Modern Approaches to Data Assimilation in Ocean Modeling
edited by P. Malanotte-Rizzoli

Combining Data and a Global Primitive Equation Ocean General Circulation Model Using the Adjoint Method

Z. Sirkes[a], E. Tziperman[b] and W. C. Thacker[c]

[a]Center for Marine Sciences, The University of Southern Mississippi, Stennis Space Center, MS 39529-5005

[b]Environmental Sciences, The Weizmann Institute of Science, Rehovot 76100, Israel

[c]Atlantic Oceanographic and Meteorological Laboratory, Miami FL 33149 USA

Abstract

A Primitive Equation Ocean General Circulation Model (PE OGCM) in a global configuration similar to that used in coupled ocean-atmosphere models is fitted to climatological data using the adjoint method. The ultimate objective is the use of data assimilation for the improvement of the ocean component of coupled models, and for the calculation of initial conditions for initializing coupled model integrations. It is argued that oceanic models that are used for coupled climate studies are an especially appropriate target for data assimilation using the adjoint method.

It is demonstrated that a successful assimilating of data into a fully complex PE OGCM critically depends on a very careful choice of the surface boundary condition formulation, on the optimization problem formulation, and on the initial guess for the optimization solution. The use of restoring rather than fixed surface-flux boundary conditions for the temperature seems to result in significantly improved model results as compared with previous studies using fixed surface-flux boundary conditions. The convergence of the optimization seems very sensitive to the cost formulation in a PE model, and a successful cost formulation is discussed and demonstrated. Finally, the use of simple, suboptimal, assimilation schemes for obtaining an initial guess for the adjoint optimization is advocated and demonstrated.

1 Introduction

Oceanographic data assimilation is a rapidly evolving field with very diverse objectives and hence many different possible methodologies to address these objectives. Two of the main purposes of combining ocean models and data are the improvement of ocean models, and the calculation of an optimal estimate of the oceanic state, based on both model dynamics and the available data (Malanotte-Rizzoli and Tziperman, Chapter 1 of this book). These two objectives are very general, and apply to a wide spectrum of

models, from high resolution to coarse, and a variety of uses can be found for the optimal ocean state estimated by data assimilation or inverse studies.

One class of ocean models for which these two objectives are especially relevant and important consists of the ocean models used in coupled ocean-atmosphere model studies. Model improvement in this context refers to the need to improve these ocean models, including their sub-grid scale parameterizations, their poorly known internal parameters such as various eddy coefficients, the surface boundary forcing fields which are often known with large uncertainty, etc. Data assimilation may be used to find those model parameters that result in a better fit of the model results to observations, and therefore in an improved performance of the model when run within a coupled ocean-atmosphere model. The state estimation problem in this context refers to the need to find "optimal" initial conditions for coupled model climate simulations. Such initial conditions, based on both the model dynamics and the oceanic observations, would hopefully result in better climate forecasts.

The combination of OGCMs and oceanographic data for the above purposes can be formulated as an optimization problem. Such an optimization would search for a set of model parameters and for an optimal ocean state which together satisfy the model equations and fit the available data as well as possible. This is done by formulating a cost function to be minimized, which measures the degree to which the model equations are satisfied, as well as the distance to the data. The minimization of this cost function is a most complex nonlinear optimization problem, requiring very efficient methodologies. A common solution for such large scale optimization problems is to use gradient-based iterative algorithms such as the conjugate gradient (c-g) algorithm. The minimization is carried out in a huge parameter space comprising of all model parameters and of the 3D model initial conditions for the temperature, salinity and velocities. The efficient estimation of the gradient of the cost function with respect to these many parameters is a crucial part of the methodology. This is done using a numerical model based on the adjoint equations of the original model equations. Thus this optimization approach is often referred to as the "adjoint method" (e.g. [1]–[4]).

The adjoint method is very efficient compared to other ways of estimating the gradient of the cost function, but is still computationally intensive. Given the power of todays computers, the adjoint method is therefore adequate primarily for medium to coarse resolution models. Due to the very high computational cost of coupled models, they are also presently limited to a fairly coarse resolution. Clearly the data assimilation problems related to coupled models are therefore an excellent match to the capabilities of the adjoint method. Moreover, it may be expected that as available computers become more powerful and allow higher resolution coupled ocean-atmosphere models, the new computational resources will also enable the use of such higher resolution models with the adjoint method.

We would like to present here a step towards the ultimate goal of using the adjoint method with the ocean component of coupled ocean-atmosphere models. We still cannot claim to having improved the model or having produced optimal initial conditions, but hopefully have made some progress. Inverting a three dimensional GCM (that is, assimilating data into a three dimensional GCM using an optimization approach) is basically a very technical problem, yet we will demonstrate here that a successful application of

the adjoint method to this problem requires a very good understanding of both the ocean circulation dynamics and of the technical issues involved. In fact, we try to emphasize here precisely those issues that require the understanding of the dynamics in order to formulate and successfully solve the inverse problem of combining ocean GCMs and data. The use of a fairly coarse resolution model here implies, of course, that we do not attempt here to produce a highly realistic simulation of the oceanic state. Rather, the above objectives are all related to the ultimate improvement of coupled ocean-atmosphere model simulations whose main tool is similar coarse-resolution models.

