Finite Difference of Adjoint or Adjoint of Finite Difference?

ZIV SIRKES

Institute of Marine Sciences, University of Southern Mississippi, Stennis Space Center, Mississipi

ELI TZIPERMAN

Environmental Sciences, Weizmann Institute of Science, Rehovot, Israel

1 October 1996 and 11 April 1997

ABSTRACT

Adjoint models are used for atmospheric and oceanic sensitivity studies in order to efficiently evaluate the sensitivity of a cost function (e.g., the temperature or pressure at some target time t_f, averaged over some region of interest) with respect to the three-dimensional model initial conditions. The time-dependent sensitivity, that is the sensitivity to initial conditions as function of the initial time t_i, may be obtained directly and most efficiently from the adjoint model solution. There are two approaches to formulating an adjoint of a given model. In the first (''finite difference of adjoint''), one derives the continuous adjoint equations from the linearized continuous forward model equations and then formulates the finite-difference implementation of the continuous adjoint equations. In the second (''adjoint of finite difference''), one derives the finite-difference adjoint equations directly from the finite difference of the forward model. It is shown here that the time-dependent sensitivity obtained by using the second approach may result in a very strong nonphysical behavior such as a large-amplitude two-time-step leapfrog computational mode, which may prevent the solution from being used for time-dependent sensitivity studies. This is an especially relevant problem now, as this second approach is the one used by automatic adjoint compilers that are becoming widely used. The two approaches are analyzed in detail using both a simple model and the adjoint of a primitive equations ocean general circulation model. It is emphasized that both approaches are valid as long as they are used for obtaining the gradient or sensitivity at a single time, as needed in data assimilation, for example. Criteria are presented for the choice of the appropriate adjoint formulation for a given problem.

1. Introduction

In atmospheric and oceanic sensitivity studies, one is often interested in finding how a change to the model initial conditions at some initial time t_i will affect some measure of the model solution (i.e., a cost function), such as an area-averaged temperature or pressure, at a later time t_f. For a given initial time t_i, this requires running the model many times and at each run varying the model initial conditions of a single model variable at a single grid point and observing its effect on the cost function at t_f. For a high-resolution model this amounts to running the model tens to hundreds of thousands of times, clearly not a feasible approach. The computational cost is yet higher if one wishes to study the time-dependent sensitivity, that is, the sensitivity to initial conditions as a function of the initial time t_i. Instead, the sensitivity as well as the time-dependent sensitivity could be obtained at a computational price equivalent to running the model only once, using an adjoint model (e.g., Hall and Cacuci 1983; Errico and Vukicevic 1992). Adjoint models are also used for four-dimensional data assimilation in general circulation models (e.g., Bennett and McIntosh 1982; Le Dimet and Talagrand 1986; Thacker and Long 1988; Tziperman and Thacker 1989; Schiller and Willebrand 1995; Harms et al. 1992).

There are two approaches that are commonly used to formulate an adjoint to a given model. The first (''finite difference of adjoint'') involves writing the adjoint equations of the linearized continuous model equations and then determining an appropriate finite-difference formulation. The second approach (''adjoint of finite difference'') derives the adjoint equations directly from the finite-difference equations of the forward model.

While using an adjoint model of a global primitive equations general circulation model derived using the adjoint of finite-difference approach (Sirkes et al. 1996; the adjoint model is based on that of Long et al. 1989), we noticed that although the gradient at time t_i computed by the adjoint model was accurate, the adjoint variables strongly oscillate between time steps (Fig. 1). This very strong numerical mode, apart from being surprising and interesting, has important implications to time-dependent sensitivity studies as defined above. Each adjoint model variable (e.g.,

Corresponding author address: Dr. Eli Tziperman, Environmental Sciences and Energy Research, Weizmann Institute of Science, Rehovot 76100, Israel.
E-mail: eli@beach.weizmann.ac.il

