Solving for the generalized inverse of the Lorenz model

Geir Evensen and Nabil Fario

Nansen Environmental and Remote Sensing Center, Bergen, Norway

In print, J. Meteor. Soc. Japan

June 5, 1995

Geir Evensen
Nansen Environmental and Remote Sensing Center
Edvard Griegsvei 3a,
N–5037 Solheimsviken
Norway
Phone: +47-55-297288
fax: +47-55-200050
e-mail: Geir.Evensen@nrsc.no

Abstract

Advanced data assimilation becomes extremely complicated and challenging when used with strongly nonlinear models. Several previous works have reported various problems when applying existing popular data assimilation techniques with strongly nonlinear dynamics. Common for these techniques is that they can all be considered as extensions to methods which have proven to work well with linear dynamics. This paper shows that a weak constraint variational formulation for the Lorenz model, where the full model state in space and time is considered as control variables, can be minimized using a gradient descent method. It is further shown that the weak constraint formulation removes some of the previous reported problems associated to the predictability limit of nonlinear models when strong constraint formulations are used. Further, by using a gradient descent method, problems associated to the use of an approximate tangent linear model when solving the Euler-Lagrange equations or when the extended Kalman filter is used, are eliminated, since the solution is found without integration of any dynamical equations. The method works well with reasonable data coverage and quality of the measurements, however, with poorer data coverage a statistical minimization method, simulated annealing, may be used to search for the global minimum.

1 Introduction

The celebrated Lorenz model has been subject of extensive studies motivated by its chaotic and strongly nonlinear nature. In the field of data assimilation the model has served as a testbed for examining the properties of various data assimilation methods when used with strongly nonlinear dynamics. The results have been used to suggest properties and possibilities of the methods for applications with oceanic and atmospheric models which may also be strongly nonlinear and chaotic.

In Gauthier (1992) the so called adjoint method, which solves a variational minimization problem where the model is imposed as a strong constraint, was tested with the Lorenz (1963) equations. He found that in a case where the model did not undergo a transition, the cost function was relatively well-behaved with respect to perturbations in the control variables, i.e. initial conditions for the three model variables. However, in a case where the model did undergo transitions he observed a very strong sensitivity of the cost function with respect to perturbations in the initial conditions, i.e. the value of the cost function depended strongly and nonlinearly on the control variables and contained local minima.

Miller *et al.* (1994) re-examined the adjoint method with the Lorenz model and proved that this behavior of the cost function was strongly dependent on the length of the assimilation time interval. While Gauthier (1992) presented an example where he used data in the time interval $t \in [0.0, T]$ with $T = 8.0$, and compared two cases, one with and one without transitions in

the reference case, Miller *et al.* (1994) examined three cases with chaotic behavior and $T = 8$, $T = 10$, and $T = 15$. These examples proved that the cost function became increasingly more sensitive with respect to small perturbations in the initial conditions as T was increased.

Stensrud and Bao (1992) compared nudging and the adjoint method with the Lorenz model and also pointed out that the adjoint method was sensitive to the length of the assimilation window. However, they used a rather short time interval with only a few phase transitions so they did not observe the extreme behavior of the cost function as was found by Miller *et al.* (1994). In their nudging experiments they found strong dependence of the results on the data density and the choice of nudging coefficients.

Miller *et al.* (1994) also gave a comprehensive discussion on applications of the extended Kalman filter with the Lorenz model. They found that the statistical linearization used in the extended Kalman filter, when deriving the error covariance evolution equation, resulted in a too simplified closure. The estimated solution was unreliable beyond $t = 11.0$. This was essentially explained by a poor prediction of error covariances resulting in insufficient gain because of a decaying mode which reflects the stability of the attractor. A generalization of the extended Kalman filter, where third and fourth order moments and evolution equations for these were included, was also examined. It was shown that this more sophisticated closure scheme provided a more consistent evolution of error statistics which also resulted in sufficient gain to keep the estimate on track. Unfortunately such an approach is not practical for a high dimensional ocean or atmospheric model, since the fourth order moment requires storage of N^4 elements (N is the number of state variables), compared to the second order moment used in the extended Kalman filter where N^2 elements must be stored.

Bürger and Cane (1994) proposed another fix for the extended Kalman filter to correct for the too low gain resulting from the decaying mode. The error evolution in the extended Kalman filter is based on the tangent linear operator of the model evaluated at the current state estimate. One can then imagine a situation where the current state is in a stable region yielding decrease of the variance, while the true state is in an unstable region, e.g. undergoing a phase transition, during which the variance should increase. By evaluating the tangent linear operator at members of an ensemble of probable states and predicting an ensemble of error covariance forecasts which then are used in a weighted average to give the new error covariance forecast, it was hoped that the growing modes should be accounted for if the state where close to the unstable regions. They found that dependent on the choice of weights and ensemble they could improve the performance of the filter. It should be noted that this method is not a closure of the error covariance equation in the extended Kalman filter. The error covariance evolution in the extended Kalman filter is based on linearized dynamics and in unstable regions it will also generate linear and unbounded error variance growth. It is only the special properties of the linearized Lorenz model with its decaying mode in the stable regions of the phase space, together with the short time period

the state is located in the unstable regions, which keeps the error variances within physically acceptable values.

In Pires *et al.* (1996) an approach was proposed for extending the assimilation interval for a strong constraint formulation. The method named quasi-static variational assimilation, was based on increasing the assimilation interval step by step, using the prior estimate as the first guess when solving the variational problem for the new increased assimilation interval. They found that this method could extend the assimilation interval by some amount but there is still an upper limit where the strong constraint formulation fails to track every transition.

Other problems related to error variance prediction have also been observed with more realistic applications of the extended Kalman filter, e.g., Evensen (1992) and Gauthier *et al.* (1993) who used the extended Kalman filter with quasi-geostrophic models. Evensen (1992) pointed out that by evaluating the model operator at an unstable, say sheared, background flow this resulted in unbounded error variance growth. So in general a more consistent closure is needed in the error covariance equation. An alternative may be to use the ensemble Kalman filter that was recently proposed by Evensen (1994a, 1994b), and applied by Evensen and van Leeuwen (1996). The ensemble Kalman filter is based on a Monte Carlo approach for solving the Kolmogorov's equation, which is the equation for evolution of the probability density function for the error statistics, and there is not any need for a closure approximation. This method has been found to work well with the Lorenz model (R. N. Miller personal communication).

This paper partly addresses whether the problems reported in these studies can be attributed to either the particular formulation used or to the solution method. We will illustrate that the strong sensitivity of the cost function with respect to initial conditions in the adjoint method is a consequence of imposing the model as a strong constraint. Obviously it must be difficult to choose an initial condition that leads to a model trajectory following a reference solution for more than a few state transitions. In the next section a weak constraint variational formulation where also the model errors are penalized is given. The error covariances used in the inverse formulation are discussed in Section 3, and in Section 4 the application of a gradient descent method is briefly discussed. Examples are given in Section 5, where also an alternative solution method based on Monte Carlo methods and simulated annealing is briefly discussed. This statistical approach can also be used to generate error estimates.

2 Generalized Inverse Formulation

The Lorenz model consists of a system of three coupled and nonlinear ordinary differential equations

$$\begin{aligned} \frac{dx}{dt} &= \sigma(y - x) + q^x, \\ \frac{dy}{dt} &= \rho x - y - xz + q^y, \\ \frac{dz}{dt} &= xy - \beta z + q^z. \end{aligned} \tag{1}$$

Here $x(t)$, $y(t)$, and $z(t)$ are the dependent variables, and we have chosen the following commonly used values for the parameters in the equation, $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$. The terms $q^x(t)$, $q^y(t)$, and $q^z(t)$ are assumed to represent the unknown model errors. Initial conditions for the model are given as

$$\begin{aligned} x(0) &= x_0 + a^x, \\ y(0) &= y_0 + a^y, \\ z(0) &= z_0 + a^z, \end{aligned} \tag{2}$$

where x_0, y_0, and z_0 are the first guess values of the initial conditions and the terms a^x, a^y, and a^z represent the errors in the first guess initial conditions. If all the error terms were known or equal to zero, these equations would formulate a well-posed problem having a unique solution in a mathematical sense.

Now a set of measurements, $\boldsymbol{d} \in \Re^M$, of the true solution are assumed given and linearly related to the model variables by the measurement equation

$$\boldsymbol{d} = \boldsymbol{\mathcal{L}}[x, y, z] + \boldsymbol{\epsilon}, \tag{3}$$

where $\mathcal{L} \in \Re^M$ is a linear measurement functional, $\boldsymbol{\epsilon} \in \Re^M$ is a vector of measurement errors, and M is the number of measurements.

If all the error terms are equal to zero the problem will be over-determined, and no solution can be found in general. However, by allowing the dynamics, the initial conditions and the measurements to contain errors, a solution can be found that minimizes these error terms in a weighted least squares sense, by minimizing the following variational integral,

$$\mathcal{J}[x, y, z] = \int_0^T dt_1 \int_0^T dt_2 \; \boldsymbol{q}(t_1)^{\mathrm{T}} \, \boldsymbol{W}_{qq}(t_1, t_2) \, \boldsymbol{q}(t_2) + \boldsymbol{a}^{\mathrm{T}} \, \boldsymbol{W}_{aa} \, \boldsymbol{a} + \boldsymbol{\epsilon}^{\mathrm{T}} \boldsymbol{w} \boldsymbol{\epsilon}. \tag{4}$$

Here the weights $\boldsymbol{W}_{qq}(t_1, t_2)$ and $\boldsymbol{W}_{aa}$ are 3×3 weight matrices, which are optimally defined as the inverses of the model error covariance matrix and the error covariance matrix for the initial conditions respectively. The $M \times M$ matrix $\boldsymbol{w}$ is the inverse of the measurement error covariance

matrix.

Note that other estimators than least squares could be defined. However, the least-squares formulation is attractive for several reasons. If the unknown errors are Gaussian, i.e., completely explained by the two first statistical moments, mean and covariance, then minimizing (4) is equivalent to finding the maximum likelihood estimate for x, y, and z. When working with methods that involve the Euler–Lagrange equations these are readily derived and the derivatives of the penalty function exist everywhere.

Since there is no integration of the model equations required in the methods used in this paper, very simple numerical discretizations based on second order centered differences for the time derivatives will be used, i.e.,

$$\begin{aligned}
\frac{x_{n+1} - x_{n-1}}{2\Delta t} &= \sigma(y_n - x_n) + q_n^x, \\
\frac{y_{n+1} - y_{n-1}}{2\Delta t} &= \rho x_n - y_n - x_n z_n + q_n^y, \\
\frac{z_{n+1} - z_{n-1}}{2\Delta t} &= x_n y_n - \beta z_n + q_n^z,
\end{aligned} \tag{5}$$

where $n = 2, N-1$ is the time-step index with N the total number of time steps. We now define the error vectors $\boldsymbol{q}_n^{\mathrm{T}} = (q_n^x, q_n^y, q_n^z)$ and $\boldsymbol{a}^{\mathrm{T}} = (a^x, a^y, a^z)$.

As a crude assumption made for convenience, the time correlation of the model errors are neglected by writing

$$\boldsymbol{W}_{qq}(t_1, t_2) = \widehat{\boldsymbol{W}}_{qq}\delta(t_1 - t_2), \tag{6}$$

where $\widehat{\boldsymbol{W}}_{qq}$ is a constant 3×3 matrix. A correlation in time of the model errors will have a time regularizing effect on the inverse estimate which is now removed in this expression. To ensure a smooth solution in time the regularization is instead accounted for by a smoothing term

$$\mathcal{J}_S[\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}] = \Delta t \sum_{n=1}^{N} \boldsymbol{\eta}_n^{\mathrm{T}} \boldsymbol{W}_{\eta\eta} \boldsymbol{\eta}_n, \tag{7}$$

where $\boldsymbol{\eta}_n^{\mathrm{T}} = (\eta_n^x, \eta_n^y, \eta_n^z)$ with,

$$\eta_n^x = \frac{x_{n+1} - 2x_n + x_{n-1}}{\Delta t^2}, \tag{8}$$

and the vector $\boldsymbol{x}^{\mathrm{T}} = (x_1, x_2, \ldots, x_N)$, and similarly for $\boldsymbol{y}$ and $\boldsymbol{z}$. The smoothing constraint and the weigths used are further discussed in the next section.

