Available online at www.sciencedirect.com
ScienceDirect
Ocean Modelling 19 (2007) 101–111
Ocean Modelling
www.elsevier.com/locate/ocemod

Cycling the representer algorithm for variational data assimilation with a nonlinear reduced gravity ocean model

Hans E. Ngodock *, Scott R. Smith, Gregg A. Jacobs

The Naval Research Laboratory, Building 1009, Stennis Space Center, MS 39529, USA

Received 15 March 2007; received in revised form 5 June 2007; accepted 6 June 2007
Available online 28 June 2007

Abstract

The representer method was used by [Ngodock, H.E., Jacobs, G.A., Chen, M., 2006. The representer method, the ensemble Kalman filter and the ensemble Kalman smoother: a comparison study using a nonlinear reduced gravity ocean model. Ocean Modelling 12, 378–400] in a comparison study with the ensemble Kalman filter and smoother involving a 1.5 nonlinear reduced gravity idealized ocean model simulating the Loop Current (LC) and the Loop Current eddies (LCE) in the Gulf of Mexico. It was reported that the representer method was more accurate than its ensemble counterparts, yet it had difficulties fitting the data in the last month of the 4-month assimilation window when the data density was significantly decreased. The authors attributed this failure to increased advective nonlinearities in the presence of an eddy shedding causing the tangent linear model (TLM) to become inaccurate. In a separate study [Ngodock, H.E., Smith, S.R., Jacobs, G.A., 2007. Cycling the representer algorithm for variational data assimilation with the Lorenz attractor. Monthly Weather Review 135 (2), 373–386] applied the cycling representer algorithm to the Lorenz attractor and demonstrated that the cycling solution was able to accurately fit the data within each cycle and beyond the range of accuracy of the TLM, once adjustments were made in the early cycles, thus overcoming the difficulties of the non-cycling solution. The cycling algorithm is used here in assimilation experiments with the nonlinear reduced gravity model. It is shown that the cycling solution overcomes the difficulties encountered by the non-cycling solution due to a limited time range of accuracy of the TLM. Thus, for variational assimilation applications where the TLM accuracy is limited in time, the cycling representer becomes a very powerful and attractive alternative, given that its computational cost is significantly lower than that of the non-cycling algorithm.
Published by Elsevier Ltd.

1. Introduction

Implementing a 4D-Var assimilation method requires the adjoint of the dynamical model in use. The adjoint is the transpose of the tangent linear approximation of the model being used. A successful assimilation experiment necessitates that the tangent linear model be stable and sufficiently accurate within the assimilation time window. In coastal ocean applications, strong nonlinearities arise from varying bathymetry, sporadic

* Corresponding author. Tel.: +1 228 688 5455.
E-mail address: Hans.Ngodock@nrlssc.navy.mil (H.E. Ngodock).

1463-5003/$ - see front matter Published by Elsevier Ltd.
doi:10.1016/j.ocemod.2007.06.004

atmospheric forcing, advection, upwelling, etc. These nonlinearities limit the time range of stability of the TLM to a few days or a few weeks depending on the forcing conditions, the circulation pattern (e.g. the boundary conditions), the model resolution, amongst others. With the TLM limitations in mind, the assimilation window must be chosen carefully: it should be a balance between the time scales of dynamical features of interest, the dynamical error correlation time scales and the time range of validity of the TLM. When the latter is smaller than the selected assimilation window, the cycling representer method therefore becomes attractive.

The cycling representer method was introduced by Xu and Daley (2000). It is an adaptation of the representer method (Bennett, 2002) for 4D-Var assimilations. The original representer method was traditionally applied to assimilating data in a rather long time window: Bennett et al. (1998, 2000, 2006); Ngodock et al. (2000). The method was also applied to a numerical weather prediction problem for hindcast and forecast, but only one cycle was used: Bennett et al. (1996). Xu and Daley (2000) proposed to cycle the representer method in time, for application in numerical weather prediction (NWP). They used a linear 1-dimensional transport model to lay down the concept, and also applied it to a linear unstable barotropic problem, Xu and Daley (2002). In these two applications there never was any issue with the TLM since the models were linear. Ngodock et al. (2007) explored the idea using the Lorenz attractor model, and made the case that cycling the representer method can be extremely beneficial in situations where the TLM is not stable for long periods of time. It was shown that cycling the representer method not only obviates the difficulties associated with an unstable TLM, it also reduces the cost of the assimilation, particularly when the outer loops are dropped. The outer loops are iterations over the linearizations of the nonlinear Euler–Lagrange problem associated with the minimization of the cost function involving a nonlinear model. Each of these iterations solves a linear data assimilation problem for which the representer method can be invoked. The latter expresses the assimilated solution or best estimate as the sum of a first guess (or background solution) and a finite linear combination of the representer functions, one per datum. In the first iteration a prescribed background field is used for linearizations purposes and, once the assimilation is completed, the corrected solution is used as background for the next iteration and so forth to formal convergence. In Ngodock et al. (2007) (hereafter NL07), the outer loops are dropped with the assumption that the TLM is accurate over the shorter cycle and at the end of the cycle the assimilation yields an accurate estimate of the solution that becomes the initial condition for the nonlinear forecast, which will be used as the background in the next cycle. The authors argue that as the system is spun-up over several cycles, the assimilation will match the data accurately. Thus, the algorithm proposed in NL07 is very suitable for coastal ocean data assimilation applications using 4D-Var methods, where short-range (a few days) hindcast assimilations are followed by short range forecasts.