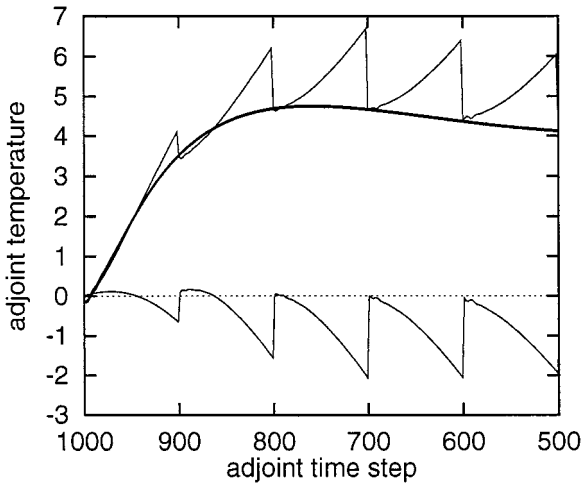

FIG. 1. Time series of the adjoint of the temperature, λ^T, at a depth of 25 m (model level 1) and horizontal location 24.45°N, 288.75°E from the adjoint of a primitive equations ocean GCM. The horizontal axis indicate time-step number, going backward in time from the $n = 1000$ step of the forward model to step $n = 500$ of the forward model. The two thin lines show the adjoint solution for the adjoint of finite-difference approach, which results in a strong computational mode. The upper thin line is the adjoint variable at odd time steps and the lower thin line is the same variable at even time steps. Note that the adjoint solution oscillates between these two curves every time step, indicating a large-amplitude two-time-step oscillation. The thick line shows the same adjoint variable for the finite difference of adjoint approach, where no computational mode is present. The dotted line indicates the zero value of the vertical axis.

the adjoint of the temperature at a given model grid point) represents the sensitivity of the cost function to the initial value of the corresponding physical variable at t_i. The adjoint solution seen in Fig. 1, however, reflects a numerical mode of the leapfrog scheme, rather than a physical time-dependent sensitivity behavior.

The purpose of this note is to show that deriving the adjoint using the adjoint of finite-difference approach may result, in certain cases, in the development of strong numerical computational modes, of which the leapfrog mode of Fig. 1 is a specific example. We emphasize here, however, that both approaches to formulating an adjoint model are perfectly valid as long as they are used for obtaining the gradient or sensitivity at a single time, as needed in data assimilation, for example. We then propose criteria for the choice of the appropriate adjoint formulation for a given problem, as well as ways of dealing with the strong nonphysical modes shown above. In the following sections, the two adjoint formulations are presented and solved for a simple advection–diffusion model equation, comparing the results to those of a global primitive equations adjoint model (section 2). We conclude in section 3.

2. The two adjoint formulations in a simple advection–diffusion model

Consider a simple one-dimensional advection–diffusion equation for a passive scalar $S(x, t)$:

$$S_t + cS_x - KS_{xx} = 0, \tag{1}$$

where c is the advection velocity and K is the diffusion coefficient. Consider the finite difference of adjoint approach for deriving the adjoint equations. Suppose that we want to study the sensitivity to initial conditions of the difference between the model solution and the data $S^{\text{data}}(x, T)$ at a time $t_f = T$. The cost function is then of the form

$$J = \frac{1}{2}\int_{x=x_0}^{L} W(x)[S(x, T) - S^{\text{data}}(x, T)]^2\, dx,$$

where $W(x)$ is some specified weight factor. The adjoint equations are formed by adding the model equations as ''hard'' constraints to the cost function (the ''soft'' constraint formulation is very similar; see Bennett and McIntosh 1982) and forming the Lagrange function

$$L = J + \int_{x=x_0}^{L}\int_{t=t_0}^{T} \lambda(x, t)\,[S_t(x, t) + cS_x(x, t) - KS_{xx}(x, t)]\, dt\, dx,$$

where $\lambda(x, t)$ is the adjoint variable (Lagrange multiplier) of the passive scalar concentration $S(x, t)$. At the minimum of the cost function, the Lagrange function is at a stationary point with respect to the model variables, and this condition is used to derive the continuous adjoint initial conditions and equations;

$$\frac{\delta L}{\delta S(x, T)} = 0 = \lambda(x, T) + W(x)[S(x, T) - S^{\text{data}}(x, T)] \tag{2}$$

and

$$\frac{\delta L}{\delta S(x, t)} = 0 = \lambda_{-t}(x, t) - c\lambda_x(x, t) - K\lambda_{xx}(x, t). \tag{3}$$