The penalty function now becomes

$$\mathcal{J}[\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}] = \Delta t \sum_{n=1}^{N} \boldsymbol{q}_n^{\mathrm{T}} \widehat{\boldsymbol{W}}_{qq} \boldsymbol{q}_n + \boldsymbol{a}^{\mathrm{T}} \boldsymbol{W}_{aa} \boldsymbol{a} + \boldsymbol{\epsilon}^{\mathrm{T}} \boldsymbol{w} \boldsymbol{\epsilon} + \Delta t \sum_{n=1}^{N} \boldsymbol{\eta}_n^{\mathrm{T}} \boldsymbol{W}_{\eta\eta} \boldsymbol{\eta}_n. \tag{9}$$

where $\boldsymbol{q}_1$, $\boldsymbol{q}_N$, $\boldsymbol{\eta}_1$ and $\boldsymbol{\eta}_N$ are defined as in (5) and (8) but using second order one-sided difference

formulas. The initial conditions and the measurements are included as before except that the measurement operator now must be considered as a matrix multiplied with the state vector consisting of all the discrete elements of $\boldsymbol{x}$, $\boldsymbol{y}$, and $\boldsymbol{z}$.

3 Error Covariances

The reference solution is calculated using a highly accurate ordinary differential equation solver. By taking the reference solution as the true state, the only contribution to the dynamical error term, $\boldsymbol{q}_n$, is the errors introduced in the discretization (5) which are due to the use of an approximate discrete time derivative.

When knowing the exact solution it is possible to calculate the numerical misfits in the model equations (5), and the error statistics of these misfits should be used to determine the weight matrices $\widehat{\boldsymbol{W}}_{qq}$ and $\boldsymbol{W}_{\eta\eta}$ used in the inverse calculation.

The centered first derivative approximation used for the time derivatives in the discrete model equations (5), is

$$\left.\frac{\partial f}{\partial t}\right|_t = \frac{f(t+\Delta t) - f(t-\Delta t)}{2\Delta t} + \left.\frac{1}{6}\frac{\partial^3 f}{\partial t^3}\right|_t \Delta t^2 + \ldots, \tag{10}$$

where the lowest order error term is included.

Given the exact reference solution we have used two methods for estimating the dynamical misfits in the discretization (5). The first method feeds the exact reference solution directly into the discrete equations to find the misfit in each time step, $\boldsymbol{q}_n$. The second method evaluates the error term in (10), given the exact reference solution. In Figure 1 these dynamical misfits are plotted for two different time steps. Clearly, the errors increase with the length of the time step and the maximum errors are located at the peaks of the reference solution. The two almost identical curves for $\Delta t = 0.033$ are generated using the two different approaches just described.

The error covariance matrix $\widehat{\boldsymbol{Q}}_{qq}$ can be calculated from a long time series of these errors, and is of course dependent of the time step used. In the experiments presented here we used a time step of $\Delta t = 0.01667$ and the corresponding error covariance matrix then becomes

$$\widehat{\boldsymbol{Q}}_{qq} = \begin{bmatrix} 0.1491 & 0.1505 & 0.0007 \\ 0.1505 & 0.9048 & 0.0014 \\ 0.0007 & 0.0014 & 0.9180 \end{bmatrix}, \tag{11}$$

where the integration has been performed for a long time interval $t \in [0, 1667]$, i.e., 100000 time steps. The inverse of this matrix is used for $\widehat{\boldsymbol{W}}_{qq}$ in the penalty function (9).

The errors are also clearly correlated in time. In Figure 2 the auto-correlation functions for the x, y, and z components of the dynamical errors are plotted. Since it is inconvenient to use a full space and time covariance matrix, we introduced the smoothing term (7), which act as a

regularization term on the minimizing solution.

It can be shown that a smoothing norm of the type

$$||\psi|| = \int_0^T \psi^2 + \gamma \psi_{tt}^2 dt \tag{12}$$

has a Fourier transform equal to

$$\hat{\psi} = (1 + \gamma f^4)^{-1}. \tag{13}$$

The limiting behavior for increasing frequency, f, is then proportional to $(\gamma f^4)^{-1}$, thus high frequencies are penalized most strongly in the smoothing norm. The ψ^2 term is added here, as a first guess penalty, for illustrational purposes. Without this term, the limiting behavior for $f \to 0$ will be singular, and the corresponding auto-correlation function becomes very flat. In the actual inverse formulation, the dynamical and initial residual will provide the first guess penalty, ensuring a well-behaved limiting behavior when $f \to 0$.

Note that at least for linear dynamics there is a unique correspondence between such a smoothing norm and an error covariance matrix, as shown by McIntosh (1990). Thus, a more consistent way of implementing the time correlation would have been to apply the smoothing constraint to the dynamical misfits instead of the actual inverse estimate. This would have ensured a smooth inverse estimate consistent with the unique error covariance corresponding to the smoothing norm choosen. These issues should be further investigated.

On the other hand will a regularization term which acts directly on the inverse estimate significantly improve the conditioning of the gradient descent method used, since searches in the space of nonsmooth functions are heavely penalized. Since the dynamical error terms are strongly dependent on the model solution, it is assumed that the auto-correlation function for the inverse estimate will be similar to the auto-correlation function found for the dynamical misfits. Thus, the weigths used for the smooting of the inverse estimate should correspond to auto-correlation functions which are similar to those for the dynamical misfits.

An inverse Fourier transform of the spectrum (13) gives an auto-correlation function which is shown in Figure 2 for two values of γ, i.e., $\gamma = 0.0008$ for curve A and $\gamma = 0.00001$ for curve B. For $\gamma = 0.0008$ the auto-correlation function has a similar half width to the auto-correlation functions of the dynamical errors. However, it turned out that for this value of γ the inverse estimate became too smooth, i.e., the peaks in the solutions were to low compared to the reference solution. We decided to use $\gamma = 0.00001$, which gave an inverse estimate more in agreement with the reference solution. Based on the time series of dynamical misfits in Figure 1, it is also clear that the errors are rather smooth for most of the time, while they have sudden changes close to the peaks of the reference solution. The computed auto-correlation function will describe an "average" smoothness of the dynamical misfits which are too smooth near the peaks in the reference solution. This can then justify the use of the smaller smoothing weight $\gamma = 0.00001$.

The smoothing weight matrix is chosen to be diagonal and given by $\boldsymbol{W}_{\eta\eta} = \gamma \boldsymbol{I}$.

The error covariance matrix $\boldsymbol{W}_{aa}^{-1}$ for the errors in the initial conditions, and the measurement error covariance matrix $\boldsymbol{w}^{-1}$ are both assumed to be diagonal and with the same error variance equal to 2.0.

4 Solution by Gradient Descent

A very simple approach for minimizing the penalty function (9) is to use a gradient descent algorithm. The gradient of $\mathcal{J}[\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}]$ with respect to the full state vector $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ in time is easily derived. When the gradient $\nabla_{(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z})}\mathcal{J}$ is known it can be used in a descent algorithm to search for the minimizing solution $(\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}, \hat{\boldsymbol{z}})$.

Gradient descent algorithms have previously been used to minimize strong constraint variational problems, using the popular adjoint method (Talagrand and Courtier, 1987). They showed that by a forward integration of the model followed by a backward integration of the adjoint equation, the adjoint variable at time $t = 0$ was also the gradient of the cost function with respect to the initial conditions. It was then possible to use this gradient in an iterative technique based on gradient descent to improve the initial conditions. When this method is used with nonlinear models one will have to store the full state in space and time since the backward equation uses the adjoint of the tangent linear operator evaluated at the current state estimate. In addition, two model integrations (one forward and one backward) for each iteration in the descent algorithm must be performed. There is also a limit on the length of the time interval for the validity of the tangent linear operator in addition to the sensitivity problems discussed in the introduction.

When using the full state space as control variables in a weak constraint formulation, the required storage will be the same as for the adjoint method. However, there is no need for any model integrations in the descent algorithm. Given a first guess estimate, the gradient of the cost function is readily evaluated, and a new state estimate can be found.

When using nonlinear dynamics, the penalty function is clearly not convex in general due to the first term in (9) containing the model residuals. However, both the measurement penalty term and the smoothing norm will give a quadratic contribution to the penalty function, and if the weights $\boldsymbol{w}$ and $\boldsymbol{W}_{\eta\eta}$ are large enough compared to the dynamical weight $\widehat{\boldsymbol{W}}_{qq}$, one can expect a nearly quadratic penalty function. On the contrary, if the model residuals are dominating in the penalty function, clearly a pure descent algorithm may get trapped in eventual local minima and the solution found may depend on the first guess in the iteration.

5 Discussion of cases

For all the cases to be discussed the initial condition for the reference case is given by $(x_0, y_0, z_0) = (1.508870, -1.531271, 25.46091)$. These are the same values as was used in Miller *et al.* (1994). The first guess solution was initially chosen as the mean of the reference solution, i.e. about $(0, 0, 23)$. We found that, using the gradient descent method, the global minimum was always found as long as the measurement density and quality was reasonable good. However, when the measurement errors became larger, or if a low number of measurements were used then the gradient descent method often converged to a local minimum. There also seemed to be a possibility for a local minima close to the zero solution where both the dynamical penalty term and the smoothing penalty vanish. It is therefore not wise to use an estimate close to the zero solution as the first guess in the descent algorithm. In all the cases below the error variances for the measurement errors and the errors in the initial conditions are 2.0. To reduce the probability of getting trapped in eventual local minima an objective analysis (OA) estimate, consistent with the measurements, was used as a first guess in the descent algorithm. It was calculated using a smoothing spline minimization algorithm, which is equivalent to objective analysis (McIntosh, 1990). This could easily be done by replacing the dynamical misfit term with a penalty of a first-guess estimate in the inverse formulation, (9).

Several examples will now be discussed and an overview is given in Table 1.

5.1 Case A

This case can be considered as a base case and is similar to the one discussed by Miller *et al.* (1994), i.e., the time interval is $t \in [0, 20]$, and the distance between the measurements are $\Delta t_{\text{obs}} = 0.25$. The gradient descent method was in this case capable of finding the global minimum when starting from the OA estimate. The minimizing solution for the three variables are given in Figures 3 and 4 together with the terms in the penalty function as a function of iteration. We find it amazing how close the inverse estimate is to the reference solution even with such large errors in the measurements. The quality of this inverse estimate is clearly superior to previous inverse calculation using the extended Kalman filter or a strong constraint formulation.

From the terms in the penalty function given in Figure 3, it is seen that the first guess is close to the measurements and rather smooth, while the dynamical residuals are large and contribute with more than 99 % of the total value of the cost function. During the iterations, the dynamical misfit is reduced while there is an increase in the smoothing and measurement terms which indicates that the final inverse solution is further from the measurements and less smooth than the first guess.

Note that by reducing the measurement error variance to 0.2, the gradient descent algorithm converged independently of the choice of first guess estimate, i.e., the penalty function is essen-

tially convex.

An error estimate has been generated using a statistical sampling method. This approach utilizes the fact that the minimum solution can be interpreted as the maximum likelihood estimate of a probability density function,

$$P_{\mathcal{J}}(\boldsymbol{\psi}) = \frac{1}{Z_{\mathcal{J}}} \exp(-\mathcal{J}[\boldsymbol{\psi}]). \tag{14}$$

The constant $Z_{\mathcal{J}}$ is chosen so that the distribution integrates to one. It is now assumed that the vector $\boldsymbol{\psi}$ contains all the elements in the state variables $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$.