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 cycling representer will be employed to assimilate data in a twin-model experiment using a 1.5 reduced-gravity nonlinear model for an idealized eddy shedding in the Gulf of Mexico. The model has previously been used in a comparison of data assimilation algorithms by Ngodock et al. (2006) (hereafter NG06). It is used here for three main reasons: firstly, it is a nonlinear model with a higher dimension than the Lorenz attractor. Secondly, it was reported in NG06 that the representer method did fit the data accurately in the entire assimilation window with a sufficient number of measurements. However when the measurements density was significantly decreased, the method had difficulties fitting the data towards the end of the assimilation window in the presence of an eddy shedding. It was hypothesized that the problem was due to inaccuracies in the TLM toward the end of the assimilation window and only a sufficient number of measurements could keep the assimilation accurate. Here we explore the ability of the cycling representer method to address this issue. Thirdly, this model is an intermediate step before applying the cycling representer method to a full three-dimensional nonlinear ocean model.

Assuming that the cycle length is chosen as mentioned earlier, it is expected that even when the TLM is stable and accurate over the entire assimilation time window, the cycling solution will still be more accurate than the solution that did not cycle, i.e. the solution obtained by assimilating all the observations in the time

window at once. Accuracy is measured here by the assimilated solution rms misfit to the data. The latter solution will be referred to as the non-cycling solution. The computational cost to achieve an accurate solution will be compared for both methods. The next section briefly describes the model and its domain, the experiments setup. Section 3 deals with the results of the experiments are while the computational cost and accuracy are discussed in Section 4. Concluding remarks follow in Section 5.

2. Model description and experiments setup

The 1.5 reduced gravity model used here is the same in NG06. It is an idealized configuration to simulate the eddy shedding in the Gulf of Mexico, see Fig. 1. The details of the model i.e. equations and parameters are not repeated here. The model parameters are tuned so that the idealized loop current eddy (LCE) sheds from the loop current (LC) with a period of 4 months. The reference initial condition has a LCE in the center of the model domain, and the remnant of a previously shed LCE at the northwestern corner of the model domain. In the next 3 months, the LCE in the middle of the domain will quasi-linearly propagate westward, interact with the western boundary and slowly move northward while dissipating. In the meantime, the loop current intrudes further into the domain and by the wake of the fourth month another LCE is about to shed from the LC. In the background initial condition, the LCE is just about to shed from the LC, and a previously shed LCE has reached the western coast of the GOM. This phase delay of approximately 2 months between the reference and the background initial conditions ensures a significant deviation of the background solution from the data.

The data are sampled from the reference solution according to 8 networks described in NG06, with 5 cm and 5 cm/sec data error for SSH and velocity respectively. Here the assimilation experiments are carried out for networks 3, 2, 1 using SSH and velocity data, and network 3 with only SSH data. The assimilation window is 4 months. In network 3, data are sampled from the reference solution every 200 km in each spatial

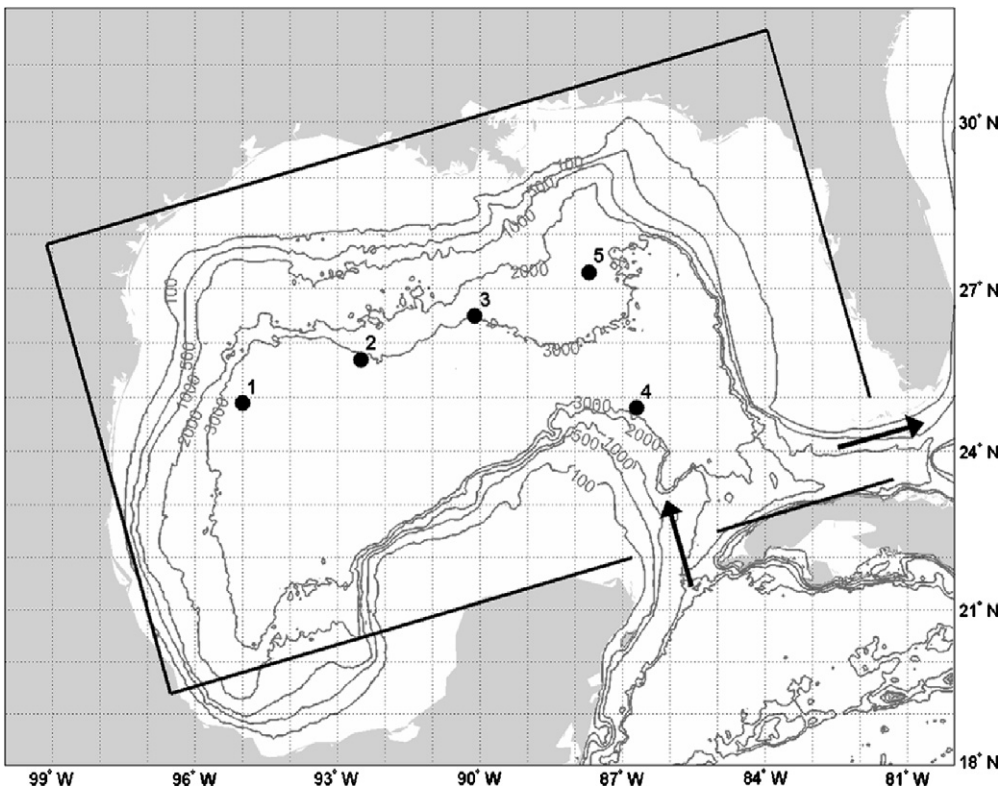

Fig. 1. Model domain and diagnostic measurement locations (black bullets).

