Cycling the Representer Algorithm for Variational Data Assimilation with the Lorenz Attractor

H. E. NGODOCK, S. R. SMITH, AND G. A. JACOBS

Naval Research Laboratory, Stennis Space Center, Mississippi

(Manuscript received 9 August 2005, in final form 14 February 2006)

ABSTRACT

Realistic dynamic systems are often strongly nonlinear, particularly those for the ocean and atmosphere. Applying variational data assimilation to these systems requires a tangent linearization of the nonlinear dynamics about a background state for the cost function minimization. The tangent linearization may be accurate for limited time scales. Here it is proposed that linearized assimilation systems may be accurate if the assimilation time period is less than the tangent linear accuracy time limit. In this paper, the cycling representer method is used to test this assumption with the Lorenz attractor. The outer loops usually required to accommodate the linear assimilation for a nonlinear problem may be dropped beyond the early cycles once the solution (and forecast used as the background in the tangent linearization) is sufficiently accurate. The combination of cycling the representer method and limiting the number of outer loops significantly lowers the cost of the overall assimilation problem. In addition, this study shows that weak constraint assimilation corrects tangent linear model inaccuracies and allows extension of the limited assimilation period. Hence, the weak constraint outperforms the strong constraint method. Assimilated solution accuracy at the first cycle end is computed as a function of the initial condition error, model parameter perturbation magnitude, and outer loops. Results indicate that at least five outer loops are needed to achieve solution accuracy in the first cycle for the selected error range. In addition, this study clearly shows that one outer loop in the first cycle does not preclude accuracy convergence in future cycles.

1. Introduction

The representer method (Bennett 1992) is a four-dimensional variational data assimilation (4DVAR) algorithm that relies on the adjoint of the dynamical model and expresses the analyzed solution as a first-guess plus a finite linear combination of representer functions, one per datum. The explicit computation and storage of all the representer functions (direct method), however, is not required since the method can be implemented indirectly using the conjugate gradient method (CGM; Amodei 1995; Egbert et al. 1994). A description of the representer methodology is provided in the appendix.

The representer method has earned an established reputation as an advanced data assimilation technique within the past decade, and has gained the attention of many potential operational users. Two primary issues, however, need to be addressed prior to implementing the representer method operationally.

The first issue that is addressed in this paper is the stability and validity of the tangent linear model (TLM) and how it impacts the assimilation accuracy. When the representer method is applied to a nonlinear model, the model must be linearized, preferably using the first-order approximation of Taylor's expansion. Traditionally, the representer method has been implemented for the assimilation of all observations in the time window considered. As with every other variational data assimilation method with nonlinear dynamics, the representer method necessitates that the TLM and its adjoint be valid and/or stable over the entire assimilation time window. The validity of the TLM is difficult to maintain over a long time period for strong nonlinear models and complex regions.

The second issue that is addressed in this paper is the computational cost of the representer method. The indirect representer method requires the integration of the adjoint and TLM within a CGM that determines the representer coefficients for the minimization of the cost function (see the appendix). This set of representer co-

Corresponding author address: Hans Ngodock, Naval Research Laboratory, Code 7320, Bldg. 1009, Stennis Space Center, MS 39529.
E-mail: ngodock@nrlssc.navy.mil

DOI: 10.1175/MWR3281.1

efficients is then used to provide a correction to the background. The number of iterations of the CGM (this is referred to as the inner loop) is typically a small fraction of the total number of measurements. Outer loops are typically required for strongly nonlinear systems. To initialize the outer loop, one would pick a first background solution around which the model is linearized. The best solution (corrected background) obtained from this assimilation would become the background for the next outer loop, and so forth until formal convergence is reached (Bennett et al. 1996; Ngodock et al. 2000; Bennett 2002; Chua and Bennett 2001; Muccino and Bennett 2002). This outer loop exacerbates the computational cost of the representer method. In this study the background that serves for linearization is also taken as the first guess.

These two issues have detracted many potential users of the representer method for operational purposes. It is possible however, to address these issues and implement the representer method at a reasonable cost for operational applications. Given a time window in which one desires to assimilate observations, it is possible to apply the representer method over cycles of subintervals. The name adopted for this approach is the cycling representer method (Xu and Daley 2000), and its associated solution is called the cycling solution. The solution that is obtained by assimilating all the observations at once in the original full time window will be called the global solution.

By using the cycling representer method, the assimilation time window is constrained to a period over which the TLM produces an accurate dynamical representation of the Lorenz attractor. Doing this reduces the need for outer loops. The outer loop is designed to solve the nonlinear Euler–Lagrange conditions associated to the assimilation problem with the nonlinear model, because the representer method solves a linear assimilation problem. In the global solution problem, the TLM may not be an accurate representation of the solution, and the adjoint would not be an accurate estimate of the derivative of the state with respect to the control variables. If the TLM is an accurate representation of the solution, the need for outer loops is removed. In the initial cycles of this assimilation approach, the first-guess or background solution may not be accurate and thus outer loops may be required. Once the system is spun up and the TLM is an accurate approximation (thanks to improved background solutions), outer loops may no longer be necessary, thus lowering the computational cost of the assimilation. There may be situations in real-world applications, however, where a few outer loops would be needed in the current cycle, even though a single outer loop sufficed in previous cycles. An example of this occurrence could be a nonlinear ocean response (advection or mixing) to a sudden, strong change in atmospheric forcing, especially in coastal areas with complex bathymetry. Nevertheless, the need of additional outer loops is assessed by the discrepancy between the assimilated solution and the data.

The idea of cycling the representer method was investigated by Xu and Daley (2000) using a 1D linear transport model with synthetic data. In that study, the error covariance of the analyzed solution was updated at the end of each cycle and used as the initial error covariance in the next cycle. Another application of the cycling representer method was performed by Xu and Daley (2002) using a 2D linear unstable barotropic model with no dynamical errors. In the 2002 study, the covariance at the end of the cycle is not updated because its computation is too costly to be practical. Updating the covariance requires the evaluation and storage of the representer functions at the final time. These two studies found that updating the covariance at the end of each cycle produced significantly more accurate analyses. However, in these two applications of the cycling representer method, only linear models were used and thus there were no need for a TLM. Most realistic applications are nonlinear and their TLM may not be stable over the time window considered. It is in this context that this study applies the cycling representer method.

A good dynamical system candidate for testing assimilation methods for strongly nonlinear models is the acclaimed Lorenz attractor model (Lorenz 1963; Fig. 1). It has been used in numerous to studies to examine the behavior of assimilation methods based on sequential filtering and variational techniques. This is done with the intent that if an algorithm performs satisfactorily well with this model, then it can be applied to atmospheric and ocean models, which is a necessary but not a sufficient condition.

Gauthier (1992) implemented the strong constraint variational assimilation method using the adjoint of the Lorenz equations and a rather sparse data network sampled every two time units. He found that the cost function was well behaved when the model did not undergo a transition. However, the solution was very sensitive to initial perturbations. This problem is closely related to the stability and validity of the TLM approximation during the different phases of the model.

Miller et al. (1994) also implemented the strong constraint variational assimilation method with the Lorenz attractor model and found that the behavior of the cost function encountered by Gauthier (1992) was strongly dependent on the length of the assimilation window.

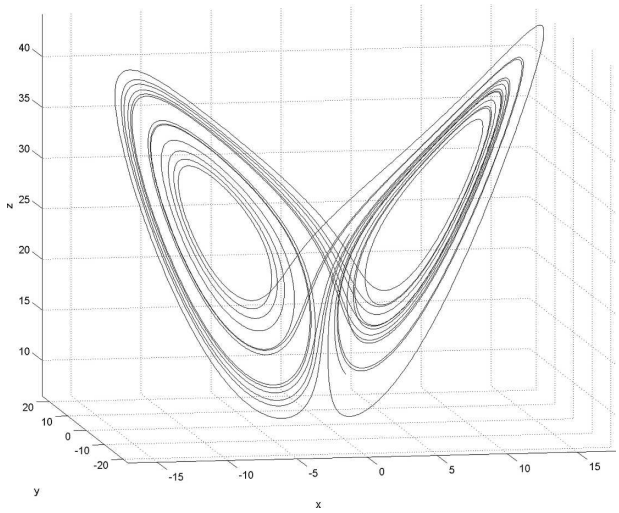

FIG. 1. The Lorenz attractor solution computed for 20 time units using the RK4 time-stepping scheme with a time step of dt = 1/600. This is the true solution from which observations are sampled every ¼ time unit.