Noting that these adjoint equations are integrated backward in time; they may be discretized using centered difference in space and a leapfrog scheme in time (Haltiner and Williams 1980). The adjoint integration is initialized, at $t = T$, using an Euler-forward time step,

$$\lambda_k^N = -W_k(S_k^N - S_k^{\text{data}})$$

$$\lambda_k^{N-1} = \lambda_k^N + \frac{c\Delta t}{2\Delta x}(\lambda_{k+1}^N - \lambda_{k-1}^N) + \frac{K\Delta t}{\Delta x^2}(\lambda_{k+1}^N - 2\lambda_k^N + \lambda_{k-1}^N) \quad (4)$$

and the following time steps are according to the leapfrog scheme

$$\lambda_k^n = \lambda_k^{n+2} + \frac{c\Delta t}{\Delta x}(\lambda_{k+1}^{n+1} - \lambda_{k-1}^{n+1}) + \frac{2K\Delta t}{\Delta x^2}(\lambda_{k+1}^{n+2} - 2\lambda_k^{n+2} + \lambda_{k-1}^{n+2}), \quad (5)$$

where

$$x = k\Delta x, \qquad k = 1, 2, \ldots, M \quad \text{and} \quad t = n\Delta t,$$

$$n = N, N-1, \ldots, 2, 1, 0 \ (N\Delta t = T).$$

The solution for the adjoint variable $\lambda(x, t)$ represents the sensitivity of the cost function to the temperature at location and time (x, t) (Thacker 1991, 1992). While the cost function is evaluated in practice using the finite-difference forward model, this first approach for deriving the adjoint solution is based on the continuous adjoint equation. The resulting sensitivity is therefore accurate only to the accuracy of the finite-difference approximations, within $O(\Delta x, \Delta t)$. [This reduced accuracy may, in principle, affect the convergence of optimizations that use an adjoint code to calculate the gradient of a cost function. However, Schiller and Willebrand (1995) used an approximate adjoint model that is even less accurate than the finite difference of adjoint formulation and still obtained a successful convergence of their optimizations.]

Next, let us consider the ''adjoint of finite difference'' approach. In the finite-difference forward model, a leapfrog scheme is used at most time steps:

$$S_k^n = S_k^{n-2} - \frac{c\Delta t}{\Delta x}(S_{k+1}^{n-1} - S_{k-1}^{n-1}) + \frac{2K\Delta t}{\Delta x^2}(S_{k+1}^{n-2} - 2S_k^{n-2} + S_{k-1}^{n-2}). \quad (6)$$

An Euler-forward step is used both for initialization ($n = 1$) and (as is common in ocean GCMs and, perhaps less so, in atmospheric GCMs) every several leapfrog steps during the run, in order to suppress the leapfrog computational mode. Denoting mixing time steps with an index m, the tracer at these time steps, S_k^m, is calculated according to

$$S_k^m = S_k^{m-1} - \frac{c\Delta t}{2\Delta x}(S_{k+1}^{m-1} - S_{k-1}^{m-1}) + \frac{K\Delta t}{\Delta x^2}(S_{k+1}^{m-1} - 2S_k^{m-1} + S_{k-1}^{m-1}). \quad (7)$$

