

# **Construction of the Adjoint MIT Ocean General Circulation Model and Application to Atlantic Heat Transport Sensitivity**

**Jochem Marotzke, Ralf Giering, Kate Q. Zhang,  
Detlef Stammer, Chris Hill and Tong Lee**

Center for Global Change Science  
Department of Earth, Atmospheric and Planetary Sciences  
MIT, Cambridge, MA 02139-4307 USA



**Report No. 63  
May 1999**

**Construction of the Adjoint MIT Ocean General Circulation Model and Application to  
Atlantic Heat Transport Sensitivity**

JOCHEM MAROTZKE\* , RALF GIERING\$, KATE Q. ZHANG# , DETLEF STAMMER,  
CHRIS HILL, and TONG LEE\$

*Center for Global Change Science  
Department of Earth, Atmospheric, and Planetary Sciences  
Massachusetts Institute of Technology  
Cambridge, MA 02139  
USA*

**April 14, 1999**

Submitted to *Journal of Geophysical Research*

---

\* Corresponding author address: Dr. Jochem Marotzke; MIT, Room 54-1514; 77 Mass. Ave.; Cambridge, MA  
02139; USA; jochem@sound.mit.edu, (617) 253-5939

\$ Current affiliation: Jet Propulsion Laboratory; Pasadena, CA 91109; USA

# Current affiliation: Department of Geophysics and Planetary Sciences; California Institute of Technology;  
Pasadena, CA 91125; USA

## Abstract

We first describe the principles and practical considerations behind the computer-generation of the adjoint to the MIT ocean general circulation model (GCM), using R. Giering's software tool Tangent-linear and Adjoint Model Compiler (TAMC). The TAMC's recipe for (FORTRAN-) line-by-line generation of adjoint code is explained by interpreting an adjoint model strictly as the operator that gives the sensitivity of the output of a model to its input.

Then, the sensitivity of 1993 annual-mean heat transport across  $29^\circ$  N in the Atlantic, to the hydrography on 1 January 1993, is calculated from a global solution of the GCM. The "kinematic sensitivity" to initial temperature variations is isolated, showing how the latter would influence heat transport if they did not affect density and hence the flow. Over one year, the heat transport at  $29^\circ$  N is influenced kinematically from regions up to  $20^\circ$  upstream in the western boundary current, and up to  $5^\circ$  upstream in the interior. In contrast, the dynamical influences of initial temperature (and salinity) perturbations spread from as far as the rim of the Labrador Sea to the  $29^\circ$  N section, along the western boundary. The sensitivities calculated with the adjoint compare excellently to those from a perturbation calculation with the dynamical model. Perturbations in initial interior salinity influence meridional overturning and heat transport when they have propagated to the western boundary and can thus influence the integrated east-west density difference. Our results support the notion that boundary monitoring of meridional mass and heat transports is feasible.

## 1. Introduction

The impending need to synthesize, basin-wide and globally, ocean data such as the entire World Ocean Circulation Experiment (WOCE) dataset including altimetry, and additionally surface forcing data obtained from weather centers and scatterometers, makes it imperative to use sophisticated ocean general circulation models (GCMs) to (1) interpolate in space and time between the observations and (2) diagnose unobservable but important quantities such as vorticity and heat transports. Conversely (and indeed prior to all these interpretations), the data stream must be used to test and improve the GCMs (the stated Goal 1 of International WOCE). A very powerful and general approach to synthesis is the use of optimization methods; we will concentrate on the particular flavor that has become known as the "adjoint approach" in meteorology and oceanography (e.g., Talagrand and Courtier, 1987; Thacker and Long, 1988; for the roots of "adjoint methods" in control theory and their relationship to sequential estimation see Wunsch, 1996).

The basic idea is quite simple: A model is defined by an algorithm (often realized through a computer code) and its independent variables, for example initial conditions, boundary conditions, or empirical parameters. A "performance index" or "cost function" measures how well a model realization matches observations; the cost function mostly is some weighted least-squares measure, but not necessarily so. The optimization determines the independent (or control) variables such that the cost function is minimized. If the model is nonlinear and large (often with  $10^5$  -  $10^6$  independent variables), iterative searches for the optimal solution are among the few practical strategies, but they need directional information (which direction is downhill?). This gradient or sensitivity of the cost function with respect to the control variables is calculated by what has become known as the "adjoint model".

Coding the "adjoint" to a complex numerical code is extremely tedious, time-consuming, and error-prone, so it is not surprising that until very recently only few adjoint ocean GCMs existed. The first effort was performed by R. B. Long, S. M. Hwang, and W. C. Thacker (Long, R. B., S. M. Hwang, and W. C. Thacker, The finite-difference equations defining the GFDL-GCM and its adjoint. Unpublished report. Atlantic Oceanographic and Meteorological Laboratory, Miami, Florida, 1989) who constructed the adjoint to Cox's (1984) version of the Geophysical Fluid Dynamics Laboratory (GFDL) GCM. It took several years to transform the prototype into a tool applicable to the inversion of hydrographic data (see the descriptions in Tziperman et al., 1992a,b; Marotzke, 1992; Marotzke and Wunsch, 1993; Bergamasco et al. 1993; Yu and Malanotte-Rizzoli, 1996). The model and its adjoint have eventually been used to estimate the time-mean and seasonally varying general circulations of the North Atlantic (Marotzke and Wunsch, 1993; Yu and Malanotte-Rizzoli, 1996, 1998), the Indian Ocean (Lee and Marotzke, 1997, 1998), and the global ocean (Sirkes et al., 1996).

The second adjoint GCM, based on the Hamburg Large-Scale Geostrophic model (LSG, Maier-Reimer et al., 1993), was constructed by R. Giering around 1990, and was documented and applied to Pacific tropical wave dynamics in Giering (1996). Schiller and Willebrand (1995) developed an approximate adjoint to the GFDL model (based only on heat and salt conservation), arguing that for optimization problems high accuracy is often not required. This model was subsequently applied to the North Atlantic mean circulation (Schiller 1995). Moreover, an adjoint to the Laboratoire d'Océanographie Dynamique et de Climatologie (LODYC) GCM has been developed (P. Delecluse, 1996, personal communication).

The above list, while incomplete, indicates that adjoint GCMs have not found very widespread use in the oceanographic community, despite their power (see, e.g., the comparison in Marotzke and Willebrand, 1996, between general circulation and simpler-dynamics inverse

models) and the popularity of the GCMs themselves. From our own experience, it is clear that a large part of the problem lies in the unyielding nature of adjoint code. An adjoint tracks the sensitivity of output with respect to input (see section 2 below); every change in the GCM therefore has to be transferred to the adjoint to correctly calculate the change in sensitivity. It became clear that the construction of an adjoint ocean model ideally would be done with a software tool (e.g., Thacker, 1991). Sensitivity of computer code results to inputs is a part of computational differentiation, which itself is a large enterprise within computational science (e.g., Griewank and Corliss, 1991; Berz et al., 1996). But until recently the general-purpose tools available (e.g., Bischof et al., 1992) were not applicable to problems with a large number of degrees of freedom - such as an ocean model for which the sensitivity to initial conditions was required (see section 4 below).

The situation changed fundamentally with the availability of automatic differentiation tools such as the Tangent-linear and Adjoint Model Compiler (TAMC, Giering, 1996, Giering and Kaminski, 1998). We document here the steps required to apply the TAMC to a fully-fledged ocean GCM ("MIT GCM", Marshall et al., 1997a,b). The MIT GCM has a modern programming structure and has been optimized for various massively parallel computers. The TAMC-created adjoint GCM retains the parallelism of the original code and is thus a prototype of complex but efficient adjoint codes. Van Oldenborgh et al. (1999) used the TAMC to construct the adjoint to the Hamburg Ocean Primitive Equation Model (HOPE, e.g., Latif and Barnett, 1994), as did Eckert (1998) and Eckert et al. (1999) with an early version of HOPE. The automatic differentiation tool *Odyssée* (Rostaing et al., 1993) was used to construct an (alternative) adjoint to the LODYC GCM (P. Delecluse, 1996, personal communication). Other efforts are undoubtedly underway; however, none of these has to our knowledge been published in the peer-reviewed literature as of writing of this manuscript. One purpose of this paper is, therefore, to make a powerful dynamical and model-data synthesis tool known and therefore available to the community at large.

Adjoint models have predominantly been applied to optimization in both meteorology (see Talagrand, 1991; Errico, 1997) and oceanography. Indeed, our first applications of the TAMC-created adjoint MIT GCM were data-synthesis studies of the global (Stammer et al., 1997) and Indian Oceans (Zhang and Marotzke, 1999). Here, we stress the "pure" sensitivity information in an adjoint solution, not its application to optimization. This creates the most direct connection between theory and actual construction of adjoint code and puts greater weight on the interpretation of the adjoint solution itself. Adjoints do not appear to have been applied to sensitivity studies using ocean GCMs, in contrast to the long history in meteorology (see, e.g., Hall et al., 1982; Errico and Vukicevic, 1992; Errico, 1997). The only oceanographic adjoint sensitivity study known to us is the one of Schröter and Wunsch, 1986, who employed a barotropic quasigeostrophic model (notice that they did not use the term "adjoint"). Here, we investigate how the heat transport across a transoceanic section depends on the initial conditions of the global model solution obtained by Stammer et al. (1997) for year 1993. In particular, we can differentiate between the purely kinematic and the dynamical consequences of temperature variations in the initial conditions. Moreover, we study how deep density perturbations lead to changes in meridional heat transport, which bears immediate connections to a potential observing system design for the meridional overturning circulation.

This paper is organized as follows. Section 2 outlines the fundamentals of adjoint sensitivity calculations and adjoint code construction, while section 3 details the concrete steps we took to obtain an efficient adjoint to the MIT GCM. Section 4 discusses the results of applying the adjoint GCM to heat transport sensitivity calculations; these sensitivities are put into a dynamical context in section 5. A few concluding remarks follow in section 6. Readers mainly interested in the physical discussion might skip sections 2 and 3. Appendix A contains a "tutorial" outlining step-by-step the calculations in a simple but nontrivial "model" and its adjoint. In appendix B, a simple example is given for calculating "kinematic sensitivity", as defined in section 4. Appendix C briefly describes the MIT GCM.

## 2. Fundamentals of adjoint sensitivity calculations and adjoint code construction

The following contains a heuristic account of how to computer-generate adjoint code; only the basic concepts will be explained. While there is overlap with the presentations by Talagrand (1991), Errico (1997), and Giering and Kaminski (1998), we differ from the former two papers in that we stress here the connection between the theoretical considerations and the actual (automatic) construction of adjoint code, a full account of which is given in Giering and Kaminski (1998). Computational differentiation in general has recently been reviewed in the conference proceedings edited by Berz et al. (1996).