Although the objective of combining 3D ocean climate models with data is of obvious interest, it is surprising to realize that there have only been very few efforts so far trying to apply the adjoint method to full complexity 3D ocean models. Tziperman et al. [5, 6] have examined the methodology using simulated data and then real North Atlantic data; Marotzke [7], and Marotzke and Wunsch [8] (hence after MW93) have considerably improved on the methodology and analyzed a North Atlantic model; Bergamasco et al. [9] used the adjoint method in the Mediterranean Sea with a full PE model, and Thacker and Raghunath [10] have examined some of the technical challenges involved in inverting a PE model. This relatively small number of studies has a simple reason: the technical difficulties in constructing an adjoint model of a full GCM are almost overwhelming. Fortunately, this difficult task was successfully tackled by Long, Huang and Thacker [11], who have generously made the results of their efforts available to others and the present study is a direct outcome of their efforts. (The adjoint code of [11] was modified here to be consistent with the global configuration and eddy parameterizations used in this study, so that the adjoint code used here is the precise adjoint of our finite difference global model). All of the above works use the the model equations as "hard" constraints. This implies that errors in the model equations are not considered explicitly. It is worthwhile noting that adjoint models can also be used for different data assimilation approaches than used here [12, 13].

Within the framework of using climate models with the adjoint method, this study has three specific objectives. First, we would like to investigate the issue of model formulation for such optimization problems, and in particular the surface boundary condition formulation. There are two commonly used surface boundary condition formulations. One is fixed-flux conditions, in which the heat flux is specified independently of the model SST. The second is restoring boundary conditions in which the heat flux is calculated by restoring the model SST to a specified temperature distribution (possibly the observed SST). Previous applications of the adjoint method to 3D GCMs used the fixed-flux formulation in an effort to calculate the surface fluxes that results in a good fit to the temperature observations. However, the optimal solution was characterized by large discrepancies, of up to 6 degrees, with the observed SST [6, 8]. Tziperman et al. [6] suggested that this discrepancy is the result of using flux boundary conditions, rather than restoring conditions that are normally used in ocean modeling. MW93 [8] suggested that this discrepancy might be a result of the use of a steady model which lacks the large seasonal signal in the SST, and that this problem might be resolved using a seasonal model. We explain and demonstrate below that using restoring boundary conditions, is better motivated physically as well as seems to eliminate the large SST discrepancies observed in previous optimization studies (section 4.2).

Our second objective is to examine various possibilities for the formulation of a cost function measuring the success of the optimization problem and their influence on the success of the optimization. Finally, we shall discuss and demonstrate methods for increasing the efficiency of the adjoint method by initializing the gradient based optimization with solutions obtained using simpler, sub-optimal, assimilation methodologies.

Ocean models presently used in coupled ocean-atmosphere studies are coarse, non-eddy-resolving, yet usually include the seasonal cycle. Faithful to our philosophy of trying to use the same models for data assimilation studies we should have used a seasonal model, and indeed work is underway to do just that. In this present work, however, we have made several steps forward going from basin to global scale, and from a simplified 3D GCM to a full PE model. These steps turned out to involve a sufficient number of new challenges, so we have decided to maintain the steady state assumption, and progress to a global PE seasonal model only at a following stage. We expect that the lessons learned from the steady state problem will be very useful at the next stage, as time dependent, presumably seasonal, models are inverted.

In the following sections we describe the model and data used in this study (section 2), discuss in detail the formulation of the optimization problem (section 3). We then present the results of the model runs carried out here (section 4), and finally discuss the lessons to be learned for future work and conclude in section 5.

2 Model and data

Ultimately, our objective is to use data to improve ocean models used in climate simulations; therefore the model used for the optimization study needs to be the same model that can be run independently in a simulation mode. This determines many of our choices concerning the model and surface boundary condition formulation.

We use the GFDL PE model, derived from the model of Bryan [14], with later modifications by Semtner [15] and Cox [16], in a coarse resolution global configuration similar to that of Bryan and Lewis [17], with the main difference being that the Arctic ocean is not included in our model. The model's geometry and resolution are also similar to those presently used by coupled ocean-atmosphere models. The model's geometry is shown in Fig. 1a. The model has 12 vertical levels, with the eddy mixing coefficients for the temperature and salinity varying with depth according to the scheme proposed by Bryan and Lewis [17]. The mixing coefficients for the temperature and salinity are given by $A_H(k) = r_H(k)2 \times 10^7 cm^2/sec$ in the horizontal direction, and $A_V(k) = r_V(k) \times 0.305 cm^2/sec$ in the vertical direction, where $r_H(k)$ and $r_V(k)$ are given in Table 1. The momentum mixing coefficients are 25×10^8, and 50 cm^2/sec in the horizontal and vertical directions correspondingly.

The choice of surface boundary condition formulation turns out to be a crucial factor in the optimization problem we have set out to solve here. We explain and demonstrate below that using restoring boundary conditions, rather than the fixed-flux formulation used previously is better motivated physically as well as eliminates the large SST discrepancy observed in previous optimization studies (section 4.2). Under restoring boundary conditions the model is driven with an implied air-sea heat flux H^{SST} that is calculated

Table 1
Model levels and horizontal and vertical mixing coefficients.

level (k)	depth (m)	horizontal mixing factor (r_H)	vertical mixing factor (r_V)
1	25.45	1.0000	1.000
2	85.10	0.8923	1.003
3	169.50	0.7794	1.007
4	295.25	0.6620	1.015
5	482.80	0.5475	1.028
6	754.60	0.4482	1.053
7	1130.65	0.3733	1.109
8	1622.40	0.3218	1.288
9	2228.35	0.2853	2.904
10	2934.75	0.2553	4.048
11	3720.90	0.2274	4.193
12	4565.55	0.2000	4.244