The finite-difference Lagrange cost function has the form

$$L = \frac{1}{2}\sum_{k=1}^{M} W_k(S_k^N - S_k^{\text{data}})^2 + \sum_{k=1}^{M}\sum_{n=0}^{N} \lambda_k^n(S_k^n - \cdots), \quad (8)$$

where the three dots denote the right-hand side of (6) and (7) inserted at the appropriate time steps. Setting the derivatives of the Lagrange function with respect to each of the variables S_k^n to zero, we obtain equations for initializing the adjoint integration:

$$\lambda_k^N = -W_k(S_k^N - S_k^{\text{data}});$$

$$\lambda_k^{N-1} = \frac{c\Delta t}{\Delta x}(\lambda_{k+1}^N - \lambda_{k-1}^N), \quad (9)$$

while the backward time stepping scheme at nonmixing time steps is as in (5). Note that the initialization (9) for λ_k^{N-1} is equivalent to the regular leapfrog step (5) evaluated at $n = N - 1$, with a dummy adjoint variable λ_k^{N+1} [which does not appear in the Lagrange function (8)] that is set to zero ($\lambda_k^{N+1} = 0$).

The adjoint of the mixing time steps in the forward model (7), derived by differentiating (8) with respect to the tracer S during the mixing steps, spreads over two time steps in the adjoint integration

$$\lambda_k^{m-1} = \lambda_k^{m+1} + \frac{c\Delta t}{2\Delta x}(\lambda_{k+1}^m - \lambda_{k-1}^m) + \frac{2K\Delta t}{\Delta x^2}(\lambda_{k+1}^{m+1} - 2\lambda_k^{m+1} + \lambda_{k-1}^{m+1}) + \lambda_k^m + \frac{K\Delta t}{\Delta x^2}(\lambda_{k+1}^m - 2\lambda_k^m + \lambda_{k-1}^m) \quad (10)$$

and

$$\lambda_k^{m-2} = \frac{c\Delta t}{\Delta x}(\lambda_{k+1}^{m-1} - \lambda_{k-1}^{m-1}). \quad (11)$$

In this second derivation of the adjoint solution, both the cost and the adjoint model are obtained using the same consistent set of finite differences, resulting in an accuracy of the sensitivity/gradient that is within the computer roundoff error. We next analyze the two approaches by solving their continuous and finite-difference equations analytically.

a. Analytic solutions of the continuous and finite-difference schemes

The analytic solution to the continuous adjoint equation 3, assuming that the adjoint initial conditions (2) reduce to a single wave, $\lambda(x, T) = Ae^{i\mu x}$, is

$$\lambda(x, t) = Ae^{i\mu(x-ct)+K\mu^2 t}, \qquad t = T, \ldots, 0, \quad (12)$$

where A is the amplitude and μ is the wavenumber. The finite difference of the adjoint to the advection–diffusion (5) is found assuming a solution of the form $\lambda_k^n =$

$B^{n\Delta t}e^{i\mu k\Delta x}$; $n = N, \ldots, 0$ (Haltiner and Williams 1980). The finite-difference solution then takes the form

$$\lambda_k^n = [Me^{-in\beta} + (-1)^n Ee^{in\beta}](1 - \delta)^{-n/2}e^{i\mu k\Delta x}, \quad (13)$$

where

$$\sigma = \frac{c\Delta t}{\Delta x}\sin(\mu\Delta x), \qquad \delta = \frac{4K\Delta t}{\Delta x^2}[1 - \cos(\mu\Delta x)],$$

$$\beta = \arcsin\left[\frac{\sigma}{(1 - \delta)^{1/2}}\right].$$

The constants M and E are determined by the initial conditions. The term whose amplitude is E oscillates with a two-time-step frequency and is the computational mode created by the leapfrog scheme (Haltiner and Williams 1980). Assuming, similarly to the continuous case, that the finite-difference adjoint initial conditions reduce to a single wave component,

$$\lambda_k^N = -W_k(S_k^N - S_k^{\text{data}}) = Ae^{i\mu k\Delta x} \quad (14)$$

yields a solution for one of the constants in the general solution for the adjoint equation

$$M = A(1 - \delta)^{N/2}e^{iN\beta} + (-1)^{N+1}Ee^{2iN\beta}. \quad (15)$$

The constant E depends on the exact adjoint initialization procedure, which in turn depends on the adjoint formulation approach used, as follows.