Denote the "state" (i.e., all prognostic variables) of a numerical model at timestep  $n$  by the vector  $\mathbf{X}_n$ ,  $0 \leq n \leq N$ . For simplicity, let us assume that the only independent or control variables are the initial conditions,  $\mathbf{X}_0$ , of the model and that the cost function  $J$  depends only on the final state,  $\mathbf{X}_N$ ; that is, one can write

$$J = f \circ \mathbf{X}_N, \quad (1)$$

where  $f$  is a scalar function that maps the state vector  $\mathbf{X}_N$  onto the real axis, and "o" stands for "operates on". The cost function is linked to the control variables by repeated application of the numerical model (once per timestep),

$$J = f \circ \Psi_N \circ \Psi_{N-1} \circ \Psi_{N-2} \circ \dots \circ \Psi_2 \circ \Psi_1 \circ \mathbf{X}_0 = J \left( f \left( \Psi_N \left( \Psi_{N-1} \left( \dots \left( \Psi_1 \left( \mathbf{X}_0 \right) \dots \right) \right) \right) \right) \right). \quad (2)$$

It is advantageous to consider the entire equation (2) as "the model", that is, the definition of the cost function is part of the "model". The sensitivity of  $J$  to the control vector is given by the chain rule



$$\frac{\partial J}{\partial \mathbf{X}_0} = 1 \circ f' \circ \Psi'_N \circ \Psi'_{N-1} \circ \Psi'_{N-2} \circ \dots \circ \Psi'_2 \circ \Psi'_1 \circ \mathbf{I}, \quad (3)$$

where the prime indicates derivative with respect to the argument, and  $\mathbf{I}$  is the unit matrix. Every  $\Psi'_n$  is the Jacobian matrix of the model at timestep  $n$ , that is, the sensitivity of the state *after* timestep  $n$  to the state *before* timestep  $n$ . Why the trivial factor  $J' \equiv dJ/df=1$  has been retained will become clear immediately. Equation (3) is a  $1 \times L$  matrix (i.e., a row vector), where  $L$  is the size of the state vector. As it stands, (3) is evaluated from the right to the left, that is,  $N$  matrix multiplications with  $L \times L$  matrices must be performed. Notice that the Jacobians never need be computed explicitly if it is possible to differentiate the *algorithm* calculating  $\mathbf{X}_n$  from  $\mathbf{X}_{n-1}$ . Every matrix multiplication corresponds to separately applying the so-called tangent linear model (TLM) to the  $L$  columns of an  $L \times L$  matrix (beginning with multiplying all  $L$  unit vectors). The last operation is the multiplication with row vector  $f'$ . This is the *forward mode* of automatic differentiation, which has been implemented in the very general software tool ADIFOR (Bischof et al., 1992).

Alternatively, one can take the transpose of (3), which gives

$$\left( \frac{\partial J}{\partial \mathbf{X}_0} \right)^T = (\Psi'_1)^T \circ (\Psi'_2)^T \circ \dots \circ (\Psi'_{N-2})^T \circ (\Psi'_{N-1})^T \circ (\Psi'_N)^T \circ (f')^T \circ 1. \quad (4)$$

Proceeding again from right to left, the number 1 is first operated upon by column vector  $(f')^T$ , followed by an application of the transpose Jacobian at timestep  $N$ , and so on. Notice that (4) involves the application of an  $L \times L$  matrix to a single column vector; more generally, if  $J$  were a vector valued function with  $K$  components,  $K$  applications would be involved. Equation (4) describes the *reverse mode* of automatic differentiation, and is equivalent to computing the solution of the adjoint model since the adjoint operator is the transpose Jacobian. Notice that  $f'$  need not formally constitute an inner product, which would be required to rigorously use the

term "adjoint", in the same sense as in the theory of differential equations (see Talagrand, 1991). In this regard, we use a more general definition of the adjoint "model", namely as the transpose of the Jacobian matrix. Taking the point of view that the adjoint model evaluates the chain rule leads to its most parsimonious definition that we are aware of, and using the term "output function" is more appropriate here than "cost function". But notice that this procedure leads to the same equations ("model") as those derived for the Lagrange multipliers in constrained optimization (e.g., Thacker and Long, 1988).

For large  $L$ , the reverse mode is clearly more computationally efficient than the forward mode, but at the expense that *all* the variables permitting the execution of the linearized model step must be available *in reverse order of the original computation*. Hence, they must be stored throughout the integration of the nonlinear model or recomputed when needed. In contrast, the TLM can be applied by running along one integration of the nonlinear model. Storage requirements will be taken up again later.

The interpretation of the  $\Psi_n$  as integrating the model over an entire timestep is by no means required. Rather, it could denote any intermediate step between input of the controls and output of the cost function. In the extreme,  $\Psi_n$  stands for the execution of a single line of computer code, which can be viewed as mapping the entire state vector plus derived variables (plus additional parameters irrelevant here) onto a single output variable. Constructing an adjoint thus means, according to (2) - (4), linearizing every code assignment and executing the transposes of all linearized assignments in reverse order, initializing with the number 1. This recipe is now demonstrated at the lowest level (we take the "atomistic" view).

Implicitly, all variables but the output variable are kept constant during an assignment, which can hence be written generally if not very efficiently,

$$\begin{pmatrix} x_1^{out} \\ \vdots \\ x_r^{out} \\ \vdots \\ x_p^{out} \end{pmatrix} = \begin{pmatrix} x_1^{in} \\ \vdots \\ g(x_r^{in}, \mathbf{a}) \\ \vdots \\ x_p^{in} \end{pmatrix}, \quad (5)$$

where  $\mathbf{a}$  is the vector of all input variables *except* the variable  $x_r^{in}$ . Since  $x_r^{in}$  and  $x_r^{out}$  allocate the same memory in the computer,  $x_r^{in}$  is overwritten by  $x_r^{out}$ . The Jacobian of the assignment or mapping (5) is

$$\begin{pmatrix} \frac{\partial x_i^{out}}{\partial x_j^{in}} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \frac{\partial g}{\partial x_1^{in}} & \dots & \frac{\partial g}{\partial x_r^{in}} & \dots & \frac{\partial g}{\partial x_p^{in}} \\ 0 & \dots & 0 & 1 & 0 \\ 0 & \dots & \dots & 0 & 1 \end{pmatrix}. \quad (6)$$

The adjoint of the assignment is defined by the transpose Jacobian, which links adjoint input and output variables (marked by an overbar) through

$$\begin{pmatrix} \frac{\partial \bar{x}_i^{out}}{\partial \bar{x}_j^{in}} \end{pmatrix} = \begin{pmatrix} 1 & 0 & \frac{\partial g}{\partial x_1^{in}} & \dots & 0 \\ 0 & 1 & \dots & \dots & 0 \\ 0 & \dots & \frac{\partial g}{\partial x_r^{in}} & \dots & 0 \\ 0 & \dots & \dots & 1 & 0 \\ 0 & \dots & \frac{\partial g}{\partial x_p^{in}} & 0 & 1 \end{pmatrix}. \quad (7)$$

Written out, this becomes

$$\bar{\mathbf{a}}^{out} = \bar{\mathbf{a}}^{in} + \frac{\partial g}{\partial \mathbf{a}} \bar{x}_r^{in}, \quad (8a)$$

$$\bar{x}_r^{out} = \frac{\partial g}{\partial x_r^{in}} \bar{x}_r^{in}. \quad (8b)$$

In the assignment corresponding to (8b), the output again overwrites the input; hence the stated order of the assignments matters. Since every assignment usually involves only a small number of input variables, most components of (8) are in practice trivial: The adjoint variable remains unchanged, since its physical variable has no impact on the output. These remarks illustrate how (2)-(4) have to be changed to include boundary conditions as control variables: Every mapping step  $\Psi_n$  then formally has as arguments not only the model state, but also the forcing functions *at all timesteps*. The mapping does not depend on the forcing except at the matching timestep; before and after, the adjoint variable to the forcing is kept constant.

A special case arises when in the assignment (5) the input variable  $x_r^{\text{in}}$  does not influence the output  $x_r^{\text{out}}$ , that is,  $x_r$  is assigned a new value but the right-hand side does not contain  $x_r$ . This means that  $\partial g / x_r^{\text{in}} = 0$  in (8) and hence  $\bar{x}_r^{\text{out}} = 0$ : The adjoint variable traces the influence on the cost function backward; since  $x_r$  is overwritten it has no effect on the later calculations, and its adjoint is set to zero.

The above exercise illustrates how the adjoint to an arbitrary assignment in program code has to be formulated. The instructions of (8) are the core of the software tool TAMC (Tangent-linear and Adjoint Model Compiler; Giering, 1996, Giering and Kaminski, 1998). They are presented here to make plausible that indeed computer-generating adjoints can be done using very general principles. To illustrate these ideas, Appendix A shows an example program that is simple in its calculations yet illustrates the power of the TAMC in dealing with nonlinear algorithms. In practice, the TAMC performs the following steps:

- 1) Parsing the code, to identify variables, syntax, etc. This step is required by every compiler. For further processing, an internal abstract representation of the code is generated.

2) Data flow analysis. This step superficially resembles what optimizing compilers do, but is performed on the entire code (globally) rather than locally (e.g., a single loop). It is needed to avoid a myriad of trivial adjoint statements that propagate zero adjoint values for constants and variables that do not influence the cost function. The data flow analysis identifies "active variables", both in the sense of what influences the cost function, and of what is influenced by the control variables.

3) Line-by-line construction of adjoint code, based on (8). Most program codes contain higher-level structures beyond assignments; their treatment is described in detail in Giering and Kaminski (1998). A loop can be considered as a formalized sequence of individual assignments. If any pass through the loop depends on the outcome of any other pass, the adjoint loop is performed in reverse order. Otherwise, the order is not important and the adjoint loop is executed in the same order as in the physical model. Conditional statements (e.g.: if statically unstable, perform convective adjustment) are not differentiable at the branch point; away from it, however, linearization is achieved by providing the adjoint with the *outcome* of the condition during the physical model integration.

4) Once a piece of adjoint code is constructed, it is known which parts of the solution of the physical model are required. An additional reverse data flow analysis is then applied to construct the sufficient but minimal recomputation of these ("program slicing"). Furthermore, the data flow analysis is needed to identify when the Jacobian depends on the solution of the physical model, the results of which then have to be stored or recomputed.

### 3. Preparation of the MIT GCM code

The overarching design principle in constructing the GCM adjoint has been that there must be no human intervention necessary in the adjoint code, because whatever change is manually inserted into the adjoint would be lost again upon a later application of the TAMC. This means, in particular, that if a particular data flow structure in the model leads to a conflict and errors in the adjoint, we *always* modify the physical model and never directly the adjoint, although the latter sometimes might be far simpler. Occasionally, errors in the adjoint suggested where the TAMC should be modified or extended.

A second principle has been to keep the optimization portion of the code, needed when the model is fitted to data, separate from the model and the adjoint. From the optimization routine's perspective, the model is merely a device that returns a single number, the cost function, if given some input vector. The adjoint delivers the gradient of the cost function with respect to the input. Hence, the optimization is oblivious to the inner workings of the model, which makes modularity of the code desirable, so that different optimization procedures and, in principle, different physical models could be readily inserted. A basic logical structure emerges in which, at the top level, an optimization is performed with respect to some dimensionless independent (or control) variables, which usually are deviations from some prior estimates, scaled by the uncertainty of the prior (see Wunsch, 1996, for a general discussion). When the model is called to perform an evaluation of the cost function, the vector of control variables is mapped onto the physical variables of the model in an interface subroutine. Its adjoint subroutine maps the solution of the adjoint model onto a dimensionless gradient vector.

The greatest obstacle to applying the adjoint technique to large-scale problems comes from the dependence of the transpose Jacobian on the physical model solution if the model is nonlinear (or, in other words, the dependence of adjoint code on the model solution). Two

extreme strategies exist: The first is recomputation of every model statement up to the point where the solution is needed, leading to a computational load proportional to the square of the "complexity" (number of operations) of the model. The second strategy stores every intermediate quantity ever computed by the code, leading to a storage requirement roughly proportional to the code's complexity. Apparently, only the latter strategy, which is infeasible for large-scale problems, was considered by the group that developed the "forward mode" automatic differentiation tool ADIFOR (Bischof et al., 1992).

In contrast, the TAMC by default performs "total" recomputation, at subroutine level, with the remaining storage requirements handled automatically. It is plausible, however, that by compromising between the recomputation and storage requirements a manageable computational load can be attained. Indeed, the example in Appendix A shows that by introducing storage directives, recalculation can be avoided. In completely general code, the efficient introduction of storage directives is not straightforward. However, most large oceanographic and meteorological computing tasks involve the time-integration of GCMs, thereby introducing a very natural guiding principle: The model state (comprised of all prognostic variables) is stored once per timestep, and all diagnostic quantities are recomputed. While by no means guaranteed to be optimal, this principle provides a very simple and useful default strategy, which can subsequently be refined.

