energy components.

The actual value of the condition number of the Hessian matrix which is found in our special example depends on the time lag between the assumed measurements of potential and kinetic energy, respectively. The variational problem turns out to be ill-posed if the time difference between the two measurements corresponds to a multiple of $T/2$ with T denoting the period of the oscillation of the integrated values of the two partial energy forms. Thacker and Lewandowicz (1994) present a comprehensive discussion of how such results can be understood as a process of information exchange between various aspects of a system during its temporal evolution.

Note, that in our specific example the distortion of the spectrum is completely symmetric. For each eigenvalue λ smaller than 1 exists a conjugate eigenvalue $2-\lambda$. Swapping the data sets, i.e., assuming h-data after 16 time steps and (u, v)-data after 32 time steps, respectively, does not affect the shape of the spectrum at all. This indicates that the observed symmetry can be explained by taking into account the reversibility of the shallow water equations. To see this interrelation, let us assume for convenience that all data values are equal to zero (the Hessian of a quadratic cost function remains unaffected by the actual values of the measurements). Further let us assume that integrating the system from some given initial model state w_0 with total energy E, the kinetic and potential energy after 16 time steps are k and P, respectively, while after 32 time steps values K and p with $k<K$ and $p<P$ are obtained. For our sample asynoptic data set the value $J_{as}(w_0)$ of the cost function will be proportional to $k+p$ and therefore be less than in the case of synoptic data with a cost function $J_s(w_0)$ being proportional to the total energy $E=k+P=K+p$ of the system. The quadratic cost function's smaller value goes along with its weaker curvature. Thus, if w_0 is chosen to be an eigenvector of the Hessian matrix, the corresponding eigenvalue must be smaller than one. Let us now integrate the system over a period of 48 time steps. We may take the result and use it as another initial state w_0^* after having changed the sign of the velocity field. Choosing the negative velocities is equivalent to inverting the time coordinate so that integrating w_0^* will produce a mirrored model trajectory with kinetic energy K after 16 time steps and potential energy P after 32 time steps. Since

$$\begin{aligned} K+P &= [(K+p)+(k+P)]-(k+p) \\ &= 2E-(k+p), \end{aligned} \tag{31}$$

the value of the cost function will be found as $J_{as}(w_0^*)=2J_s(w_0)-J_{as}(w_0)$. Therefore w_0^* represents an eigenvector with eigenvalue $2-\lambda$ being the symmetric counterpart of the eigenvector w_0 with eigenvalue λ.

The dashed line in Fig. 1 represents the eigenvalue spectrum if u-data are available at the initial

Table 1. *Common sampling schedule assumed for Figs. 2–13; number and locations of the data are specified for each experiment individually*

Time (s)	Kind of data
8000	u-observations
9000	v-observations
11000	h-observations
12000	u-observations
13000	v-observations
14000	h-observations
16000	u-observations
17000	v-observations
19000	h-observations
21000	u-observations
23000	v-observations
24000	h-observations
25000	u-observations
27000	v-observations
29000	h-observations
29000	u-observations
30000	v-observations
32000	h-observations

time t_0 while v- and h-data are specified as before. In this case only 36 eigenvalues remain equal to one, for the related flow patterns the integral values of u'^2 and v'^2 are individually conserved. Though the spectrum has now become more structured, the average value 1 of the spectrum is maintained. However, the structure of the eigenvalue spectrum for data given at three different times is no longer symmetric.

From now on, we will always presume that the number of available data is much larger than the number of independent variables. Table 1 specifies the temporal distribution of observations that is underlying all subsequent experiments. 18 data sets are supposed to exist at different time levels, but each of these sets involves only one model variable. It turns out that for irregular sampling intervals the eigenvalue spectrum is qualitatively insensitive to the actual sampling schedule. The overall length of the assimilation interval is again 32,000 s so that the product of Courant number and number of time steps amounts to approximately 3.

3.2. *Spacially displaced or sparse data*

This subsection addresses the question of how spatial location and density of the observational network influence the solution of a variational data assimilation problem.

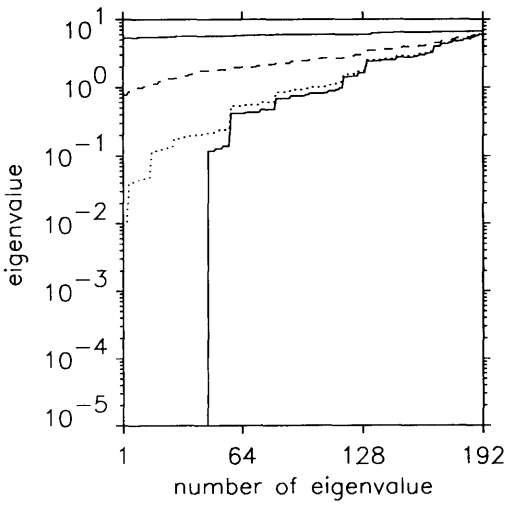

Fig. 2. Eigenvalues of the nondimensional Hessian matrix for asynoptic data sampled according to Table 1. Observation points are located according to shift parameters (cf., text) $\xi = 0$ (upper solid line), $\xi = 0.2$ (dashed), $\xi = 0.4$ (dotted) and $\xi = 0.5$ (lower solid line). The model equations are integrated using the time-centered implicit method with time step 1000 s (Courant number 0.093).

For our first experiment let us assume that the 18 fields of the data set (cf., Table 1) provide measurements at each grid point. This means that to optimize 192 independent initial values, $18 \times 64 = 1152$ measurements are available. For model integration a time increment of 1000 s (Courant number 0.093) is chosen to avoid temporal interpolation when calculating the data-counterparts. We shall investigate the effect of shifting the location of each measurement from point (i, j) along a line, connecting (i, j) and $(i+1, j+1)$. The shift is quantified by a parameter ξ: $\xi = 0$ corresponds to no shift while $\xi = 0.5$ corresponds to observations in the center of every grid cell. The bi-linear scheme is used to interpolate calculated model values to the observational points.

Fig. 2 shows that displacing the observation points relative to the computational grid in a regular manner increases the condition number of the Hessian. For $\xi = 0.5$ the solution of the inverse problem is no longer unique since 45 eigenvalues vanish. The results demonstrate that the use of

interpolation to calculate the cost function may turn a well-posed problem into an ill-posed one even if the number of available data is much larger than the number of independent variables. Visual examination of the eigenvectors with zero eigenvalues (not shown) indicates that the structures which are uncontrolled by the data are all related to wave lengths of twice the grid spacing (2Δ-waves).

Appendix A addresses the mathematical problem of reversing the distortion of the eigenvalue spectrum by means of preconditioning, considering an example of spatial interpolation in one dimension. Such a correction of the ill-shaped objective function could be done in principal for $\xi \neq 0.5$; it must fail, however, for $\xi = 0.5$.

The problem with interpolating model fields into the center of grid cells is that several distinguishable model states may be mapped onto identical interpolated values. A similar difficulty occurs if the observational grid coincides with the numerical grid but measurements are not available at every grid point. In this situation specific modes may propagate through the model domain exhibiting non-zero values only away from observation points and thus being uncontrolled by the data. Our next experiment serves to illustrate this point.

Let us assume that the fields of our data set provide measurements at every second grid point (where the sum of the grid indices $i + j$ is an even number), so that the total number of data is now $18 \times 32 = 576$. All model parameters are kept unchanged. The dashed line in Fig. 3 depicts the resulting eigenvalue spectrum. Comparing these eigenvalues with those originating from dense data (upper solid line, copied from Fig. 2), shows an increased slope of the spectrum together with 36 eigenvalues that vanish. Again uncontrolled 2Δ-waves make the optimization problem ill-posed.

Regarding the structure of those initial state components which cannot be recovered from the sparse data, one suspects that it is the regular spacing of the data gaps that causes the problems with the Hessian matrix. This suspicion is confirmed by the result of an experiment with a slightly modified setting. Again the data sets are assumed to be diluted, but now the respective location of each measurement is not at the grid point itself but at some random position in the grid cell next to it. The lower solid line in Fig. 3 shows indeed that randomizing the data positions in this way is a remedy to remove the problem of zero eigenvalues, i.e., of uncontrolled patterns in the intial model fields.