In the finite difference of adjoint, using the Euler-forward scheme (4) for the first time step, we obtain

$$E = A\frac{[(1 - \sigma^2 - \delta)^{1/2} - (1 - \delta/2)](-1)^{-N}(1 - \delta)^{N/2}e^{-iN\beta}}{2(1 - \sigma^2 - \delta)^{1/2}}. \quad (16)$$

For $(\Delta x, \Delta t) \to 0$, [while keeping $c\Delta t/\Delta x = O(1)$ due to the Courant–Friedrichs–Levy (CFL) criterion], the amplitude of the physical mode in the general solution for the finite-difference adjoint equation (13) converges to A, while the amplitude of the computational mode converges to zero.

In the adjoint of finite difference, one initializes the adjoint model with a leapfrog step and sets $\lambda_k^{N+1} = 0$ (9), leading to

$$E = A\frac{[(1 - \sigma^2 - \delta)^{1/2} - i\sigma](-1)^{-N}(1 - \delta)^{N/2}e^{-iN\beta}}{2(1 - \sigma^2 - \delta)^{1/2}}. \quad (17)$$

The amplitudes of both terms in the general finite-difference solution (13) now remain finite, of order A, for $(\Delta x, \Delta t) \to 0$. This implies that the numerical mode is as strong as the physical mode and the total solution is very different from that of the continuous adjoint equation (12). Note that the two-step numerical mode is excited in the adjoint of finite difference formulation both by the initialization procedure and by the adjoint of the forward model mixing steps.

b. Numerical examples

In order to compare the results of the two adjoint formulations, their equations were solved over a domain divided into 100 grid points, using cyclic boundary conditions in x and setting the adjoint initial conditions to a sine wave (14). The results are plotted from the last time step ($n = 125$) to $n = 0$. The adjoint solution in the finite difference of adjoint approach (Fig. 2, thick line) is smooth and, for this simple model, is indistinguishable from the analytic solution (12) of the continuous adjoint equation. In the adjoint of finite difference approach (Fig. 2, thin line) the adjoint solution contains a large-amplitude numerical mode, as anticipated above and seen also for the adjoint of the primitive equation ocean GCM (Fig. 1). The solutions for the two adjoint formulations coincide, within $O(\Delta x, \Delta t)$, every 50 time steps, between the two parts (10) and (11) of the adjoint of the mixing time step (circles in Fig. 2), as well as at the final step of the adjoint model, when the value of

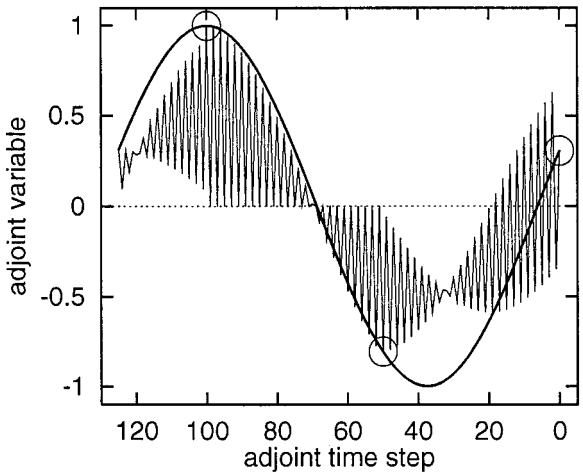

FIG. 2. Time series of the adjoint variable λ_k^n, from the adjoint of the simple advection–diffusion model, at $k = 20$. The thick line shows the finite difference of adjoint formulation, with no computational mode. The thin line shows the adjoint of finite-difference formulation that results in a strong leapfrog computational mode. Note that the two formulations converge at mixing time steps of the forward model, every 50 time steps (at $n = 100, 50$), as well as at the final time step ($n = 0$) that is used to calculate the gradient of the cost function. These convergence points are denoted by circles in the figure.