Storing the model state at every timestep still would make memory requirements prohibitive for even a moderate-sized GCM, but the TAMC supports an elegant strategy called checkpointing, which was proposed by Griewank (1992) and later (manually) implemented into a quasigeostrophic model and its adjoint by Restrepo et al. (1995). A good summary of the method was provided by Hersbach (1998), who applied the predecessor of the TAMC to a surface wave model. The integration period of  $N$  timesteps is subdivided into  $M$  segments each of length  $N/M$ ; the result of the forward model is stored at the end of each segment but not at

other timesteps. The forward model is re-run over the last segment and its history stored at every timestep; then the adjoint is run backwards over the last segment. This is repeated for all segments, moving backwards in time (from the last to the first segment). Resulting is the exact adjoint, at the cost of one extra forward integration.

For a given number  $N$  of timesteps and assuming that the "cost" of every type of storage is the same, the optimum number  $M$  of segments is calculated from minimizing the total storage, which is proportional to

$$R = N / M + M . \tag{9}$$

At any given time, an entire segment of length  $N/M$  and the number  $M$  of segments must be held in memory or on disk. Minimal total storage is achieved for  $M_{opt} = N^{1/2}$  and is proportional to  $R_{min} = 2N^{1/2}$ , which for  $N \sim 10^4$  (i.e., roughly one model year at a timestep of 1 hour) is a reduction by a factor of 50. Many refinements to this simplest checkpointing strategy are possible; for example, early segments can be made longer than later segments, and further levels of recomputation can be added. Griewank (1992) proved that storage can be made to grow no faster than logarithmically, at the cost of logarithmically increasing runtime. Here, we use only one or two levels of recomputation, to limit not only storage but also runtime cost.

Checkpointing is implemented in the following way. The timestepping loop is split into two nested loops, contained in two subroutines, `outer` and `inner`. Subroutine `outer` runs over  $M$  segments, defined above. Subroutine `inner` performs  $N/M$  timesteps each time it is called. By placing specific directives before and at the very beginning of these loops, the TAMC generates the checkpointing adjoint code. Two different storage devices ("tapes") are used; in `inner` the prognostic and some diagnostic variables are stored in auxiliary COMMON blocks, while in `outer` the prognostic variables are stored in a direct-access file



on disk. This technique provides the most efficient adjoint code on a vector computer; on a parallel computer with large core memory the outer tape can also be stored in memory, providing higher performance since input/output to disk is rather slow.

An extra saving is achieved by recognizing that in minimization problems, the adjoint is used only to provide the downhill direction, which is useful even if only approximate (Schiller and Willebrand, 1995). Marotzke and Wunsch (1993) implemented the very efficient shortcut of storing the model history only at certain intervals, so the adjoint operator is approximated; the permissible length of storage interval is problem dependent. Through explicitly specified storage locations, the TAMC generates code to reuse storage; this trick can be combined with the checkpointing strategy and implemented into the adjoint code without human intervention. Even under rapidly varying circumstances, additional storage savings by a factor of 5 have been possible. Storing with reduced precision further decreases the memory requirements.

#### **4. Application: Sensitivity of Atlantic heat transport**

The TAMC-generated adjoint GCM has already been used for very complex tasks, which justifies the considerable investment in development. Stammer et al. (1997) performed a global assimilation with a 2°-resolution model for the year 1993, using TOPEX/POSEIDON altimetry and 10-day averages of synoptic surface fluxes. Although global assimilation experiments with adjoint GCMs have been performed earlier (e.g., Sirkes et al. 1996), Stammer et al.'s (1997) application appears to be the first with synoptic forcing and data. Zhang and Marotzke (1999) estimated the Indian Ocean general circulation from climatology and a basin-scale ocean model; for the first time with an ocean GCM, they estimated the parameters of open boundary conditions. Both these developments were enormously simplified through the flexible construction and modification of the adjoint model.

We now describe in some depth the application of the adjoint model to a sensitivity analysis of our ocean GCM. Although straightforward in principle and extensively done with atmospheric models (see Introduction, Hall et al., 1982, and Errico, 1997), such a sensitivity calculation appears never to have been done with an ocean GCM. Previously, the solution of an adjoint ocean GCM has mainly been used as directional information in a minimization procedure but otherwise discarded. Here, we calculate the sensitivity of an important function of the model solution, the heat transport across 29°N in the Atlantic, to the initial conditions in temperature and salinity. More specifically, the heat transport is the annual mean for 1993 and taken from the solution of Stammer et al. (1997), while the initial conditions are the temperatures and salinities estimated for 1 January 1993. The latitude of 29° N is chosen because it is close to the maximum northward heat transport. Hence, the output function is

$$Q = P^{-1} c_p \rho_0 \iiint v T dx dz dt, \quad (10)$$

where  $P$  is the averaging period,  $c_p$  specific heat,  $\rho_0$  a reference density,  $v$  meridional velocity, and  $T$  potential temperature.

Figure 1a shows the sensitivity of the annual-mean heat transport at 29° N,  $Q$ , to the sea surface temperature of 1 January, 1993. The large positive contribution near the western boundary and south of 29° N is readily understood from the surface velocity field (Fig. 1b). The considerable northward boundary current contributes a larger temperature flux if its temperature is higher; downstream of the section, there is no contribution from this effect. With a northward velocity component of 10 cm/s, a temperature anomaly can be transported in the boundary current over 3000 km in one year, roughly consistent with the southward extent of the region of positive sensitivity in Fig. 1a. In contrast, simple passive advection of temperature anomalies cannot explain the negative sensitivity away from the western boundary

and right to the north of the section, because the surface flow across the section is weakly northward over the entire year. Neither can passive advection account for the large negative sensitivities along the western boundary to the north of the section, or the large positive sensitivity off Africa.

To gain further insight, we have to consider the effect of a temperature perturbation more broadly. A change in temperature has a dynamically active component, through its effect on density and hence thermal wind shear, and a dynamically passive component (temperature change on an isopycnal, contributing to the combination of temperature and salinity that is locally orthogonal to the density change and was called "veronicity" by Munk, 1981, p. 282). For small amplitudes, the "active" component of a temperature perturbation has the same consequences as a salinity perturbation, scaled by  $(-\alpha/\beta)$ , where  $\alpha$  and  $\beta$  are respectively the thermal and haline expansion coefficients,

$$\alpha \equiv -\frac{1}{\rho} \left( \frac{\partial \rho}{\partial T} \right)_S; \quad \beta \equiv \frac{1}{\rho} \left( \frac{\partial \rho}{\partial S} \right)_T. \quad (11)$$

Hence, we can identify the sensitivity of heat transport to the dynamically active part of the initial temperature by considering the sensitivity of heat transport across 29° N to initial salinity perturbations, which is shown in Fig. 1c. In the western boundary current and off Africa, the salinity contributions are of the same structure and opposite sign as the sensitivity to temperature; over the interior, there is a contribution of the same sign but to the south of the section.

The following heuristic reasoning suggests a way to combine the heat transport sensitivities to temperature and salinity changes such that only the kinematic (dynamically inactive) sensitivity to temperature variations remains. This decomposition is reminiscent of the one by Bindoff and McDougall (1994), who analyzed changes in hydrography as being

caused by pure heave of isopycnals (or neutral surfaces), by heating on isopycnals, or by freshening on isopycnals. In some sense, we identify a passive tracer in our model, without having to timestep an additional equation. Write the northward heat transport,  $Q$ , symbolically as a function of initial temperature and salinity,

$$Q = Q(\rho(T, S), T), \quad (12)$$

where the grouping of arguments indicates that we consider heave and temperature changes on isopycnals as the independent processes influencing heat transport. Implicitly, (12) invokes the thermal wind relationship. The sensitivity of  $Q$  to initial temperature variations is shown in Fig. 1a and given by

$$\left(\frac{\partial Q}{\partial T}\right)_s = \left(\frac{\partial Q}{\partial \rho}\right)_T \left(\frac{\partial \rho}{\partial T}\right)_s + \left(\frac{\partial Q}{\partial T}\right)_\rho = -\alpha \rho \left(\frac{\partial Q}{\partial \rho}\right)_T + \left(\frac{\partial Q}{\partial T}\right)_\rho, \quad (13)$$

using (11), where the last term on the right-hand side is the "kinematic sensitivity". The sensitivity of  $Q$  to initial salinity variations is shown in Fig. 1c and given by

$$\left(\frac{\partial Q}{\partial S}\right)_T = \left(\frac{\partial Q}{\partial \rho}\right)_T \left(\frac{\partial \rho}{\partial S}\right)_T = \beta \rho \left(\frac{\partial Q}{\partial \rho}\right)_T. \quad (14)$$

We can now isolate the kinematic sensitivity of heat transport to initial temperature variations by rearranging (13) and using (14), to obtain

$$\left(\frac{\partial Q}{\partial T}\right)_\rho = \left(\frac{\partial Q}{\partial T}\right)_s + \alpha \rho \left(\frac{\partial Q}{\partial \rho}\right)_T = \left(\frac{\partial Q}{\partial T}\right)_s + \frac{\alpha}{\beta} \left(\frac{\partial Q}{\partial S}\right)_T. \quad (15)$$

Appendix B contains a worked example of (15), where the result is demonstrated for a box model that is "run" over one timestep. Figure 1d shows the "sum" of Figs. 1a and 1c, weighted according to (15). There is near-perfect cancellation of the dynamical contributions

to heat transport sensitivity in the western boundary current to the north of the section, and off Africa. Near the western boundary, an increase in surface density leads to a decrease in sea level and hence a stronger baroclinic shear in the Gulf Stream. This perturbation would propagate southward, as a mix of Kelvin and topographic waves (e.g., Döscher et al., 1994), and influence heat transport across  $29^\circ$  N from a significant distance. The nearly antisymmetric contributions off Florida and off Africa, just north and south of the section, respectively (Fig. 1c), lead to increased northward thermal wind shear and hence surface flow; since they are located right at the boundary, they do not lead to compensating southward flow as an isolated mid-ocean anomaly would.

The isolation of the kinematic sensitivity helps explain, tentatively, the narrow negative bands north and south of the section, in the heat transport sensitivities to temperature and salinity variations, respectively. The kinematic sensitivity to initial SST is negative over the ocean interior and shows considerable symmetry about the section, the latter reflecting the model's finite-difference scheme (heat transport is calculated at a velocity gridpoint, so the two neighboring temperatures to the north and south must be averaged to form the product of velocity and temperature). The minima in Fig. 1d, across  $29^\circ$  N, reflect the banded structure of generally southward velocity at level 2 (Fig. 1e), except west of  $70^\circ$  W. Temperatures in the top two layers are strongly coupled in winter (the mixed layer is ca. 100 m deep; figure not shown), so in the interior, where the meridional surface velocities are small, the flow at level 2 dominates the advection of a temperature anomaly. Overall, the kinematic sensitivity displayed in Fig. 1d reflects the classical picture of a subtropical gyre, with a strong poleward western boundary current and broad equatorward return flow. It is the ability to isolate the kinematic sensitivity that allows us to state, loosely, that we have identified a passive tracer from temperature and salinity without having to run a passive-tracer equation or an ensemble of Lagrangian drifters.

No such simple picture emerges from the dynamical sensitivity (Fig. 1c), but a tentative explanation can be given as follows. A positive salinity anomaly in the center of the subtropical gyre causes a positive density and hence negative sea level anomaly. If this anomaly can spread westward as a barotropic Rossby wave (without, however, reaching the western boundary), the net result is a weakening of the subtropical gyre and its associated heat transport. A positive temperature anomaly has the opposite effect; notice that only a narrow latitudinal range is involved. The near-total compensation between dynamical and kinematic temperature sensitivities south of the section is conspicuous but appears accidental, judging from an analogous sensitivity calculation of heat transport across  $33^\circ$  N (not shown).