dimension and every 10 days, while networks 2 and 1 sample the reference solution every 300 km (in both x and y directions) and every 5 and 10 days respectively. This produces a data density that increases with the network number. The covariances for the data, model and initial errors are the same as in NG06: the data error covariance is assumed diagonal with a variance of 25 cm^2 for SSH and 25 $cm^2\ s^{-2}$ for both components of velocity; the model errors are allowed only in the momentum equations following Jacobs and Ngodock (2003), and have spatial correlation scales 100 km in both x and y directions, a standard deviation $10^{-4}\ m^2\ s^{-2}$ (obtained by accounting for a typical stress of 0.1 Nm^{-2} which in turn is divided by a typical density of 1000 kg m^{-3}), and a time correlation scale of 10 days. The results from the non-cycling assimilation experiments are available from the experiments reported in the same reference. Only the cycling assimilation experiments are carried out here and compared to the corresponding non-cycling solution obtained with 6 outer loops. It should be noted that the initial error covariance at the beginning of a new cycle is not updated as the posterior error covariance from the previous cycle. This procedure is computationally expensive and is avoided here. The original initial error covariance is used in every cycle. A set of 5 diagnostic stations is used for evaluation in this study. The station locations are shown in Fig. 1. They are selected in such a way that they are common to all the sampling networks; locations 1–3 are distributed along the path of the LCE, location 4 is in the region where the LCE sheds, and location 5 is north of the LCE shedding region.

3. Cycling assimilation experiments

The first cycling representer assimilation experiments are carried out for network 1 using 4 cycles of 1 month each and 3 outer loops in each cycle. A cycle length of 1 month is chosen to allow (i) a stable TLM, (ii) time distributed data within each cycle (especially when the data is sampled every 10 days e.g. Networks 1 and 3), (iii) and the propagation of the data influence in time through the model dynamics and the model error covariance function. Fig. 2 shows the difference between the reference and the assimilated solutions for both the non-cycling and the cycling at the end of each month. This figure shows that although both solutions have comparable discrepancies in velocity and sea surface height with the reference solution at the end of the first month, the discrepancies decrease rapidly in the cycling solution and by the end of the assimilation window they are greatly reduced relative to the non-cycling solution. It is not the case with the non-cycling solution; the discrepancies persist and are mostly located around the region where the LCE sheds from the LC, i.e. where advective nonlinearities are strongest. This indicated that the failure of the non-cycling solution is associated with an inaccurate TLM as suggested in NG06. It is also worth mentioning here that the cycling solution is obtained with 3 outer loops in each cycle, which is half the computational cost of the non-cycling solution computed with 6 outer loops as reported in NG06.

In the second set of cycling representer experiments, data is assimilated for each networks 3, 2 and 1 using 4 1-month cycles in two cases: in the first case 3 outer loops are used in each cycle, and in the second case 3 outer loops are used only in the first cycle and 1 outer loop in the remaining cycles. To compare these two sets of solutions for each network against the corresponding non-cycling solution, the rms of the difference between the reference and the assimilated solutions is computed at the five diagnostic locations (Fig. 1). This rms is computed over the entire assimilation window and shown in Fig. 3. Results indicate that the cycling solutions have almost always the lower rms except at the first 2 locations for both velocity components, and at the second location for SSH, even though the cycling rms and the difference between the cycling and non-cycling rms never exceed the data error by a standard deviation. In contrast, at all other locations, the cycling rms is always lower than the data error for all networks while the non-cycling rms and the difference between the cycling and the non-cycling rms sometimes exceed the data error by more than a standard deviation, e.g. v-velocity at locations 3, 4 and 5. This figure also shows that the cycling solution obtained with 3 outer loops in the first cycle only and 1 outer loop in subsequent cycles is as accurate as the one with 3 outer loops in each cycle, except at location 1 for all three variables. The latter solution had already reduced the computational cost of the non-cycling roughly by 50%, so that the comparable accuracy of the former indicates that the cost can further be reduced roughly by another 50%.

The rms is also computed over the last two months of the assimilation and is shown in Fig. 4. It can be seen that all the rms values are slightly lower compared to in Fig. 3, indicating that the accuracy of both the cycling and non-cycling solutions is improving over time.