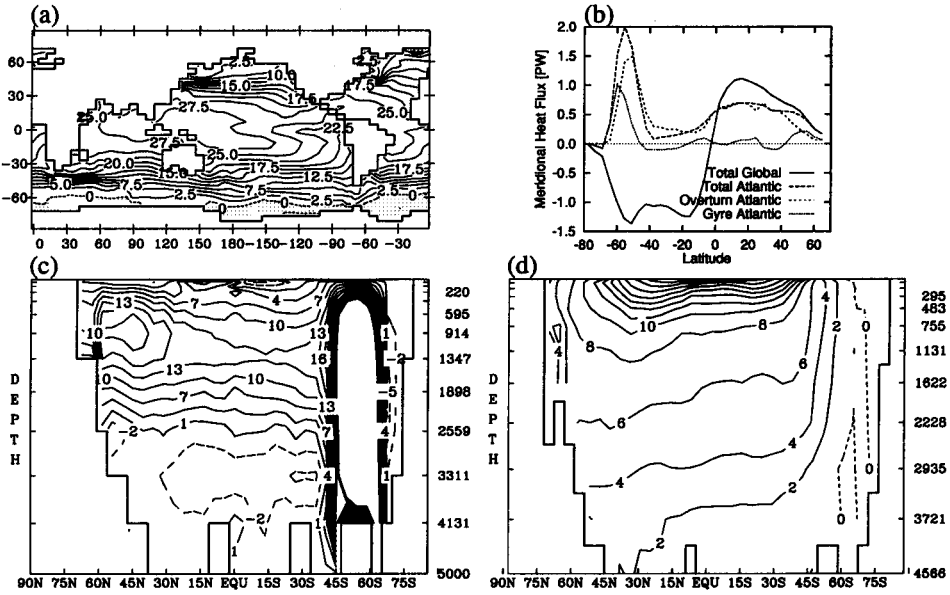

Figure 1: The steady state model solution for the surface temperature obtained without the use of interior data: (a) Model geometry and the sea surface temperature at steady state. Contour intervals are 2.5°C. Negative areas are dotted. (b) Total meridional heat flux for the global ocean (solid), for the Atlantic ocean (dash), overturning circulation contribution to the meridional heat flux (short-dash) and gyre contribution of the meridional heat flux (dot). (c) North Atlantic meridional stream function. (d) Temperature section through the North Atlantic model sector solution.

at time step n from the model upper level temperature, $T^n_{i,j,k=1}$, and the temperature data at this depth, $T^d_{i,j,k=1}$, (where the indices i, j denote horizontal grid point location, and k vertical level) as follows

$$H^{SST,n}_{ij} = \rho_0 C_p \gamma^T \Delta z_1 (T^d_{i,j,k=1} - T^n_{i,j,k=1}). \tag{1}$$

The restoring coefficient γ^T has units of one over time, C_p is the heat capacity of sea water, ρ_0 is a constant reference density, and Δz_1 denotes the thickness of the upper model level. Similarly, an implied fresh water flux is calculated from the difference of the model surface salinity and the surface salinity data,

$$[E\text{–}P]^{SSS,n}_{ij} = \gamma^S \Delta z_1 (S^d_{i,j,k=1} - S^n_{i,j,k=1})/S_0, \tag{2}$$

where S_0 is a constant reference salinity used to convert the virtual salt flux to an implied fresh water flux. In our runs, where $\Delta z_1 = 50m$, we use $\gamma^T = 1/30\text{days}^{-1}$ and $\gamma^S = 1/120\text{days}^{-1}$. Following Hirst and Cai [21], we restore our model surface temperature and salinity to (-1.9° C, 34.84ppt) in the North Atlantic portion of our model, at two grid points only, located at (68.9N; 7.5W and 11.25W) using restoring coefficients that are 10 times larger than those used elsewhere. This results in an improved simulation of the NADW formation and spreading. Finally, as the Mediterranean Sea is not included in our model, a sponge layer is used at two grid points near the Mediterranean outflow region, in which model temperature and salinity are restored to the Levitus data at all depths.

The steady state model results obtained by integrating the model for about 1500 years (without data assimilation) are shown in Fig. 1 Depicted are the surface temperature field (Fig. 1a), the global and North Atlantic meridional heat flux (Fig. 1b, see [17, 18] for the meridional heat flux decomposition used here) the North Atlantic overturning circulation (Fig. 1c) and a temperature section through the North Atlantic ocean (Fig. 1d). Note that the overturning circulation is about 16Sv at 30N, close to the commonly assumed value of about 18 Sv there. This is due to the strong restoring at the two northern surface grid points mentioned above, without which the overturning at 30N reduces by about 25%.

The "data" used in this study are the annually averaged temperature and salinity analysis of Levitus [22]; the annually averaged climatologies of heat flux from Esbensen and Kushnir [20], of fresh water flux ($[E\text{–}P]$) from Baumgartner and Reichel [23] and of winds from Hellerman and Rosenstein [24]. All of these are, in fact, gridded analyses rather than raw data. While it is clearly more convenient to use such analyses, future applications of the adjoint method may use the raw data instead. The use of the raw observations, together with detailed error information, may result in more reliable results and better error statistics for the model solution than is possible here.