Moments of $P_{\mathcal{J}}(\boldsymbol{\psi})$ could be estimated using standard numerical integration based on Monte Carlo methods using points selected at random from some distribution, however this would be extremely inefficient due to the huge state space associated with the oceanic and atmospheric models. Instead a method based on the Metropolis algorithm (Metropolis *et al.*, 1953), is used. Duane *et al.* (1987) proposed the use of guidance Hamiltonian to take the physical properties of the system into account when calculating the next trial. Their remarkable result is that detailed balance is obtained even with the use of a guidance Hamiltonian. An extensive discussion of the hybrid Monte Carlo method is given by Neal (1992, 1993).

By using the hybrid Monte Carlo method to generate a Markov chain that samples this function, a statistical variance estimate can be generated. Note that this method may be used to generate error estimates independently of the minimization technique used to solve the weak constraint problem. The estimated standard deviations of the errors and the true differences between the estimate and the reference solution is given in Figure 5 for the x-component. The largest errors appear around the peaks of the solution and the statistical and true errors are similar. An attempt of generating error estimates by sampling the posterior distribution was also made by Bennett and Chua (1994). However, they did not use the hybrid Monte Carlo method when sampling the posterior distribution and the error estimates vere not very realistic.

5.2 Case B

Here, we extended the time interval to $T = 60$, to test the sensitivity of the inverse estimate with respect to a longer time interval. The number of observations is increased by a factor of 3 to give the same data density as in case A. Note that the value of the cost function is also increased by about a factor of 3. This case behaves similarly to case A, with convergence to the global minimum with a similar rate as in case A. In Figure 6 the x-component of the solution is given together with the terms in the penalty function.

An important conclusion from this example is that by using a weak constraint variational formulation for the inverse, the strong sensitivity with respect to perturbations in initial conditions which is observed for strong constraint variational formulations, is completely removed.

The weak constraint formulation allows the dynamical model to "forget" very past and future information, and the convergence of the inverse calculation therefore has a "local" behavior, where the current estimate at two distant locations have vanishing influence on each other. This suggests that a gradient descent method might work well also for a very large state vector, e.g., the state vector of a primitive equation model, as long as proper preconditioning is used.

5.3 Case C

This case is also similar to case A, except that the number of observations is reduced by 25 %, so the distance between the measurements is now $\Delta t_{\text{obs}} = 0.33$. The solution for the x-component is given in Figure 7, and even with this coarser measurement density the gradient descent method finds the global minimum, although the convergence is slower than in case A.

5.4 Case D

When the distance between the measurements is increased to $\Delta t_{\text{obs}} = 0.40$, the solution found using the gradient descent method is a local minimum. The solution for the x-component is given in Figure 8 together with the terms in the penalty function. That this actually is a local minimum, and not an artifact of the variational formulation will be clear from the next example where the same problem is solved using a statistical method, and a solution which gives a lower value for the penalty function is found. The convergence to a local minimum illustrates that the gradient descent method is not capable of tracking every transition in the Lorenz model with this lower data density.

5.5 Case D1

This case is similar to case D, but now the hybrid Monte Carlo method is used in combination with simulated annealing (see Kirkpatrick *et al.* 1983) for minimizing the penalty function. When working with a penalty function which has many local minima, the so-called simulated annealing technique may be used to improve the convergence to the stationary distribution, based on the method's capability of escaping local minima. A similar approach was used by Bennett and Chua (1994) to solve for the inverse of a nonlinear shallow water model having ill-posed open boundary conditions. The use of the hybrid Monte Carlo method in combination with simulated annealing has been extensively discussed by Neal (1992, 1993), in the context of Bayesian training of back-propagation networks. Clearly, an approach which is based on statistical sampling will be significantly more expensive in terms of CPU-time, than the gradient descent method.

The minimizing solution in this case is given in Figure 9. Note that the number of iterations required for convergence in this case is higher than in other cases. This is due to perturbations caused by the annealing process that allows uphill moves to migrate out of local minima. A local

minima was still found in this case too, where one of the transitions where missed. However, the solution is significantly better than in the previous case where three transitions where missed. There is a possibility that with better choices of the parameters in the annealing schedule the global minimum could be found. The method used here is actually not proper annealing but should be denoted quenching, since the system is cooled too fast to guarantee that the global minimum will be found.

5.6 Case E

In a final case which again is similar to case A, only the x-component of the solution is measured. This implies that the solutions for y and z will be entirely determined by the choice of model error covariance matrix and interactions between the dynamical equations. In Figure 10 the solution for the x-component is given. The solution is at all times following the reference solution, however, the accuracy of the solution is a little poorer than in case A near $T = 20$. This is a result of poor conditioning and slower convergence of the gradient descent method in this case. The slow convergence can be expected since the quadratic contribution from the measurement term is lower when only the x-component of the solution is measured. The problems near $T = 20$ did not occur when the true solution was used as the first guess in the gradient descent method, and with a sufficient number of iterations a better result can be expected.

The y and z-components are given in Figure 11. They are clearly as good as the estimate for the x component. The use of consistent covariances for the model errors was believed to be crucial for a successful inversion. However, it turned out that using a diagonal error covariance matrix for the dynamical errors, gave essentially the same result for all the three components. Thus the dynamics imposed by the model was in this case sufficient for generating the correct inverse estimate for the y and z component, however this is not believed to be a very general result.

6 Discussion

A weak constraint variational formulation for the Lorenz model have been minimized using a gradient descent method. It has been illustrated that by imposing the dynamical model as a weak constraint by allowing the dynamics to contain errors, this leads to a better posed problem than the strong constraint formulation. The extreme sensitivity of the penalty function with respect to initial conditions in strong constraint applications is completely removed. Further, there are no limitations on the length of the assimilation interval.

Here, the inverse was calculated using a gradient descent method using the full state in "space" and time as control variables. The huge state space associated with such a formulation is the main objection against using a gradient descent method for a weak constraint inverse calculation.

It could be compared to the mathematically very appealing representer method (Bennett, 1992), where the solution is searched for in a space with dimension equal to the number of measurements. On the other hand, using a gradient descent approach there is no need to integrate any dynamical equations, since a new candidate for the solution in space and time is substituted every step in the iteration procedure. This gives rise to the notation substitution methods, where the essential point is how the candidates in the iteration are chosen. A gradient descent method will always provide a solution, however it may be a local minimum. Statistical methods based on simulated annealing in combination with a hybrid Monte Carlo method for generating the candidates are much more expensive than a gradient descent approach but has a higher probability of finding the global minimum.

It should be noted that with reasonable good measurement coverage the penalty function is essentially convex, but, when either the number of measurements are decreased or with poorer quality of the measurements, the quadratic contribution to the penalty function from the measurement term has less influence and nonlinearities in the dynamics give raise to local minima. Of course, all data assimilation methods will have problems tracking the true state with too little information from measurements.

When using a gradient descent method to solve a weak constraint problem the full state in space and time must be stored simultaneously. This may be a huge task for some applications, however, it is still comparable to using the adjoint technique, where the tangent linear operator is evaluated at the current state estimate from the forward integration. The main advantage of using a substitution method is recognized when the dynamics is strongly nonlinear, like in the Lorenz model. For linear or weakly nonlinear dynamics, the weak constraint problem may be more easily solved using representer calculations (see Bennett, 1992).

It should also be pointed out that the gradient descent does not directly provide error estimates for the minimizing solution. However, if the gradient descent method is first used to find the solution then the hybrid Monte Carlo method can be used to sample from the posterior distribution function and error variance estimates can be calculated.

The methods discussed here should also be capable of handling a nonlinear measurement functional.

References

Bennett, A. F., *Inverse Methods in Physical Oceanography*, Cambridge University Press, 1992.

Bennett, A. F., and B. S. Chua, Open-ocean modeling as an inverse problem: The primitive equations, *Mon. Weather Rev.*, *122*, 1326–1336, 1994.

Bürger, G., and M. A. Cane, Interactive Kalman filtering, *J. Geophys. Res.*, *99*(C4), 8015–8031, 1994.

Duane, S., A. D. Kennedy, B. J. Pendleton, and D. Roweth, Hybrid Monte Carlo, *Phys. Lett. B.*, *195*,

216–222, 1987.

Evensen, G., Using the extended Kalman filter with a multilayer quasi-geostrophic ocean model, *J. Geophys. Res.*, *97*(C11), 17,905–17,924, 1992.

Evensen, G., Inverse methods and data assimilation in nonlinear ocean models, *Physica D*, *77*, 108–129, 1994a.

Evensen, G., Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *J. Geophys. Res.*, *99*(C5), 10,143–10,162, 1994b.

Evensen, G., and P. J. van Leeuwen, Assimilation of Geosat altimeter data for the Agulhas current using the ensemble kalman filter with a quasi-geostrophic model, *Mon. Weather Rev.*, *124*, 85–96, 1996.

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

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

Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, *Science*, *220*(4598), 671–680, 1983.

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

McIntosh, P. C., Oceanic data interpolation: Objective analysis and splines, *J. Geophys. Res.*, *95*(C8), 13,529–13,541, 1990.

Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.*, *21*(6), 1087–1092, 1953.

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

Neal, R. M., Bayesian training of backpropagation networks by the Hybrid Monte Carlo method, Technical Report CRG-TR-92-1, Department of Computer Science, University of Toronto, 1992.

Neal, R. M., Probabilistic inference using Markov chain Monte Carlo methods, Technical Report CRG-TR-93-1, Department of Computer Science, University of Toronto, 1993.

Pires, C., R. Vautard, and O. Talagrand, On extending the limits of variational assimilation in nonlinear chaotic systems, *Tellus, Ser. A*, To appear, 1996.

Stensrud, D. J., and J. wen Bao, Behaviors of variational and nudging assimilation techniques with a chaotic low-order model, *Mon. Weather Rev.*, *120*, 3016–3028, 1992.

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

Description of Cases

Case	T	Δt_{obs}	Measured	$\mathcal{J}_{\text{global}}$	$\hat{\mathcal{J}}_{\text{est}}$
A	20	0.25	x, y, z	$0.87\,10^3$	$0.87\,10^3$
B	60	0.25	x, y, z	$0.28\,10^4$	$0.28\,10^4$
C	20	0.33	x, y, z	$0.74\,10^3$	$0.74\,10^3$
D	20	0.40	x, y, z	$0.73\,10^3$	$0.22\,10^4$
D1	20	0.40	x, y, z	$0.73\,10^3$	$0.13\,10^4$
E	20	0.25	x	$0.71\,10^3$	$0.73\,10^3$

Table 1: T is the length of the assimilation time interval, Δt_{obs} is the time interval in between measurements, then the measured variables are given, and the last two columns contain the value of the cost function at the global minimum and at the estimate.