The sensitivity of heat transport across  $29^\circ$  N to initial conditions deeper down differs significantly from those at the surface (Fig. 2). First, and somewhat trivially, a unit perturbation in deep temperature or salinity creates a larger heat or salt *content* change, due to the larger layer thickness at greater depth. But in addition there is a considerably larger dynamical contribution of initial temperature changes to the heat transport; correspondingly, there is a larger contribution from initial salinity (Fig. 2). As a consequence, the kinematic sensitivity to a deep initial temperature change (again reflecting the horizontal velocity field at the same depth; figure not shown) plays a negligible role; the cancellation between the dynamical sensitivities is almost complete. At 1160 m, the sensitivity to temperature shows a pattern with negative contributions north of the section, concentrated near the western boundary, while south of the section the contributions are positive and arise along the entire section. Additionally, positive contributions arise along the eastern boundary, from as far south of the section as  $10^\circ$  N. Most remarkable is perhaps that both shallow and deep temperature and salinity anomalies at the rim of the Labrador Sea can influence, within one year, the heat transport across a subtropical section. The next section will put these considerations into the larger context of the dynamics of the meridional overturning circulation.

## 5. Dynamics of heat transport sensitivity to salinity perturbations

We now change our perspective and consider how a deep salinity anomaly influences meridional heat transport over a considerable distance. Apart from the dynamical information this procedure yields, it provides an independent test of the adjoint sensitivity calculation. A salinity perturbation is added to the initial conditions of the optimized run of Stammer et al. (1997), in that salinity at 1160 m depth is perturbed by 0.01, between latitudes  $28^{\circ}$  N and  $38^{\circ}$  N, and between longitudes  $68^{\circ}$  W and  $58^{\circ}$  W (a total of 25 grid cells).

Figure 3 shows the resulting perturbations in horizontal velocity at 160 m depth, at 60-day intervals, beginning with day 30. The salinity anomaly at 1160 m depth creates a response at all depths (counterclockwise vortex above, clockwise vortex below), which migrates westward with the speed of first-mode Rossby waves. The perimeter of the vortex reaches the western boundary after about 90 days. The subsequent evolution south of the anomaly resembles that of a shoaling Kelvin wave, with a northward current emerging along the western boundary; when the wave reaches the equator, it "draws" water from the east. The behavior in Fig. 3 is entirely consistent with the Kelvin wave patterns described in Kawase (1987). Notice that within one year, there is only a weak response south of the equator.

Figure 4 shows the perturbation in meridional overturning stream function, at the same times as the velocity plots of Fig. 3. A coherent positive pattern (indicating clockwise anomalous transport) is established roughly at the latitudes of the perturbation within 90 days, which is when the anomaly vortex has reached the western boundary. This pattern intensifies until day 270, when its southern edge begins to extend southward across the equator. In contrast, the core of the stream function anomaly migrates southward more slowly (from about  $30^{\circ}$  N to  $20^{\circ}$  N, between days 90 and 330). Below about 3000 m depth, the southern edge migrates southward more slowly than above.

Figure 5 shows the time series of the resulting heat transport anomalies, at a number of latitudes. To the north of the perturbation, there is virtually no response. At 35° N, the initial response of 1.5 Terawatts (TW, 1 TW  $\equiv 10^{12}$  W) decays over 100 days, but regains strength to end up at about 2 TW. At both 24° N and 29° N, the heat transport anomaly starts very small but grows steadily to a maximum of about 4 and 3.5 TW, respectively, between days 200 and 250, roughly coincident with the maximum in meridional overturning. Finally, the perturbation heat transport at 9° N begins to increase considerably only after day 100. The evolution of the meridional overturning anomaly well explains that of the heat transport anomaly: There is no northward migration at all, while to the south of the perturbation the heat transport evolves in lockstep with the meridional overturning; notice that positive stream function values mean northward flow near the surface and southward flow at depth, causing positive heat transport anomalies.

The annual mean of the heat transport anomaly at 29° N is 2.3 TW. This compares excellently with that implied by the adjoint sensitivity calculation,

$$\delta Q = N_{\delta S} \left( \overline{\frac{\partial Q}{\partial S}} \right)_T \delta S = 25 \times 1.0 \times 10^{13} \frac{W}{psu} \times 0.01 psu = 2.5 TW, \quad (16)$$

where  $N_{\delta S}$  is the number of perturbed gridpoints,  $\delta S$  the magnitude of the perturbation, and  $\left( \overline{\frac{\partial Q}{\partial S}} \right)_T$  the patch-average of the adjoint sensitivities of Fig. 2b. We have confirmed the linearity of our perturbation run by scaling the salinity perturbation down by an order of magnitude; the response is nearly identical except for the same order-of-magnitude reduction in amplitude.

We have experimented with other deep salinity perturbations (figures not shown), and the picture that emerges explains why the pattern of positive influence is wider at lower



latitudes than at higher latitudes (Fig. 2b). When the anomaly is placed right at the boundary, it sooner reaches latitudes farther south; the Kelvin wave is set off immediately, rather than following the westward migration. Over one year, hence, a salinity perturbation from near 50° N can reach the "output latitude" of 29° N only if the perturbation starts right at the boundary. Anomalies further south, in contrast, have enough time to migrate westward prior to setting off the Kelvin wave. Conversely, in the adjoint, the "anti-Kelvin wave" induces *eastward* propagation, which has more time to migrate at lower latitudes (e.g., 40° N), because it is excited earlier. But notice that the wave picture is complicated considerably by bottom topography; the time evolution of the adjoint solution (figure not shown) clearly shows that the maximum near 38° N in Fig. 2b is due to trapping by topography and local recirculation.

The above considerations have demonstrated the importance of density anomalies near the boundaries, for meridional overturning and heat transport. Thus, a suggestion arises for the monitoring of meridional overturning and heat transport, by measuring density changes near the ocean margins as part of a climate observing system. We now place this connection into a more general dynamical context. If meridional flow is everywhere in thermal wind balance,

$$\partial_z v = -\frac{g}{f\rho_0} \partial_x \rho, \quad (17)$$

zonal integration across an ocean basin yields

$$L_x \overline{\partial_z v} = -\frac{g}{f\rho_0} (\rho_E - \rho_W), \quad (18)$$

where  $L_x$  is zonal extent of the basin,  $\rho_E$  and  $\rho_W$  density at eastern and western walls, respectively, and the overbar marks a zonal average. Other notation is standard; Cartesian coordinates are used for simplicity. Equations (17) and (18) reflect the well-known fact that the geostrophic *transport* between any two points is proportional to the pressure difference

between these points, irrespective of their distance (provided that topography is nowhere intersected).

Assuming vertical sidewalls (i.e.,  $L_x$  independent of depth) and the existence of a meridional streamfunction  $\psi$  (guaranteed if no mass enters or leaves through the zonal boundaries), leads to

$$\frac{f\rho_0}{g}\partial_{zz}\psi \equiv -\frac{f\rho_0}{g}\partial_z L_x \bar{v} = \rho_E - \rho_W, \quad (19)$$

meaning that under the simplest of all circumstances, the curvature of the meridional overturning streamfunction, with respect to the vertical, is proportional to the east-west density difference. This latter quantity took center stage in the theory of the purely buoyancy-driven meridional overturning circulation developed by Marotzke (1997) and Marotzke and Klinger (1999), but it does not appear to have found widespread use elsewhere.

The situation is considerably more complex if wind forcing and irregular bottom topography are admitted, which lead to contributions to the meridional overturning streamfunction that are not in thermal wind balance. Specifically, the external mode (vertical average) projects onto the meridional overturning in the presence of bottom topography (e.g., Robbins and Toole, 1997; Lee and Marotzke, 1998). Moreover, meridional overturning variability on timescales of seasonal and shorter is dominated by the varying Ekman transport and its depth-independent return flow (Lee and Marotzke, 1998; Jayne, 1999). None of these contributions to the meridional overturning has a straightforward relation to the density at the zonal boundaries, but as we show now, (19) can be a good approximation, even under non-idealized circumstances.

Figure 6 shows the left-hand and right-hand sides of (19), evaluated from the model solution obtained by Stammer et al. (1997) for the year 1993. For simplicity, we only show regions of the North Atlantic where in the analysis we could easily avoid cutting through bottom topography. Between 500 m and 1700 m depth, the correspondence between the curvature of the meridional overturning streamfunction and the east-west density difference is quite good; the difference plot between the two shows considerably smaller contributions (Fig. 6c). Largely, density is lower at the eastern boundary than at the western boundary, above 1700 m depth, meaning positive meridional transport shear, in turn consistent with northward mass transport in the thermocline. In contrast, the fit is poor above 500 m, presumably because of the reasons outlined above. A deeper investigation would go beyond the scope of this paper and will be left for future study.

## **6. Concluding remarks**

We have presented the foundations and a novel application of a powerful recent development in ocean modeling. A software tool has been used to computer-generate the adjoint of a fully fledged ocean GCM, without human manipulation of the adjoint code itself to make it correct or efficient. Possible applications of the adjoint model are:

1) Oceanographic inverse or optimization problems (synthesizing observations); the model described here has been applied in this mode by Stammer et al. (1997) and Zhang and Marotzke (1999).

2) The calculation of fastest-growing structures in dynamical instabilities (e.g., Farrell and Moore, 1992; Eckert, 1998; Eckert et al., 1999); we have not yet applied the adjoint MIT model to a problem of this class.

3) The sensitivity of parts of the model solution to independent parameters; only this application has been the focus of this paper.

Our results must in part be considered preliminary since the model solution obtained by Stammer et al. (1997) is only the first application of a global state estimation system. A subsequent experiment comprising a longer integration time and model improvements is being carried out; this run will be used for a more in-depth application of the ideas about sensitivity studies presented here. Moreover, in order to draw a conclusion concerning the importance of exact initialization at various locations, one must multiply the sensitivities shown here by typical variabilities; this would reduce the apparent influence of the deep temperatures and salinities. On the other hand, the isolation of kinematic and dynamical influences of temperature anomalies is independent of the particular model realization; it shows a way of disentangling individual processes from a complex superposition of effects.

Our sensitivity analysis lays out a strategy for designing field experiments aimed at climate monitoring, for example the Atlantic Climate Variability Experiment (ACVE) currently discussed as part of the Climate Variability and Predictability Program (CLIVAR). We see that deep-ocean properties far removed from the section of interest influence the heat transport; hence, confining observations to the near-surface ocean would be detrimental to the ability to dynamically link temperature measurements and heat transport estimates. We also obtain a feel for the influence, in a state estimation system, of a single hydrographic observation on the estimate of a large-scale climate quantity (although, strictly speaking, in that case we consider the reverse problem). Indeed this point is at the heart of the impending synthesis of the WOCE hydrography (using a GCM and station, rather than climatological, data). The maps shown in Figs. 1 and 2 here are encouraging, in that they show that there is a considerable nonlocal, large-scale influence of temperature perturbations. Finally, the highly localized regions of large

dynamical influence of temperature changes on heat transport, visible in Fig. 1a at both ends of the heat transport line, indicate that "boundary monitoring" of the meridional overturning may be a feasible observational strategy. This last point warrants further study, because the thermal wind relation is not guaranteed to hold in the vicinity of the boundary and the presence of small-scale variability. It remains to be shown whether spatial averaging would alleviate these concerns.

### **Acknowledgements.**

We are indebted to Carl Wunsch and John Marshall for their support and encouragement throughout this project. Christian Eckert and Carl Wunsch made helpful comments on earlier versions of the manuscript. This work was supported by NSF grants OCE-930135 and OCE-9617570 (JM, KQZ, and TL), contracts 958125 with JPL and NAG5-3724 with NASA (DS and RG), contract NAG5-7162 with NASA (DS), and a grant from the American Automobile Manufacturers Association (CH).