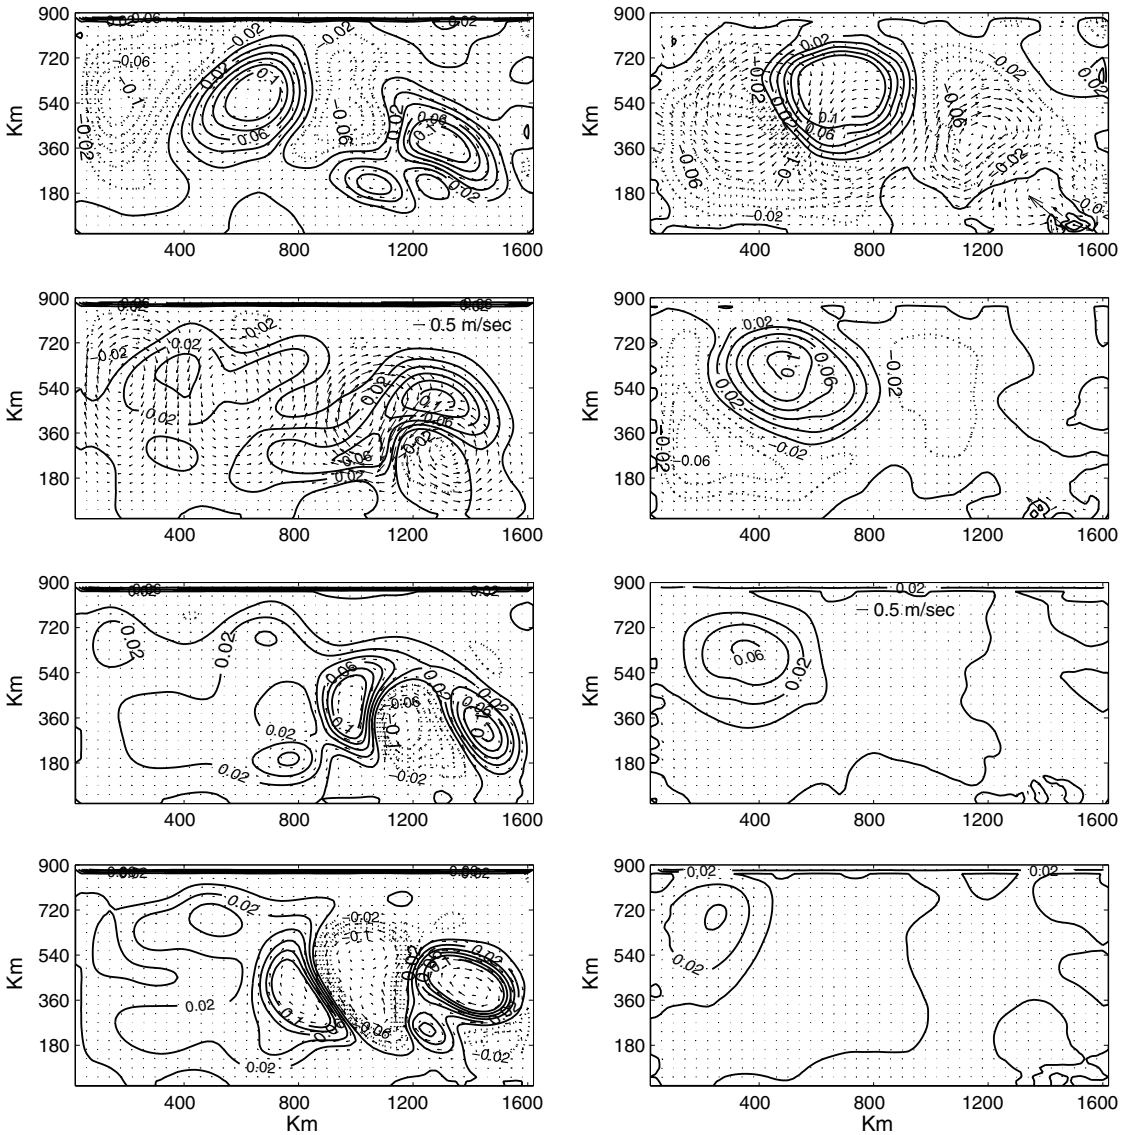

Fig. 2. The difference between the reference and the assimilated solutions obtained from the non-cycling (left column) and the cycling (right column) representer algorithms for network 1. The differences are shown at the end the first month (top row), second month (second row), third month (third row) and fourth month (fourth row). Arrows represent the velocity and the contour lines represent the sea surface height, with a contour line of 0.01 m (1 cm).

There is only one occurrence of the non-cycling solution being significantly more accurate than the cycling solutions, at location 1 and network 1 for all variables (SSH and both velocity components). Location 2 all three solutions are comparably accurate across the networks and model variables. However, at locations 3–5, for all networks and all variables, the two cycling solutions clearly outperform the non-cycling solution. The latter sometimes exceeds the data error by more than 2 standard deviations; e.g. v-velocity at locations 3 and 5 for networks 1 and 3. The rms values at these three locations (3–5) simply reflect what was already shown above in Fig. 2 for network 1, i.e. the discrepancy between the non-cycling solution and the reference solutions persists in the area where the LCE sheds from the LC, an area where these locations are contained.

The difficulties of the non-cycling solution in this area are not limited to network 1 which contains the least number of data, neither are they point wise as Figs. 3 and 4 suggest. Similar to Fig. 2, Fig. 5 shows the discrepancies to the reference solution are computed for the non-cycling and the cycling solutions at the