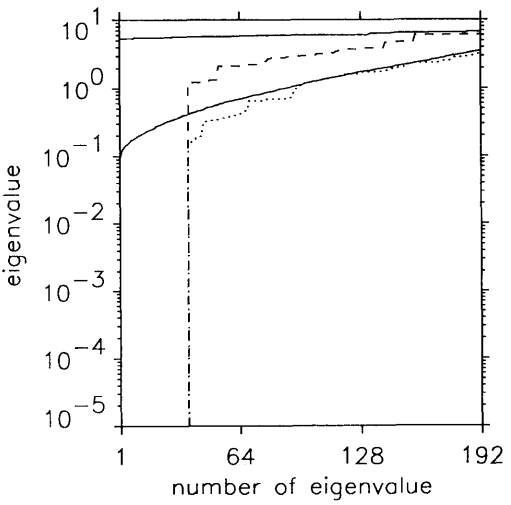

Fig. 3. The effect of using sparse data fields. Observation schedule and model specifications are the same as in Fig. 2. Eigenvalues of the nondimensional Hessian are shown for (a) data at every grid point (upper solid line, copied from Fig. 2), (b) data at every second grid point (dashed), (c) data in every second grid cell with $\xi = 0.25$ (dotted), (d) irregularly located data in every second grid cell (lower solid line). Condition numbers are 1.3 in case (a) and 38.3 in case (d).

Note that for randomized data points, the large eigenvalues exhibit a similar distribution as if the whole data fields would have been displaced by a common distance ξ (cf., Fig. 2). Fig. 3 depicts the spectrum which is obtained for $\xi = 0.25$ with data in every second grid cell (dotted line). Compared to this, the right hand part of the spectrum with random data positions is smoothed by removing degenerecies of eigenvalues.

3.3. *Choice of the numerical scheme*

So far, we have found two kinds of singularities which were related to the spatial data distribution. In the first instance, observations were located in the center of the numerical grid cells so that irreversible interpolation had to be employed in order to calculate the data counterparts m_ν. In the second instance, the data points coincided with the numerical grid but observations were available only at every second grid point. In both experi-

ments zero eigenvalues of the Hessian matrix indicated that the evaluation of the m_v did not depend on all components of the initial state fields.

In the latter case the detailed relation $m_v(w'_0)$ exclusively depends on the numerical algorithm used for time-integration of the model. Therefore it can be expected that the observed singularities may at least partly be removed by application of another integration algorithm, which introduces a stronger coupling between neighboured grid values. Lax-Wendroff's scheme (e.g., Anderson et al., 1984) with its explicit diffusion term provides an example for such a scheme. Fig. 4 compares the eigenvalue spectrum obtained with Lax-Wendroff's method (solid line) to the spectrum obtained with the time-centered implicit scheme (dashed line, copied from Fig. 3). Employing Lax-Wendroff's scheme reduces the number of zero eigenvalues from 45 to 3, though the magnitude of the other 42 eigenvalues remains very small.

However, the diffusion term in Lax-Wendroff's numerical approximation (cf., Appendix B), which seems to be responsible for the shrinkage in the space of uncontrolled intial state patterns, turns out to cause problems itself when the Courant number is increased. This is shown in Fig. 5. Let us consider the experiment with irregularly located data in every second grid cell (cf., lower solid line in Fig. 3). The result of varying the Courant number is hardly noticable when the time centered implicit scheme is used (not shown). The left panel of Fig. 5 depicts results of the same experiment, except that now the time centered scheme has been replaced by a fully implicit Euler scheme. The influence of time step variations on the eigenvalue spectrum still remains weak. The right panel of Fig. 5 depicts the corresponding results when Lax-Wendroff's scheme is employed. Now, increasing the Courant number from 0.1 to 0.3 changes the condition number of the nondimensional Hessian from 76 to 871. Zero eigenvalues are found for two Courant numbers c: for $c = 0.5$ one eigenvalue vanishes, for $c = 1/\sqrt{2} \approx 0.707$ (limit of linear stability) there are 18 eigenvalues equal to zero.

A more detailed examination of the Hessian's eigenvectors with eigenvalue zero reveals that they are all related to 2Δ-waves. For $c = 1/2$ and $c = 1/\sqrt{2}$ the diffusion terms of Lax-Wendroff's scheme map these small scale structures onto zero fields (cf., Appendix B), i.e., the 2Δ-waves

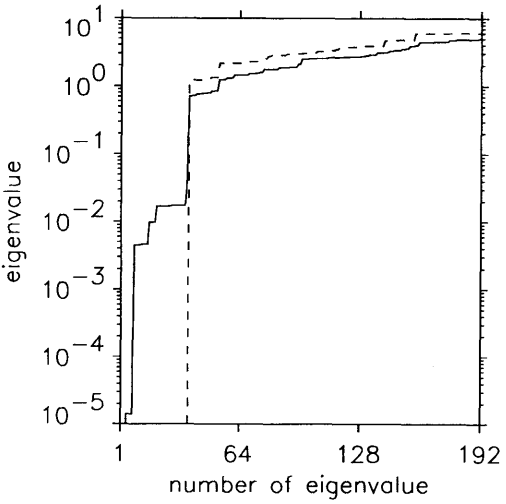

Fig. 4. The eigenvalues of the nondimensional Hessian matrix for different numerical integration schemes. Data are sampled according to Table 1 at every second grid point. Results are shown for using (a) the time centered implicit scheme (dashed line, copied from Fig. 3), (b) Lax-Wendroff's scheme; both schemes are applied with time step 1000 s.

are eigenvectors of Lax-Wendroff's scheme with corresponding eigenvalues equal zero. The mutual interrelation between vanishing eigenvalues of both the Hessian and the numerical scheme becomes obvious considering (5).

Note that the singular spectra of eigenvalues depicted in Fig. 5 depend on the fact that inside the time interval covered by the first time step no observations are available. Otherwise the modelled counterparts of such data would be directly influenced by the model's initial fields (via temporal interpolation) and therefore be partly independent of the choice of the numerical scheme.

Finally, it should be mentioned that the Courant numbers used in Fig. 5 are all related to time steps which are not suitable to meet the temporal spacing of the data. In all runs temporal interpolation is needed to evaluate the m_v. However, this temporal interpolation does not essentially spoil the condition of the Hessian.

4. Error amplification when solving an ill-posed problem

In the previous sections, examples have been considered which elucidate some aspects of how

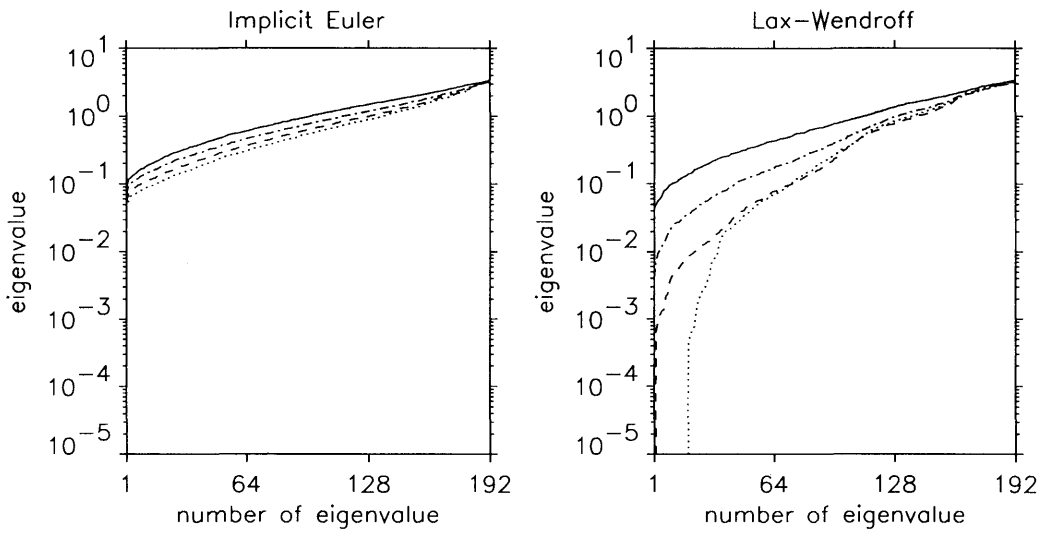

Fig. 5. The eigenvalues of the nondimensional Hessian matrix as a function of Courant number c. Results are shown for $c = 0.1$ (solid lines), $c = 0.3$ (chain-dotted lines), $c = 0.5$ (dashed lines) and $c = 1/\sqrt{2}$ (dotted lines); the related time steps are 1071 s, 3213 s, 5354 s and 7572 s, respectively. The experiment compares Euler's fully implicit method (left panel) and Lax-Wendroff's method (right panel). Asynoptic data (cf., Table 1) in every second grid cell are assumed to be irregularly located.

the Hessian of a variational data assimilation problem may become ill-conditioned due to various details of the problem's design. Let us now consider the consequences encountered when trying to solve an ill-posed data assimilation problem.