## Appendix A. A nontrivial example of automatic adjoint code construction

To illustrate the ideas outlined in section 2, an example program is shown that is simple in its calculations yet illustrates the power of the TAMC in dealing with nonlinear algorithms. The FORTRAN code (Fig. A1) and its TAMC-generated adjoint (Fig. A2) are shown with minimal editing, essentially taking out the declarative statements to make the codes compact enough for illustration. Moreover, the numerical results are inserted step-by-step as comments. Before the subroutine `model` is called, the number `nc` of control variables is set to 3, and the vector `x0` initialized to (1,3,3). The quantity `fc` is the cost function and is the only output of `model`. The state vector `x` is initially set to `x0`, and the algorithm then performs nonlinear (quadratic) operations on the `x`, followed by "convective adjustment". The comments to the right of the program code show the values of the pertinent variables *after* the assignment has been carried out. With the given values of the control variables, the "model" returns the cost function value 40.25. Notice the storage directives (`CADJ STORE`), which are converted into FORTRAN statements by the TAMC, storing the value of the vector `x` in an array or into a file.

The adjoint `admodel` is constructed with the TAMC, and is called with the number and the values of the control variables, the cost function value, and the initialization of the adjoint to the cost function variable, `adfc=1`, as input [in the case considered here, the latter is the "1" at the far right of (4) defining the sensitivity in reverse mode. Under more general conditions, the subroutine under consideration is one of many, and would have as input arguments `adfc` and the adjoints to the control variables that resulted from previous adjoint subroutines.] The adjoint statements are executed in reverse order, as are the `do` loops, where necessary. Repeatedly, the solution of the model is restored from an array or file and then used in the adjoints to (linearized) nonlinear statements. Most notable is the adjoint to the convective adjustment. The "static stability" of the physical model solution is checked and the convective adjustment performed on the adjoint variables. The outcome of this check is

predetermined once the physical model has been run, which is tantamount to the statement that we have performed a linearization of the branching statement (this is, of course, not possible right at the branch point). After the  $i=2$  step,  $adx(2)$  and  $adx(3)$  are completely mixed, while after the  $i=1$  step,  $adx(1)=adx(2)$ , both in direct analogy to the convective adjustment in the model. The physical interpretation is that since the properties of two adjacent "boxes" are mixed, each box contributes equally much to the value after mixing; hence the adjoint variables are equal, too.

```

subroutine model(nc,x0,fc)
...
CADJ INIT tape = common,2      ! Tape 'tape' initialized as COMMON block
CADJ INIT sfile = 'storefile' ! Tape 'sfile' initialized as file 'storefile'
  do i=1,3
    x(i)=x0(i)
  enddo                          ! x=(1,3,3)
c
  y=x(1)*x(1)                    ! y=1, x=(1,3,3)
  do i=1,3
    x(i)=y+x(i)*x(i)
  enddo                          ! x=(2,10,10)
c
c Convective adjustment:
  do i=1,2
CADJ STORE x = tape            ! Store: i=1: x1=x=(2,10,10) ! i=2: x2=x=(6,6,10)
    if (x(i).lt.x(i+1)) then ! i=1: 2<10? Yes: ! i=2: 6<10? Yes:
      x(i)=0.5*(x(i)+x(i+1)) ! i=1: x(1)=6 ! i=2: x(2)=8
      x(i+1)=x(i) ! i=1: x(2)=6 ! i=2: x(3)=8
    endif
  enddo                          ! x=(6,8,8)
c
CADJ STORE x = sfile          ! Store x3 = x = (6,8,8)
  fc=(x(1)-5.5)**2+2.*x(2)+3.*x(3) ! fc = 40.25
end

```

Fig. A1: Example code. Included as comments are the values of the variables *after* the statements have been executed.

To provide an independent test of the outcome of the adjoint model, we now perturb the control variables by "tracers" such that

$$\mathbf{x0} = (1 + \varepsilon_1, 3 + \varepsilon_2, 3 + \varepsilon_3) \quad (\text{A1})$$

This leads to the intermediate results of model (see program code for definitions; perturbations are carried to first order only)

$$\mathbf{x}_1 = (2 + 4\varepsilon_1, 10 + 2\varepsilon_1 + 6\varepsilon_2, 10 + 2\varepsilon_1 + 6\varepsilon_3) \quad (\text{A2})$$

$$\mathbf{x}_2 = (6 + 3\varepsilon_1 + 3\varepsilon_2, 6 + 3\varepsilon_1 + 3\varepsilon_2, 10 + 2\varepsilon_1 + 6\varepsilon_3) \quad (\text{A3})$$

and, after convective adjustment is complete, to

$$\mathbf{x}_3 = (6 + 3\varepsilon_1 + 3\varepsilon_2, 8 + 2.5\varepsilon_1 + 1.5\varepsilon_2 + 3\varepsilon_3, 8 + 2.5\varepsilon_1 + 1.5\varepsilon_2 + 3\varepsilon_3). \quad (\text{A4})$$

Hence,

$$fc = 40.25 + 15.5\varepsilon_1 + 10.5\varepsilon_2 + 15\varepsilon_3, \quad (\text{A5})$$

so that, in accord with the adjoint result,

$$\frac{\partial}{\partial \mathbf{x0}} fc(\mathbf{x0} = (1, 3, 3)) = \frac{\partial}{\partial \varepsilon} fc(\varepsilon = \mathbf{0}) = (15.5, 10.5, 15). \quad (\text{A6})$$

---

Figure A2 (next page). TAMC-generated adjoint of the code displayed in Fig. A1. Included as comments are the values of the variables *after* the statements have been executed.



```

      subroutine admodel( nc, x0, fc, adx0, adfc )
C*****
C** This routine was generated by the **
C** Tangent linear and Adjoint Model Compiler, TAMC 5.0.10 **
C*****
...
C RESET LOCAL ADJOINT VARIABLES
      do ip1 = 1, 3
        adx(ip1) = 0.
      end do
      ady = 0.          ! Nonlocal adjoint variables: adfc=1 adx0=(0,0,0)
C OPEN FILES OF TAPE: storefile
      open(60,file='storefile_1_model_x',ACCESS='DIRECT',RECL=1*3*8)
C ADJOINT COMPUTATIONS
      read(60,REC=1) x          ! Restore from file: x = x3 = (6,8,8)
      adx(3) = adx(3)+3*adfc      ! adx(3)=3
      adx(2) = adx(2)+2*adfc      ! adx(2)=2
      adx(1) = adx(1)+2*adfc*(x(1)-5.5) ! adx(1)=1
      adfc = 0.
c Adjoint of convective adjustment:
      do i = 2, 1, -1
        do ip1 = 1, 3          ! Restore from COMMON block:
          x(ip1) = xh(ip1,1)
        end do                ! i=2: x=x2=(6,6,10) ! i=1:x=x1=(2,10,10)
        if (x(i) .lt. x(i+1)) then ! i=2: 6<10? Yes: ! i=1: 2<10? Yes:
          adx(i) = adx(i)+adx(i+1) ! i=2: adx(2)=5 ! i=1: adx(1)=3.5
          adx(i+1) = 0.          ! i=2: adx(3)=0 ! i=1: adx(2)=0
          adx(i+1) = adx(i+1)+0.5*adx(i) ! i=2: adx(3)=2.5 ! i=1:adx(2)=1.75
          adx(i) = 0.5*adx(i)      ! i=2: adx(2)=2.5 ! i=1:adx(1)=1.75
        endif
      end do                  ! adx = (1.75, 1.75, 2.5)
      do i = 1, 3
        x(i) = x0(i)
      end do                  ! Reset: x = x0 = (1,3,3)
      do i = 1, 3
        ady = ady+adx(i)      !i=1: ady=1.75 !i=2: ady= 3.5 !i=3: ady= 6
        adx(i) = 2*adx(i)*x(i)
      end do                  ! ady= 6 adx=(3.5,10.5,15)
      adx(1) = adx(1)+2*ady*x(1) ! adx(1) = 15.5
      ady = 0.
      do i = 1, 3
        adx0(i) = adx0(i)+adx(i)
        adx(i) = 0.
      end do                  ! adx0 = (15.5,10.5,15)
C CLOSE FILES OF TAPE: storefile
      close(60)
      end

```

## Appendix B: Derivation of kinematic sensitivity in a simple example

As the simplest example of a nonlinear ocean model, with the flow dependent on temperature and salinity and a spatial "extent", consider the 2-box model of the thermohaline circulation originally devised by Stommel (1961) and later used extensively for studying multiple equilibria and atmospheric interactions of the thermohaline circulation (e.g., Marotzke, 1996). The model consists of two ocean boxes, which are well mixed and have constant depth; box 1 represents the high-latitude ocean and box 2 the low-latitude ocean (Fig. B1).  $H_1$  and  $H_2$  are ocean heat gain through the surface, and  $H_S$  is a virtual surface salinity flux, which can be simply related to a surface freshwater flux. All surface fluxes are assumed constant. The conservation equations for the ocean are

$$\dot{T}_1 = H_1 + |q|(T_2 - T_1), \quad (\text{B1})$$

$$\dot{T}_2 = H_2 - |q|(T_2 - T_1), \quad (\text{B2})$$

$$\dot{S}_1 = -H_S + |q|(S_2 - S_1), \quad (\text{B3})$$

$$\dot{S}_2 = H_S - |q|(S_2 - S_1), \quad (\text{B4})$$

The flow strength,  $q$ , is related to the meridional density gradient by a linear law,

$$q = k[\alpha(T_2 - T_1) - \beta(S_2 - S_1)], \quad (\text{B5})$$

where a linear equation of state has been assumed, and  $\alpha$  and  $\beta$  are, respectively, the thermal and haline expansion coefficients. Positive  $q$  means poleward surface flow while negative  $q$  means equatorward surface flow. Here, we can assume  $q$  positive without loss of generality and hence can omit the absolute magnitude signs in the equations. We denote initial temperatures by  $\theta_1$  and  $\theta_2$ , and initial salinities by  $\Gamma_1$  and  $\Gamma_2$ . The initial flow field is

$$q_0 = k[\alpha(\theta_2 - \theta_1) - \beta(\Gamma_2 - \Gamma_1)]. \quad (\text{B6})$$

The model is stepped forward by one timestep according to

$$T_{1/2} = \theta_{1/2} + \Delta t H_{1/2} \pm \Delta t q_0 (\theta_2 - \theta_1), \quad (\text{B7})$$

similarly for salinities. The flow after one timestep is

$$q = k[\alpha(T_2 - T_1) - \beta(S_2 - S_1)]. \quad (\text{B8})$$

Notice that in contrast to (B5), the variables in (B8) are discretized in time. Analogous to the sensitivity analysis in the main text, the output function is chosen to be the heat transport,

$$J \equiv q(T_2 - T_1). \quad (\text{B9})$$




Fig. B1: Geometry of the simple nonlinear example model

Despite the simplicity of the model, analyzing the sensitivity of  $J$  with respect to the initial conditions is fairly involved, owing to the multiple dependencies of the variables. The complete algebra is rather tedious, and we will not present it here; instead, we sketch the development by deriving representative portions of the procedure.