3 Optimization Problem

One of the main lessons that have been learned over the past few years while trying to combine 3D ocean models and data, is that the correct formulation of the inverse problem is of crucial importance to the success of the optimization. Much thought and understanding of the dynamics should enter the process of posing the optimization problem. This

process includes the choice of a cost function that measures the optimization success and that needs to be minimized, the specification of the initial guess for the optimization solution from which the iterative minimization should begin, and the choice of control variables which are varied in the optimization. We now examine each of these steps in some detail. The results of an optimization formulated according to the ideas presented in this section are shown and discussed in section 4.

3.1 Cost Function

Once the data and model have been specified, the next stage in the formulation of the inverse problem is to specify a measure for the success of the optimization, i.e., the cost function to be minimized. The cost function measures both the fit of the model results to the data, and the degree to which the dynamical constraints are satisfied. A given dynamical constraint can be formulated in many different ways. It has been shown for simpler GCMs that the ability of the optimization to minimize the cost function critically depends on the precise form of the cost function [7]. We find that a Primitive Equations model is even more sensitive to the precise cost formulation.

Let us consider the various dynamical and data constraints and the possibilities of specifying them within a cost function to be minimized. Begin from the dynamical constraints, which in our case are the requirement for the solution to be as close as possible to a steady state of the model equations. This condition may be obtained by minimizing a measure of the deviation of the model from a steady state solution. Tziperman and Thacker [4] and then Tziperman et al [5, 6] have suggested to minimize the finite difference form of $(\partial T/\partial t)^2$, obtained by stepping the model from the initial conditions $T_{ijk}^{n=0}$ a single time step to $T_{ijk}^{n=1}$, and minimizing the sum of terms such as $(T_{ijk}^{n=1} - T_{ijk}^{n=0})^2$. This seems reasonable, and worked for a QG model [4], yet encountered major difficulties when applied to a 3D model [5, 6]. Marotzke [7], in an important contribution, suggested to use instead $(T_{ijk}^{n=N} - T_{ijk}^{n=0})^2$, such that the model integration time $N\Delta t$ corresponds to the time scale of physically relevant processes in the model (e.g. O(10 years) for a problem involving the upper ocean, longer time scales for the deeper ocean, etc). Marotzke's suggestion resulted in most significantly improved convergence of the optimization, as presented in both Marotzke [7] and MW93 [8].

A useful perspective for evaluating the usefulness of a given formulation of the dynamical constraints in the cost function is the conditioning of the resulting optimization problem. The cost surface in parameter space near the cost minimum is of a bowl shape. The bowl may be nearly flat in some directions and very steep in others. If such a discrepancy occurs, the optimization is said to be ill conditioned [25]. An ill conditioned optimization may stall and not progress towards the minimum even after many iterations of the minimization algorithm. If the steepness of the cost surface is nearly even in all directions, the optimization is said to be well conditioned, and the solution is found within a few iterations. The conditioning issue was discussed in detail in Tziperman et al. [6], where the analysis pointed out to some possible ways of improving the conditioning using various formulations for the cost function. The conditioning of the steady penalties of temperature and salinity for the PE model used here is examined in section 3.1.1. For a primitive equation model such as used here, there are additional considerations con-

cerning the form of the dynamical constraints for the velocity field which turns out to be most crucial for the success of the optimization, and these are discussed in section 3.1.2. Finally, the cost formulation for the penalties requiring the model heat flux (and fresh water flux) to be close to the observations is discussed in section 3.1.3.

3.1.1 Dynamical constraints for temperature and salinity

In order to evaluate the conditioning of the dynamical constraints, we have plotted them together with the data penalties along a somewhat arbitrary section between two points in parameter space. The two points correspond to two choices for the 3D temperature, salinity, velocity and stream function initial conditions. The two points were obtained by running a few iterations of the optimization algorithm once from the steady state solution and once from a robust diagnostic solution ([27]; see below for details). The plotted cost function is of the form

$$\begin{aligned} J(T^0, S^0, u^0, v^0, \psi^0) &= \sum_{ijk} \left[\overline{W}_k^T (T_{ijk}^{n=N} - T_{ijk}^{n=0})^2 + W_k^T (T_{ijk}^d - T_{ijk}^{n=0})^2 \right] \\ &+ \sum_{ijk} \left[\overline{W}_k^S (S_{ijk}^{n=N} - S_{ijk}^{n=0})^2 + W_k^S (S_{ijk}^d - S_{ijk}^{n=0})^2 \right] \end{aligned} \tag{3}$$

where $T^0 = T^{n=0}$ is the initial condition for temperature, and similarly for S^0, u^0, v^0, ψ^0. The precise choice of the weights is discussed below. Let the two points in parameter space be $\mathbf{x_1}, \mathbf{x_2}$. Then the various terms of the cost function were evaluated and plotted along the straight line in parameter space connecting these two points at $\mathbf{x} = \mathbf{x_1} + r(\mathbf{x_2} - \mathbf{x_1})$, with r varying from $r = -0.6$ to $r = 1.6$ at intervals of $\Delta r = 0.1$, using an integration time of $N\Delta t = 2$ years. The results are shown in Fig. 2.

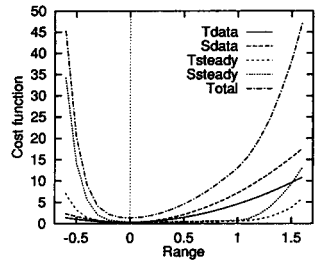

Figure 2: Cost function along a section in parameter space. Shown are the steady temperature penalties (short-dash); steady salinity penalties (dot); data temperature penalties (solid); data salinity penalties (dash) and the total cost (dash-dot).