Our experiment simulates a specific assimilation process, referring to the same model-data configuration that underlies the eigenvalue spectra in the right panel of Fig. 5. However, while the foregoing investigations of the Hessian matrix were independent of measured data values, we now have to specify some hypothetical data set. We do this by integrating an initial model state consisting of the following periodic height field pattern,

$$\begin{array}{ccccccc} & \vdots & \vdots & \vdots & \vdots & \vdots & \\ \cdots & 0 & 0 & 0 & 0 & 0 & \cdots \\ \cdots & 0 & 2.5 & 5 & 2.5 & 0 & \cdots \\ \cdots & 0 & 5 & 10 & 5 & 0 & \cdots \\ \cdots & 0 & 2.5 & 5 & 2.5 & 0 & \cdots \\ \cdots & 0 & 0 & 0 & 0 & 0 & \cdots \\ & \vdots & \vdots & \vdots & \vdots & \vdots & \end{array} \qquad (32)$$

together with a zero velocity field. Lax-Wendroff's scheme is applied with Courant number 0.6. Choosing this Courant number results in a large condition number 31030 (cf., Fig. 8) of the Hessian but avoids the two singular spectra depicted in Fig. 5.

Fig. 6 shows the projections of the initial state (32) onto the complete system of orthogonal eigenvectors of the nondimensional Hessian. In all further investigations we will refer to these eigenvectors as a basis in order to specify a certain model state. Note that components of pattern (32) being parallel to eigenvectors with very small eigenvalues are negligable. Since the eigenvalue spectrum for Courant number 0.6 contains no zeros, the inverse problem of deriving initial model values from given data can be solved in principle. Assuming correct data (here: those data, which have been produced by the model itself), the height field component of the solution of the inverse problem will prove to be identical with pattern (32).

However, the situation of having perfect data is not realistic. Inconsistencies between observed data and model calculations will inevitably be involved in any realistic data assimilation problem, and one has to judge carefully the possible influence of such model-data discrepan-

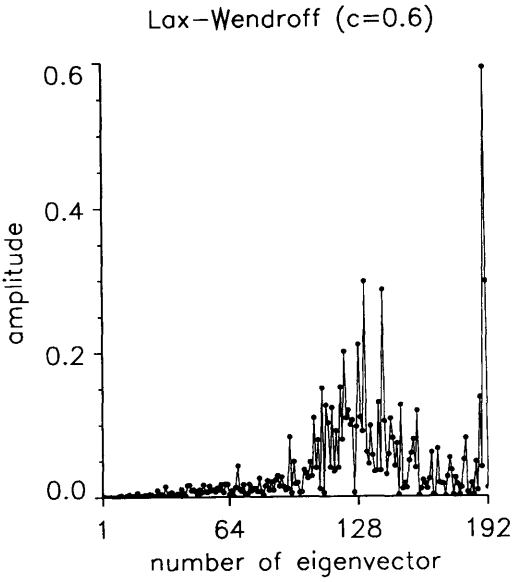

Fig. 6. Projections of the sample model state (cf., text, eq. (32)) onto the eigenvectors of the nondimensional Hessian resulting from using Lax-Wendroff's scheme with Courant number 0.6. The data configuration is the same as in Fig. 5.

cies on the reliability of the variational problem's solution.

The differences between data and their modelled counterparts may be random or systematic. In our example we concentrate on the latter case which occurs if the model conception is unable to accurately describe the physical processes. To simulate such a situation, whose investigation has been the issue of Wergen (1992), let the hypothetical true nature be represented by a model run with Courant number 0.3 (performed to generate the data). We may then compare the "true" initial state (32) with the optimum state determined on the basis of "approximated" model equations that will be simulated by running the model with Courant number 0.6.

Fig. 7 shows the result of this experiment in spectral representation as well as in terms of some representative cross section through the optimized height field. Comparing Figs. 6, 7 reveals that the retrieved optimum state differs from the reference model state in having large components in a subdomain that is spanned by eigenvectors of the Hessian related to small eigenvalues. Considering the deviation of the recovered height field from the reference field used for data generation (left part of

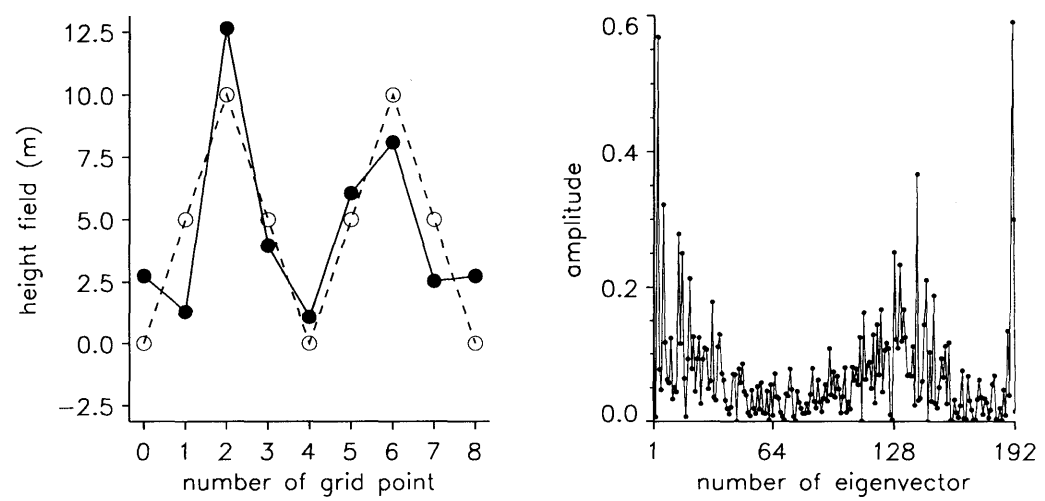

Fig. 7. Right panel: the same as in Fig. 6 except that the analyzed model state has been retrieved from inconsistent artificial data produced by running the model with Courant number 0.3. Left panel: characteristic cross section through the height field of the retrieved model state (solid line). The dashed line shows the related cross section through the reference state used for data generation.

Fig. 7) one realizes indeed a strong short wave component. Note that the difference between the two fields is non-symmetric, although no random error distribution has been introduced. The irregular structure of the optimum model state is exclusively due to irregular spacing of our hypothetical data points. This irregular spacing (being necessary to obtain a well defined variational problem for sparse data sets) of exact data is sufficient to cause irregular short wave patterns in the optimum initial state, which look very similar to those in solutions inferred from noisy data being regularly spaced (not shown).

It is obvious that the result depicted in Fig. 7 would be of little value if it occurred in practical applications. Those components of the retrieved model state that project onto eigenvectors of the Hessian matrix with very small eigenvalues are unreliable and should not be considered as significant. Therefore one may try to use prior knowledge of the optimum model state's desirable properties in order to formulate a modified variational problem which does not allow for solutions with large amplitude model components being uncontrolled by data. The eigenvalue spectrum of the Hessian matrix needs to be changed in order to achieve this.

In all operational implementations of data assimilation schemes the cost function includes some term measuring the deviation of the independent fields from background information provided by a numerical weather prediction model. Such prior information prevents the appearance of underdetermined model states. However, in many nonoperational applications background information will not be available. Then, some penalty constraint may be introduced for the sake of regularization suppressing unrealistic large error amplitudes of flow patterns that correspond to eigenvectors of the Hessian with small eigenvalues.