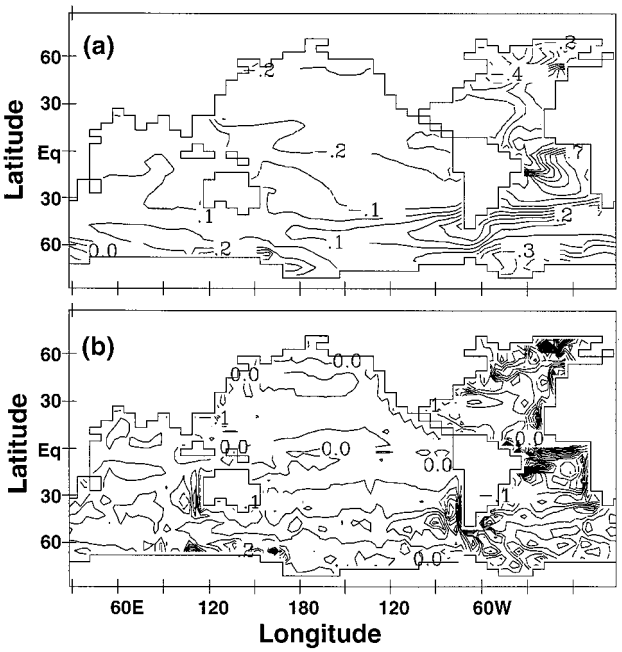

FIG. 3. A contour plot of the gradient field $\lambda^T(x, y, t = 200 \text{ yr})$ in the North Atlantic Ocean, from the global adjoint model of Sirkes et al. (1996). (a) The exact gradient obtained from $\lambda^T(x, y, t = 200 \text{ yr})$. Contours are scaled by 0.663E+02. (b) The same adjoint variable one time step before the time (a) was drawn, includes a large component of computational mode. Contours scaled by 0.126E+06.

the adjoint variable is used to calculate the sensitivity/gradient of the cost function. This convergence of the two adjoint formulations at mixing time steps can also be shown from the analytic solutions of the previous section, by a somewhat tedious manipulation of the adjoint solution (17) and the mixing-step equation (10).

The above simple model results of the leapfrog computational mode and the convergence of the two approaches at mixing steps occur also in the complex adjoint GCM of Sirkes et al. (1996), as seen in Fig. 1. Furthermore, Fig. 3 shows that the entire three-dimensional adjoint temperature–salinity fields in that model go through a large-amplitude two-step oscillation before producing a perfectly accurate gradient at the last time step of the adjoint model integration. Both the amplitude (note different scaling of contours in Fig. 3a,b) and the spatial structure of the adjoint solution are significantly different between consecutive time steps. The spatial structure at time steps that include a contribution from the numerical mode (Fig. 3b) is much noisier than the physical sensitivity (Fig. 3a).

3. Conclusions

We have considered the two possible formulations of an adjoint model for a given forward model. The first formulation, finite difference of adjoint, results in a sensitivity that is accurate to only within $O(\Delta x, \Delta t)$, but the adjoint solution does not contain any numerical modes. The second formulation, adjoint of finite difference, calculates a gradient at the final time step of the adjoint integration that is accurate within computer roundoff error, but the adjoint solution may include strong computational modes such as the leapfrog computational mode analyzed above. Note that the adjoint leapfrog computational mode seen in Figs. 1 and 2 makes sense from a numerical point of view, as it represents the actual model sensitivity to initial conditions. Because of the leapfrog scheme, the forward model solution at a given time step is sensitive in a different way to previous model states at odd and even time steps, as reflected in the adjoint computational mode. But for time-dependent sensitivity studies, we are interested in the sensitivity of the physical system underlying the model numerics, and that, of course, should not contain any two-step oscillations.

It should be clear that the leapfrog computational mode analyzed here is only an example of a more general problem. The adjoint of finite-difference formulation does not take care explicitly of the numerical stability of the adjoint solution. Thus other schemes than leapfrog schemes used in the forward model may result in other numerical artifacts in the adjoint solution. While we have offered a solution for the specific leapfrog scheme problem, of extracting the physical time-dependent sensitivity at mixing time steps, this solution may clearly not be relevant to other adjoint/numerical problems. The more general take-home message of this note, therefore, is that in order to avoid numerical artifacts in the adjoint solution, one needs to follow one of two alternative routes. The first is to use the finite difference of adjoint formulation, where the numerical stability of the adjoint solution is explicitly taken care of. The second alternative is to develop numerical schemes that are stable and whose adjoint of finite difference is numerically stable as well. This should eliminate the problems reported here, and we see this as an important future research direction. Finally, when only the gradient at the initial time is needed, one may certainly use the adjoint of finite difference formulation.