The data penalties along the section are clearly simple parabolas. The dynamical constraints for the temperature and salinity, however, have a very nonlinear character, reflecting the nonlinearity of the model equations used to obtain $T^{n=N}$ from $T^{n=0}$. These terms of the cost function are nearly flat between the two points ($r = 0$ and $r = 1$), and then rise very rapidly outside of the interval. In particular, going from the minimum point

at $r \approx 0$, corresponding to the optimization started at the robust diagnostic solution, to $r = 1$, the data penalties increase significantly, indicating a very significant change in the temperature and salinity fields (Fig. 2). Yet the steady penalties hardly change. This seems to indicate a possible ill conditioning of the dynamical constraints, so that they are not well constraining the optimization which would feel mostly the variation of the data penalties along this section. As these dynamical constraints were evaluated using a 2 year integration time, they are presumably much better conditioned than using a single time step or other short integration time. It seems likely, however, that a more thoughtful formulation of the steady penalties may result in an even better conditioned form of the dynamical constraints.

While there is probably room for improvement in the cost formulation, we wish to emphasize that an optimization problem formulated using a cost function similar to the above is, in fact, successfully solved below (section 4).

3.1.2 Dynamical constraints for velocities and barotropic stream function

Under the primitive equation approximation, there are 5 prognostic fields: temperature, salinity, two horizontal baroclinic velocities and the barotropic stream function. In principle, each of these needs to be required to be at a steady state if such a model solution is desired. We have attempted to do this by adding to the cost function 3 terms such as

$$J_{steady\ u,v,\psi} = \sum_{ijk} \overline{W}^U_k \left((u^{n=N}_{ijk} - u^{n=0}_{ijk})^2 + (v^{n=N}_{ijk} - v^{n=0}_{ijk})^2 \right) + \sum_{ij} \overline{W}^\psi (\psi^{n=N}_{ij} - \psi^{n=0}_{ij})^2. \quad (4)$$

Several optimizations were performed using this formulation, starting from the data, from the steady state or from a robust diagnostic solution (see next section). In all cases, the optimization efficiently reduced the steady penalties for the velocities and stream function using minute changes to the temperature and salinity, leaving the steady and data penalties for the temperature and salinity nearly unchanged. This could, of course, be due to a poor choice of the cost weights, although we feel that we have come up with a reasonable choice for them (see Table 2 and discussion below).

Note that given the density stratification, the velocity field in a rotating fluid must adjust to the density stratification within a few pendulum days. Therefore, there seems to be no point in penalizing the velocity field separately from the temperature and salinity fields. Once the temperature and salinity penalties are minimized by the optimization, the velocity field just adjusts to the optimal stratification. Indeed, removing the velocity and stream function penalties resulted in an immediate improvement of the convergence of the optimization, and the steady velocity penalties are therefore not used in this study.

It is interesting to note that this problem did not arise in previous studies such as [5]-[8], because they were all using a simpler GCM in which the momentum equations were diagnostic, and therefore did not require separate steady velocity penalties. The issue of dynamical constraints for the velocity field in a PE model is one of the new insights we seem to have gained by going to a full PE model in the present study.

3.1.3 Constraints for surface flux data

In all previous applications of the adjoint method to a 3D GCM, the model was formulated using fixed surface-flux boundary conditions for the temperature and salinity. Then an optimal flux which minimizes the cost function was sought using the optimization algorithm. This involved penalizing the deviations of the optimal heat flux, H, from the heat-flux data, H^d as follows:

$$J(H) = \sum_{ij} \left[W^H (H^d_{i,j} - H_{i,j})^2 \right]. \tag{5}$$

Note that the cost function in this case is an explicit function of the heat flux H which is used as a control variable to be directly calculated in the optimization. In previous applications of the adjoint method, this formulation resulted in very large discrepancies between the model surface temperature and the observed one, in spite of the data penalties in the cost function.

In this study, we wish to examine the suggestion of Tziperman et al. [6] that restoring conditions may resolve the problem of large SST discrepancies, by using a cost function of the form

$$J(SST) = \sum_{ij} \left[W^H \left(H^{SST,n=0}_{ij} - H^d_{i,j} \right)^2 \right]. \tag{6}$$

where $H^{SST,n=0}$ is the restoring conditions heat flux (1) at the beginning of the run, and the control variable is the surface temperature, rather than the flux itself.

Let us now write the complete cost function (selected parts of this cost function are used in the optimization presented below):

$$\begin{aligned}
J(T^{n=0}, S^{n=0}, u^{n=0}, v^{n=0}, \psi^{n=0}) \quad &= \quad \sum_{ij} \left[\overline{W}^{\psi} (\psi^{n=N}_{ij} - \psi^{n=0}_{ij})^2 \right] \\
&+ \sum_{ijk} \left[\overline{W}^T_k (T^{n=N}_{ijk} - T^{n=0}_{ijk})^2 + \overline{W}^S_k (S^{n=N}_{ijk} - S^{n=0}_{ijk})^2 \right] \\
&+ \sum_{ijk} \left[W^T_k (T^d_{ijk} - T^{n=0}_{ijk})^2 + W^S_k (S^d_{ijk} - S^{n=0}_{ijk})^2 \right] \\
&+ \sum_{ijk} \left[\overline{W}^U_k (u^{n=N}_{ijk} - u^{n=0}_{ijk})^2 + \overline{W}^U_k (v^{n=N}_{ijk} - v^{n=0}_{ijk})^2 \right] \\
&+ \sum_{ij} \left[W^H \left(H^{SST,n=0}_{ij} - H^d_{i,j} \right)^2 + W^{[E-P]} \left([E-P]^{SSS,n=0}_{ij} - [E-P]^d_{i,j} \right)^2 \right].
\end{aligned} \tag{7}$$