The easiest way to remove extremely small eigenvalues of the Hessian matrix while leaving large eigenvalues untouched is to add a small constant α to the whole spectrum. This is what happens if we extend the cost function J by a penalty term of the following form,

$$J(w'_0) = J_D(w'_0) + \tfrac{1}{2}\alpha \, \|w'_0\|^2, \tag{33}$$

which we will refer to as a zero order penalty (penalyzing the zero order derivative). Differentiating the penalty term twice with respect to w'_0 results in α times the identity matrix. Thus, the above penalty term does not change the eigenvectors of the Hessian matrix which served as a reference system in Figs. 6, 7. This makes formulation (33) especially suited for demonstrations of how the penalty term works.

We want to investigate how the optimum model state given in Fig. 7 will be modified if we assume definition (33) with some non-zero coefficient α. From (30) follows the general expression for the model state minimizing the cost function,

$$w'_{0,\mathrm{opt}} = H'^{-1} g'_0 \qquad \text{with} \quad g'_0 = -\left.\frac{\mathrm{d}J}{\mathrm{d}w'_0}\right|_{w'_0 = 0}. \tag{34}$$

Let the columns of the matrix E contain the eigenvectors of the nondimensional Hessian H'. Then the spectral representation of $w'_{0,\mathrm{opt}}$ arises from multiplication of (34) with the transpose of E,

$$E^{\mathrm{T}} w'_{0,\mathrm{opt}} = E^{\mathrm{T}} H'^{-1} g'_0 = D^{-1} E^{\mathrm{T}} g'_0, \tag{35}$$

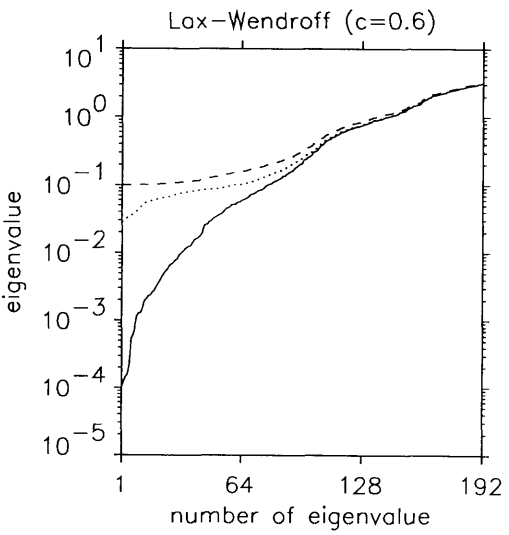

Fig. 8. The eigenvalues of the nondimensional Hessian matrix resulting from using Lax-Wendroff's scheme with Courant number 0.6 and irregularly located asynoptic data in every second grid cell (cf., right panel of Fig. 5). Eigenvalues are shown for a cost function (a) without penalty (solid line), (b) including a zero order penalty constraint with $\alpha = 0.1$ (dashed line), (c) including a second order penalty constraint with $\alpha = 0.1/64 = 1.56 \times 10^{-3}$ (dotted line).

with diagonal matrix D^{-1} containing the inverse eigenvalues of H' (note that $E^{\mathrm{T}}E = I$ for an orthonormal system of eigenvectors):

$$D^{-1} = E^{\mathrm{T}} H'^{-1} E. \tag{36}$$

Each element of D^{-1} has the form $1/(\lambda_i + \alpha)$ with λ_i denoting the ith eigenvalue of H' for $\alpha = 0$. Both vector g'_0 and matrix E are independent of α. Thus, it becomes obvious from (35) that increasing the coefficient α will result in a reduction of all spectral components of the retrieved initial model state $w'_{0,\mathrm{opt}}$. However, this reduction will be large for components related to small eigenvalues $\lambda_i \ll \alpha$ and negligable for modes with eigenvalues $\lambda_i \gg \alpha$. One may suspect that an appropriate choice of the coefficient α can turn an ill-posed variational problem into a reasonable well-posed one.

The solid line in Fig. 8 shows the nondimensional Hessian's eigenvalue spectrum for the unpenalized problem. Comparing this spectrum with Fig. 6 reveals that the reference field (32) hardly projects on eigenvectors with eigenvalues smaller than 0.1. Therefore 0.1 appears to be a reasonable choice for α in (33). The dashed line in Fig. 8 depicts the modified eigenvalue spectrum and Fig. 9 shows the solution of the related optimization problem. As expected, including the penalty constraint into the definition of the cost function has suppressed all components of the solution which project onto the left hand part of the spectrum with non-smooth eigenvectors. The same conclusion can be drawn considering the cross section through the obtained height field.

On the other hand, comparing the spectral representations of the unpenalized (Fig. 7) and the penalized (Fig. 9) solution, respectively, reveals hardly any differences in the details of the amplitudes in the domain of eigenvectors with large eigenvalues. The penalty term has left those components unchanged, being unable to suppress systematic smooth error structures which have been introduced by using different Courant numbers for data production and solution of the inverse problem, respectively. The existence of such smooth error components becomes evident from comparing Fig. 7 with the spectral representation of our reference initial state (Fig. 6).

Regarding the very simple structure of the zero order penalty constraint applied so far, one may ask if a higher order penalty term performs much better by affecting different spatial scales selectively. Wahba and Wendelberger (1980) describe a method to find systematically the order of the penalized derivative and the coefficient α most consistent with the data. Their criterium for a good choice is the ability to predict model values where data are withheld (a method called cross validation). We do not intend to specify the penalty constraint which performs best for our present example. But let us shortly consider what happens if the order of the penalty constraint is increased.

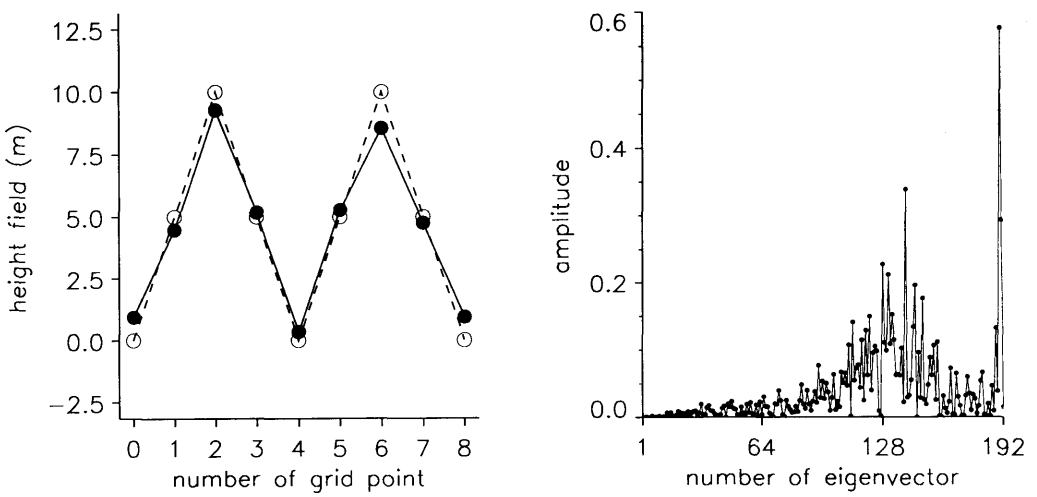

Fig. 9. The same as Fig. 7 except with zero order penalty constraint ($\alpha = 0.1$).