The variation of  $J$  is

$$\delta J = (T_2 - T_1)\delta q + q\delta T_2 - q\delta T_1; \quad (\text{B10})$$

with

$$\delta q = -k\alpha\delta T_1 + k\alpha\delta T_2 - k\beta\delta S_1 + k\beta\delta S_2, \quad (\text{B11})$$

this becomes

$$\begin{aligned} \delta J = & -[q + k\alpha(T_2 - T_1)]\delta T_1 + [q + k\alpha(T_2 - T_1)]\delta T_2 \\ & + k\beta(T_2 - T_1)\delta S_1 - k\beta(T_2 - T_1)\delta S_2. \end{aligned} \quad (\text{B12})$$

Now, we use the timestepping equation (B7) and the variation of the equation (B6) for the initial flow,  $q_0$ , to obtain

$$\begin{aligned} \delta T_1 = & \delta\theta_1 + \Delta t\delta H_1 + \\ & \Delta t\{q_0\delta\theta_2 - q_0\delta\theta_1 + (\theta_2 - \theta_1)[-k\alpha\delta\theta_1 + k\alpha\delta\theta_2 + k\beta\delta\Gamma_1 - k\beta\delta\Gamma_2]\}. \end{aligned} \quad (\text{B13})$$

Equation (B13) and three more of the same style are used to substitute for  $\delta T_1$ ,  $\delta T_2$ ,  $\delta S_1$ , and  $\delta S_2$  in (B12). Then, terms in  $\delta\theta_1$ ,  $\delta\theta_2$ ,  $\delta\Gamma_1$ , and  $\delta\Gamma_2$  are collected to obtain, respectively,  $\partial J / \partial\theta_1$ ,  $\partial J / \partial\theta_2$ ,  $\partial J / \partial\Gamma_1$ , and  $\partial J / \partial\Gamma_2$ , as

$$\begin{aligned} \frac{\partial J}{\partial\theta_1} = & 2k\alpha\Delta t(\theta_2 - \theta_1)[q + k\alpha(T_2 - T_1)] - k\alpha(T_2 - T_1)[1 - 2q_0\Delta t + 2k\beta\Delta t(\Gamma_2 - \Gamma_1)] \\ & - q(1 - 2q_0\Delta t), \end{aligned} \quad (\text{B14})$$

$$\frac{\partial J}{\partial \Gamma_1} = -2k\beta\Delta t(\theta_2 - \theta_1)[q + k\alpha(T_2 - T_1)] + k\beta(T_2 - T_1)[1 - 2q_0\Delta t + 2k\beta\Delta t(\Gamma_2 - \Gamma_1)], \quad (\text{B15})$$

$$\frac{\partial J}{\partial \theta_2} = -\frac{\partial J}{\partial \theta_1}, \quad (\text{B16})$$

$$\frac{\partial J}{\partial \Gamma_2} = -\frac{\partial J}{\partial \Gamma_1}. \quad (\text{B17})$$

The last two equations follow from the antisymmetry of the problem. When (B14) and (B15) are combined according to (15) from the main text, one obtains

$$\frac{\partial J}{\partial \theta_1} + \frac{\alpha}{\beta} \frac{\partial J}{\partial \Gamma_1} = -q(1 - 2q_0\Delta t). \quad (\text{B18})$$

This is indeed the sensitivity of the transport of a purely passive tracer, after one timestep, to the initial condition, as can be seen by inserting the time-stepping equation (B7) into the transport definition (B9) and assuming that  $q$  and  $q_0$  are known a priori.

## Appendix C: The MIT general circulation model

The MIT GCM has been documented extensively in Marshall et al. (1997a,b), so a brief description suffices here. The model is used in its hydrostatic version, which is based on the primitive equations on a sphere under the Boussinesq approximation. Spatial coordinates are longitude, latitude, and height. Conservation of horizontal and vertical momentum, volume, heat, and salt, and an equation of state are written in the form

$$\partial_t \mathbf{u} = -\nabla p / \rho_0 + \mathbf{G}_u \quad (\text{C1})$$

$$\partial_z p = -g\rho \quad (\text{C2})$$

$$-\mathbf{v} + \partial_z w = 0 \quad (\text{C3})$$

$$\partial_t T = G_T \quad (\text{C4})$$

$$\partial_t S = G_S \quad (\text{C5})$$

$$\rho = \rho(S, T, p), \quad (\text{C6})$$

where  $\mathbf{u}$  and  $\nabla$  are the horizontal components of velocity and the gradient operator, respectively,  $T$  is potential temperature, and  $p$  is the deviation of the pressure from that of a resting ocean of constant reference density  $\rho_0$ . The  $G$  stand for all contributions to the tendency terms except pressure gradient forces in (C1), i.e., for advective, Coriolis, metric, forcing, dissipation, and mixing terms including a convective adjustment. The nonhydrostatic version of the model will not be considered here, and the quasi-hydrostatic approximation (Marshall et al. 1997a) would produce only minor modifications in (C2). The model equations are solved on a staggered grid ("C"-grid, Arakawa and Lamb, 1977) using the usual boundary conditions [insulating side walls and bottom, surface fluxes of momentum, heat, and (equivalent) salinity, and either no-slip or free-slip conditions tangential to side-walls or the bottom]. Horizontal velocity, temperature, and salinity are marched forward in time using (C1), (C4), and (C5) in

an Adams-Bashforth discretization. Vertical velocity, density, and pressure are diagnostic variables; the first two of them are readily calculated from (C3) and (C6), respectively.

The computation of pressure is less straightforward and hence roughly sketched here, in particular since this is important for the construction of the adjoint. The deviation of pressure from the reference state is separated into "surface" and "hydrostatic" contributions according to

$$p(\lambda, \phi, z) = p_s(\lambda, \phi) + p_{HY}(\lambda, \phi, z), \quad (C7)$$

where  $p_s$  is the pressure at  $z=0$  (the "surface") and  $p_{HY}$  is calculated from the hydrostatic relationship (C2) using the integration constant  $p_{HY}(z=0)=0$ . (The calculation of  $p_s$  and  $p_{HY}$  does not change even in the nonhydrostatic case, Marshall et al., 1997a,b). An equation for the surface pressure is obtained by integrating the continuity equation (C3) vertically over the entire depth. As shown in Appendix 2 of Marshall et al. (1997b), the end result is an elliptic equation for the surface pressure, which is solved through an iterative conjugate gradient algorithm. The discrete elliptic operator is a symmetric (and hence self-adjoint) matrix, meaning that the adjoint of the procedure calculating the surface pressure is identical to the procedure itself. This leads to considerable simplification, because constructing the adjoint to the iterative procedure would lead to either excessive recomputation or storage, because nonlinear operations (e.g., scalar products and ratios, Press et al., 1992) are involved. However, since the equation for surface pressure and hence its solver are self-adjoint, the subroutine itself is called from the adjoint, with input and output variables appropriately interchanged by the TAMC.

## References

- Arakawa, A., and V. Lamb, Computational design of the basic dynamical processes of the UCLA general circulation model, *Methods in Computational Physics*, 17, 174-267, 1977.
- Bergamasco, A., P. Malanotte-Rizzoli, W. C. Thacker, and R. B. Long, The seasonal steady circulation of the Eastern Mediterranean determined with the adjoint method, *Deep-Sea Res.*, 40, 1269-1294, 1993.
- Berz, M., C. Bischof, G. Corliss, and A. Griewank (Eds.), *Computational Differentiation: Techniques, Applications, and Tools*, 421pp, SIAM, 1996.
- Bindoff, N. L., and T. J. McDougall, Diagnosing climate change and ocean ventilation using hydrographic data, *J. Phys. Oceanogr.*, 24, 1137-1152, 1994.
- Bischof, C., A. Carle, G. Corliss, A. Griewank, and P. Hovland, ADIFOR: Generating derivative codes from FORTRAN programs, *Scientific Programming*, 1, 1-29, 1992.
- Cox, M. D., A primitive equation, 3-dimensional model of the ocean, GFDL Ocean Group Tech. Rep. No. 1, GFDL/Princeton University, 1984.
- Errico, R. M., What is an adjoint model? *Bull. Am. Met. Soc.*, 78, 2577-2591, 1997.
- Errico, R. M., and Vukicevic, Sensitivity analysis using an adjoint of the PSU-NCAR mesoscale model, *Mon. Wea. Rev.*, 120, 1644-1660, 1992.
- Eckert, C., *On Predictability Limits of ENSO*, Ph. D. thesis, 76pp, MPI Hamburg, 1998.
- Eckert, C., R. Giering, and M. Latif, Optimal perturbations of a hybrid coupled model of El Nino, Submitted to *Q. J. R. Meteorol. Soc.*, 1999.
- Farrell, B. F., and A. M. Moore, An adjoint method for obtaining the most rapidly growing perturbation to oceanic flows, *J. Phys. Oceanogr.*, 22, 338-349, 1992.
- Giering, R., *Erstellung eines adjungierten Modells zur Assimilierung von Daten in ein Modell der globalen ozeanischen Zirkulation*, Ph.D. thesis, MPI Hamburg (in German), 1996.
- Giering, R., and T. Kaminski, Recipes for adjoint code construction, *Association for Computing Machinery Transactions on Mathematical Software*, Volume 24, Number 4, 437-474, 1998.



- Griewank, A., Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation, *Optimization Methods and Software*, 1, 35-54, 1992.
- Griewank, A., and G. F. Corliss, (Eds.), *Automatic Differentiation of Algorithms: Theory, Implementation, and Application*, SIAM, 1991.
- Hall, M. C. G., D. G. Cacuci, and M. E. Schlesinger, Sensitivity analysis of a radiative-convective model by the adjoint method, *J. Atmos. Sci.*, 39, 2038-2050, 1982.
- Hersbach, H., Application of the adjoint of the WAM model to inverse wave modeling, *J. Geophys. Res.*, 103, 10,469-10,489, 1998.
- Jayne, S. R., *The Dynamics of Global Ocean Heat Transport Variability*, Ph.D. thesis, 169 pp, MIT-Woods Hole Joint Program in Oceanography, 1999.
- Latif, M., and T. P. Barnett, Causes of decadal climate variability in the North Pacific/North American sector, *Science*, 266, 634-637, 1994.
- Kawase, M., Establishment of deep ocean circulation driven by deep-water production, *J. Phys. Oceanogr.*, 17, 2294-2317, 1987.
- Lee, T., and J. Marotzke, Inferring meridional mass and heat transports of the Indian Ocean by fitting a general circulation model to climatological data, *J. Geophys. Res.*, 102, 10,585-10,602, 1997.
- Lee, T., and J. Marotzke, Seasonal cycles of meridional overturning and heat transport of the Indian Ocean, *J. Phys. Oceanogr.*, 28, 923-943, 1998.
- Maier-Reimer, E., U. Mikolajewicz, and K. Hasselmann, Mean circulation of the Hamburg LSG OGCM and its sensitivity to the thermohaline surface forcing, *J. Phys. Oceanogr.*, 23, 731-757, 1993.
- Marotzke, J., The role of integration time in determining a steady state through data assimilation, *J. Phys. Oceanogr.*, 22, 1556-1567, 1992.
- Marotzke, J., Analysis of thermohaline feedbacks, in *Decadal Climate Variability: Dynamics and Predictability*, edited by D. L. T. Anderson and J. Willebrand, NATO ASI series, 333-378, 1996.

- Marotzke, J., Boundary mixing and the dynamics of three-dimensional thermohaline circulations, *J. Phys. Oceanogr.*, 27, 1713-1728, 1997.
- Marotzke, J., and B. A. Klinger, The dynamics of equatorially asymmetric thermohaline circulations, Submitted to *J. Phys. Oceanogr.*, 1999.
- Marotzke, J., and J. Willebrand, The North Atlantic mean circulation: Combining data and dynamics, in *The Warm Water Sphere of the North Atlantic Ocean*, edited by W. Krauss, 55-90, 1996.
- Marotzke, J., and C. Wunsch, Finding the steady state of a general circulation model through data assimilation: Application to the North Atlantic ocean, *J. Geophys. Res.*, 98, 20,149-20,167, 1993.
- Marshall, J., C. Hill, L. Perelman, and A. Adcroft, Hydrostatic, quasi-hydrostatic and non-hydrostatic ocean modeling, *J. Geophys. Res.*, 102, 5733-5752, 1997a.
- Marshall, J., A. Adcroft, C. Hill, L. Perelman, and C. Heisey, A finite volume, incompressible Navier Stokes model for studies of the ocean on parallel computers, *J. Geophys. Res.*, 102, 5753-5766, 1997b.
- Munk, W., Internal waves, in *Evolution of Physical Oceanography, Scientific Surveys in Honor of Henry Stommel*, B. A. Warren and C. Wunsch (eds.), The MIT Press, Cambridge, MA, 264-291, 1981.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in FORTRAN* (2nd ed.), 963pp, Cambridge University Press, 1992.
- Restrepo, J. M., G. K. Leaf, and A. Griewank, Circumventing storage limitations in variational data assimilation studies, *SIAM J. Sci. Comput.*, 1, 1-16, 1995.
- Robbins, P., and J. Toole, The dissolved silica budget as a constraint on the meridional overturning circulation of the Indian Ocean, *Deep-Sea Res.*, 44, 879-906, 1997.
- Rostaing, N., S. Dalmas, and A. Galligo, Automatic differentiation in Odyssee, *Tellus*, 45A, 558-568, 1993.