However, contrary to Gauthier (1992), Miller et al. (1994) did not derive the adjoint equations for the minimization of the cost function. They computed the gradient of the cost function directly and invoked a conjugate gradient routine to minimize this cost function. The assimilated solution computed by Miller et al. (1994), fitted the data (sampled every 0.25 time units) for 7 time units. The forecast following the assimilation was able to track the data up to three additional time units before errors increased significantly.

Evensen and Fario (1997) implemented the weak constraint variational assimilation method using the Lorenz attractor equations with the same configuration as Miller et al. (1994). By including a dynamic error and directly computing the gradient of the cost function with respect to all state variables (in space and time), they were able to use a gradient descent algorithm to accurately assimilate data (sampled every 0.25 time unit) for longer time periods (20 and 60 time units) with no unexpected behavior of the cost function. However, when the data density decreased, the same method produced only a local minimum of the cost function that resulted in three missed transitions within an interval of 20 time units. This suggests that the good fit that Miller et al. (1994) obtained for seven time units in the strong constraint approach depends on the rather dense data network. Note that in Evensen and Fario (1997), there is no need of the TLM or the adjoint.

Realistic applications, however, will not always have the luxury of an exact hand-computed gradient of the cost function. The computation of the gradient typically relies on the adjoint, which in turn relies on the TLM. Because of the chaotic behavior of realistic models, any linearization is subject to error growth therefore limiting the application of the TLM approximation to a time scale for which errors are lower than a prescribed threshold. Our approach is different from the studies mentioned above in that we compute the gradient of the cost function by means of the adjoint and implement the weak constraint variational assimilation using the cycling representer method. In this paper, the global and cycling solutions are compared and the ability to remove the outer loops is examined. In these experiments a significance test is not performed. This would turn the assimilation problem into a search for suitable prior assumptions about the data, initial condition, and dynamical errors, and hence cloud the issue at hand. This study specifies prior error covariances similar to those used in previous studies. The Lorenz model and the variational assimilation problem are described in section 2, and the TLM stability and accuracy are discussed in section 3. Section 4 deals with the assimilation experiments. A discussion and concluding remarks follow in sections 5 and 6, respectively.

2. The Lorenz model and the variational assimilation problem

The Lorenz model is a coupled system of three nonlinear ordinary differential equations:

$$\frac{dx}{dt} = \sigma(y - x) + q^x,$$

$$\frac{dy}{dt} = \rho x - y - xz + q^y, \quad \text{and}$$

$$\frac{dz}{dt} = xy - \beta z + q^z, \tag{1}$$

where x, y, and z are the dependent variables. The commonly used time invariant coefficients are $\sigma = 28$, $\rho = 10$, and $\beta = 8/3$ and the model errors are represented by q^x, q^y, and q^z. The initial conditions for Eq. (1) are

$$x(0) = x_0 + i^x,$$

$$y(0) = y_0 + i^y, \quad \text{and}$$

$$z(0) = z_0 + i^z, \tag{2}$$

where $x_0 = 1.508\,87$, $y_0 = -1.531\,271$, and $z_0 = 25.460\,91$ are the first guess of the initial conditions. These are the same values that are used in the data assimilation studies by Miller et al. (1994), Evensen (1997), Evensen and Fario (1997), Miller et al. (1999), and Evensen and Van Leeuwen (2000). The initial condition errors are

represented by i^x, i^y, and i^z. By setting the model and initial condition errors in Eqs. (1) and (2) to zero, the solution to the Lorenz attractor is computed for the time interval [0, 20] using the fourth-order Runge–Kutta (RK4) discretization scheme with a time step of $dt = 1/600$ (Fig. 1). This solution is labeled as the true solution, since using time steps smaller than $dt = 1/600$ does not significantly alter the solution within the specified time period. Thus, the numerical representation has converged.

In the time interval [0, 20] there is a set of M observations $\mathbf{d} \in \Re^M$ such that

$$\mathbf{d} = \mathbf{H}(x, y, z) + \boldsymbol{\varepsilon}, \tag{3}$$

where $\mathbf{H}$ is a linear measurement functional ($M \times 3$ matrix) and $\boldsymbol{\varepsilon} \in \Re^M$ is the vector of measurement errors. The data used for all assimilation experiments are sampled from the true solution (Fig. 1) with a frequency of 0.25 time units. Since there are three position variables, there are a total of $M = 237$ measurements. Purposefully, data is not sampled at $T = 0$ or $T = 20$ in order to demonstrate the capability of the assimilation method to propagate the information from future (past) observations to correct the initial (final) conditions. The measurement error is assumed to be $\varepsilon = 0.002$, and its covariance matrix is assumed to be diagonal.

In general, the errors in Eqs. (1), (2), and (3) are not known. Therefore solving Eq. (1) directly is an ill-posed problem. By using generalized inversion, a solution to Eq. (1) is obtained that minimizes the errors in Eqs. (1)–(3) in a weighted least squares sense using a cost function:

$$J(x, y, z) = \int_0^T \int_0^T \mathbf{q}(t_1)\mathbf{W}_{qq}(t_1, t_2)\mathbf{q}(t_2)\, dt_1\, dt_2 + \mathbf{i}^{\mathrm{T}}\mathbf{W}_{ii}\mathbf{i} + \boldsymbol{\varepsilon}^{\mathrm{T}}\mathbf{w}\boldsymbol{\varepsilon}. \tag{4}$$

The weights ($\mathbf{W}_{qq}$, $\mathbf{W}_{ii}$, and $\mathbf{w}$) in Eq. (4) are defined as the inverses of the respective error covariance matrices $\mathbf{C}_{qq}$, $\mathbf{C}_{ii}$, and $\mathbf{C}_{\varepsilon}$ for the model, initial condition, and measurement errors, respectively.

For the sake of clarity, Eqs. (1)–(2) are rewritten here in vector notation:

$$\frac{d\mathbf{x}}{dt} = \mathbf{F}(\mathbf{x}) + \mathbf{q}$$
$$\mathbf{x}(0) = \mathbf{x}_0 + \mathbf{i}, \tag{5}$$

where $\mathbf{x}$ is the state vector consisting of the three independent variables (x, y, z) at each time step and $\mathbf{F}(\mathbf{x})$ is the model dynamics. The observations may be written in a similar notation:

$$\mathbf{d} = \mathbf{H}\mathbf{x} + \boldsymbol{\varepsilon}. \tag{6}$$

In real-world problems, the observation operator $\mathbf{H}$ may be nonlinear for some types of measurements (e.g., radiances in the atmosphere, acoustic tomography in the ocean) therefore requiring a tangent linear approximation so that the representer method may be used. The Euler–Lagrange (EL) conditions for a local extremum of the cost function are derived by setting the first variation of Eq. (4) to zero. These conditions are

$$-\frac{d\boldsymbol{\lambda}}{dt} = \left[\frac{d\mathbf{F}(\hat{\mathbf{x}})}{d\mathbf{x}}\right]^{\mathrm{T}}\boldsymbol{\lambda} - \mathbf{H}^{\mathrm{T}}\mathbf{w}(\mathbf{d} - \mathbf{H}\hat{\mathbf{x}})$$
$$\boldsymbol{\lambda}(T) = 0 \tag{7}$$

and

$$\frac{d\hat{\mathbf{x}}}{dt} = \mathbf{F}(\hat{\mathbf{x}}) + \mathbf{C}_{qq} \cdot \boldsymbol{\lambda}$$
$$\hat{\mathbf{x}}(0) = \mathbf{x}_0 + \mathbf{C}_{ii}\boldsymbol{\lambda}(0). \tag{8}$$

In the above equations, $\hat{\mathbf{x}}$ is the optimal solution and $\boldsymbol{\lambda}$ is the weighted residual or adjoint variable defined as

$$\boldsymbol{\lambda}(t) = \int_0^T \mathbf{W}_{qq}(t, t_2)\mathbf{q}(t_2)\, dt_2$$
$$(\mathbf{C}_{qq} \cdot \boldsymbol{\lambda})(t) = \int_0^T \mathbf{C}_{qq}(t, t_2)\boldsymbol{\lambda}(t_2)\, dt_2. \tag{9}$$

The initial condition error that is used to perturb Eq. (2) is specified to be 10% of the standard deviation of each state variable of the true solution ($i^x = 0.784$, $i^y = 0.897$, and $i^z = 0.870$). The initial condition error covariance ($\mathbf{C}_{ii}$) is simply a 3×3 diagonal matrix with values equal to the square of the RMS of these initial condition errors.