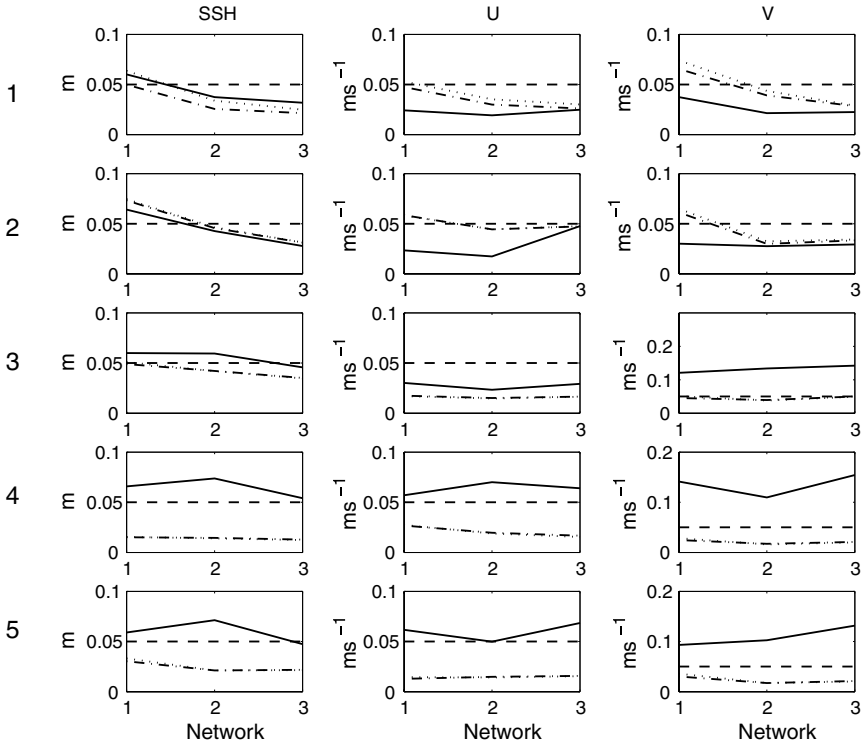

Fig. 3. Comparison of rms misfits to the data for the non-cycling solution (solid line), the cycling solution with 3 outer loops in each cycle (dash-dotted line) and the cycling solution with 3 outer loops only in the first cycle and one outer loop in subsequent cycles (dotted line). The rms is computed over the entire assimilation time interval. The dashed line is the data error. Numbers to the left (1–5) indicate the diagnostic location.

end of the third month for all networks, including an experiment where only SSH data from network 3 is assimilated. This figure shows that the errors in the non-cycling solution are consistent for all networks. One might have expected increasing errors as the data coverage decreases from network 3 to network 1. Such is the case for the cycling solution and not for the non-cycling. One can hypothesize that the errors in the non-cycling solution are dominated by systematic errors in the TLM. Fortunately, the cycling solution is able to fit the data properly because the TLM errors are inhibited by a limited assimilation interval and a more accurate background provided by the previous cycle nonlinear forecast.

A final experiment is carried out with the assimilation of only SSH data from network 3. As in NG06 for the non-cycling solution, we test the ability of the cycling algorithm to infer the velocity field through the model dynamics by assimilating only SSH measurements. The non-cycling and the cycling solutions accuracy is evaluated through the rms error to the reference solution at the selected locations. Results in Table 1 show that the non-cycling solution is able to accurately fit the SSH data at all locations (except for location 4 where the rms exceeds 2 standard deviations) and the velocity only at the first two locations. At the remaining and critical locations 3–5, the non-cycling solution miserably fails to correct the velocity components with rms values sometimes exceeding 5–10 standard deviations. In contrast, the cycling solution accurately fits the SSH data and the inferred velocity accurately matches the non-assimilated velocity data within expected errors.

4. Accuracy and cost comparison

The computational cost of the representer method applied to nonlinear models has understandably deterred many potential users. The algorithm requires a sequence of linear iterates of the nonlinear Euler–Lagrange system. Each linear iterate can be solved using the representer method, the solution of which becomes the

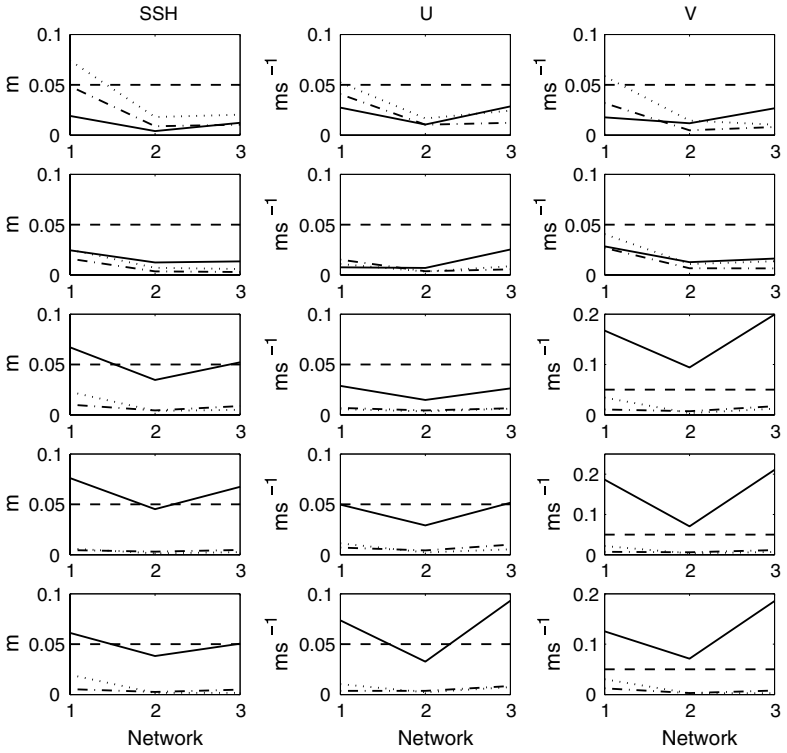

Fig. 4. Same as Fig. 3, except the rms is computed over the last 2 months of the assimilation.

background for linearization in the next iterate until formal convergence. These linear iterations, also called outer loops, are what exacerbates the computational cost of the algorithm; each iterate is a multiplicative factor to the cost of a linear assimilation in a given time window. Because the outer loops are an attempt to solve the nonlinear Euler–Lagrange system, an inaccurate linearization, which could also result from the accumulation of linearization errors over time, would cause the algorithm to require more outer loops. The non-cycling solution in this study used 6 outer loops. Two cycling solutions were computed using 4 cycles of 1-month each. The first solution used 3 outer loops in each cycle and the second solution used 3 outer loops only for the first cycle and one outer loop for the remaining cycles. If C denotes the computational cost for 1 outer loop for the entire four-month assimilation window, then the cost of the non-cycling solution is 6C. The cost of one outer loop in one cycle is C/4 at worst. This assumes that the minimization of the cost function converges in the same number of conjugate gradient iterations, which is a very conservative assumption since the number of data is significantly lower within the cycle and the minimization problem is eventually better conditioned than in the non-cycling solution. Even in this worse case scenario, the first cycling solution costs only 3C, and is by far more accurate than the non-cycling solution as we have seen above. The second cycling solution was shown to be equally accurate as the first, yet it costs only 1.5C i.e. half the cost of the first cycling solution. This represents a huge gain in computational cost, added to an already improved solution. Although the gains here are specific to the chosen application, the general conclusion is that the cycling solution will be more accurate and computationally less expensive in the case where the TLM accuracy is limited.