The data weights for the temperature, salinity and velocities are the inverse square error in the temperature data as estimated in Table 2, normalized by the number of model's grid points, M. The steady penalties require that the drift in temperature (or salinity) during a period of 15 years is equal to the assumed data error. The integration time of 2 years used to evaluate the steady penalties dictates the following choice for the steady penalties [6, 7]:

$$\overline{W}_k^T = \frac{1}{M}\left(\frac{2 \text{ years} \times \epsilon_k(T)}{15 \text{ years}}\right)^{-2}. \tag{8}$$

The steady penalties for the velocities and stream function are similarly calculated from $\epsilon_k(U)$ and $\epsilon(\psi)$ given in Table 2. The errors in the flux data were assumed to be 50Watts/m^2 for the climatological heat flux and 50cm/yr for the evaporation minus precipitation data [18].

The above choice of weights implied uncorrelated error statistics. For correlated errors, non diagonal weight matrices must be used. The errors in oceanic observations are not only correlated, but the correlation distances are, in fact, variable. This necessitates the use of non diagonal, inhomogeneous and non-isotropic error statistics. The use of horizontally uniform diagonal weights here is due to both the simplicity of this formulation and to the lack of reliable information about error statistics in oceanic observations.

Table 2
Error estimates used to calculate the cost function weights.

level	$\epsilon_k(T)$ (°C)	$\epsilon_k(S)$ (ppt)	$\epsilon_k(U)$ (*cm/sec*)
1	2.000	0.2500	5.000
2	1.858	0.2323	4.677
3	1.675	0.2095	4.258
4	1.436	0.1796	3.712
5	1.142	0.1429	3.041
6	0.8218	0.1029	2.309
7	0.5249	0.06580	1.630
8	0.2976	0.03742	1.111
9	0.1555	0.01967	0.7866
10	0.08189	0.01048	0.6185
11	0.04942	0.006425	0.5444
12	0.03676	0.004844	0.5154

With the above choice for the cost weights, a given constraint can be said to be consistent with the assumed error level if the corresponding term in the cost function is less than one. Larger value of the temperature data penalties, for example, would indicate that the solution is not consistent with the requirement that the solution is near the Levitus analysis. A large steady penalty contribution indicates that the solution is not consistent with the steady state model equations. An optimal solution should have all terms, representing dynamical constraints as as well as data constraints, smaller than one.

3.2 Initial guess

The minimization of a cost function based on the equations of a complex OGCM as constraints is a highly nonlinear optimization problem. If started too far from the absolute minimum of the cost function, the gradient based optimization could lead to a local minimum of the cost function which does not represent the optimal combination of dynamics

and data. Tziperman et al. [6] found evidence for such local minima and MW93 [8] also found that when starting their optimization directly from the data it seemed to converge to a different solution than the one they felt reflects the optimal state.

It is clearly important, therefore, to initialize the optimization with a good initial guess for the optimization solution. This can reduce the possibility of falling into a local minimum, as well as save much of the effort of minimizing the cost function through the expensive conjugate gradient iterations.

The initial guess for the optimization solution can be obtained by using simpler assimilation methods that are not optimal in the least square sense, yet have been shown to produce a very good approximation for the optimal solution. Let us briefly consider two such methods and demonstrate them using the present global model.

Suppose that our cost function consists of steady and data penalties for the temperature,

$$J(T) = \sum_{i,j,k} \left[\overline{W}^T_{ijk} \left(\mathbf{u}\nabla T - K_h \nabla_H^2 T - K_v T_{zz} \right)^2 + W^T_{ijk} \left(\hat{T} - T \right)^2 \right], \tag{9}$$

(the steady penalty here is simply the square of the steady state model equations). Because each term in the cost function is weighted by its expected error, we expect that at the optimal solution the total contribution of the steady penalties over the entire model domain should be roughly of the same order as that of the data penalties [6]. Assuming (with no rigorous justification) that this global condition may be applied locally, we have

$$\left(\mathbf{u}\nabla T - K_h \nabla_H^2 T - K_v T_{zz} \right)^2 \approx \left[\frac{W^T_{ijk}}{\overline{W}^T_{ijk}} \right] \left(\hat{T} - T \right)^2, \tag{10}$$

which is exactly the robust diagnostic equation [27] for the temperature

$$\frac{\partial T}{\partial t} + \left(\mathbf{u}\nabla T - K_h \nabla_H^2 T - K_v T_{zz} \right) = \gamma \left(\hat{T} - T \right), \tag{11}$$

at a steady state, with the restoring coefficient set to [6]

$$\gamma = \left[\frac{W^T_{ijk}}{\overline{W}^T_{ijk}} \right]^{\frac{1}{2}}. \tag{12}$$

In order to demonstrate the efficiency of the robust diagnostics approach, when used in the above fashion, to produce a good guess of the optimal solution, we show in Table 3 the cost parts obtained from the points in parameter space representing the Levitus data [entry (a)], the steady state model solution [entry (b)], and the robust diagnostic solution [entry (c)]. As may be expected, the point representing the Levitus data is characterized by large steady penalties and zero value for the data penalties; the steady state has vanishingly small values for the steady penalties but relatively large values for the data penalties, indicating that the steady state is not consistent with the data. Finally, the robust diagnostic solution has a well balanced distribution of steady and data penalties such that they are all small, and has therefore produced a near-optimal solution of our inverse problem, as anticipated in the above discussion.