The results of this work are especially relevant at this time due to the ''automatic adjoint compilers'' (e.g., Bischof et al. 1992) that are becoming widely used now in an effort to reduce the time-consuming adjoint code development. These compilers derive an adjoint computer code directly from the computer code (e.g., a Fortran code) of the forward model. They follow the adjoint of finite difference approach and may, therefore, suffer all of the potential difficulties discussed above.

We wish to end by emphasizing that both approaches, the adjoint of finite difference and finite difference of adjoint are perfectly valid as long as they are used carefully and for obtaining the gradient or sensitivity at a single time, as needed in data assimilation, for example. We hope that the present note will help future studies in choosing the appropriate approach for a given problem that requires the use of an adjoint model.

Acknowledgments. We thank Long, Huang, and Thacker for generously providing us their adjoint of the GFDL model. We gratefully acknowledge the most useful comments on an earlier draft from Jeff Anderson, Carlisle Thacker, and three anonymous reviewers.

REFERENCES

Bennett, A. F., and P. C. McIntosh, 1982: Open ocean modeling as an inverse problem: Tidal theory. *J. Phys. Oceanogr.,* **12,** 1004–1018.

Bischof, C., A. Carle, G. Corliss, A. Griewank, and P. Hovland, 1991: ADIFOR—Generating derivative codes from Fortran programs. *Sci. Programm.,* **1,** 1–29.

Errico, R. M., and T. Vukicevic, 1992: Sensitivity analysis using an adjoint of the PSU-NCAR mesoscale model. *Mon. Wea. Rev.,* **120,** 1644–1660.

Hall, M. C. G., and D. G. Cacuci, 1983: Physical interpretation of the adjoint functions for sensitivity analysis of atmospheric models. *J. Atmos. Sci.,* **40,** 2537–2546.

Haltiner, G. J., and R. T. Williams, 1980: *Numerical Prediction and Dynamical Meteorology*. John Wiley and Sons, 470 pp.

Harms, D. E., S. Raman, and R. V. Madala 1992: An examination of four-dimensional data-assimilation techniques for numerical weather prediction *Bull. Amer. Meteor. Soc.,* **73,** 425–440.

Le Dimet, F., and O. Talagrand, 1986: Variational algorithm for analysis and assimilation of meteorological observations: Theoretical aspects. *Tellus,* **38A,** 97–110.

Long, R. B., S.-M. Huang, and W. C. Thacker 1989: The finite-difference equations defining the GFDL-GCM and its adjoint. Atlantic Oceanographic and Meteorological Laboratory Rep., Miami, FL, 28 pp. [Available from AOML, Miami, FL 33149.]

Schiller, A., and J. Willebrand, 1995: A technique for the determination of surface heat and freshwater fluxes from hydrographic observations, using an approximate adjoint ocean circulation model. *J. Mar. Res.,* **53,** 433–451.

Sirkes, Z., E. Tziperman, and C. W. Thacker, 1996: Combining data and a global primitive equation ocean general circulation model using the adjoint method. *Modern Approaches to Data Assimilation in Ocean Modeling,* P. Malanotte-Rizzoli, Ed., Elsevier, 119–145.

Thacker, W. C., 1991: Automatic differentiation from an oceanographer's perspective. *Automatic Differentiation of Algorithms,* SIAM, 191–201.

——, 1992: Oceanographic inverse problems. *Physica D,* **60,** 16–37.

——, and R. B. Long, 1988: Fitting dynamics to data. *J. Geophys. Res.,* **93,** 1227–1240.

Tziperman, E., and W. C. Thacker, 1989: An optimal control/adjoint equations approach to studying the oceanic general circulation. *J. Phys. Oceanogr.* **19,** 1471–1485.