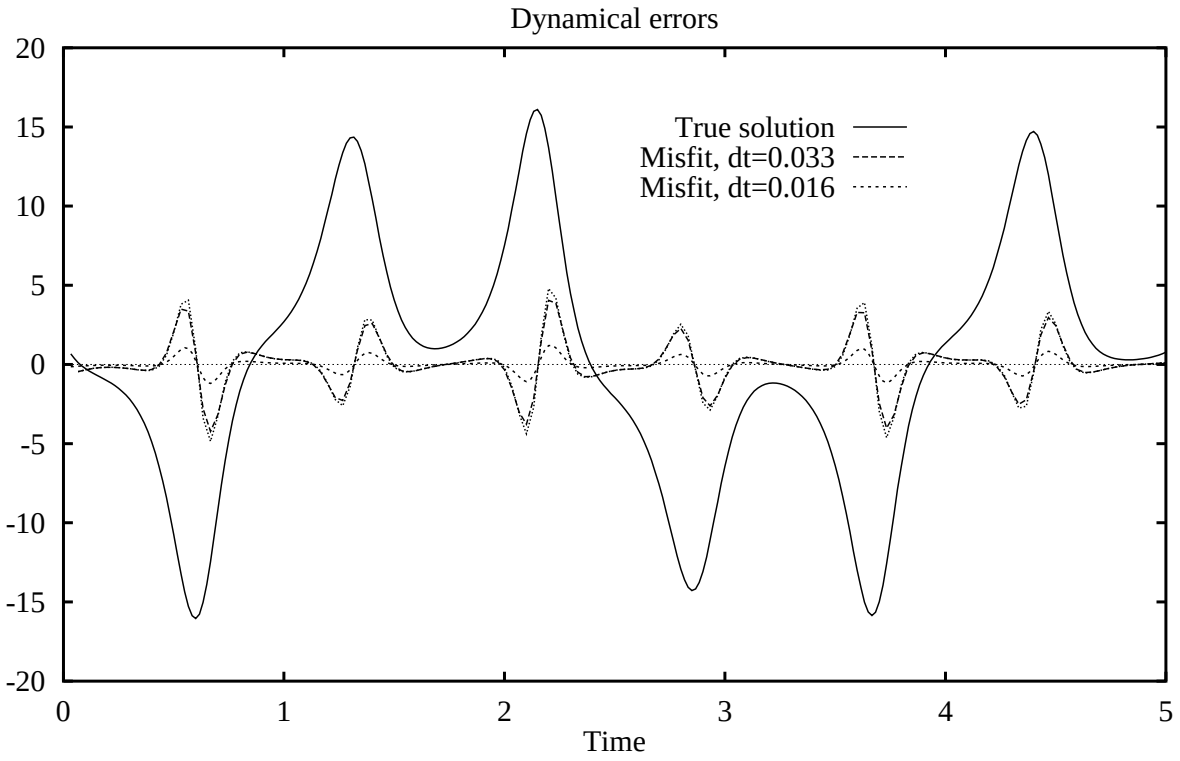

Figure 1: Errors in the difference approximation used for the time derivative, plotted together with the reference solution used in the calculation of the errors. The two similar curves for $\Delta t = 0.033$ are comparing the actual calculated misfits and the lowest-order error term in the discrete time derivative.

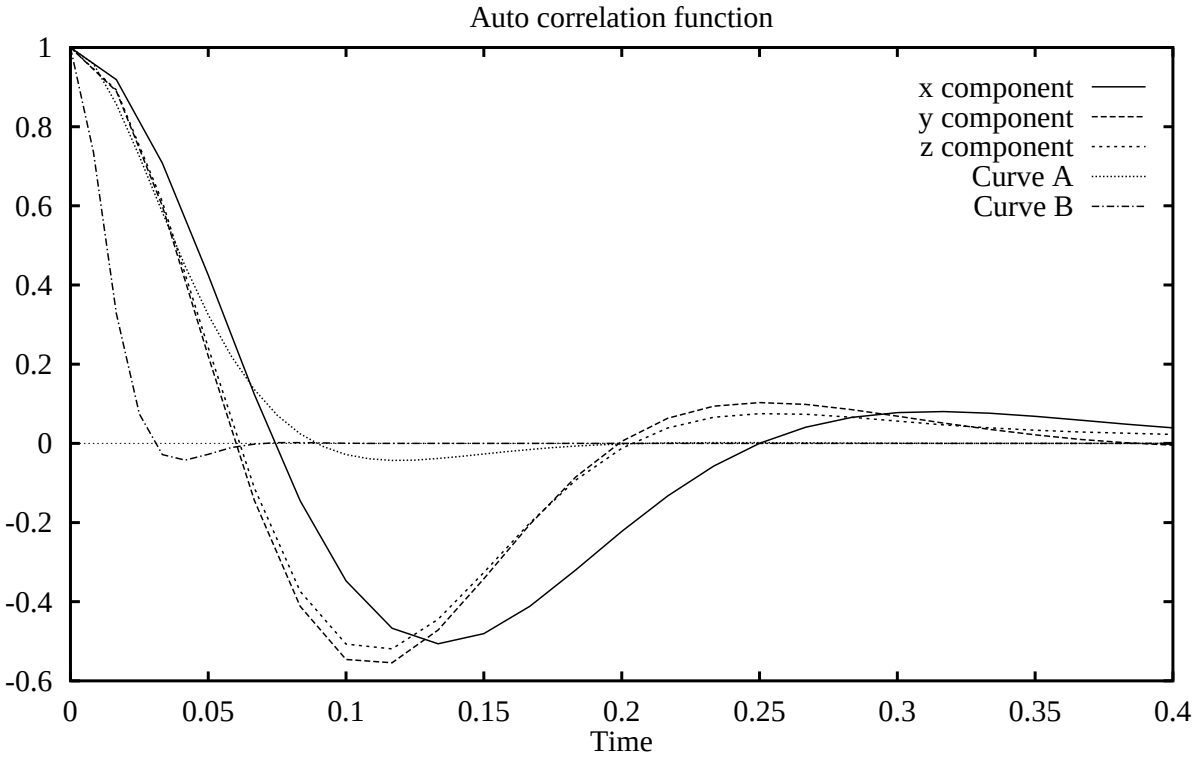

Figure 2: Auto-correlation functions calculated for the computed dynamical misfits for the x, y, and z component of the solution, and two auto-correlation functions corresponding to the smoothing norm with $\gamma = 0.0008$ (curve A) and $\gamma = 0.00001$ (curve B).

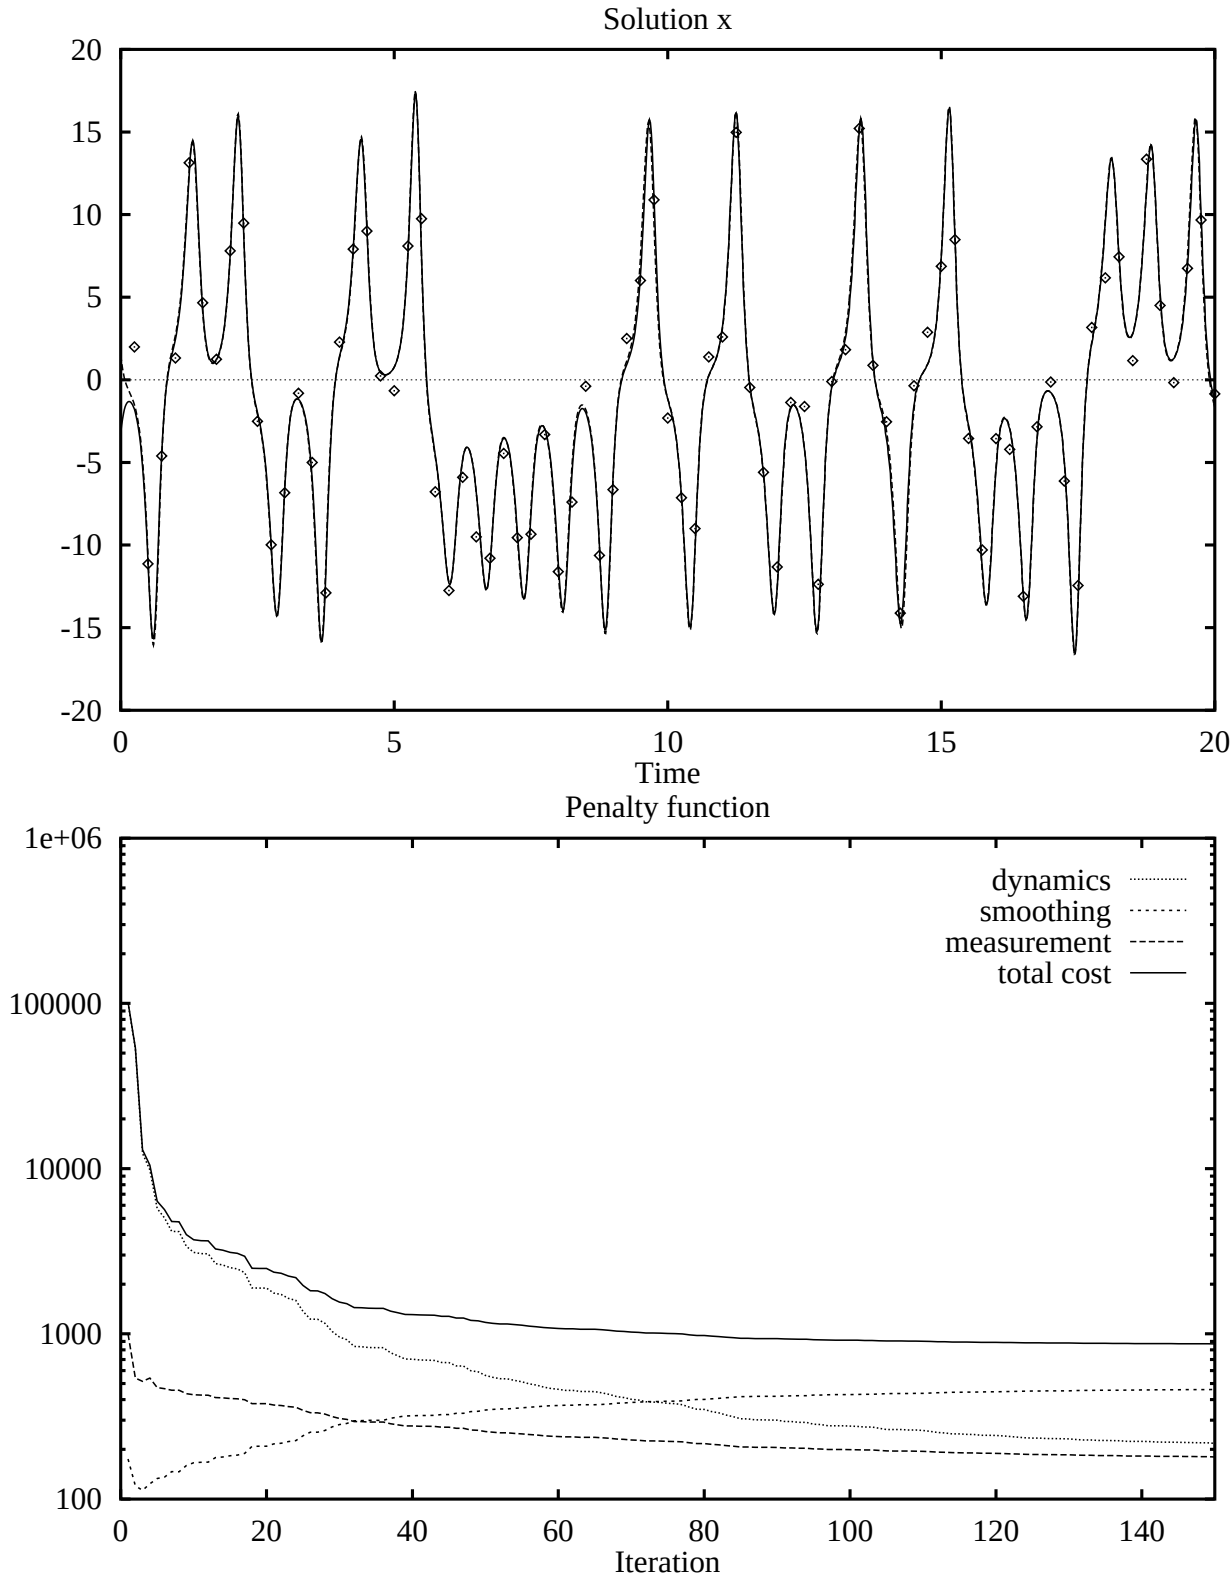

Figure 3: Case A: The inverse estimate for x (top) and the terms in the penalty function (bottom). The estimated solution is given by the solid line. The dashed line is the true reference solution, and the diamonds shows the simulated observations. The same line types will be used also in the following figures.