The model errors used to perturb Eq. (1) have unique values for each dependent variable at each time step of the model integration. The statistics of these model errors are estimated by first computing a degraded solution using a time step of $dt = 1/60$. The difference between this degraded solution and the truth is examined up to the point where the solutions diverge significantly (5.4 time units). The RMS of the differences between the true and degraded solutions prior to this time is divided by the number of time steps during this time period (5.4 time units) to obtain an estimate for the standard deviation (STD) of the model error [$\mathrm{STD}(q) = 3.69 \times 10^{-3}$]. By examining the temporal

correlation of these differences, a decorrelation time scale of $\tau = 0.25$ time units is estimated for the model error. Therefore, the model errors are created by first generating three time series of Gaussian distributed random numbers (one for each dependent variable). A temporal convolution of these random errors is then conducted with the prescribed exponential function $\exp\{-[(t - t')/\tau]^2\}$. This array of correlated errors is finally scaled to have the estimated STD (3.69×10^{-3}). The spatial covariance of these dynamic errors (i.e., the stationary part of $\mathbf{C}_{qq}$) is

$$\tilde{\mathbf{C}}_{qq} = \begin{bmatrix} 1.36 \times 10^{-5} & 5.99 \times 10^{-7} & -1.56 \times 10^{-6} \\ 5.99 \times 10^{-7} & 1.36 \times 10^{-5} & -2.07 \times 10^{-6} \\ -1.56 \times 10^{-6} & -2.07 \times 10^{-6} & 1.36 \times 10^{-5} \end{bmatrix}. \tag{10}$$

3. The TLM

The dynamical model in Eq. (5) is tangent linearized around a first-guess trajectory $\mathbf{x}^f(t)$ using the first-order approximation of Taylor's expansion:

$$\frac{d\mathbf{x}}{dt} = \mathbf{F}(\mathbf{x}^f) + \frac{d\mathbf{F}(\mathbf{x}^f)}{d\mathbf{x}}(\mathbf{x} - \mathbf{x}^f) + \mathbf{q}$$

$$\mathbf{x}(0) = \mathbf{x}_0^f + \delta\mathbf{x}_0 + \mathbf{i}. \tag{11}$$

In the remainder of this paper, $\mathbf{x}^f$ is referred to as the background solution. The assimilation problem is then solved with the linearized model [Eq. (11)] using the representer method. The representer expansion can be written for the linear EL system as follows (also see appendix):

$$\hat{\mathbf{x}}(t) = \mathbf{x}^f(t) + \sum_{m=1}^{M} \alpha_m \mathbf{r}_m(t). \tag{12}$$

By choosing $\mathbf{x}^f$ for both the background (for linearization) and the first guess (for assimilation), one can formulate the assimilation problem as a search for the optimal correction to $\mathbf{x}^f$ (Uboldi and Kamachi 2000; Jacobs and Ngodock 2003), that is, the second term in the right-hand side of Eq. (12). The solution $\hat{\mathbf{x}}$ from the assimilation is taken as the new trajectory around which the model is linearized (i.e., $\hat{\mathbf{x}}$ replaces $\mathbf{x}^f$ in the next iteration of the outer loop).

There are three problems that are associated with this approach. First, the convergence of the linear iterations (outer loops) depends on the background, and the background may not always be sufficiently accurate. Second, the TLM around the background may not always be valid and accurate over the entire assimilation window, especially when we are dealing with strongly nonlinear dynamical models. This is because the accuracy/validity of the TLM is limited by the growth of errors. Finally, the outer loop amplifies the cost of the assimilation associated with a linearized model, and renders the cost of the representer method expensive for operational purposes.

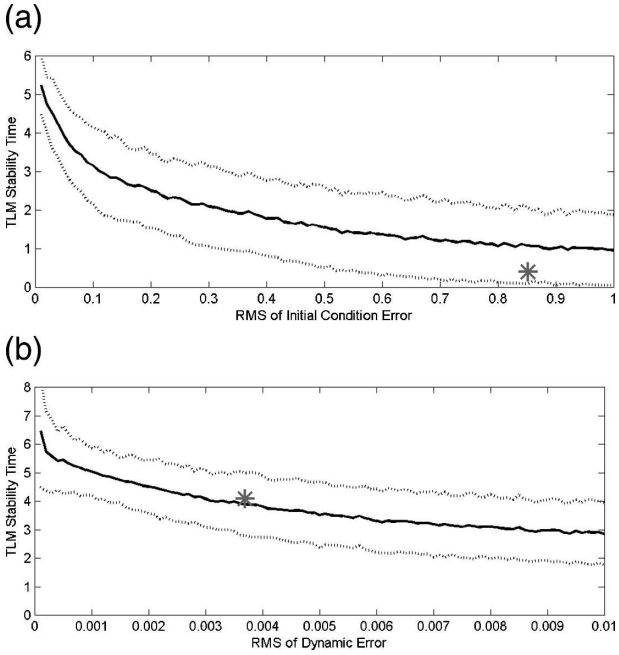

FIG. 2. Stability of the TLM relative to the RMS of the (a) initial condition and (b) model error. For each RMS error, the TLM was computed 1000 times, each time perturbed by a new set of randomly generated numbers scaled by the specified RMS error. The solid line is the mean of these realizations and the dotted lines are ±1 STD. The gray asterisk depicts the errors that are used in the experiments presented in this paper.

The accuracy of the TLM depends on the initial condition and model errors. The estimation of these errors is described in section 2. For completeness, the TLM accuracy limit for a range of error levels in both the initial conditions and model are separately examined. The experiments for each type of error contain 100 different values for the RMS error ranging from 0 to 1 (initial condition error) and 0 to 0.01 (model error). For each RMS error, the TLM was computed 1000 times (for statistical significance), and in each run a different random realization of the errors was used. The accuracy criterion for the TLM is based on the 1000 member ensembles and is chosen as the time when the difference between the TLM and the background exceeds the standard deviation of the background. In Fig. 2 the mean and ±1 STD of the 1000-member TLM accuracy time is shown along with the accuracy time obtained with our choice of errors in section 2 (gray asterisks).

This realization (gray asterisks) is the only one used for the experiments throughout the remainder of this paper. Assimilations using significantly large ensembles would be prohibitive and the results difficult to display.

In section 2, the RMS of the initial condition error is estimated to be 0.852. Using one realization of this random initial condition error, the TLM accuracy is 0.4 time unit, which is less than the mean accuracy time of the 1000 member ensemble. This rather short accuracy time given by the initial condition error indicates potential difficulties for the assimilation in the first cycle. In the previous section, the RMS of the model error is estimated to be 0.003 69 and using one realization of this model error, the TLM accuracy is 4.077 time units, which is very close to the mean of the 1000 member ensemble. This implies that for the choice of our one realization of initial and model errors, once the initial condition errors have been corrected in the first cycle, the TLM will have a longer time of validity in subsequent cycles, allowing for a possible extension of the cycle length.

The initial and model errors selected for this paper are applied to the TLM [Eq. (11)], where the background $\mathbf{x}^f$ is the nonlinear solution obtained with the initial conditions (x_0, y_0, z_0) using the RK4 time stepping with $dt = 1/60$. The time series of the background and TLM trajectories are shown in Fig. 3. The TLM strays from the background in less than half a time unit. This deviation of the TLM is caused by the magnitude of errors and their growth, as well as the early transition that the background undergoes around $t = 0.5$. The TLM follows this transition with a slight delay and wrong amplitude. After this transition, the errors continue to grow unbounded and the TLM solution does not recover.