- Schiller, A., The mean circulation of the Atlantic Ocean north of 30°S determined by the adjoint method, *J. Mar. Res.*, 53, 433-451, 1995.
- Schiller, A., and J. Willebrand, 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, 453-497, 1995.
- Schröter, J., and C. Wunsch, Solution of nonlinear finite difference ocean models by optimization methods with sensitivity and observational strategy analysis, *J. Phys. Oceanogr.*, 16, 1855-1874, 1986.
- Sirkes, Z., E. Tziperman, and W. C. Thacker, Combining data and a global primitive equation ocean general circulation model using the adjoint method, in *Modern Approaches to Data Assimilation in Ocean Modeling*, edited by P. Malanotte-Rizzoli, Elsevier, 119-145, 1996.
- Stommel, H., Thermohaline convection with two stable regimes of flow, *Tellus*, 13, 224-230, 1961.
- Stammer, D., C. Wunsch, R. Giering, Q. Zhang, J. Marotzke, J. Marshall, and C. Hill, The global ocean circulation estimated from TOPEX/POSEIDON altimetry and a general circulation model, Center for Global Change Science Report No. 49, 40pp, MIT, 1997.
- Talagrand, O., Adjoint models, in *Numerical Methods in Atmospheric Models*, Vol. II, European Centre for Medium-Range Weather Forecasts, 73-91, 1991.
- Talagrand, O., and P. Courtier, Variational assimilation of meteorological observations with the adjoint vorticity equation. I: Theory, *Q. J. R. Meteorol. Soc.*, 113, 1311-1328, 1987.
- Thacker, W. C., Automatic differentiation from an oceanographer's perspective, in *Automatic Differentiation of Algorithms: Theory, Implementation, and Application*, edited by A. Griewank and G. F. Corliss, SIAM, 191-201, 1991.
- Thacker, W. C., and R. B. Long, Fitting dynamics to data, *J. Geophys. Res.*, 93, 1227-1240, 1988.
- Van Oldenborgh, G. J., G. Burgers, S. Venzke, C. Eckert, and R. Giering, Tracking down the delayed ENSO oscillator with an adjoint OGCM, *Mon. Wea. Rev.*, in press, 1999.
- Wunsch, C., *The Ocean Circulation Inverse Problem*, 442 pp, Cambridge University Press, 1996.

- Yu, L., and P. Malanotte-Rizzoli, Analysis of the North Atlantic climatologies using a combined OGCM/adjoint approach, *J. Mar. Res.*, 54, 867-913, 1996.
- Yu, L., and P. Malanotte-Rizzoli, Inverse modeling of seasonal variations in the North Atlantic Ocean, *J. Phys. Oceanogr.*, 28, 902-922, 1998.
- Zhang, K. Q., and J. Marotzke, The importance of open-boundary estimation for an Indian Ocean GCM-data synthesis, *J. Mar. Res.*, in press, 1999.


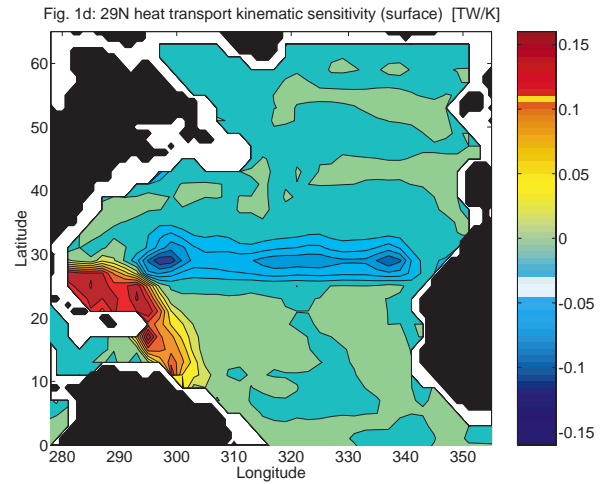

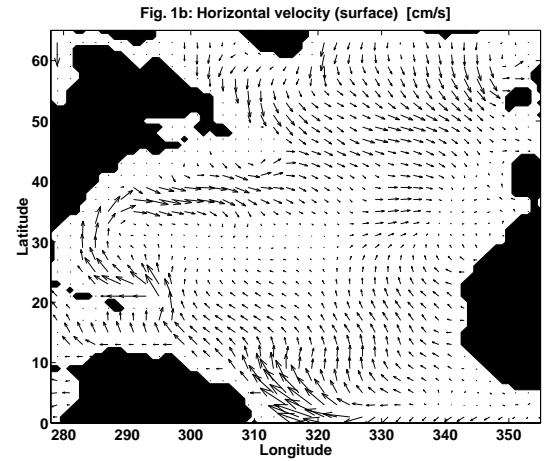
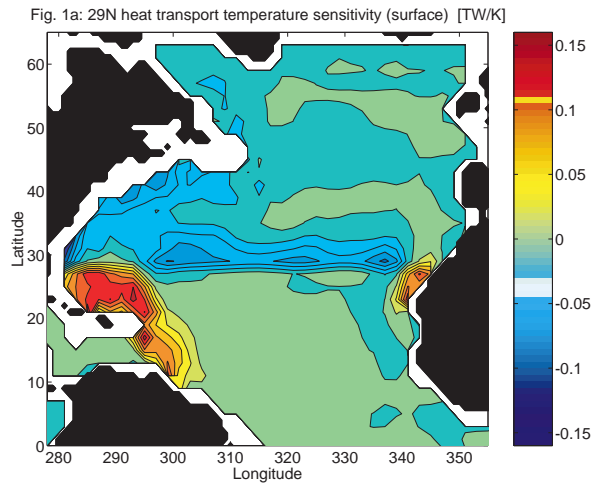


Figure 1.

- (a) Sensitivity of Atlantic heat transport across 29°N (1993 mean), to surface temperature on 1 January 1993. Contour interval is  $0.02 \times 10^{12}$  W/K.
- (b) Annual-mean surface velocity. The reference vector in the upper-left corner denotes 20 cm/s.
- (c) Sensitivity of Atlantic heat transport across 29°N (1993 mean), to surface salinity on 1 January 1993. Contour interval is  $0.04 \times 10^{12}$  W/psu.
- (d) As in (a) but for the kinematic sensitivity.
- (e) Annual-mean velocity for level 2; reference vector denotes 20 cm/s.

Fig. 2a: 29N heat transport temperature sensitivity ( 1160m) [TW/K]




Fig. 2b: 29N heat transport salinity sensitivity ( 1160m) [TW/psu]

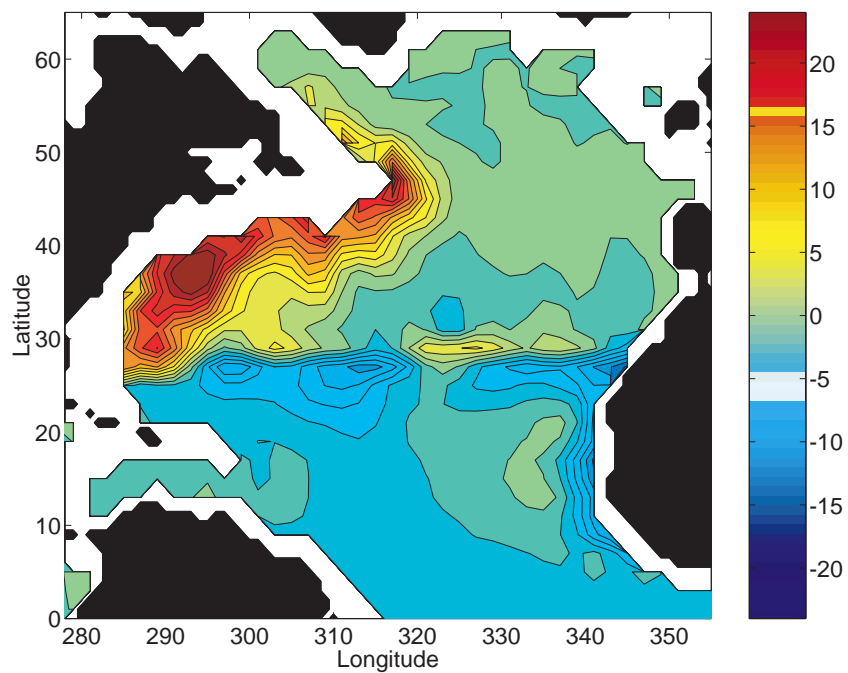


Figure 2.

(a) Sensitivity of Atlantic heat transport across 29°N (1993 mean), to temperature at 1160 m, on 1 January 1993. Contour interval is  $0.5 \times 10^{12}$  W/K.

(b) Sensitivity of Atlantic heat transport across 29°N (1993 mean), to salinity at 1160 m, on 1 January 1993. Contour interval is  $2 \times 10^{12}$  W/psu.