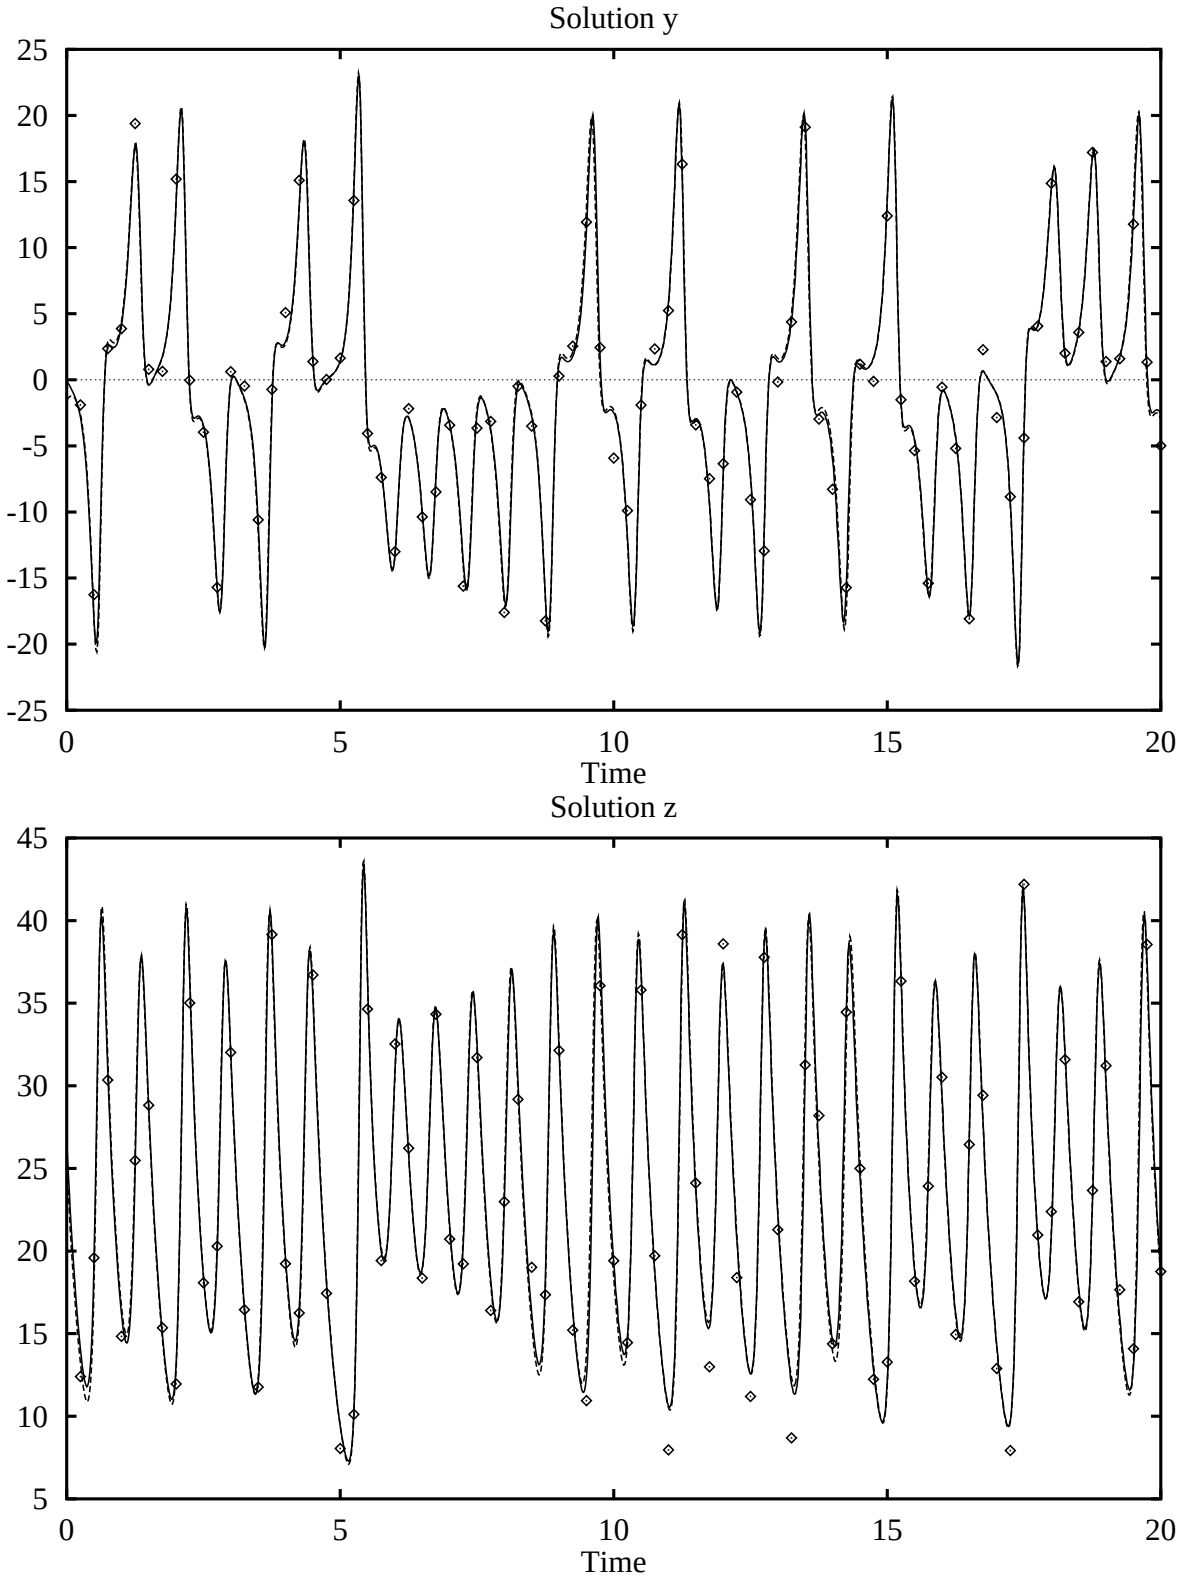

Figure 4: Case A: The inverse estimate for y (top) and z (bottom).

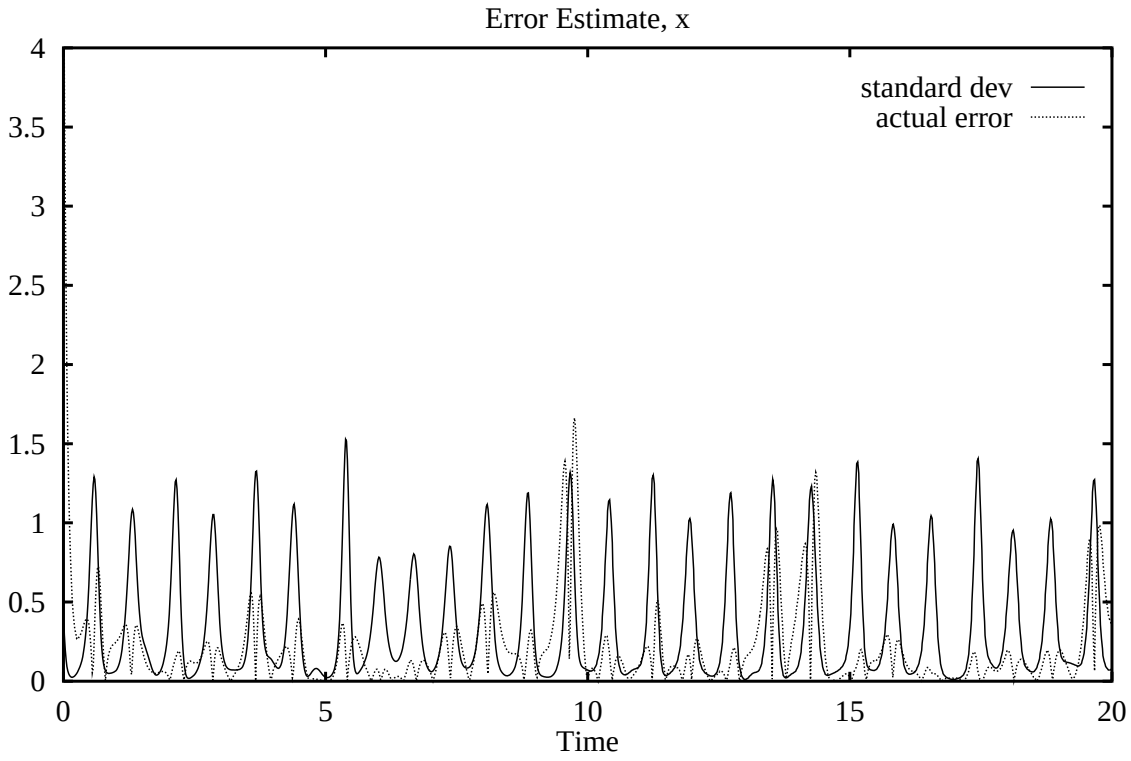

Figure 5: Case A: Statistical error estimates (standard deviations) for x together with the absolute value of the actual errors.