We may look for an initial model state whose fields exhibit a weak curvature, minimizing the following extended cost function,

$$J(w'_0) = J_D(w'_0) + \tfrac{1}{2}\alpha \, \|\Delta w'_0\|^2. \tag{37}$$

The symbol Δ represents the Laplacian on a grid with unit spacing, i.e.,

$$(\Delta w')_{i,j} = \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix} w'_{i,j}. \tag{38}$$

We would like to specify the coefficient α in (37) such that for small scale structures the strength of the second order penalty term compares with that of the former zero order constraint. In other words, the induced gradients of the penalty terms should be of similar size, due to similar eigenvalues of the related Hessian matrices. The gradient of the second order penalty term has the following components:

$$\frac{1}{2}\alpha \frac{\mathrm{d}\,\|\Delta w'_0\|^2}{\mathrm{d}(w'_0)_{i,j}} = \frac{1}{2}\alpha \frac{\mathrm{d}\sum_{k,l}(\Delta w'_0)^2_{k,l}}{\mathrm{d}(w'_0)_{i,j}}$$

$$= \alpha \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 2 & -8 & 2 & 0 \\ 1 & -8 & 20 & -8 & 1 \\ 0 & 2 & -8 & 2 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix} (w'_0)_{i,j}. \tag{39}$$

It can readily be verified that a checkerboard structure consisting of alternating values 1 and -1 is an eigenvector of the penalty term's Hessian represented by the numerical stencil in (39). The eigenvalue is found to be 64α. The related eigenvalue arising from a zero order penalty term would be α. Since we chose $\alpha = 0.1$ in (33), it appears reasonable to define $\alpha = 0.1/64$ in the context of formulation (37).

The resulting eigenvalue spectrum of the cost function is depicted in Fig. 8 (dotted line). It can be noticed that modifications of the large eigenvalues are even less than with the zero order penalty. In the domain of small eigenvalues it becomes obvious from the distorted spectrum that the new penalty term does not just add a constant. A direct mode-by-mode comparison with the zero-order penalized problem is not possible since the structure of the eigenvectors becomes modified by a second order penalty term. However, comparing the spectral representations of the model state minimizing (37) (Fig. 10) and the unpenalized solution (Fig. 7), respectively, reveals that the amplitudes of smooth structures on the right hand side of the spectrum remained untouched while the non-smooth structures represented on the left hand side of the spectrum obviously experienced essential modifications.

The left panel of Fig. 10 indicates that the solution of (37) fits remarkably well the reference height field. It is interesting to compare the residuals of the cost functions' data terms for the

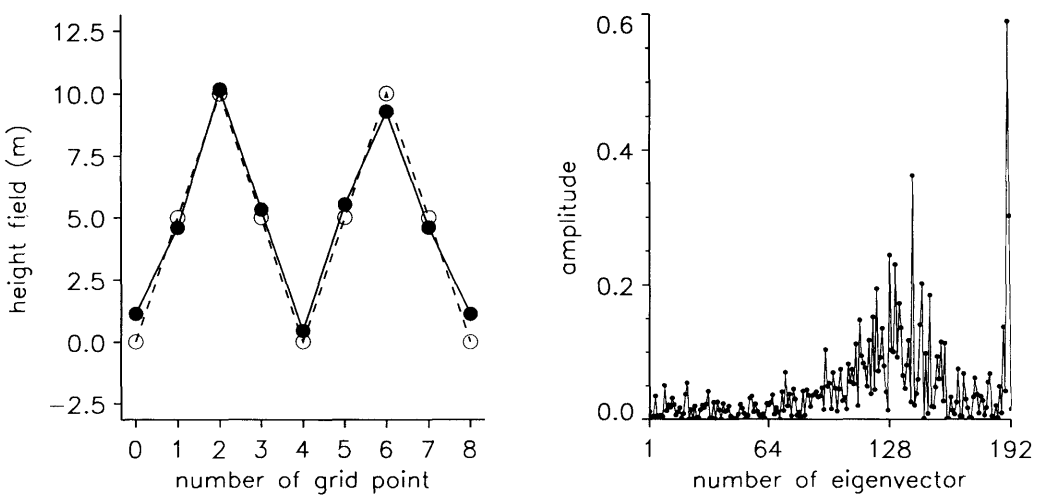

Fig. 10. The same as Fig. 7 except with second order penalty constraint ($\alpha = 1.56 \times 10^{-3}$).

initial model states minimizing (33) and (37), respectively. It turns out that the data misfit is some factor 0.6 smaller in the latter case. What is more, using the second order penalty with the above value for α results in balanced magnitudes of the data component and the penalty component of the cost function.

5. Performance of gradient based minimization algorithms for an ill-conditioned variational problem

The previous section was meant to illustrate how the shape of the nondimensional Hessian reflects the sensitivity of an optimized initial model state with respect to erroneous data or systematic model errors. Let us now consider how the eigenvalue spectrum of the Hessian affects the performance of minimization algorithms based on gradient information.

If gradient information is available, the simplest strategy for minimizing a functional is the steepest descend technique (cf., Gill et al., 1981). This method tries to find the minimum of the functional by a series of successive line searchs along the respective direction of the local negative gradient (direction of steepest descent). Knowing the details of the Hessian matrix it is possible to formulate the amount of relative error reduction for a single step of this algorithm. Let us assume that after m iterations the current approximation for the initial model state is w_0^m (primes are omitted in the following for the sake of convenience). Projecting the remaining error onto the eigenvectors of the Hessian matrix results in its spectral representation F^m:

$$F^m = E^{\mathrm{T}}(w_0^m - w_{0,\mathrm{opt}}). \tag{40}$$

Starting from w_0^m, the steepest descent algorithm will produce an improved approximation w_0^{m+1} according to

$$\begin{aligned} w_0^{m+1} &= w_0^m - \beta^m \frac{\mathrm{d}J}{\mathrm{d}w_0^m} \\ &= w_0^m - \beta^m H(w_0^m - w_{0,\mathrm{opt}}), \end{aligned} \tag{41}$$

with some step size β^m. The new approximation will contain the following error:

$$\begin{aligned} F^{m+1} &= E^{\mathrm{T}}(w_0^{m+1} - w_{0,\mathrm{opt}}) \\ &= E^{\mathrm{T}}(w_0^m - w_{0,\mathrm{opt}}) - \beta^m E^{\mathrm{T}} H(w_0^m - w_{0,\mathrm{opt}}) \\ &= F^m - \beta^m E^{\mathrm{T}} H E E^{\mathrm{T}}(w_0^m - w_{0,\mathrm{opt}}) \\ &= (I - \beta^m D)\, F^m, \end{aligned} \tag{42}$$

D being defined according to (36). According to (42), the updated error spectrum results from the old one after multiplication with a diagonal matrix whose entries are essentially determined by the eigenvalues of the Hessian matrix H. If $1/\beta^m$ happens to be an eigenvalue of H, the error component projecting onto the corresponding eigenvector will be removed completely. Error components projecting onto eigenvectors with eigenvalues $\lambda_i \neq 1/\beta^m$ will be reduced less effectively or may even be amplified. Only if all eigenvalues have the same value (condition number one), the algorithm will be able to find the solution after one single step.

The main problem with the steepest descent algorithm is that for many applications the sequence of step sizes β^m tends to converge towards a series of only two alternating values. This means that the spectral structure of relative

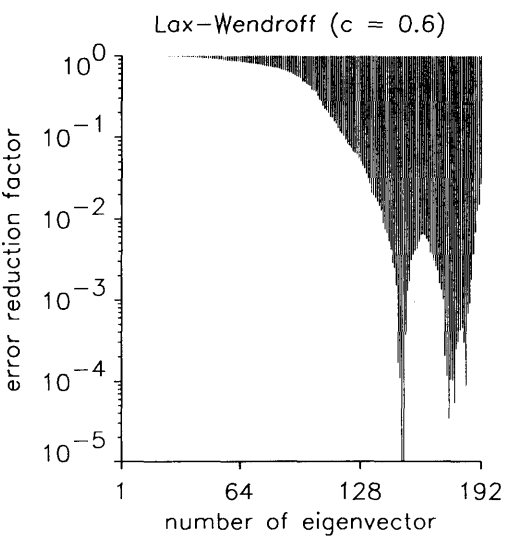

Fig. 11. Performance of the steepest descent algorithm when solving the unpenalized problem whose solution is shown in Fig. 7. The picture shows the ratio of error amplitudes contained in the approximation after 5 steepest descent cycles and in the first guess (zero fields), respectively.