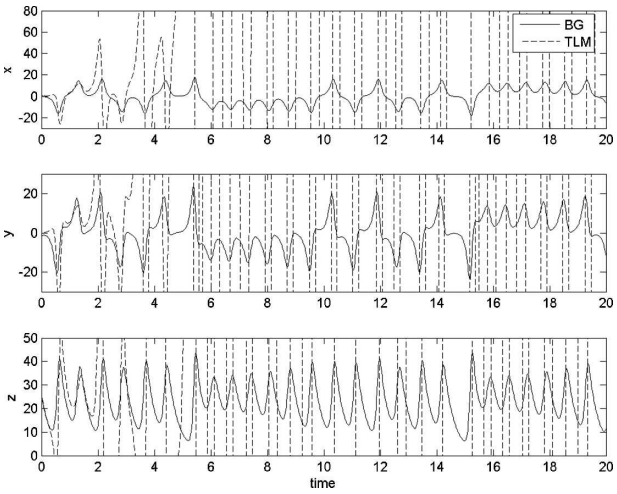

FIG. 3. Comparison between the background solution (BG) and the TLM linearized about this background. The background is the Lorenz attractor computed using the RK4 time-stepping scheme with a time step of $dt = 1/60$, and the coefficients and initial conditions stated in the text. The TLM initial conditions and dynamics are perturbed from the background by the initial condition and model error specified in the text and represented by the gray asterisks in Fig. 2. The result shows that the TLM diverges from the background solution after about 0.4 time units.

4. The experiments

The previous section demonstrates how the TLM accuracy is limited in time given the expected initial condition and model errors. For the choice of errors in section 2, the TLM is accurate for only 0.4 time units. It is clear that under these conditions, we cannot expect a satisfactory assimilation solution for the entire assimilation window [0, 20]. The background trajectory used to compute the global solution is obtained by integrating the nonlinear model, and adding the initial condition and model errors described in section 2. Note that although the role of the background for linearization purposes is the same in the assimilation problem and in the TLM, the background need not be the same in both. In the TLM, the background and the perturbation errors that force the TLM are needed for testing purposes. Thus, an unperturbed nonlinear solution was used as the background. For the assimilation experiments discussed in this section, the background also serves as the field to be corrected. In this case, the background is obtained by solving the nonlinear model and adding the selected errors as perturbations in order to simulate a solution that is significantly different from the data. An additional experiment will introduce a model error through the perturbation of the dynamical model parameters.

a. The inaccuracy of the global solution

Figure 4 compares the global solution to the background and true solutions (note that the data are sampled from the true solution every 0.25 time units except at $t = 0$ and $t = 20$). The global solution attempts to correct the perturbed background by assimilating all data over the entire time frame [0, 20] using the direct representer method with one outer loop. Even though the time frame of assimilation is far greater than the stability of the TLM, the solution is able to track the data somewhat for about seven time units. One should notice that whenever the discrepancy between the background and the data increases, the global solution follows the background due to the relatively small covariance of the model error. It can be seen from Fig. 5

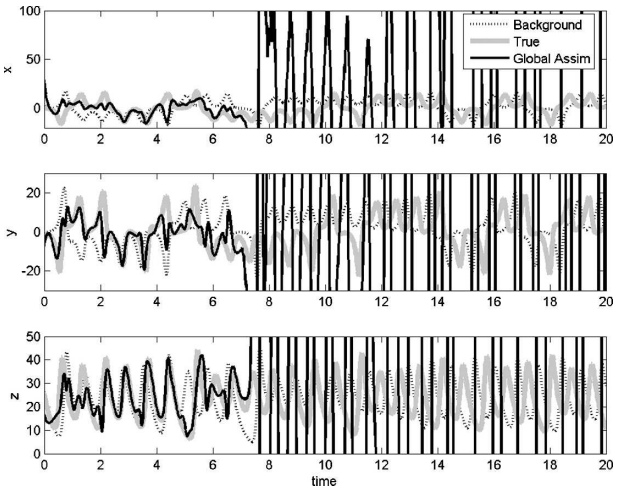

FIG. 4. The assimilated solution (solid black line) computed using the direct representer method over the entire time period (20 time units) goes unstable after about 7 time units. The data are sampled from the true solution (gray line) every ¼ time unit and are assimilated in an attempt to correct the wrong background (dotted line). The background is computed in the same fashion as in Fig. 3, except instead of the TLM being perturbed, the background is perturbed by the specified initial condition and model error.

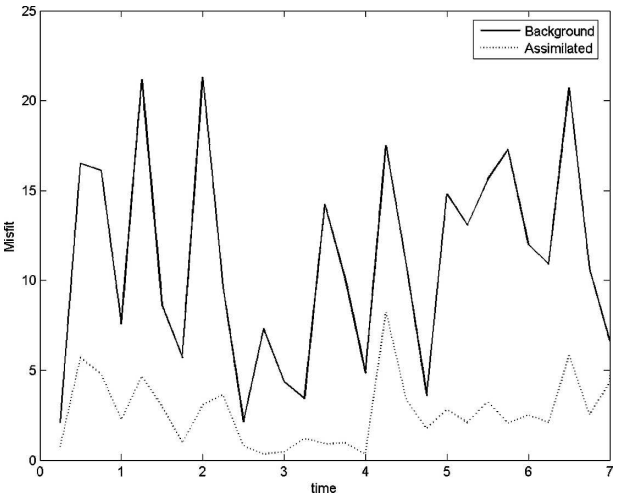

FIG. 5. RMS misfits between the data and the background (solid line) and assimilated (dotted line) solutions for the first 7 time units in Fig. 4. This plot reveals that even though the TLM is only reliable for about 0.4 time units (see Fig. 3), the assimilated solution is stable for about seven time units and is correcting the background toward the data during this time period.

that the global solution is able to reduce the prior misfits significantly (even beyond the time range of accuracy of the TLM) before losing track of the data. Beyond seven time units, the misfit between assimilated solutions and data grows rapidly and can be attributed to the increasing errors in the TLM approximation. One can therefore conjecture that the error growth in the TLM can be limited by reducing the length of the assimilation window. This provides the incentive for reducing the assimilation window and using the cycling representer method.

b. The advantage of the cycling representer approach

In this section, the cycling representer method is implemented for the same assimilation problem by subdividing the assimilation window into cycles of equal length. This is motivated by the reduced range of accuracy of the TLM and the ability of the assimilation method to adjust to the data beyond that range. There are three clear advantages that one can foresee in this approach: (i) a shorter assimilation window will limit the growth of errors in the TLM, (ii) the background for the next cycle will be improved, and (iii) the overall computational cost is reduced. It is assumed that the assimilation in the current cycle will improve the estimate of the state at the final time. The ensuing forecast (the solution of the nonlinear model propagated from the final state) is a better background for the next cycle than the corresponding portion of the background used in the global solution. The same random error is added to the dynamics of the nonlinear forecast of each cycle as was added to the corresponding portion of the background used in the global assimilation. Except for the first cycle, the initial condition of each forecast is perturbed by the model error instead of the initial condition error. The chaotic behavior of the model and the random error will cause a divergence (that grows with the cycle length) between the observations and the background. All experiments that follow use the same model error specifications as used in the global assimilation, and unless noted otherwise, the direct representer method is applied with a single outer loop, which is exactly what was done for the global solution.

The first cycling representer experiment uses 2 cycles of 10 time units each. The results in Fig. 6 show that the solution is drastically improved compared to the global solution in Fig. 4. The first thing one observes is the disappearance of the large values of the global solution in the last 2/3 of the assimilation window. The analyzed solution in both cycles clearly benefits from a shorter assimilation window limiting error growth in the TLM, and the second cycle has the extra advantage of using a better estimate of the background. The misfits to the data are decreased, although they still exceed the observation standard deviation in many portions of the trajectory.

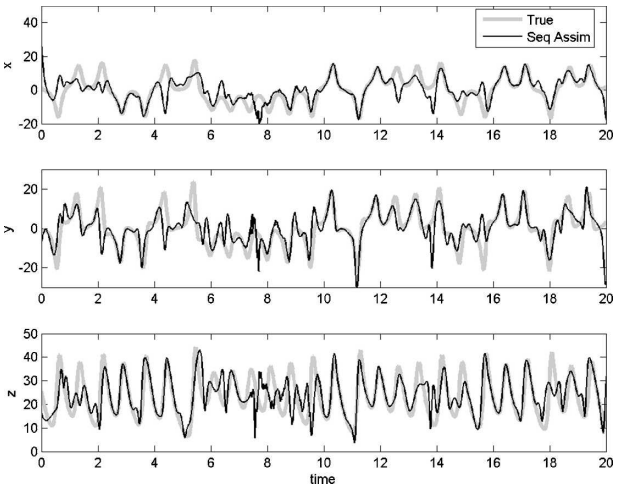

FIG. 6. Same as in Fig. 4 except that here the assimilation is separated into two cycles, where the final assimilated solution of the first cycle (at $t = 10$ time units) is used as the initial conditions for the computation of the background in the second cycle. The background for the first cycle is perturbed with both initial condition and model error, whereas for the second cycle the background is only perturbed with model error. By performing two cycles of the direct representer method instead of one (Fig. 4), the assimilated solution is now stable for the entire time period.