5. Weaker constraints

Compared to the non-cycling, the cycling approach introduces additional controls that are not present in the original assimilation problem. These additional controls are the initial conditions for each cycle after the first. Thus, even though the same error covariances are used in both solutions, the cycling assimilation problem has ‘weaker’ constraints than the non-cycling. To demonstrate this, comparisons are made between the

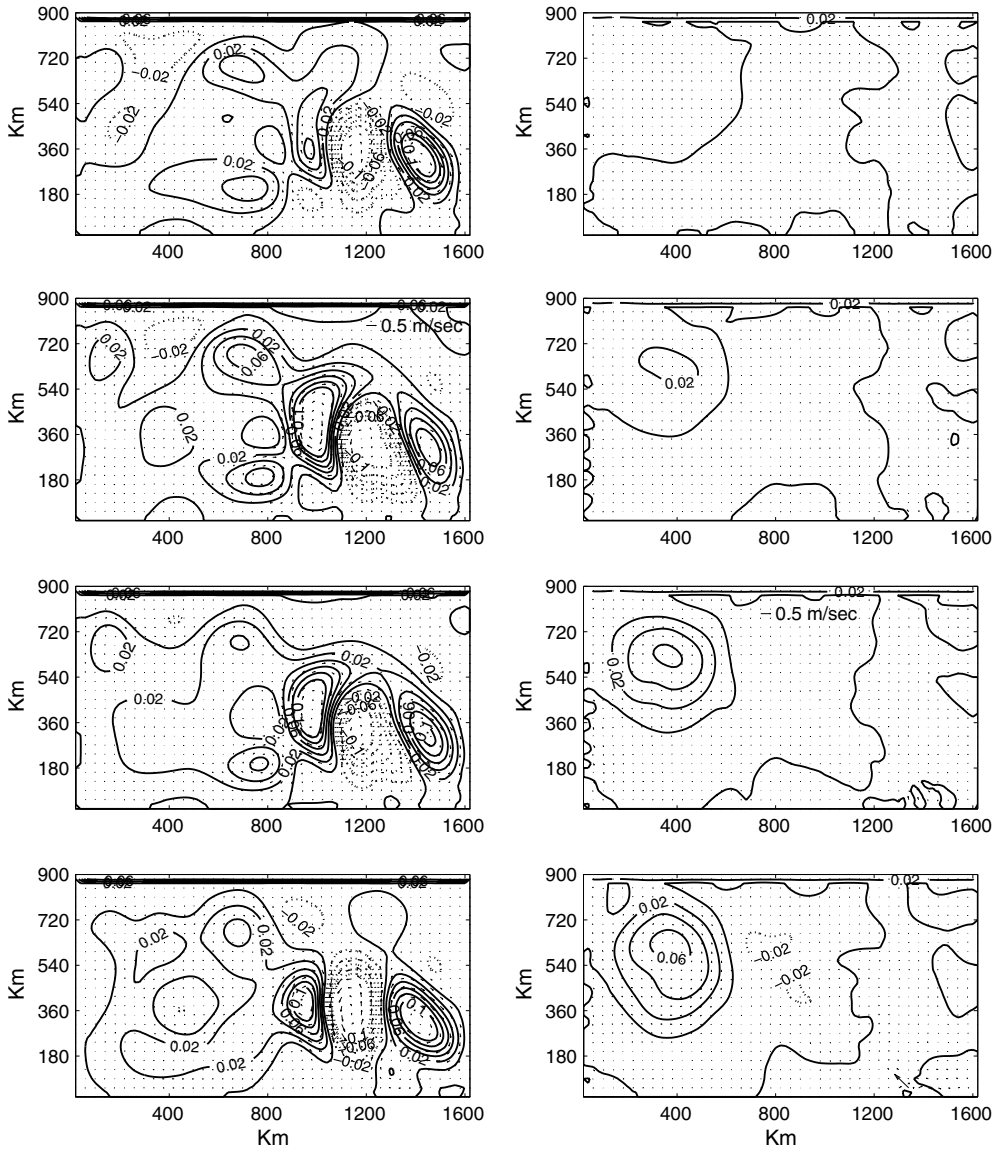

Fig. 5. Comparison of the difference between the reference and the assimilated solution using the non-cycling (left column) and the cycling (right column) algorithms at the end of the third month for networks 3 (first row), 2 (second row), 1 (third row) and network 1 with only SSH data assimilated (last row).

Table 1
RMS of non-cycling and cycling solutions at the five diagnostic locations for network 3 assimilating only SSH data

Location	SSH		U		V	
	Non-cycling	Cycling	Non-cycling	Cycling	Non-cycling	Cycling
1	0.0160	0.0619	0.0871	0.0353	0.0307	0.0658
2	0.0253	0.0330	0.0521	0.0211	0.0416	0.0670
3	0.0679	0.0173	0.1073	0.0164	0.4772	0.0170
4	0.1354	0.0060	0.1795	0.0094	0.3671	0.0212
5	0.0963	0.0075	0.2926	0.0156	0.5844	0.0244

cycling approach using the original covariance and the non-cycling approach with weaker constraints. The non-cycling constraint is weakened by multiplying the model error variance by 10.