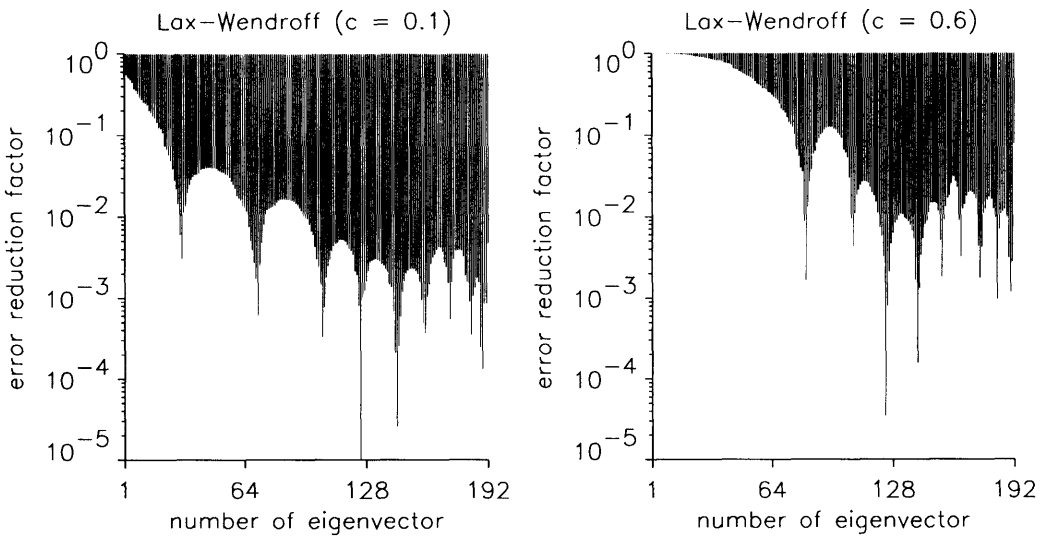

Fig. 12. Right panel: the same as Fig. 11 except with 10 conjugate gradient cycles. Left panel: the same as right panel except with using Courant number 0.1 (instead of 0.6) for inverse modelling.

error reduction is approximately the same after every second iteration making the scheme very inefficient when minimizing a cost function with a wide spectrum of curvatures in different directions. Fig. 11 illustrates the "bi-polar" behaviour of the steepest descent algorithm depicting the spectral error reduction after 5 steepest descent cycles to solve the unpenalized problem whose solution is shown in Fig. 7.

The shortcoming of the steepest descent algorithm to concentrate on error reduction in a confined spectral domain, is partly overcome by the class of conjugate gradient algorithms which determine the new search direction as a weighted combination of the former search direction and the actual gradient. Fig. 12 illustrates the ability of conjugate gradient algorithms to cover with an increasing number of iterations larger parts of the spectrum. However, conjugate gradient algorithms too depend on the condition number of the Hessian matrix. For inverse modelling with Lax-Wendroff's scheme and Courant number 0.6 (right panel, to be compared with Fig. 11) the errors in the very left part of the spectrum are still hardly modified after 10 iterations. Using Courant number 0.1 (left panel) the overall error reduction is much better. In both cases Courant number 0.3 has been chosen to calculate the artificial data.

Interrupting a minimization algorithm after a certain number of iterations one may take advantage of its selective efficiency as an indirect penalty constraint. Courtier and Talagrand (1987) mention this possibility (cf., also the numerical experiments of Li and Droegemeier, 1993).

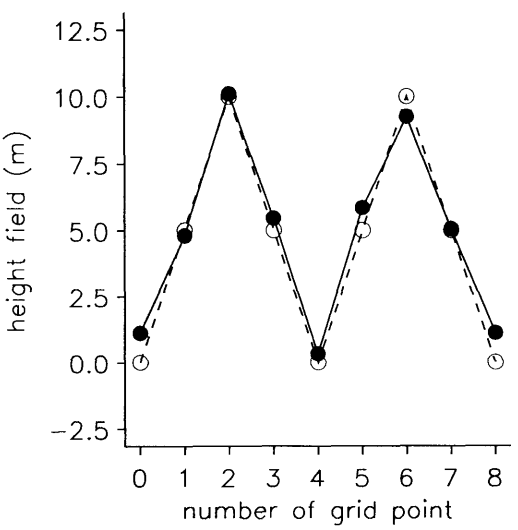

Fig. 13. The same as the left panel of Fig. 7 except that the optimum initial state has been approximated by the result after 10 conjugate gradient cycles starting from first guess zero. The corresponding error reduction is depicted in the right panel of Fig. 12.

A suitable number of iterations would reduce the error components corresponding to large eigenvalues of the Hessian matrix while being insufficient to attack the errors corresponding to small eigenvalues. If one takes care that the first guess prescribes zero amplitudes in the domain of those modes being poorly controlled by data (e.g., first guess identical zero), the remaining error will represent those components of the true solution which we want to suppress. In our example, leaving these error components untouched means to produce smooth fields. To demonstrate the efficiency of the method, Fig. 13 depicts the approximate optimum height field obtained after ten conjugate gradient cycles. The related error reduction is shown in the right panel of Fig. 12. Comparing Figs. 13, 10 shows little difference between the solutions obtained by introducing an explicit second order penalty term or by using a truncated minimization algorithm, respectively.

6. Conclusions

A simple hyperbolic system has been used as an example to elucidate the dependence of variational data assimilation on various ingredients of the entire problem: data distribution, numerical integration scheme and the minimization algorithm. Concerning all these aspects the Hessian matrix of the cost function is the tool at hand to allow the assessment of variational results as well as a systematic analysis of the minimization procedure.

The Hessian matrix encapsulates all information about the sensitivity of the cost function with respect to variations of the control parameters. However, the numerical value for the curvature of the cost function along different directions in the space of independent variables depends on the choice of the underlying metric being intrinsically related to scaling of the model variables.

The purpose of proper scaling can be twofold: either round-off errors during numerical model integration are to be minimized or minimization algorithms shall be accelerated. Both intentions are independent of each other. While dynamic scaling is easily done considering the governing equations, optimum scaling of a variational problem, if at all possible, must always refer to a specific data configuration. Since the level of information about different aspects of a model, reflected by the model error covariance matrix, varies in time, the amount of information, which can be gained from a certain type of measurement, is time dependent, too.

Of course, changing the metric (or scaling) is irrelevant from the physical point of view in the sense that it does not change the amount of information which can be extracted from a given set of observations.

If the observed quantities are not direct counterparts of the model variables, a separate model has to be employed (e.g., a radiation model calculating radiances as functions of the temperature distribution) in order to map the model values onto data counterparts. The sensitivity of this interface model's output with respect to its input may be different for various spectral components. Our studies including spatial interpolation may be seen as a very simple example of this situation. Assuming an observational grid being shifted relative to the computational one, interpolation reduces the amplitudes of non-smooth field components while leaving the amplitudes of large scale structures approximately unmodified. This spectral characteristic of the interpolation procedure entails a widened range of eigenvalues of the Hessian matrix.

Dealing with hyperbolic equations, the reference system can be chosen in order to obtain a skew-symmetric representation of the dynamic matrix, which reflects a conservation law. However, discretizing the model equations for the sake of a numerical treatment may spoil the hyperbolic character of the equations via introduction of some symmetric matrix component. It has been shown that using a non-conservative scheme can modify the eigenvalue spectrum of the nondimensional Hessian in an essential way. Considering the special example of Lax-Wendroff's scheme, a strong dependence of the Hessian matrix on the Courant number could be observed. For large Courant numbers small scale components in the initial state fields will dissipate after a few time steps and thereafter not affect the calculated counterparts of observations. The conclusion is that, when performing inverse calculations on the basis of Lax-Wendroff type schemes, one should be especially careful concerning the appearance of unrealistic rough fields of the control variables.

The general remedy for ill-conditioned problems is to introduce some prior knowledge about the control variables (Lorenc, 1986). In operational

implementations of data assimilation schemes a numerical weather prediction model will provide accurate and dense background information. For many other studies such information will not be available and one has to derive prior knowledge from other sources. If small eigenvalues of the nondimensional Hessian are clearly related to small scale structures near the maximum resolution of the numerical grid, a penalty term may be chosen which acts scale dependent (Courtier and Talagrand, 1987; Wergen, 1992). An example would be the squared Laplacian. If the magnitude of the penalty is appropriately adjusted, the constraint will produce a regularized solution which is smooth with nearly unmodified large scale components. The solution will be stable with respect to small modifications of the data input.