The results in Fig. 7 show the RMS error between the truth and the assimilated solution with respect to time for cycle lengths of 1, 2, 5, and 10 time units. It is shown that the RMS error increases with the cycle length. This is to be expected since longer cycles violate the TLM accuracy criterion. In other words, the steady decrease of RMS error with respect to the cycle length indicate that as the latter approaches the TLM stability time for the range of perturbations given by the adjoint model, the assimilation algorithm is better able to fit the data.

c. *The cycling indirect representer approach*

In realistic atmospheric and ocean models it is not possible to compute all the representer functions and implement the direct representer algorithm. In this section the indirect representer algorithm is implemented (Amodei 1995; Egbert et al. 1994; Chua and Bennett 2001; also see the appendix). In the first attempt to apply the indirect approach a cycle length of one time unit is used. Unlike the direct representer approach where an exact matrix inversion is used to compute the representer coefficients, the indirect approach uses an iterative solver (i.e., the CGM). The CGM does not require the entire matrix to be inverted; rather it iterates through search directions in data space. The stopping criterion of the CGM is either (i) when the relative norm of the residual is less than 10^{-3}, or (ii) when the number of iterations equals the number of measurements within the cycle. For well-conditioned systems, the first stopping criterion is typically reached first therefore requiring less computation time compared to the direct method.

The first remark in this experiment is that the representer matrix for the first cycle assimilation is ill conditioned (condition number $= 2.3 \times 10^9$) due to the rather high variance of the initial condition errors. As a consequence, the CGM converges very slowly and does not reach the first stopping criterion. Since proper convergence is scarcely reached within the first cycle, outer loops are required in order to achieve an acceptable solution. In this experiment, four outer loops are applied to each cycle (Fig. 8b). This assimilated solution tracks the data well and compares fairly well with the direct approach (Fig. 8a), albeit for a higher computational cost. Since the indirect method seems to only have difficulties in the first cycle, an experiment is performed where four outer loops are only applied in the first cycle, and a single outer loop is applied in subsequent cycles (Fig. 8c). This solution is less expensive than the previous and its accuracy and cost is comparable to the direct method. The reason that this approach performs well is because it is important to perform an accurate assimilation in the first cycle in order to obtain a good fit to the data and correct the background containing the initial condition error. This improved state estimate provides a good initial condition and thus background for the next cycle and so forth. It should be emphasized here that the proposed method does not require or rely on the ability to obtain a good analysis in the first cycle, although it is a desirable outcome. The additional outer loops in the first cycle were used in the context of the indirect representer approach, in order to overcome the slow convergence of the conjugate gradient method due to the poor conditioning of the representer matrix. If the first cycle fails to give a good analysis, it still will minimize to some extent the discrepancy between the background and the data at the final time, which will yield a better background for the next cycle. As this process is repeated, the system will fit the data, as illustrated in Figs. 7 and 8. The lesson learned here is that in cycling the indirect representer method, only the first few cycles may necessitate outer loops.

It should also be mentioned that when the length of the cycle is reduced to 0.5 time units, which is roughly the stability time for the TLM in the first cycle, the CGM converges rather well and there is no need for multiple outer loops. This result is shown in Fig. 9 and implies that the indirect method works satisfactorily well without outer loops when the cycle length is at

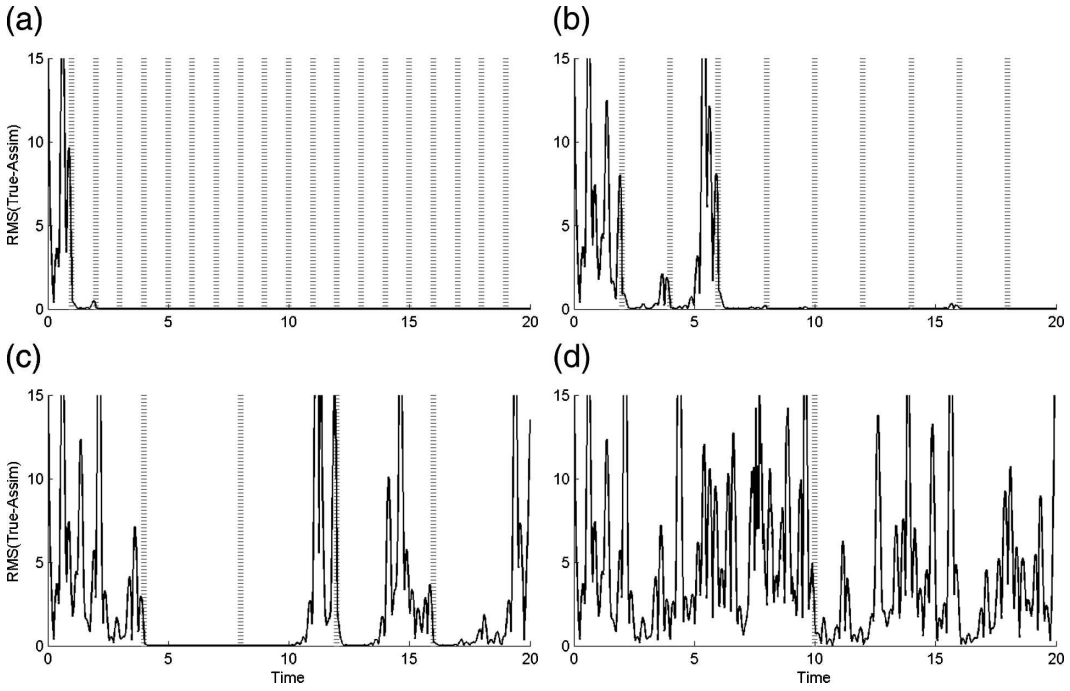

FIG. 7. RMS of the misfit between assimilated and true solutions using different numbers of cycles: (a) 20, (b) 10, (c) 5, and (d) 2 cycles. The cycle boundaries are depicted by dashed lines. By increasing the number of cycles from 2 (Fig. 7d is the misfit of Fig. 6) to 20 (Fig. 7a), a significant improvement in the assimilated solution is achieved.

about the TLM accuracy range. Of course, this conclusion would have to be tested with realistic models.

d. The cost

One major reason why the representer method is not widely implemented is the perceived computational cost. The biggest reduction in cost is achieved by limiting the outer loops to one, as was mentioned above. Further gains in computational cost are obtained by cycling the representer method. Assume that the matrix inversion in the direct method is performed with a cost of $O(M\log M)$ for computing M representer coefficients, where M is the number of measurements. The cycling approach total cost will be $N_{cy} \times O(M_{cy}\log M_{cy})$, where N_{cy} is the number of cycles and M_{cy} is the number of measurements within each cycle (assuming that the measurements are uniformly distributed in the assimilation interval). Although $N_{cy} \times M_{cy} = M$, $\log M_{cy}$ gets exponentially smaller with increasing N_{cy}, thus, decreasing the computational cost as illustrated in Table 1. However, there is a drawback to reducing the cycle length. The data influence is extended beyond the cycle interval only through an improved initial condition for the next cycle. Future data contained in subsequent cycles will not contribute to the assimilation in the current and past cycles. One should keep this in mind, as well as the time decorrelation scale of the model errors, in choosing the appropriate cycle length.