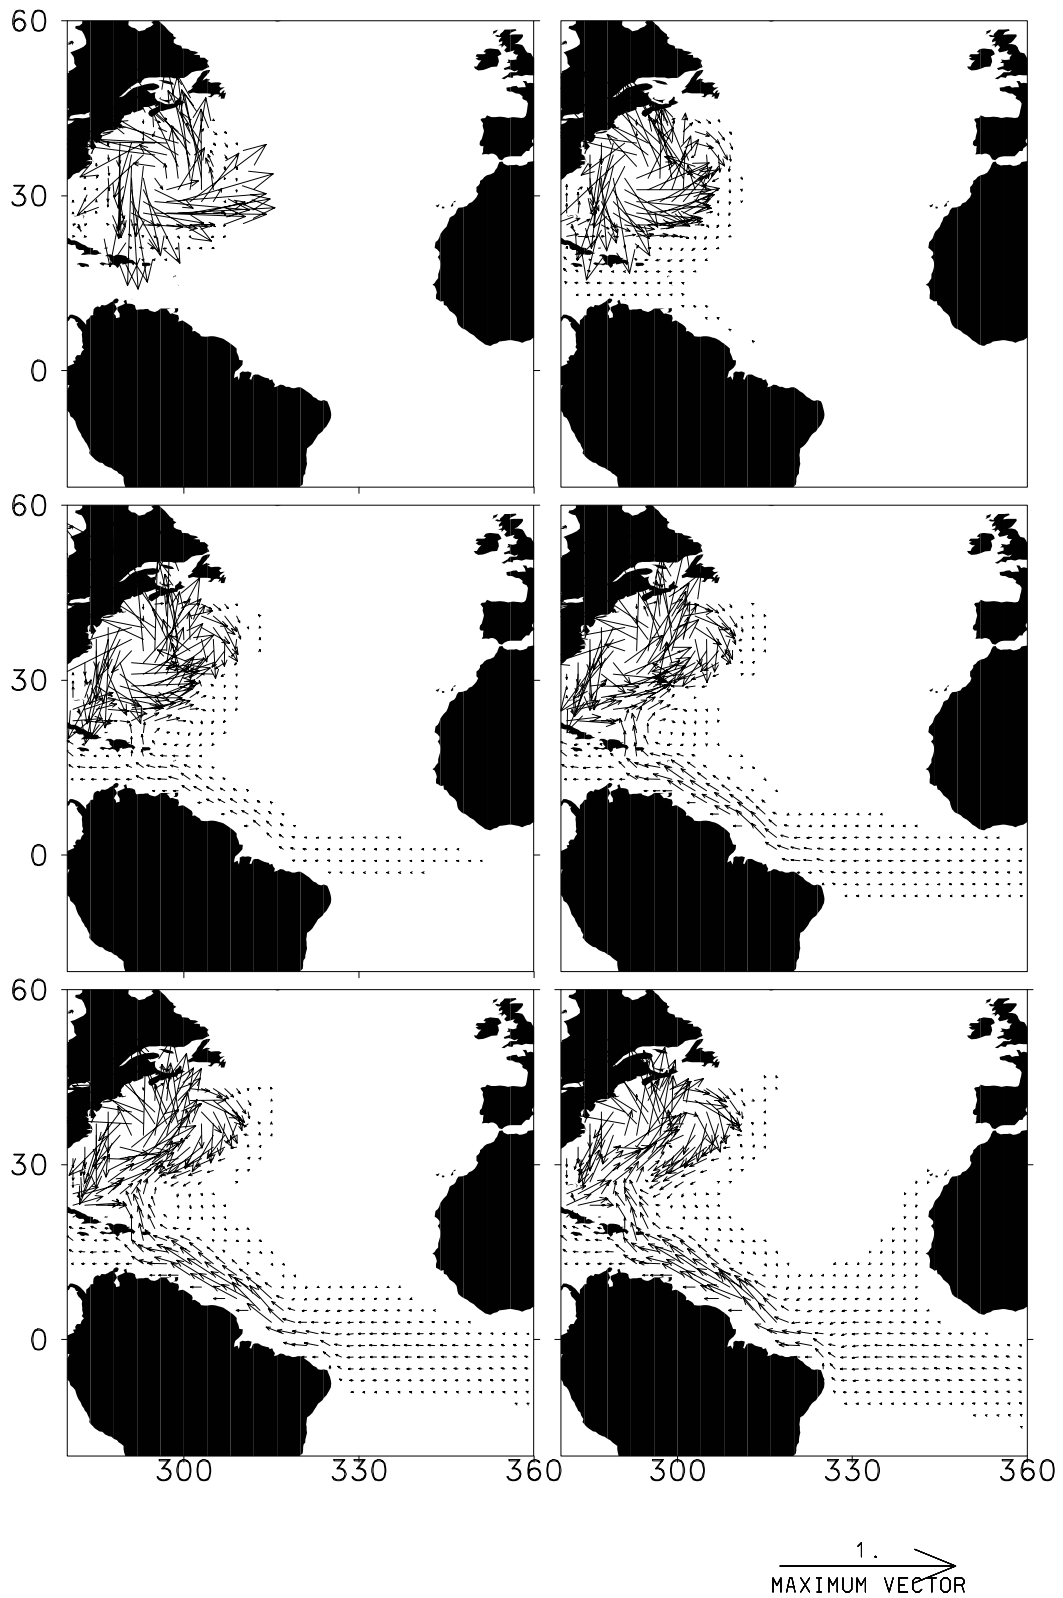


Figure 3. Velocity anomaly at 160 m, at 60-day intervals, beginning with day 30, following a salinity perturbation of amplitude 0.01 at 1160 m, between latitudes 28°N and 38°N, and between longitudes 68°W and 58°W. Reference vector denotes 1 cm/s.

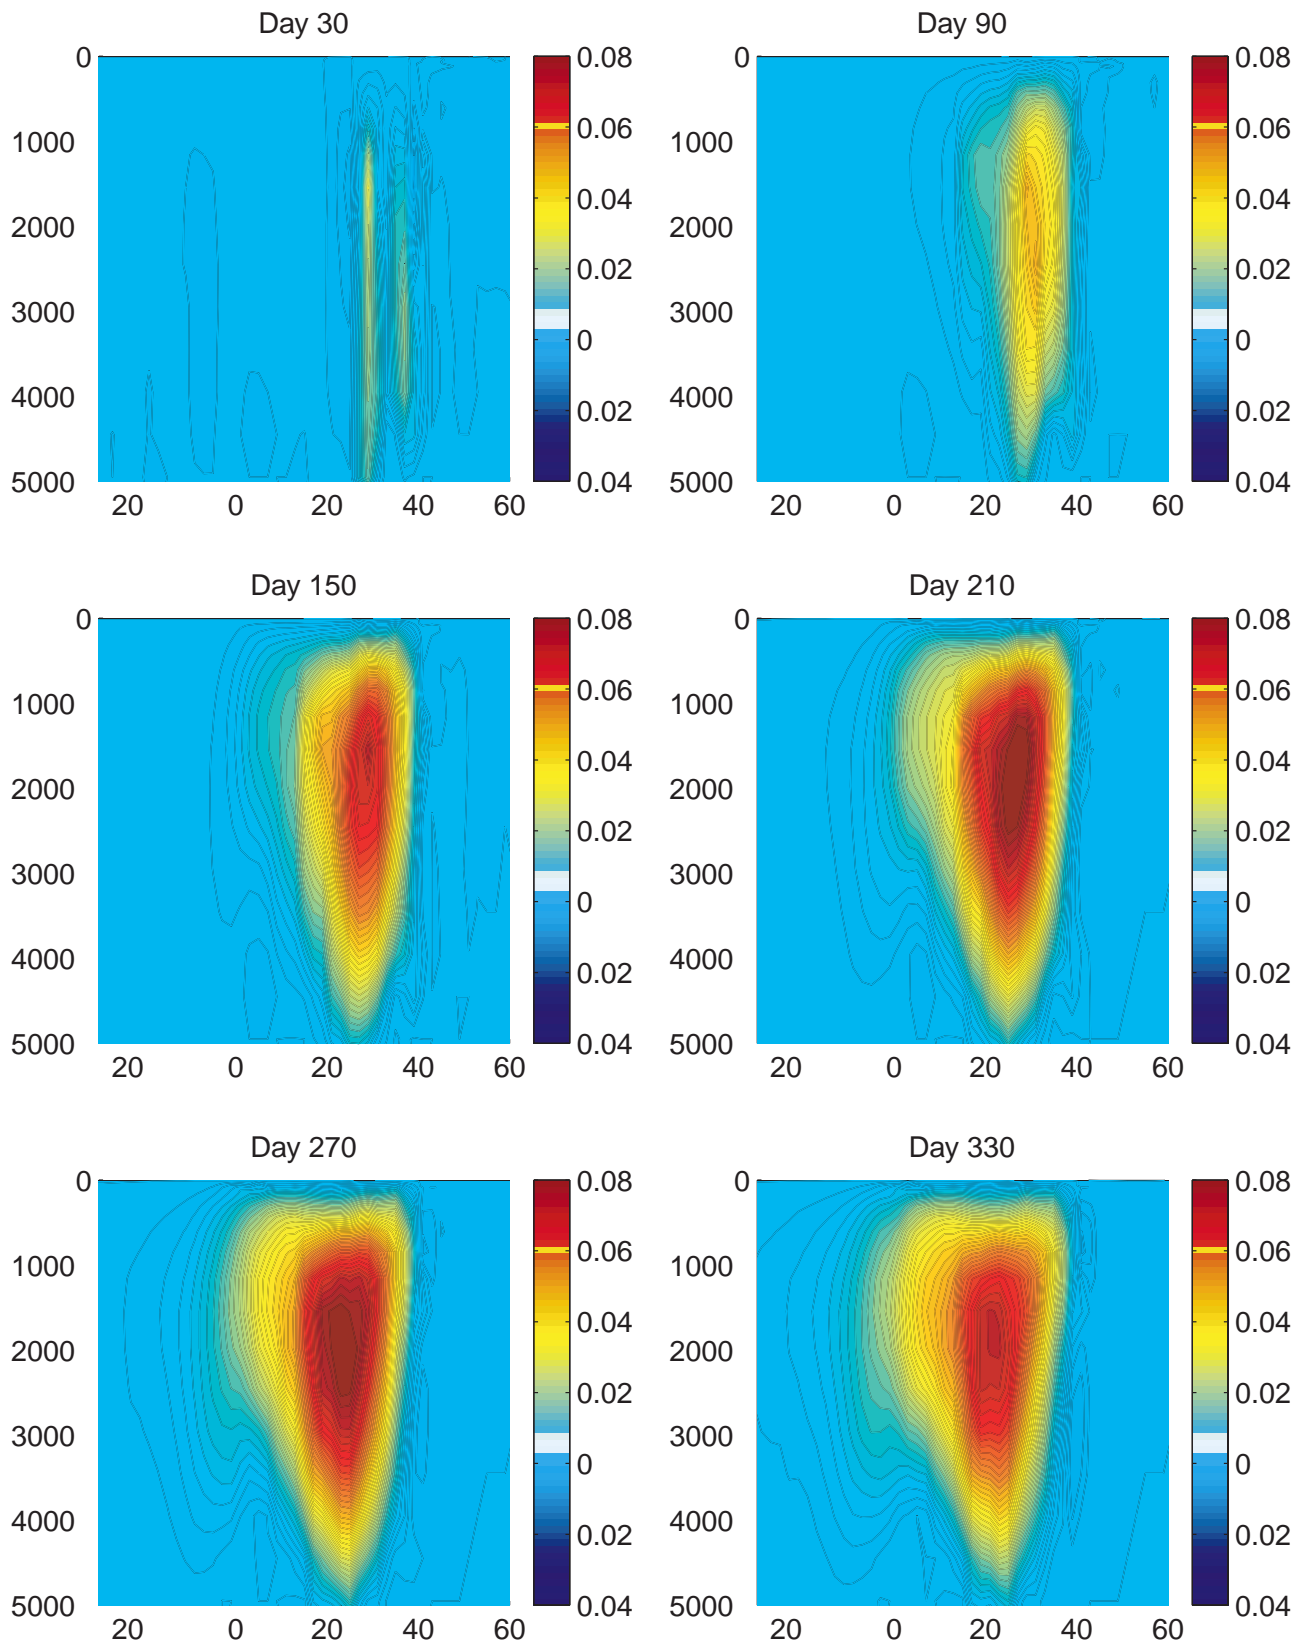


Figure 4. Perturbation in Atlantic meridional overturning stream function (in Sv), at 60-day intervals, beginning with day 30, following a salinity perturbation of amplitude 0.01 at 1160 m, between latitudes 28°N and 38°N, and between longitudes 68°W and 58°W. Positive values indicate clockwise anomalous transport.






Figure 5. Time series of Atlantic heat transport anomalies, at 35°N (solid), 29°N (dashed), 24°N (dash-dotted), and 9°N (dotted), following a salinity perturbation of amplitude 0.01 at 1160 m, between latitudes 28°N and 38°N, and between longitudes 68°W and 58°W.

Fig. 6a: North Atlantic,  $f\rho_0/g\sigma_{\theta}$ Psicurvature




Fig. 6b: NA EW insitu density difference




Fig. 6c: North Atlantic, EW delrho minus  $f\rho_0/g\sigma_{\theta}$ Psicurvature




Figure 6. Diagnostics from the global solution of Stammer et al. (1997) over the North Atlantic sector.

(a) Left-hand side of Eq. (19),

$$f\rho_0/g\sigma_{\theta}$$

(b) Right-hand side of Eq. (19),

$$\rho_E - \rho_W$$

(c) Residual of Eq. (19),

$$(\rho_E - \rho_W) - f\rho_0/g\sigma_{\theta}$$

The contour interval is  $0.1 \text{ kg m}^{-3}$ , and the zero contour is drawn in red.