An analysis of the performance of a conjugate gradient algorithm for an ill-posed variational problem reveals that the undesired poorly controlled structures remain untouched during the first few cycles of the minimization process. This observation offers the opportunity to use the number of minimization steps as a control parameter to suppress unrealistically sensitive modes of the initial fields. After a limited number of iterations the adjusted control parameters do hardly depend on whether a penalty term has been used or not. However, without an appropriate penalty constraint the actual approximation will just represent the true solution's projection onto some subdomain in the space of independent variables. The use of a penalty term makes the obtained result stable with respect to further minimization steps as well as with respect to modifications of the first guess.

This study presented various examples which indicate that in any realistic assimilation problem one has to be aware of poorly controlled initial field components, even when the data outnumber the control parameters. Therefore use of a penalty term or other background information in the formulation of the cost function is deemed indispensible. Though the reason for specific ill-conditioning will stay unknown in most practical applications, the examples studied in the present paper convey some idea of situations which should be avoided generally. First, the probability of encountering an ill-posed problem will be high if data are sampled according to a regular schedule. This point has been discussed in detail by Thacker and Lewandowicz (1994). Second, an observational grid should never be shifted systematically with respect to the computational grid. In the extreme case of data being centered between computational grid points, the shortest waves on the computational grid are unvisible on the observational grid so that their amplitude remains uncontrolled by the data. Such a situation might possibly occur when feeding data into a model formulated on a staggered grid.

It is interesting to note that interpolation of data onto the computational grid may be seen as penalizing rough structures. Interpolation of smooth data will hardly modify them. However, the situation is different for small scale information. To illustrate this, consider the hypothetical situation where observations very close to (but not exactly at) the center of each computational grid cell reveal regular alternating values $+1/-1$, i.e., a pure 2Δ-wave. It is possible to define grid values such that the measurements are reproduced. However, the retrieval produces an unrealistically large amplitude approaching infinity if the observation points approach the cell centers. Let us now presume, that before starting the assimilation cycle the raw data have been interpolated onto the computational grid. Interpolation will extinguish a good deal of the observed signal, and we will end up with low-amplitude preprocessed information on the computational grid controlling the smale scale model components effectively. Using the raw data in combination with a suitable penalty constraint would produce comparable results. However, the formulation of an extra penalty term is always somewhat artificial, and it may sometimes appear advantageous to apply interpolation to cure the basic problem in a direct and obvious way.

7. Acknowledgements

Numerous clarifying discussions with W. C. Thacker and R. Lewandowicz are greatfully appreciated as well as the anonymous reviewer's comments. One of us (D.P.E.) was partially supported by NATO travel grant 0487/87.

8. Appendix A

Preconditioning for a 1-dimensional problem involving spatial interpolation

A data assimilation problem is considered where the vector m of data counterparts arises from the vector w of independent variables by simple spatial interpolation. The model vector w is assumed to consist of four independent components being alligned with one space dimension. Assuming periodic boundaries, linear interpolation is described by the following matrix multiplication,

$$\begin{pmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \end{pmatrix} = G \begin{pmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \end{pmatrix}$$

with

$$G = \begin{pmatrix} (1-\varepsilon) & \varepsilon & 0 & 0 \\ 0 & (1-\varepsilon) & \varepsilon & 0 \\ 0 & 0 & (1-\varepsilon) & \varepsilon \\ \varepsilon & 0 & 0 & (1-\varepsilon) \end{pmatrix}. \quad \text{(A.1)}$$

The parameter ε may adopt values between 0 and 1. We wish to specify the four components w_i of the model vector such that the resulting interpolated values m_j agree with four observations d_j. The model state w we are looking for will minimize the following objective function,

$$J(w) = \tfrac{1}{2}(Gw-d)^{\mathrm{T}}\,\Sigma^{-1}(Gw-d), \quad \text{(A.2)}$$

where Σ denotes some positive symmetric matrix decribing the covariances of the data uncertainties. The Hessian matrix of this cost function reads,

$$H = G^{\mathrm{T}}\Sigma^{-1}G. \quad \text{(A.3)}$$

Suppose that, in order to update our current approximation of the optimum model state, we use some gradient method. We try to improve our model fit searching along some direction ϕ,

$$\phi = P\frac{\mathrm{d}J}{\mathrm{d}w} = PH(w-w_{\mathrm{opt}}), \quad \text{(A.4)}$$

which results from the gradient of the cost function after multiplication with some preconditioning matrix P. Note that the product PH equals the curvature matrix $\hat{H}$ when P is chosen to equal the metric tensor C_{W} (cf., (9)). Therefore the direction ϕ defined by (A.4) may be considered as the direction of steepest descent with respect to some metric based on the inverse of the preconditioning matrix P (Tarantola, 1987).

In order to obtain an efficient optimization algorithm, $PH \simeq I$ should hold, i.e., the preconditioning matrix should approximate the inverse of the Hessian matrix. For $\Sigma \simeq I$, a reasonable choice of P would be such that the product PH becomes independent of the actual position of the given data. This can be achieved by defining

$$P = \Sigma^{-1}(G^{\mathrm{T}}\Sigma^{-1}G)^{-1}. \quad \text{(A.5)}$$

Let us restrict ourselves to the simple situation with $\Sigma = I$. Then the preconditioning matrix P can be calculated explicitly. One obtains

$$P = (G^{\mathrm{T}}G)^{-1} = \frac{1}{(1-2\varepsilon)^2} \times \begin{pmatrix} p_1 & p_2 & p_3 & p_2 \\ p_2 & p_1 & p_2 & p_3 \\ p_3 & p_2 & p_1 & p_2 \\ p_2 & p_3 & p_2 & p_1 \end{pmatrix}, \quad \text{(A.6)}$$

with three independent matrix elements being defined as

$$p_1 = \frac{1-2\varepsilon(1-\varepsilon)(\varepsilon^2-\varepsilon+2)}{1-2\varepsilon(1-\varepsilon)},$$
$$p_2 = -\varepsilon(1-\varepsilon), \quad \text{(A.7)}$$
$$p_3 = \frac{2\varepsilon^2(1-\varepsilon)^2}{1-2\varepsilon(1-\varepsilon)}.$$

For $\varepsilon = 0$ or $\varepsilon = 1$ (i.e., no interpolation) the preconditioning matrix P equals the identity, i.e., $p_2 = p_3 = 0$. However, the general expression for P is rather complex even when considering our very simple example. If ε approaches 0.5, the off-diagonal elements of P become more and more important. The absolute values of all matrix elements with alternating signs tend to infinity while differences between them remain small. This means that numerical calculations involving the preconditioning matrix will become more and more inaccurate for ε approaching 0.5. For $\varepsilon = 0.5$ the interpolation matrix G is singular and expression (A.5) becomes undefined.

9. Appendix B

Eigenvectors of Lax-Wendroff's scheme with eigenvalue zero

Fig. 5 shows that some eigenvalues of the Hessian may vanish if we apply Lax-Wendroff's method with Courant number $1/2$ or $1/\sqrt{2}$ to our problem. A vanishing eigenvalue of the Hessian matrix means that, if the corresponding eigenvector is added to some current guess of the initial model state, this will not modify the calculated values m_v of data counterparts. The reason for this is that Lax-Wendroff's algorithm will map certain small scale structures onto zero, i.e., the scheme has eigenvectors with eigenvalue zero. This phenomenon will be elucidated in the following.