The cost reduction in the indirect method is less obvious. However, for well-conditioned representer matrices, the indirect method typically converges in a number of iterations that is a fraction of the total number of measurements. That is, the cost of the indirect method is a fraction of the cost of the direct method. Therefore, since cycling and avoiding the outer loops reduce the cost of the representer method in the direct method, they will further reduce the cost in the indirect method. Given the difficulties encountered by the indirect method in the first cycle of one time unit, longer cycles were not implemented. The computational cost of the 20-cycle experiment with the conjugate gradient using 4 outer loops in the first cycle and 1 outer loop in the subsequent cycles is 1.27 s, just slightly higher than the cost of the direct method, the excess cost being attributed to the outer loops in the first cycle.

e. Strong constraint cycling versus weak constraint cycling

The assimilation problem can be solved using the strong constraint representer approach. Here the strong and weak constraint solutions are compared within the same setting. From the known sensitivity of

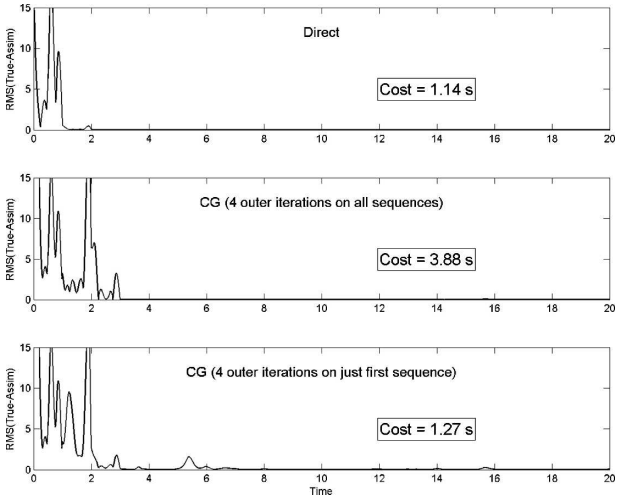

FIG. 8. RMS of the misfit between assimilated and true solutions using 20 cycles. Assimilated solutions are computed using (a) the direct representer method (same as in Fig. 7a), (b) the indirect conjugate gradient method using four outer loops for each cycle, and (c) the indirect conjugate gradient method using four outer loops only for the first cycle. Also displayed in this figure are the times required for each computation. The indirect approach (Fig. 8c) is a more applicable method and has about the same cost as the direct method (Fig. 8a). The indirect approach, however, takes about twice as long for the assimilated solution to reach an accurate level.

the Lorenz model to initial conditions, the question one tries to address by the strong constraint approach is whether it is possible to accurately fit the data within the cycle by correcting/controlling only the initial condition. How short does a cycle need to be to achieve an accurate solution? The strong constraint solution is obtained by the same procedure as the weak, except that the model error covariance is set to zero. In comparison to Fig. 7, results in Fig. 10 reveal that the weak constraint solution is not only more accurate, but can also afford longer cycles than the strong constraint. The

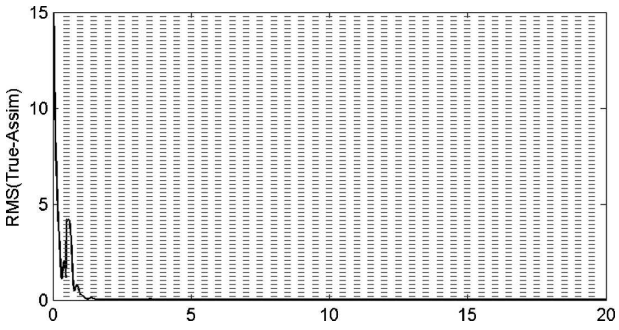

FIG. 9. RMS of the misfit between assimilated and true solutions using 40 equal-length cycles. The assimilated solution is computed using the indirect conjugate gradient method with a single outer loop applied to all cycles.

TABLE 1. Computational cost of the global and cycling solutions using the direct method with a single outer loop.

	Global	2 cycles	4 cycles	5 cycles	10 cycles	20 cycles
Time (s)	21.37	9.33	4.49	3.59	1.90	1.09

strong constraint is almost confined to the TLM validity time, and needs quite a few cycles to start matching the data. In the experiment with cycles of two time units, the weak constraint accurately fits the data after three cycles, but the strong constraint never does. When the cycle length is decreased to one time unit, the weak constraint fits the data in the second cycle and afterward, while the strong constraint starts fitting the data only in the 15th cycle. The computations here show the superior performance of the weak constraint approach. The weak constraint approach, however, comes with a higher computation cost. For these experiments, the computational cost of the strong constraint is about half of that for the weak constraint. The reason for this large difference in cost is that the number of operations to propagate the Lorenz attractor is relatively small and is about equal to the number of operations required to convolve the model errors in space and time. For larger operational models, this computational cost difference would be significantly smaller. Both approaches could perform well in operational settings if they are given sufficient spinup cycles, less for the weak constraint.

f. Assimilating only the x variable

It is unlikely that in realistic applications all components of a dynamical model can be observed or sampled. For example, the majority of measurements taken in the ocean are from the surface. To mimic this scenario, an experiment is carried out where only the x component of the Lorenz model is sampled at the same frequency of the 0.25 time unit. The cycling representer algorithm with cycles of one time unit and one outer loop is used to assimilate these data. Results in Fig. 11 show that the method is able to fit the data satisfactorily, albeit for a longer adjustment or spinup time. Compared to Fig. 7a, it takes five additional cycles for the assimilation to start matching the data. Therefore, a decrease in the total number of assimilated observations and the fact that only one state variable is sampled are not limiting factors for the proposed algorithm.

g. First cycle accuracy as a function of initial condition error variance and outer loops

As mentioned above, the proposed algorithm does not solely rely on the ability to obtain a good analysis in

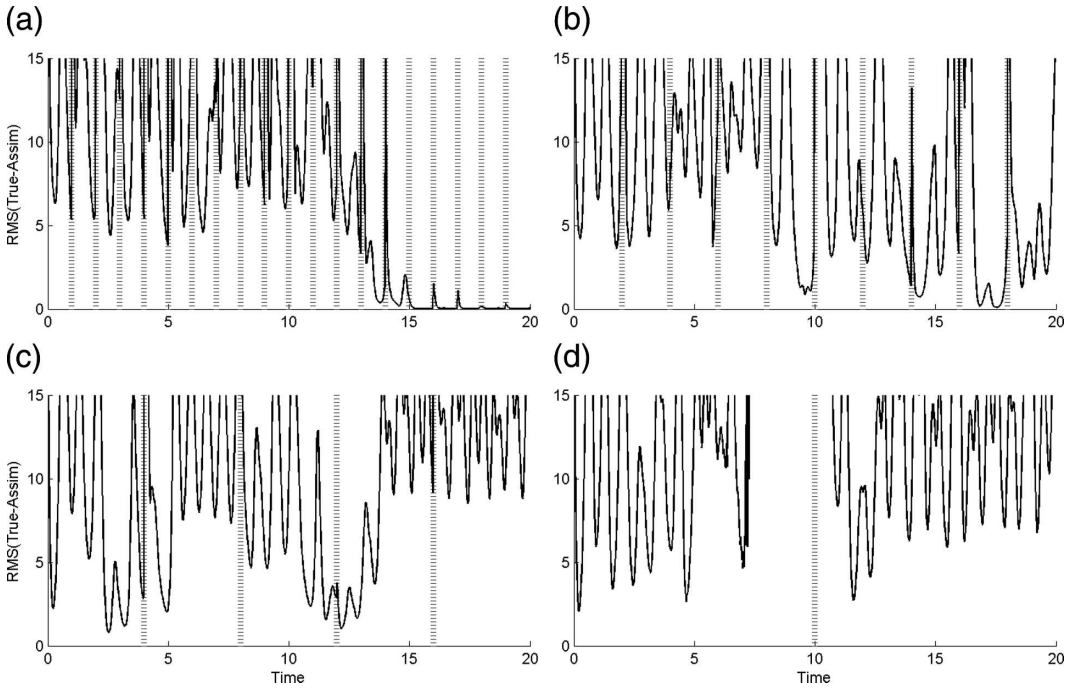

FIG. 10. Same as in Fig. 7, except that the strong constraint representer method is used instead of the weak constraint. Comparison of this figure with Fig. 7 shows the benefit of using weak constraint to assimilate data with strongly nonlinear dynamic systems.

the first cycle. Since the Lorenz model is extremely sensitive to initial perturbations, a significantly large error was prescribed to the initial condition to test the robustness of the algorithm. It was shown in Figs. 8b,c that comparable accuracy can be obtained using either outer loops in every cycle or just in the first cycle. An experiment is presented here that displays the accuracy of the assimilation after one time unit as a function of the magnitude of the initial perturbations and the number of outer loops. The results in Fig. 12 show that the accuracy of the assimilated solution at the end of the first cycle decreases as the RMS of the initial condition error increases, which is to be expected, and that the same accuracy increases with the number of outer loops used.