Table 3
Summary of model runs and assimilations used in this study.

Run	Cost Parts				Comments
	data T/S	steady T/S	steady $u,v/\psi$	data $H/[E-P]$	
(a)	0.00 / 0.00	9.18 / 8.98	12.00 / 3221.	0.15 / 1.87	data
(b)	19.41 / 61.3	0.01 / 0.02	0.00 / 0.00	0.25 / 1.92	steady state
(c)	0.31 / 0.32	0.51 / 0.49	0.06 / 1.47	0.15 / 1.81	robust (rest. b.c)
(d)	0.62 / 0.74	1.11 / 1.24	0.08 / 1.67	0.00 / 0.00	robust (flux b.c)
(e)	0.31 / 0.34	0.51 / 0.50	0.06 / 1.48	0.10 / 0.66	extended robust
(f)	0.31 / 0.32	0.32 / 0.42	*0.03 / *2.17	*0.15 / *1.81	optimization

Terms marked by "*" were not part of the cost function used in the optimization and are only given for comparison with the other runs.

A second example of using a simple assimilation technique to obtain a good approximation of a complex optimization problem involves the optimal combination of heat-flux data and SST data [18]. Given the SST data, an estimated implied heat-flux field H^{SST} may be obtained using the restoring conditions formulation (1). Given also a climatological flux estimate, H^d, we can formulate an optimization problem in order to calculate an optimal heat flux H which is based on both estimates H^d and H^{SST}. The appropriate cost function is of the form:

$$J(SST,H) = \sum_{ij} \left[W^{SST} \left(H_{ij}^{SST} - H_{i,j} \right)^2 + W^H (H_{i,j}^d - H_{i,j})^2 \right]. \tag{13}$$

To obtain an approximate solution to the optimization problem posed by the above cost function, we simply write the model heat flux at every time step as a weighted average of the implied fluxes obtained from the restoring boundary conditions, and the climatological flux data

$$H^n = \alpha^T H^d + (1-\alpha^T) H^{SST,n}, \tag{14}$$

Integrating the model to a steady state using this heat flux, we obtain a solution for H which serves as the approximated solution to the above optimization problem. To derive an expression for α^T, we again use the expectation that at the minimum of the cost function, the different cost terms have roughly the same magnitude,

$$\sum_{ij} \left[W^{SST} \left(H_{ij}^{SST} - H_{i,j} \right)^2 \right] \approx \sum_{ij} \left[W^H (H_{i,j}^d - H_{i,j})^2 \right]. \tag{15}$$

assuming this holds locally and taking the square root, we have

$$\left[\frac{W^{SST}}{W^H} \right]^{1/2} \left(H_{ij}^{SST} - H_{i,j} \right) \approx (H_{i,j} - H_{i,j}^d). \tag{16}$$

A final manipulation of (16) brings us to the form postulated before in (14) and the relation between the weights in the cost function (13) and the coefficient α^T is found to be [18]

$$\alpha^T = \left[1 + \sqrt{W^{SST}/W^H}\right]^{-1}. \tag{17}$$

The runs in Table 3 demonstrate how the above scheme, which we term "extended robust" serves to minimize the heat-flux penalties in the cost function. The heat flux and fresh water flux penalties in entries (a-c) in Table 3, reflecting the data, steady state and robust diagnostics, are relatively large. Entry (e) represents the solution obtained using the robust diagnostics scheme (11) in the ocean interior plus the extended robust diagnostics scheme (14) at the surface. The extended robust scheme can be seen to be very efficient in reducing the value of the flux terms in the cost function, demonstrating again that simpler assimilation methods, when used wisely, can most efficiently calculate a near-optimal solution of most complex nonlinear optimization problems.

Both of the simple assimilation schemes used above can be shown to be equivalent to a corresponding optimization problem and give the same results under certain simplifying assumptions such as linearity, a single time step in evaluating the steady penalties etc. Thus the success of the simpler methods is not surprising. It is important to note however, that these simpler methods cannot replace the optimization approach for its ultimate objectives of parameter estimation and 4D data assimilation, both of which are still not tackled here.

3.3 Control Variables for a PE optimization

A primitive equation ocean model such as we use here requires the specification of temperature, salinity, horizontal baroclinic velocity field and the barotropic stream function as initial conditions. This multiplicity of initial conditions that must be calculated by the optimization algorithm poses two potential difficulties. First, the parameter space is significantly larger due to the addition of the baroclinic velocities and stream function as control variables. In general, the larger the parameter space, the more iterations are required to locate the cost minimum. Second, the additional control variables are very different from the temperature and salinity initial conditions, and thus pose new conditioning problems. Some of the complexities of using the baroclinic velocities and barotropic stream function as control variables, and the resulting ill conditioning were carefully examined by Thacker and Raghunath [10]. These potential difficulties with the velocity initial conditions lead Tziperman et al. [5, 6] to develop and use a model with diagnostic momentum equations for which only temperature and salinity initial conditions needed to be specified. However, in the present work we are faced with an optimization based on a full PE model, with more than double the number of initial conditions (per a given model resolution) than in Tziperman et al. [5, 6].