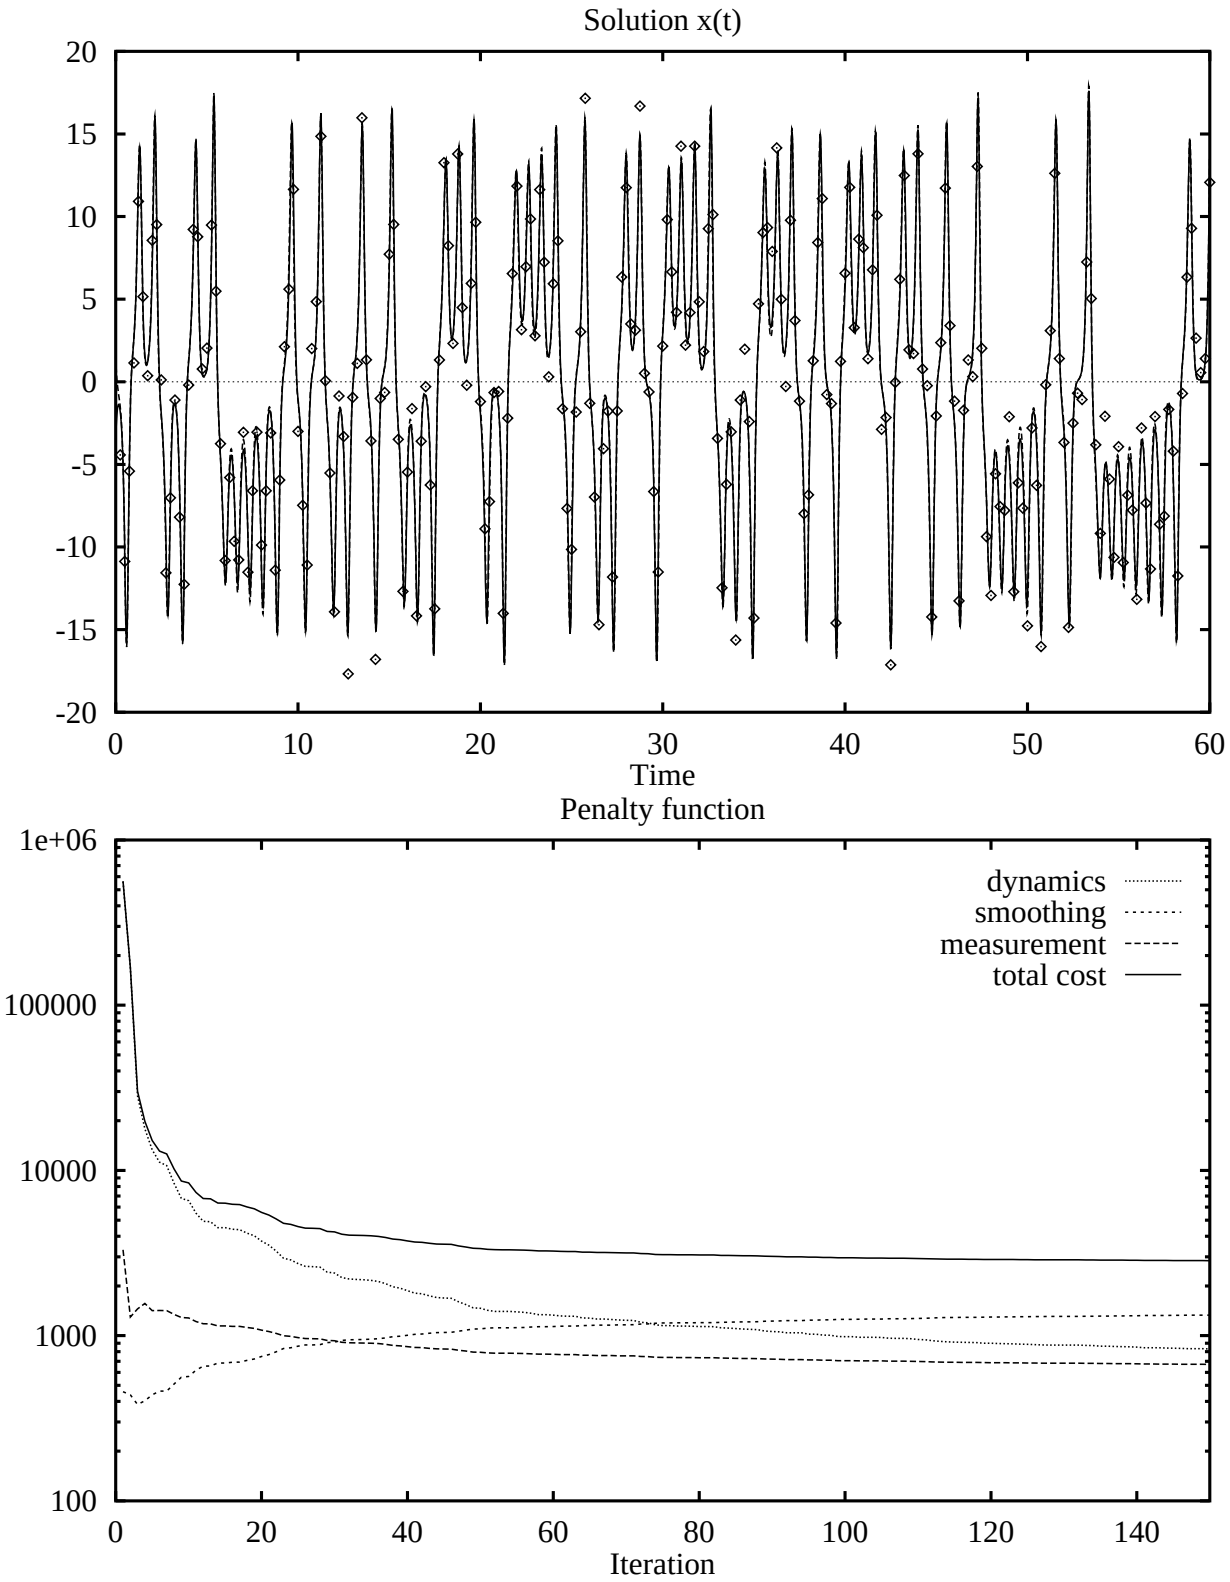

Figure 6: Case B: The inverse estimate for x (top) and the penalty function (bottom).

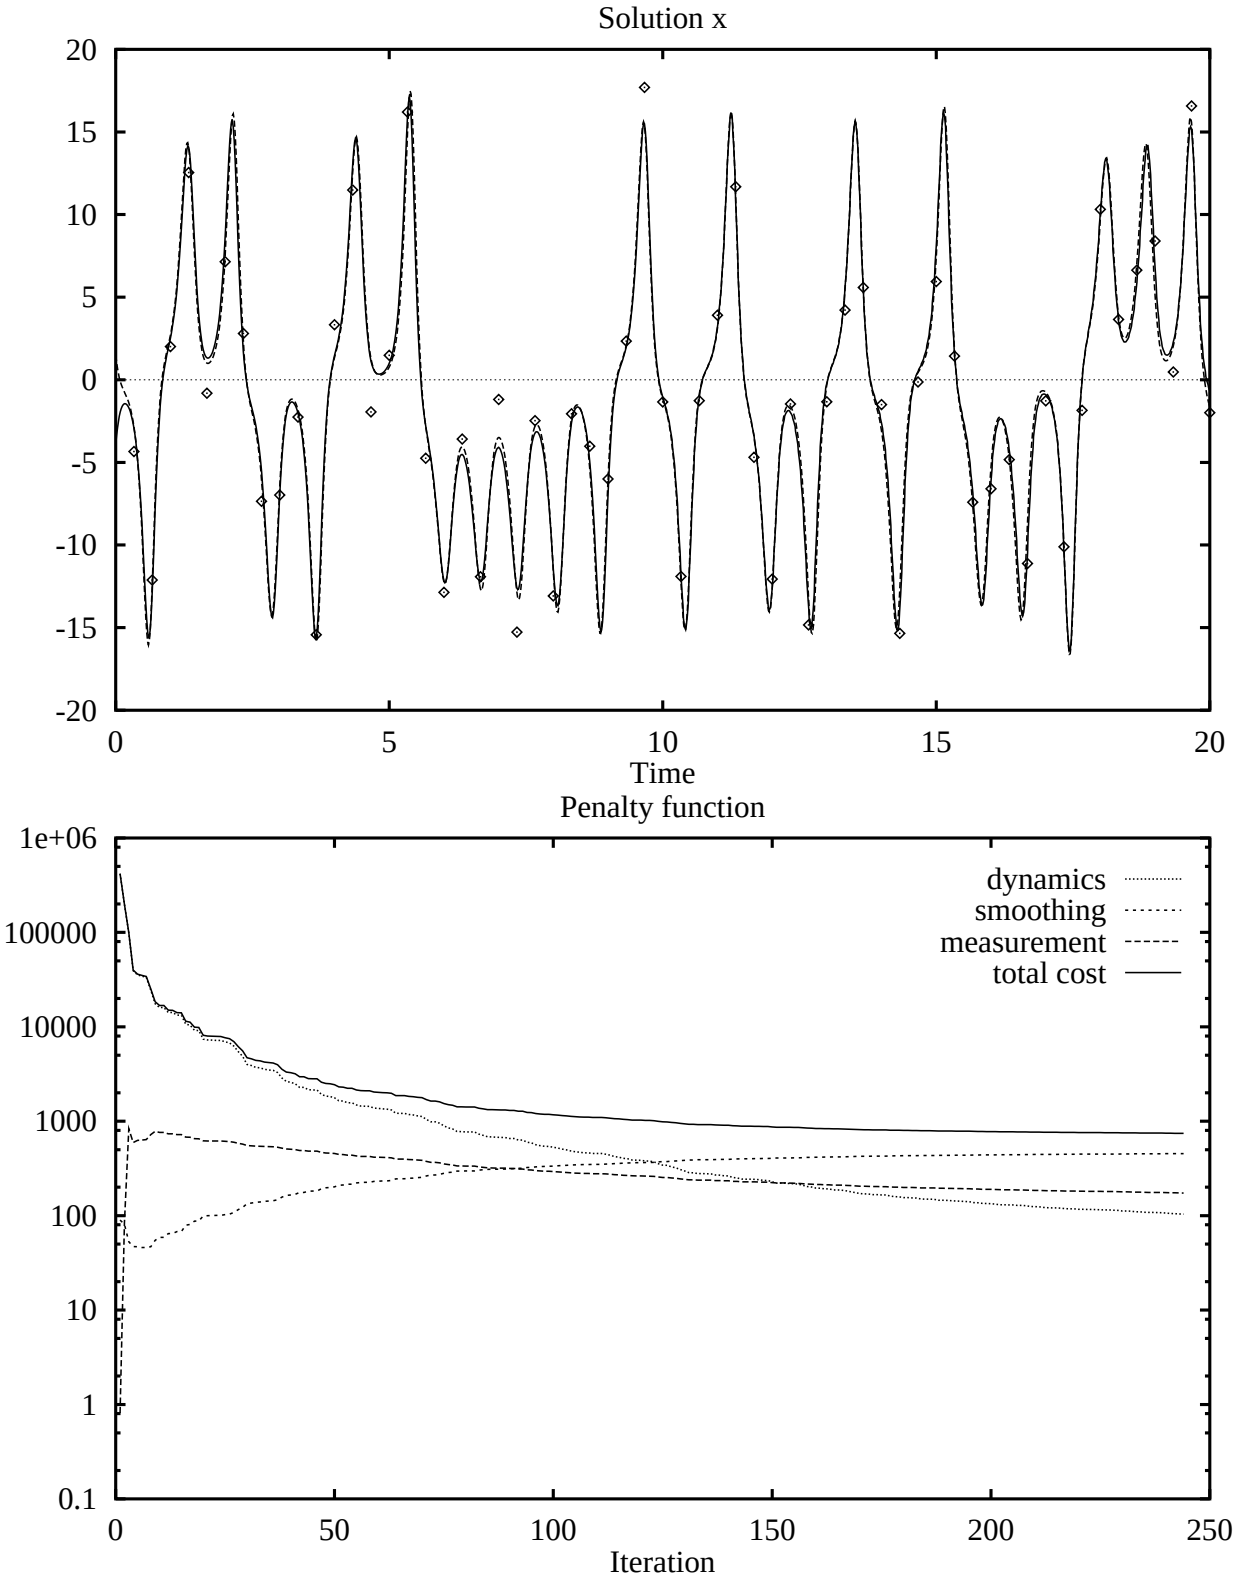

Figure 7: Case C: The inverse estimate for x (top) and the penalty function (bottom).