For initial condition RMS errors below 20%, one outer loop is sufficient to achieve an accurate solution at the end of the first cycle (i.e., the data misfit is less than the data error). When the initial condition RMS error is between 20% and 40%, two outer loops are necessary. Three outer loops are needed when the RMS is between 40% and 45%, and beyond 45% at least five outer loops are needed. The initial condition errors selected in section 2 and used in the assimilation experiments (Fig. 2a) have an RMS error of 85%, which according to Fig. 12, would require at least five outer loops to obtain an accurate solution at the end of the first cycle. However, using these initial condition errors with one outer loop (Fig. 7a) produced an accurate assimilated solution after two cycles. Therefore, one must determine if it is worth the computational cost to perform additional outer loops to improve the accuracy of the first cycle, knowing that it may not significantly influence the accuracy of subsequent cycles.

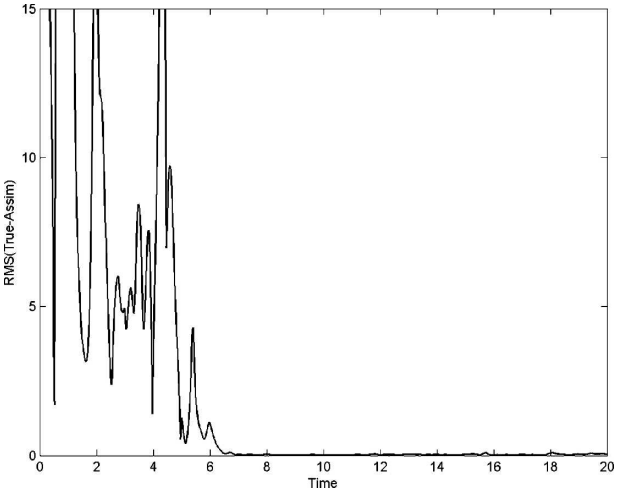

FIG. 11. RMS of the misfit between assimilated and true solutions using only the x observations sampled every 0.25 time unit and the direct cycling representer algorithm with 20 cycles and one outer loop.

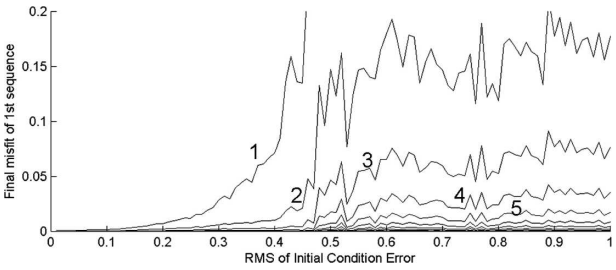

FIG. 12. RMS of posterior misfit at the end of the first cycle as a function of outer loops and the initial condition RMS error.

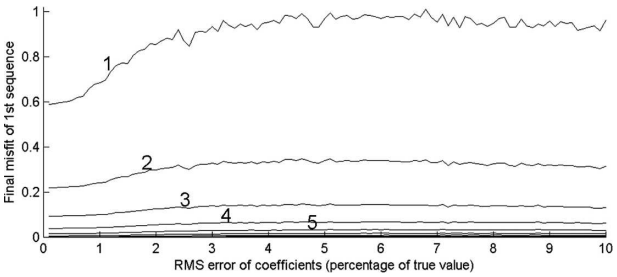

FIG. 13. RMS of posterior misfit at the end of the first cycle as a function of the parameter error (in terms of percentage of true value) and outer loops.

h. Model error through perturbed parameters

The model error described in section 2 consists of perturbations of the right-hand side of the model while keeping the same dynamics that generated the observations. Another possible method is to generate the model error by perturbing the parameters in the dynamical model. Contrary to the previous experiments, perturbing the parameters in the model ensures that the dynamics used in the assimilation are different from the dynamics that generated the data. The perturbation of the model parameters is similar to that of the initial conditions described in section 2. First the RMS of parameters is computed, and the perturbation is a percentage of that RMS multiplied by a normally distributed random number. For each selection of the perturbation magnitude, 1000 realizations of the random number are used in separate assimilation experiments, generating a 1000-member ensemble of assimilated solutions. In Fig. 13, the mean of this ensemble of posterior misfits at the end of the first cycle is shown as a function of the perturbation magnitude and the outer loops. As the parameter perturbation magnitude varies from 0.1% to 10%, there is only a slight increase in the posterior misfits at the end of the first cycle in the first outer loop. Figure 13 indicates that the misfits decrease as more outer loops are used, with the biggest decrease in the second and third outer loops. However, at least five outer loops are needed to obtain an accurate solution at the end of the first cycle. After the first outer loop the misfits are fairly constant for perturbations between 2% and 10% in each subsequent outer loop.

5. Discussion

The implications of the proposed algorithm to the real world of oceanography and meteorology are immediate. The dimensionless time (τ) in the Lorenz model is related to a simplified one-layer atmospheric model time (t) by $\tau = \pi^2 H^{-2}(1 + a^2)\kappa t$, where $a^2 = 0.5$, H is the depth of the fluid, and κ is the conductivity. Therefore, for a fluid depth of 500 m and a conductivity of 25×10^{-3}, a time unit in the Lorenz model corresponds to 7.818 days. The TLM stability range depends on the resolved processes and scales, the model resolution, and of course, the nonlinear interactions within the model dynamics. For example, the TLM of a general circulation ocean model forced by climatological or monthly means atmospheric fluxes and designed to resolve seasonal to interannual variability in the tropical Pacific will be stable for months (Ngodock et al. 2000), whereas the TLM of a fine resolution coastal circulation ocean model forced with synoptic atmospheric fluxes with complex bottom topography may not be stable beyond a week or two. In the latter case any hindcast experiment using the representer method for time ranges longer than a month will benefit from the algorithm proposed in this study. Additionally, the outer loops may not be restricted to the first cycle as is the case here with the Lorenz model. They can be turned on as needed. In real-world applications, even though the assimilation system has been spun up over several cycles and only one outer loop sufficed for the previous cycle, there may arise a situation where nonlinearities become stronger than usual, causing the background in the current cycle (i.e., the forecast from the previous cycle) to deviate significantly far from the data.

The proposed algorithm will therefore be suitable both for reanalysis computations and cycles of analysis and predictions (e.g., in coastal ocean monitoring).

6. Conclusions

The impact of TLM accuracy in a cycling representer implementation is examined using the Lorenz attractor. Errors in the initial conditions and dynamics of the Lorenz attractor are assumed and added to the background. The error estimates used here indicate that the TLM about this background is accurate up to 0.4 time

units. For the Lorenz equations, weak constraint assimilation globally over a 20 time unit period is not possible because of the inaccuracy and instability of the TLM over such a long time period. Experiments using successively shorter time intervals indicate that a weak constraint cycling representer assimilation is able to provide an accurate solution with cycle intervals of two time units or less.

For the problems examined here, it is found that once a sufficiently accurate initial condition is provided for a cycle, the TLM is sufficiently accurate in subsequent cycles therefore neglecting the need for outer loops to perform the nonlinear minimization. For the Lorenz attractor system, outer loops are required for the first cycle and not for subsequent cycles.

The representer method is used for the 4DVAR problem with the Lorenz model in various settings: the direct approach for both global and cycling assimilations, and the indirect approach for the cycling. The first guess is taken as the nonlinear background solution that serves for the linearization in the TLM. The cycling approach is motivated by the limited stability of the TLM and the potential operational applications. For the selected set of initial and model errors, the global assimilation solution could track the data for about seven time units. After this time, the solution rapidly diverges from the data owing to unbounded error growth in the TLM. The cycling solution is more apt to fit the data, especially with small cycle lengths, because the error growth in the TLM is limited. The iterated indirect representer method was then implemented via the CGM. A slow convergence of the CGM was observed in the first cycle, due to the poor conditioning of the representer matrix. By applying four outer loops to this first cycle, a sufficiently accurate assimilation was achieved. When the cycle length was reduced to 0.5 time units, which is about the time range of the TLM accuracy, the indirect approach converged rather well without outer loops. In addition, these weak constraint assimilation experiments were shown to significantly outperform their strong constraint counterparts.

This paper shows that cycling the representer method for strongly nonlinear problems is a valuable assimilation tool. The cycle length could be slightly longer than the time range of accuracy of the TLM. The indirect representer algorithm may or may not require outer loops in the early cycles depending on the length of the cycle. Once the assimilated solution becomes sufficiently accurate in the estimation of the final state for the early cycles, the need of outer loops is removed. This significantly reduces the cost of the weak constraint variational assimilation, in addition to the reduction obtained by cycling.