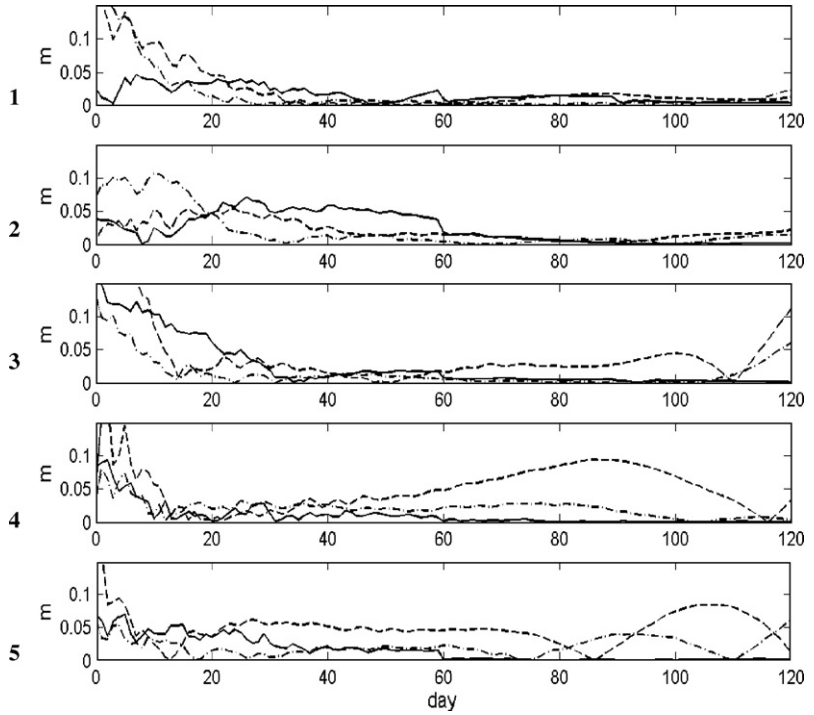

Fig. 6. Absolute difference between the reference solution and the assimilated solutions by cycling (solid line), non-cycling (dashed line) and non-cycling with larger (factor of 10) variance (dash-dotted line) for SSH at the five diagnostic stations.

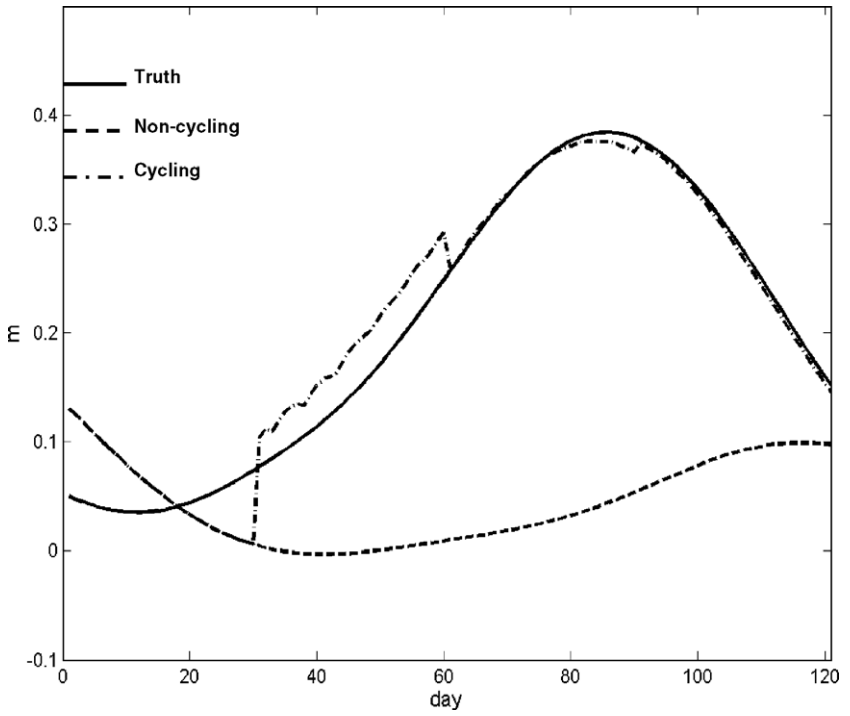

Fig. 7. Timeseries of the SSH at location 5 for the truth or reference solution (solid line), the background in the non-cycling solution (dashed line) and the background in the cycling solution (dash-dotted line).

The absolute difference between the reference solution and the assimilated solutions obtained by cycling, non-cycling and non-cycling with ten times the original model error variance is shown in Fig. 6 for SSH at the five diagnostic stations. It can be seen from Fig. 6 that the non-cycling solution accuracy improves when the model error variance is increased, i.e. the weaker constraints enable the non-cycling assimilation to better fit the data to a level of accuracy that compares well with the cycling solution even though the non-cycling solution deviates from the reference solution. Finally, the weaker non-cycling solution appears to be smoother in time than the cycling solution because of the time correlation function that applies to the entire assimilation window. The cycling solution error, however, generally decreases in time as the assimilation moves from one cycle to the next. This is a direct benefit from using the previous cycle forecast as background for the current cycle assimilation.

It is arguable that the accuracy of the cycled solution is solely due to weaker constraints as seen above. Fig. 7 shows the immediate improvements that the cycling solution applies to the background of the subsequent cycle. This improvement reduces the magnitude of the innovations and thus enables the tangent linear approximation and the assimilation to be more accurate. In contrast, the non-cycling solution has to overcome larger innovations to fit the data, which will require more inner and outer iterations for the process to converge. Starting from the same background as the non-cycling solution, it is shown in Fig. 7 that although the assimilation has not been carried out yet in the second cycle, the background is already more accurate than the corresponding section from the original background. Together with a shorter time window, this improved background will inhibit error growth in the TLM and yield an accurate assimilated solution.