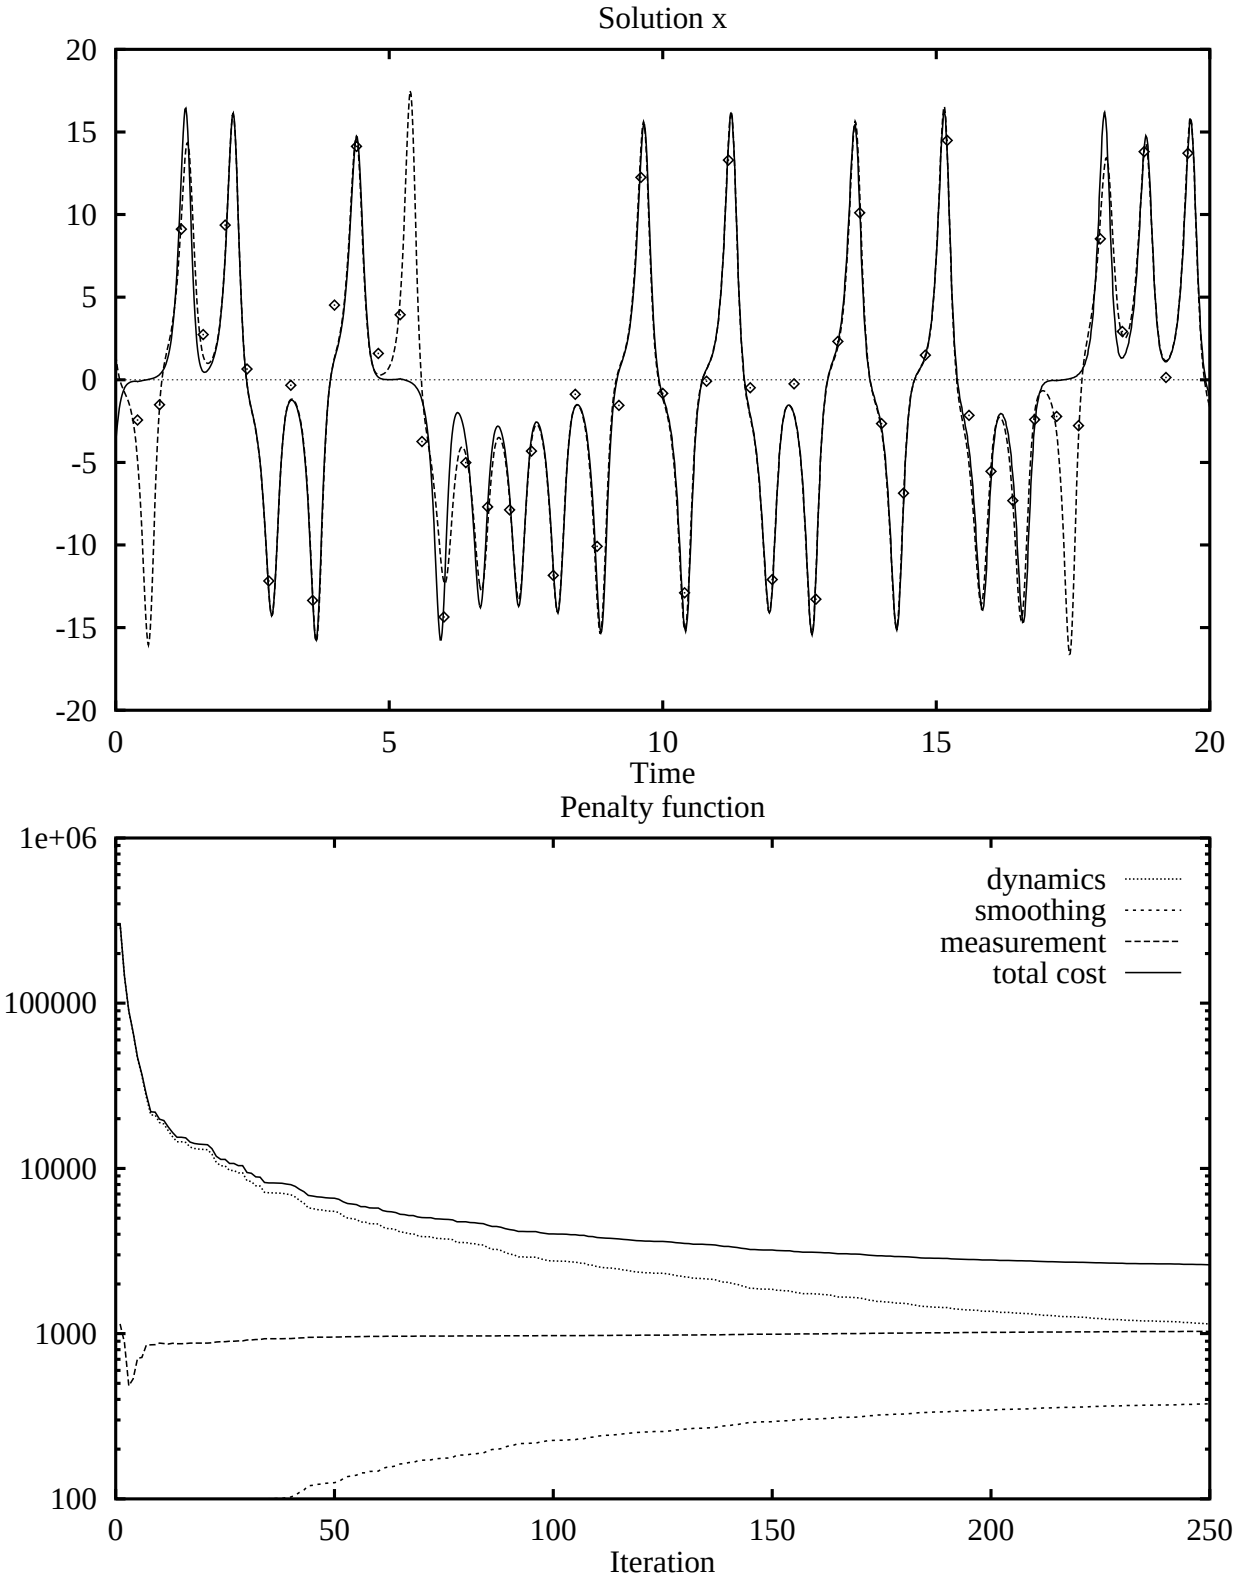

Figure 8: Case D: The inverse estimate for x (top) and the penalty function (bottom).

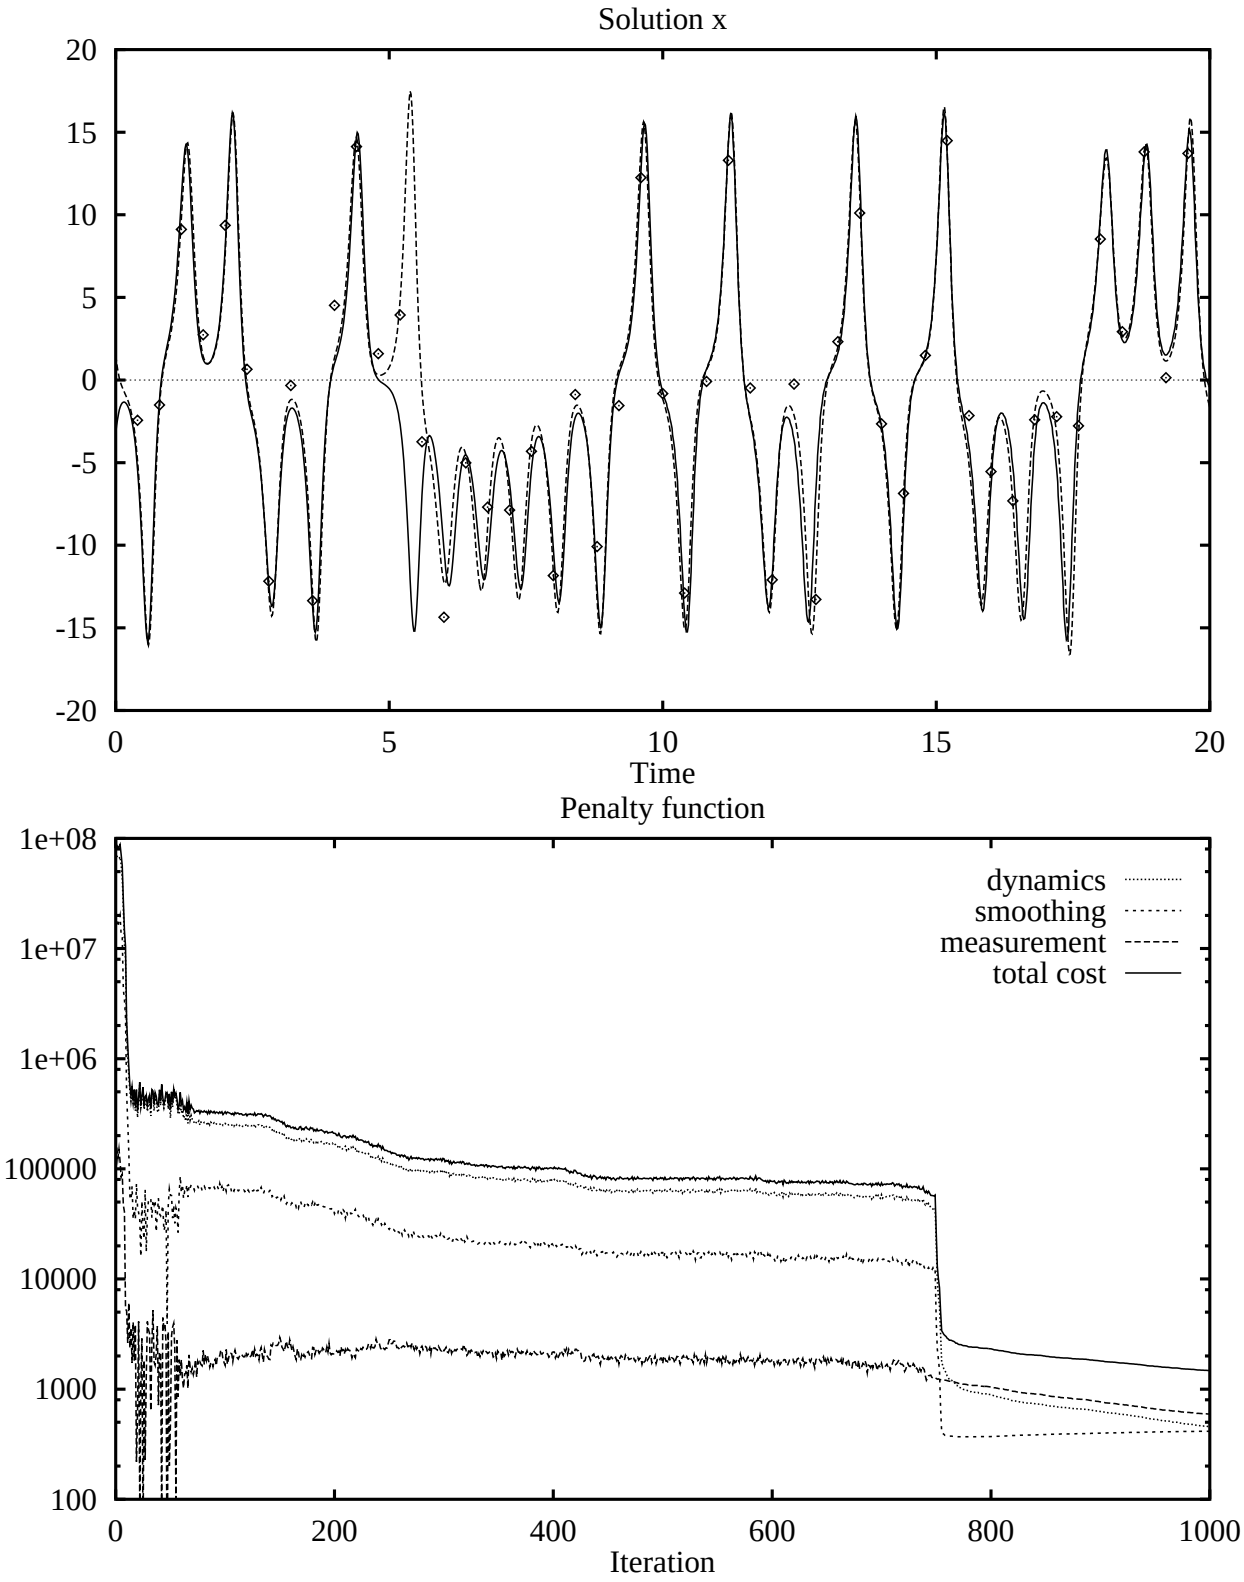

Figure 9: Case D1: The inverse estimate for x (top) and the penalty function (bottom). The sudden decrease in the cost function after about 750 iteration was caused by a change to a gradient descent method after the equlibrium distribution had been reached.