The robustness of the proposed algorithm is demonstrated by its ability to recover the true solution after few spinup cycles in a series of separate experiments: the number of measurements was decreased and only the x component of the Lorenz model was sampled, the initial condition error magnitude was varied, and finally the model error was introduced by perturbing the parameters in the dynamical model. In these settings additional outer loops were necessary to achieve an accurate solution at the end of the first cycle. However, this study clearly shows that dropping the outer loops was not detrimental to the algorithm, even with the rather pessimistic choice of initial condition and model errors.

Acknowledgments. This work was sponsored by the Naval Research Laboratory as part of the project "Slope to Shelf Energetics and Exchange Dynamics" (NRL/JA/7320-05-5286). The authors are thankful to two anonymous reviewers whose comments helped to improve the quality of this paper.

APPENDIX

Solving the Linear EL System Using the Representer Method

Given a background field $\mathbf{x}^f$, the linear EL to be solved is

$$-\frac{d\boldsymbol{\lambda}}{dt} = \left[\frac{d\mathbf{F}(\mathbf{x}^f)}{d\mathbf{x}}\right]^{\mathrm{T}} \boldsymbol{\lambda} - \mathbf{H}^{\mathrm{T}}\mathbf{w}(\mathbf{d} - \mathbf{H}\hat{\mathbf{x}})$$
$$\boldsymbol{\lambda}(T) = 0 \tag{A1}$$

and

$$\frac{d\hat{\mathbf{x}}}{dt} = \mathbf{F}(\mathbf{x}^f) + \frac{d\mathbf{F}(\mathbf{x}^f)}{d\mathbf{x}}(\hat{\mathbf{x}} - \mathbf{x}^f) + \mathbf{C}_{qq} \cdot \boldsymbol{\lambda}$$
$$\hat{\mathbf{x}}(0) = \mathbf{x}_0^f + \mathbf{C}_{ii}\boldsymbol{\lambda}(0). \tag{A2}$$

The representer expansion for uncoupling Eqs. (A1)–(A2) is

$$\hat{\mathbf{x}}(t) = \mathbf{x}^f(t) + \sum_{m=1}^{M} \alpha_m \mathbf{r}_m(t). \tag{A3}$$

Here the background (i.e., the trajectory around which the model is linearized) is also taken as the first guess (the solution that the assimilation will correct). The representer functions $\mathbf{r}_m$, $m = 1, \ldots M$ are computed from

$$-\frac{d\boldsymbol{\lambda}_m}{dt} = \left[\frac{d\mathbf{F}(\mathbf{x}^f)}{d\mathbf{x}}\right]^{\mathrm{T}} \boldsymbol{\lambda}_m - \mathbf{H}^{\mathrm{T}}\delta(t - t_m)$$

$$\boldsymbol{\lambda}(T) = 0 \tag{A4}$$

and

$$\frac{d\mathbf{r}_m}{dt} = \frac{d\mathbf{F}(\mathbf{x}^f)}{d\mathbf{x}}(\mathbf{r}_m) + \mathbf{C}_{qq} \cdot \boldsymbol{\lambda}_m$$

$$\mathbf{r}_m(0) = \mathbf{C}_{ii}\boldsymbol{\lambda}_m(0). \tag{A5}$$

It may be shown (e.g., Bennett 2002) that the representer coefficients α_m, $m = 1, \ldots M$, in Eq. (A3) are the solution of the linear system:

$$[\mathbf{R}^e + \mathbf{w}^{-1}]\boldsymbol{\alpha} = \mathbf{d} - \mathbf{H}\mathbf{x}^f, \tag{A6}$$

where $\mathbf{R}^e$ is the representer matrix, obtained by evaluating the representer functions at the measurements sites (i.e., the mth column of $\mathbf{R}^e$ is $\mathbf{H}\mathbf{r}_m$). In practice, solving Eq. (A6) does not require the computation of the entire representer matrix. An iterative method such as the conjugate gradient may be invoked, as long as the matrix-vector product on the left-hand side of Eq. (A6) can be computed for any vector in the data space. This is made possible through the indirect representer algorithm (Amodei 1995; Egbert et al. 1994), which is also used to assemble the right-hand side of Eq. (A3) without the explicit computation and storage of the representer functions. Specifically, given a vector $\mathbf{y}$ in the data space, the product $\mathbf{R}^e\mathbf{y}$ is obtained by solving Eqs. (A4)–(A5) with $\mathbf{y}$ replacing the Dirac delta in the right-hand side of Eq. (A4), then applying the observation operator $\mathbf{H}$ to the resulting $\mathbf{r}$. Once the representer coefficients $\boldsymbol{\alpha}$ are obtained, the optimal residuals are computed by solving Eq. (A4), where the single Dirac delta function is now replaced by the linear combination $\Sigma_{m=1}^{M}\alpha_m\delta(t - t_m)$. These residuals are then used in the right-hand side of Eq. (A5) to compute the optimal correction to the first guess $\mathbf{x}^f$.

Given a background trajectory $x^f(t)$, $y^f(t)$, and $z^f(t)$, [Eq. (1)] the system is tangent linearized as

$$\frac{dx}{dt} = \sigma(y - x)$$

$$\frac{dy}{dt} = \rho x - y - x^f z - (x - x^f)z^f$$

$$\frac{dz}{dt} = x^f y - (x - x^f)y^f - \beta z, \tag{A7}$$

and the adjoint of the Lorenz model is given by

$$-\frac{d\lambda^x}{dt} = -\sigma\lambda^x + \rho\lambda^y - z^f\lambda^y - y^f\lambda^z$$

$$-\frac{d\lambda^y}{dt} = \sigma\lambda^x - \lambda^y + x^f\lambda^z$$

$$-\frac{d\lambda^z}{dt} = -x^f\lambda^y - \beta\lambda^z. \tag{A8}$$

REFERENCES

Amodei, L., 1995: Solution approchée pour un problème d'assimilation de données avec prise en compte de l'erreur du modèle. *Comp. Rend. Acad. Sci.,* **321** (IIa), 1087–1094.

Bennett, A. F., 1992: *Inverse Methods in Physical Oceanography*. Cambridge University Press, 347 pp.

——, 2002: *Inverse Modeling of the Ocean and Atmosphere*. Cambridge University Press, 320 pp.

——, B. S. Chua, and L. M. Leslie, 1996: Generalized inversion of a global numerical weather prediction model. *Meteor. Atmos. Phys.,* **60,** 165–178.

Chua, B. S., and A. F. Bennett, 2001: An inverse ocean modeling system. *Ocean Modell.,* **3,** 137–165.

Egbert, G. D., A. F. Bennett, and M. G. G. Foreman, 1994: TOPEX/POSEIDON tides estimated using a global inverse method. *J. Geophys. Res.,* **99,** 24 821–24 852.

Evensen, G., 1997: Advanced data assimilation for strongly nonlinear dynamics. *Mon. Wea. Rev.,* **125,** 1342–1354.

——, and N. Fario, 1997: solving for the generalized inverse of the Lorenz model. *J. Meteor. Soc. Japan,* **75** (1B), 229–243.

——, and P. J. Van Leeuwen, 2000: An ensemble Kalman smoother for nonlinear dynamics. *Mon. Wea. Rev.,* **128,** 1852–1867.

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

Jacobs, G. A., and H. E. Ngodock, 2003: The maintenance of conservative physical laws within data assimilation systems. *Mon. Wea. Rev.,* **131,** 2595–2607.

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

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

——, E. F. Carter, and S. T. Blue, 1999: Data assimilation into nonlinear stochastic models. *Tellus,* **51A,** 167–194.

Muccino, J. C., and A. F. Bennett, 2002: Generalized inversion of the Korteweg-De Vries equation. *Dyn. Atmos. Oceans,* **35** (3), 227–263.

Ngodock, H. E., B. S. Chua, and A. F. Bennett, 2000: Generalized inversion of a reduced gravity primitive equation ocean model and tropical atmosphere ocean data. *Mon. Wea. Rev.,* **128,** 1757–1777.

Uboldi, F., and M. Kamachi, 2000: Time-space weak-constraint data assimilation for nonlinear models. *Tellus,* **52A,** 412–421.

Xu, L., and R. Daley, 2000: Towards a true 4-dimensional data assimilation algorithm: Application of a cycling representer algorithm to a simple transport problem. *Tellus,* **52A,** 109–128.

——, and ——, 2002: Data assimilation with a barotropically unstable shallow water system using representer algorithms. *Tellus,* **54A,** 125–137.