According to Lax-Wendroff's scheme, starting from some model state w^n at time level n, the model state w^{n+1} one time step Δt later is predicted as (e.g., Anderson et al., 1984):

$$w^{n+1} = w^n + \Delta t \frac{\partial w^n}{\partial t} + \frac{1}{2}(\Delta t)^2 \frac{\partial^2 w^n}{\partial t^2}. \tag{B.1}$$

Using the model equations (23)–(24) and replacing spatial derivatives with central differences leads to the following formulas holding at each grid point i, j (for $\Delta x = \Delta y$),

$$\begin{aligned} u_{i,j}^{n+1} = u_{i,j}^n &+ \alpha \begin{bmatrix} 0 & 0 & 0 \\ 1 & -2 & 1 \\ 0 & 0 & 0 \end{bmatrix} u_{i,j}^n \\ &+ \beta \begin{bmatrix} -1 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & -1 \end{bmatrix} v_{i,j}^n \\ &+ \gamma \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & -1 \\ 0 & 0 & 0 \end{bmatrix} h_{i,j}^n, \end{aligned} \tag{B.2}$$

$$\begin{aligned} v_{i,j}^{n+1} = v_{i,j}^n &+ \beta \begin{bmatrix} -1 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & -1 \end{bmatrix} u_{i,j}^n \\ &+ \alpha \begin{bmatrix} 0 & 1 & 0 \\ 0 & -2 & 0 \\ 0 & 1 & 0 \end{bmatrix} v_{i,j}^n \\ &+ \gamma \begin{bmatrix} 0 & -1 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} h_{i,j}^n, \end{aligned} \tag{B.3}$$

$$\begin{aligned} h_{i,j}^{n+1} = h_{i,j}^n &+ \delta \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & -1 \\ 0 & 0 & 0 \end{bmatrix} u_{i,j}^n \\ &+ \delta \begin{bmatrix} 0 & -1 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} v_{i,j}^n \\ &+ \alpha \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix} h_{i,j}^n, \end{aligned} \tag{B.4}$$

where $\alpha = c^2/2$, $\beta = c^2/8$, $\gamma = (g\,\Delta t)/(2\,\Delta x)$ and $\delta = (h_r\,\Delta t)/(2\,\Delta x)$ (for unscaled equations) with Courant number $c = \sqrt{gh_r}\,\Delta t/\Delta x$.

Examination of those eigenvectors of the Hessian matrix which are related to the zero eigenvalues depicted in Fig. 5 reveals that all these model states exhibit 2Δ-structures. We wish to demonstrate that the diffusion terms in (B.2–B.4) being proportional to α are those which extinguish such small scale structures for certain Courant numbers.

Let us first consider an initial model state with zero velocity fields. For the height field we presume a general pattern representing the shortest waves that can be resolved on our grid,

$$\begin{array}{c} \vdots\ \vdots\ \vdots\ \vdots \\ \cdots\ a\ b\ a\ b\ \cdots \\ \cdots\ c\ d\ c\ d\ \cdots \\ \cdots\ a\ b\ a\ b\ \cdots \\ \cdots\ c\ d\ c\ d\ \cdots \\ \vdots\ \vdots\ \vdots\ \vdots \end{array}, \tag{B.5}$$

where a, b, c and d are arbitrary numbers. Note, that according to (B.2, B.3), advancing this height field will not stir a velocity field. For certain values

of the Courant number, the height field arising from integrating the model a single time step will vanish, too. From (B.4) we get the following condition for a, b, c and d to be mapped onto zero:

$$\begin{pmatrix} (1-4\alpha) & 2\alpha & 2\alpha & 0 \\ 2\alpha & (1-4\alpha) & 0 & 2\alpha \\ 2\alpha & 0 & (1-4\alpha) & 2\alpha \\ 0 & 2\alpha & 2\alpha & (1-4\alpha) \end{pmatrix} \times \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \tag{B.6}$$

In order to allow for a non-zero solution (a, b, c, d) of (B.6), the determinant of the matrix on the left hand side of (B.6) must vanish,

$$(1-8\alpha)(1-4\alpha)^2 = 0. \tag{B.7}$$

With $\alpha = c^2/2$ the solutions of (B.7) determine the critical values c_1 and $c_{2/3}$ of the Courant number:

$$c_1 = \frac{1}{2}, \qquad c_{2/3} = \frac{1}{\sqrt{2}}. \tag{B.8}$$

The specific structure of the related height fields can be specified from (B.6). Using $\alpha_1 = 1/8$ one obtains

$$(a, b, c, d)_1 \sim (1, -1, -1, 1), \tag{B.9}$$

which corresponds to a checkerboard pattern. For $\alpha_{2/3} = 1/4$ two orthogonal eigenvectors with zero eigenvalues are found,

$$\begin{aligned} (a, b, c, d)_2 &= (1, 1, -1, -1), \\ (a, b, c, d)_3 &= (1, -1, 1, -1), \end{aligned} \tag{B.10}$$

which, according to pattern (B.5), represent 1-dimensional 2Δ-waves in x- and y-direction, respectively.

A similar analysis can be done considering hypothetical initial model states whose only non-zero components are small scale u or v fields, respectively. This analysis will be simplified by the fact that the diffusion operators in (B.2, B.3) are 1-dimensional. For both equations $c = 1/\sqrt{2}$ is retained as a critical Courant number. The related patterns that are extinguished by the scheme consist of alternating values $a, -a, a, -a, \ldots$ in one spatial direction. Since the diffusion operator for u (or v) in (B.2, B.3) does not provide any coupling between different rows (or columns) of the field matrix, the parameter a of the alternating pattern can be chosen for each row (column) independently. This leads to a subspace of singular u-fields (v-fields) being proportional to the number N_I (N_J) of rows (columns) of the numerical grid. Therefore, for Courant number $1/\sqrt{2}$ the total number N of eigenvectors with zero eigenvalues amounts to $N = 2 + N_I + N_J$, including the two singular height patterns. With $N_I = N_J = 8$ we obtain 18 zero eigenvalues in accordance with Fig. 5.

REFERENCES

Anderson, D. A., Tannehil, J. C. and Pletcher, R. H. 1984. *Computational fluid mechanics and heat transfer.* McGraw-Hill, New York, 599 pp.

Bennet, A. F. and Miller, R. N. 1991. Weighting initial conditions in variational assimilation schemes. *Mon. Wea. Rev.* **119**, 1098–1102.

Browning, G. L., Holland, W. R., Kreiss, H.-O. and Worley, S. J. 1990. An accurate hyperbolic system for approximately hydrostatic and incompressible oceanographic flows. *Dynamics of Atmospheres and Oceans* **14**, 303–332.

Courtier, P. and Talagrand, O. 1987. Variational assimilation of meteorological observations with the adjoint vorticity equation. II: Numerical results. *Q. J. Roy. Meteorol. Soc.* **113**, 1329–1347.

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

Ghil, M. and Malanotte-Rizzoli, P. 1991. Data assimilation in meteorology and oceanography. *Adv. in Geophys.* **33**, 141–265.

Kapitza, H. 1991. Numerical experiments with the adjoint of a nonhydrostatic mesoscale model. *Mon. Wea. Rev.* **119**, 2993–3011.

Li, Y. and Droegemeier, K. K. 1993. The influence of diffusion and associated errors on the adjoint data assimilation technique. *Tellus* **45A**, 435–448.

Long, R. B. and Thacker, W. C. 1989. Data assimilation into a numerical equatorial ocean model, Part I. The model and the assimilation algorithm. *Dyn. Atmos. Oceans* **13**, 379–412.

Lorenc, A. 1986. Analysis methods for numerical weather prediction. *Q. J. Roy. Meteorol. Soc.* **112**, 1177–1194.

Talagrand, O. 1981. A study of the dynamics of four-dimensional data assimilation. *Tellus* **33**, 43–60.

Tarantola, A. 1987. *Inverse problem theory; methods for data fitting and model parameter estimation.* Elsevier Science Publishers, Amsterdam, 613 pp.

Thacker, W. C. 1989. The role of the Hessian matrix in fitting models to measurements. *J. Geophys. Res.* **94**, 6177–6196.

Thacker, W. C. and Lewandowicz, R. 1994. Dynamics of information and uncertainty. *Tellus* **46A**, 651–670.

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

Tikhonov, A. N. and Arsenin, V. Y. 1977. *Solutions of ill-posed problems.* Halsted Press, New York, 258 pp.

Wahba, G. and Wendelberger, J. 1980. Some new mathematical methods for variational objective analysis using splines and cross validation. *Mon. Wea. Rev.* **108**, 1122–1143.

Wergen, W. 1992. The effect of model errors in variational assimilation. *Tellus* **44A**, 297–313.