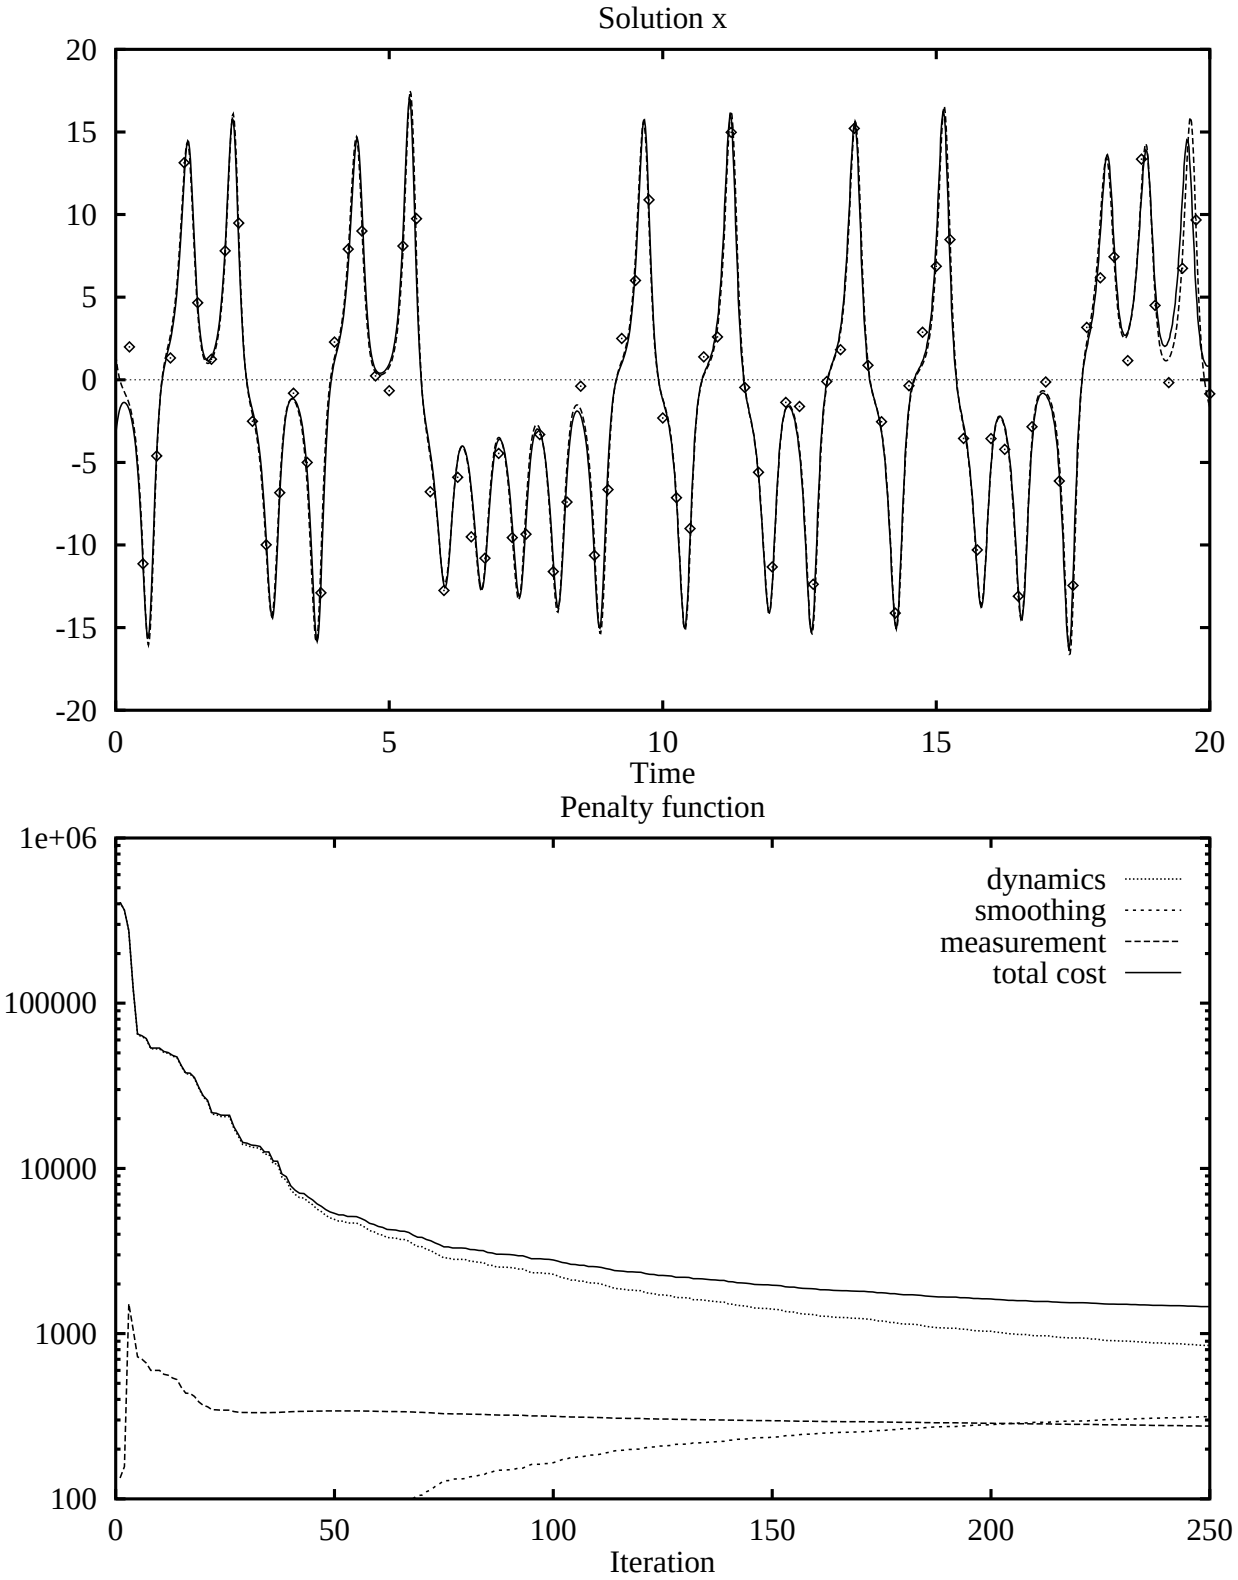

Figure 10: Case E: The inverse estimate for x (top) and the penalty function (bottom).

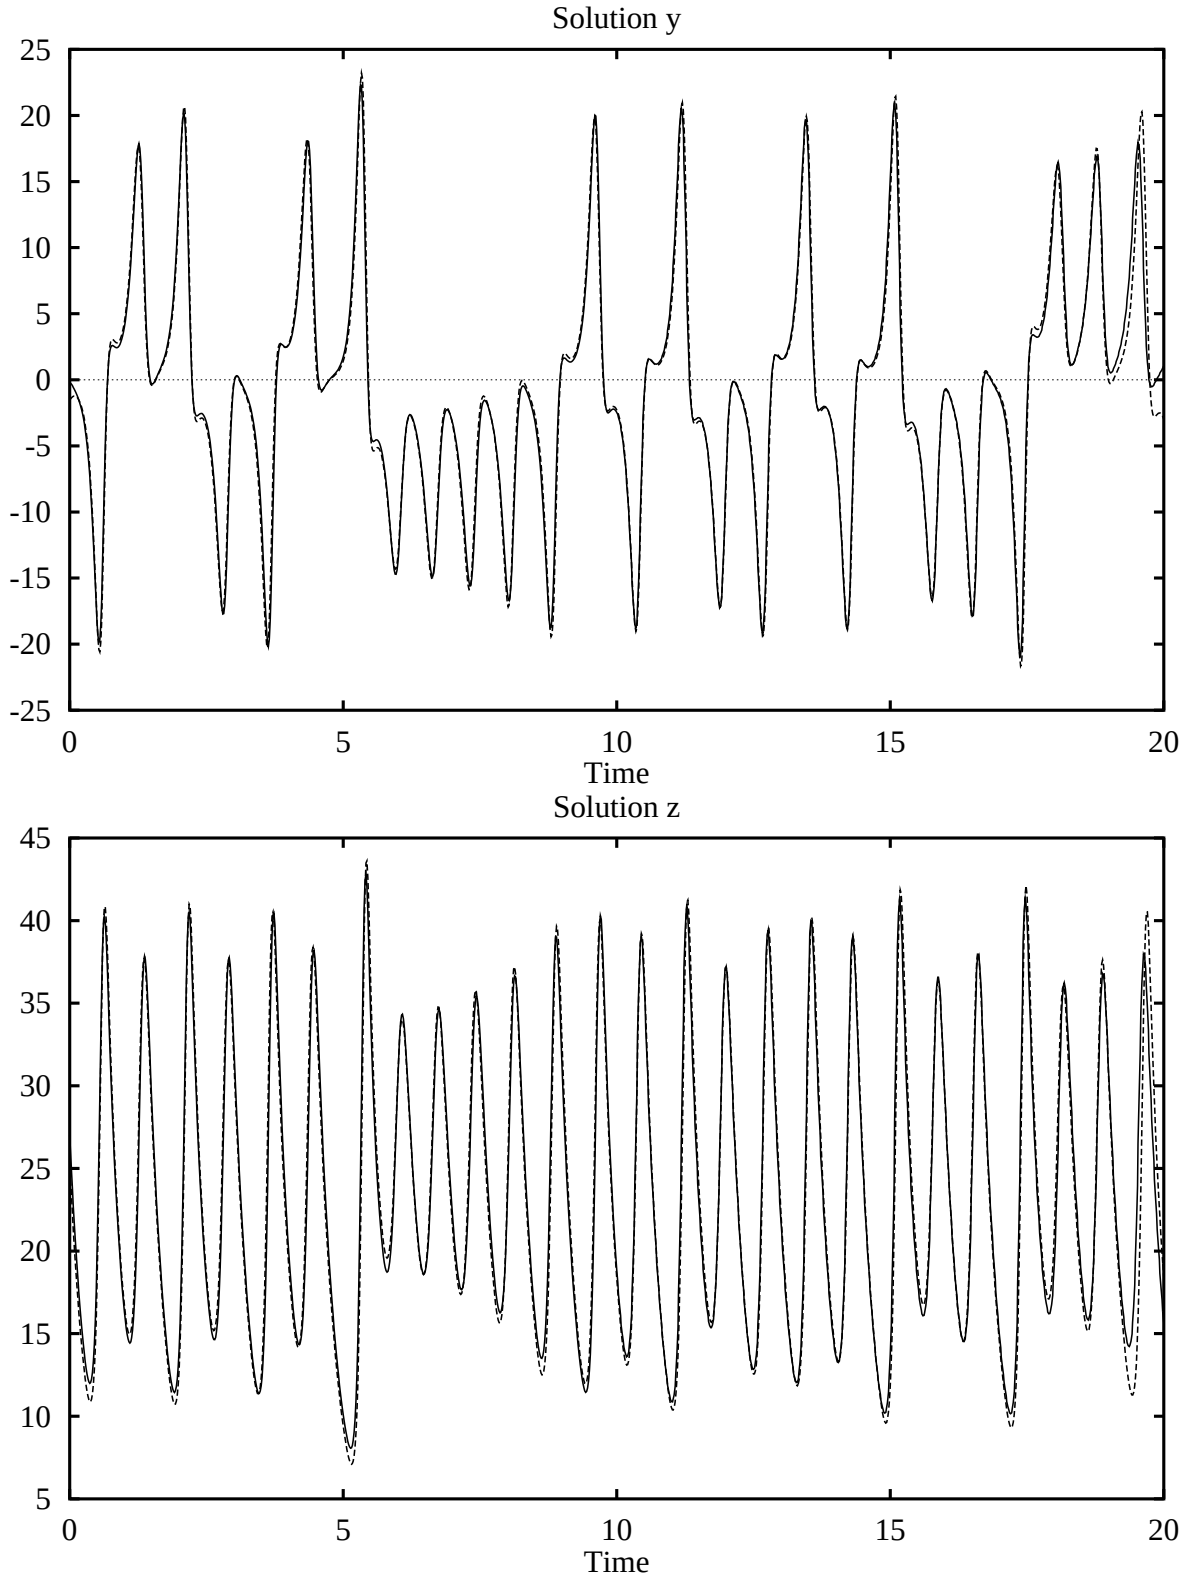

Figure 11: Case E: The inverse estimate for the unobserved variables y (top) and z (bottom).

List of Figures

List of Tables