6. Discussion and conclusion

In an earlier study the cycling representer algorithm was applied for the first time to the highly nonlinear yet low dimension Lorenz attractor model (NL07). Here it is applied in a twin model assimilation experiments to a 1.5 nonlinear reduced gravity ocean model for an idealized LCE shedding from the LC in the Gulf of Mexico. The model is chosen as an intermediate step between the low dimension Lorenz attractor and a full multi-layer nonlinear ocean model. Non-cycling and cycling assimilation solutions were compared in terms of accuracy (lower rms error to the reference solution) and computational cost. It was found that for the same data sets the cycling solution is by far more accurate than the non-cycling solution, and it is also computationally less expensive. The deficiencies of the non-cycling solution may be associated with error growth in the TLM over a long period of time. Taking the difference between the assimilated and reference solutions showed that errors in the non-cycling solution were largest in the region with strong nonlinearities. The cycling solution overcomes these deficiencies by using a sequence of shorter assimilation windows over which the growth of errors in the TLM is inhibited, i.e. the validity of the tangent linear approximation is assured. For the example at hand it was shown that the cost of the non-cycling solution can be reduced by 50% (with half the number of outer loops in each cycle) to 75% (with half the number of outer loops only in the first cycle and 1 outer loop in the remaining cycles). The need of outer loops should be guided by the desired accuracy in the assimilation. A situation may arise where more outer loops are needed in the current cycle whereas the previous cycle was accurate with only one outer loop. This may be the case for a strongly nonlinear response due to stronger than usual atmospheric forcing in coastal oceans with complex bathymetry. The strength of the cycling algorithm lies in the limitation of the error growth in the TLM due to a reduced time interval and a faster inversion for the minimization of the cost function thanks to a smaller number of observations and a better conditioning of the assimilation problem. Also, the nonlinear background, which is the nonlinear forecast from the previous cycle's final condition, eliminates the linearization errors that would have been introduced by using a linearized solution as a first guess in the assimilation.

The cycling approach introduces additional controls that are not present in the original assimilation problem. These additional controls are the initial conditions for each cycle after the first. They render the cycling assimilation 'weaker' than the non-cycling. As an attempt to weaken the non-cycling assimilation, an experiment was carried out with the model error variance multiplied by 10. This weaker non-cycling solution was found to be more accurate than the non-cycling solution of the original assimilation problem and roughly as accurate as the cycling solution (albeit for some deviations from the reference solution). Nonetheless, the accu-

racy of the cycling solution was superior and more consistent, albeit for temporal discontinuities at the beginning of each new cycle.

One should be very careful about further weakening the constraints in the non-cycling assimilation in order to obtain a more accurate solution; a weaker model will yield a better fit to the data. However, the assimilation will generate larger residuals that may dominate the term balances in the model equations. Thus, for a given assimilation problem, in which errors are prescribed adequately, the cycling approach will yield the most accurate solution at a lower cost.

Acknowledgements

This work was sponsored by the Office of Naval Research (Program Element number 0601153N) as part of the project 'Shelf to Slope Energetics and Exchange Dynamics'. This paper is NRL paper contribution number NRL/JA/7320-07-7103. The authors acknowledge the constructive comments of the anonymous reviewers that helped improve the manuscript.

References

Bennett, A.F., 2002. Inverse Modeling of the Ocean and Atmosphere. Cambridge University Press.

Bennett, A.F., Chua, B.S., Leslie, L.M., 1996. Generalized inversion of a global numerical weather prediction model. Meteorology and Atmosphere Physics 60, 165–178.

Bennett, A.F., Chua, B.S., Harrison, D.Ed., McPhaden, M.J., 1998. Generalized inversion of tropical atmosphere-ocean data and a coupled model of the tropical pacific. Journal of Climate 11, 1768–1792.

Bennett, A.F., Chua, B.S., Harrison, D.Ed., McPhaden, M.J., 2000. Generalized inversion of tropical atmosphere-ocean data and a coupled model of the tropical pacific. II. The 1995-96 La Nina and 1997-1998 El Nino. Journal of Climate 11, 2770–2785.

Bennett, A.F., Chua, B.S., Ngodock, H.-E., Harrison, D.E., McPhaden, M.J., 2006. Generalized inversion of the gent-cane model of the tropical pacific with tropical atmosphere-ocean (TAO) data. Journal of Marine Research 64 (1), 1–42.

Jacobs, A.G., Ngodock, H.E., 2003. The maintenance of conservative physical laws within data assimilation systems. Monthly Weather Review 131 (11), 2595–2607.

Ngodock, H.E., Chua, B.S., Bennett, A.F., 2000. Generalized inversion of a reduced gravity primitive equation ocean model and tropical atmosphere ocean data. Monthly Weather Review 128, 1757–1777.

Ngodock, H.E., Jacobs, G.A., Chen, M., 2006. The representer method, the ensemble Kalman filter and the ensemble Kalman smoother: a comparison study using a nonlinear reduced gravity ocean model. Ocean Modelling 12, 378–400.

Ngodock, H.E., Smith, S.R., Jacobs, G.A., 2007. Cycling the representer algorithm for variational data assimilation with the Lorenz attractor. Monthly Weather Review 135 (2), 373–386.

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

Xu, L., Daley, R., 2002. Data assimilation with a barotropically unstable shallow water system using representer algorithms. Tellus 54A, 125–137.