As before, we can use our knowledge of the physics to formulate the optimization problem in a way that is more likely to result in an efficient solution. It is known, and this fact has been used above to formulate the steady cost penalties, that given the density stratification, the velocity field in a rotating fluid must adjust to the density stratification within a few pendulum days. It seems most reasonable, therefore, that one would not need to calculate initial conditions for the velocities, and restrict the optimization problem to finding only the optimal temperature and salinity. The optimal velocity field will be found by the model after a very short initial adjustment period that should not have a

significant effect on the cost function that is based on the difference in temperature and salinity over an integration period of years. Every several iterations, the initial conditions for u, v, ψ may be updated by integrating the model for a few days starting from the last initial conditions for the temperature and salinity calculated by the c-g optimization and saving the results for the adjusted velocities and stream function and other models variables to be used as the new starting point for the optimization. Because of the short integration period, the temperature and salinity hardly change from their value calculated by the optimization.

This procedure should result in a better conditioning of the optimization problem due to the significantly reduced number of control variables. In Fig. 3 we show the reduction of the cost function for the optimization (run (f) in Table 3) that was started from a robust diagnostic solution. The optimization procedure was able to reduce the value of the cost, but eventually stalled after about 17 iterations. It seems that the optimization has converged to a local or global minimum solution; however after restarting the optimization with only T, S as control variables, additional progress was obtained, indicating that the stalling was more likely due to ill conditioning. Note that if the solution found at iteration 17 (Fig. 3) was indeed a minimum solution in the full parameter space spanned by T, S, u, v, ψ, then it is also a minimum in the subspace of T, S, and no further progress should have been obtained.

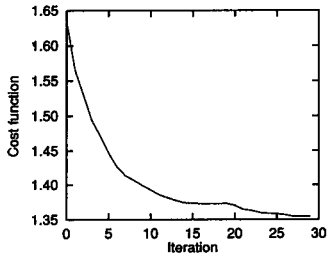

Figure 3: Cost value as function of iteration number for the optimization (run (f) in Table 3) beginning from the extended robust diagnostic solution (run (e) in Table 3).

Another issue related to the choice of control variables for the optimization is that of preconditioning. Preconditioning refers to a transformation of the control variables in order to improve the conditioning of the optimization. The control variables may be measured in various units and have very different typical numerical magnitudes. This may result in a badly conditioned optimization and therefore in the optimization stalling and not progressing towards the minimum of the cost function. The simplest remedy is to scale the control variables so that they all have similar numerical ranges. This may be improved upon by scaling the variables by the diagonal of the Hessian matrix if it can be estimated. The control variables may also be scaled by a non-diagonal transformation if a reasonably efficient transformation is available (see, e.g., [25, 26, 10]). Although somewhat neglected in the discussion here, the issue of preconditioning is a most important one.

4 Results

So far we have discussed in detail the issues of correctly formulating the optimization problem, and trying to guarantee its successful solution by starting from a good initial approximation of the optimization solution. We now wish to describe the results of a few model runs in some more details. We begin in section 4.1 by describing and analyzing the solution of optimization (f) in Table 3. We then analyze model solutions obtained under restoring conditions and under flux conditions in section 4.2.

4.1 The optimization solution

One of the advantages of nonlinear optimization is that it can be used to re-map the data in a way that is consistent with the model equations. Fig. 4 shows the horizontal temperature and salinity fields at model levels 2 and 7, as obtained from the optimization (run (f) in Table 3), as well as the Levitus data at the same levels.

The data residuals at levels 2 and 7 for the temperature and salinity (Fig. 5a) are quite small over most of the ocean volume, as indicated by the fact that the global measure of the data penalties (see Table 3) is less than one for both the temperature and the salinity. But there are some regions, most notably the western boundary regions in the North Atlantic and North Pacific, as well as the equatorial Pacific region, in which the deviations from the data are systematic and larger than the errors specified by the cost function weights (Table 2). In these regions, the optimization has clearly modified some features of the Levitus analysis quite substantially [See for example (Fig. 4) the temperature field in the tropical Pacific at level 2, or the smoother salinity contours created by the optimization at level 7]. In some cases the changes made by the optimization could be considered improvements, in others they are certainly a reflection of model deficiencies. Considering the coarse model we use here, we do not wish to claim to have improved on the Levitus analysis. But the temperature and salinity distributions we find are clearly more consistent with the model dynamics and therefore more appropriate for starting a coupled model integration using an ocean model similar to ours than is the original Levitus analysis.

The steady residuals at levels 2 and 7 for the temperature and salinity are shown in Fig. 5b. The quantity plotted is the temperature after two year integration from the optimal state, minus the optimal state, multiplied by 7.5, to get the extrapolated drift expected in a 15 year period, as it appears in the cost function. The projected temperature drift is quite small at level 2, except in the Pacific sector of the southern ocean, where a strong convection creates some numerical noise of no physical significance. At level 7 one notices systematic warming in the north west Atlantic, probably due to the inability of the model to create the NADW at the right level and to have it spread southward correctly. In the north east Atlantic, the cooling trend is related to the Mediterranean tongue outflow that while simulated fairly reasonably thanks to the Mediterranean sponge layer, is still not sufficiently consistent with the data in that region. The steady salinity residuals reflect basically the same model problems indicated by their temperature counterparts.

It is important to understand that while the optimization results suffer some obvious deficiencies as indicated above, they still provide a significant improvement over both the steady state model solution obtained without data assimilation and the Levitus analysis.

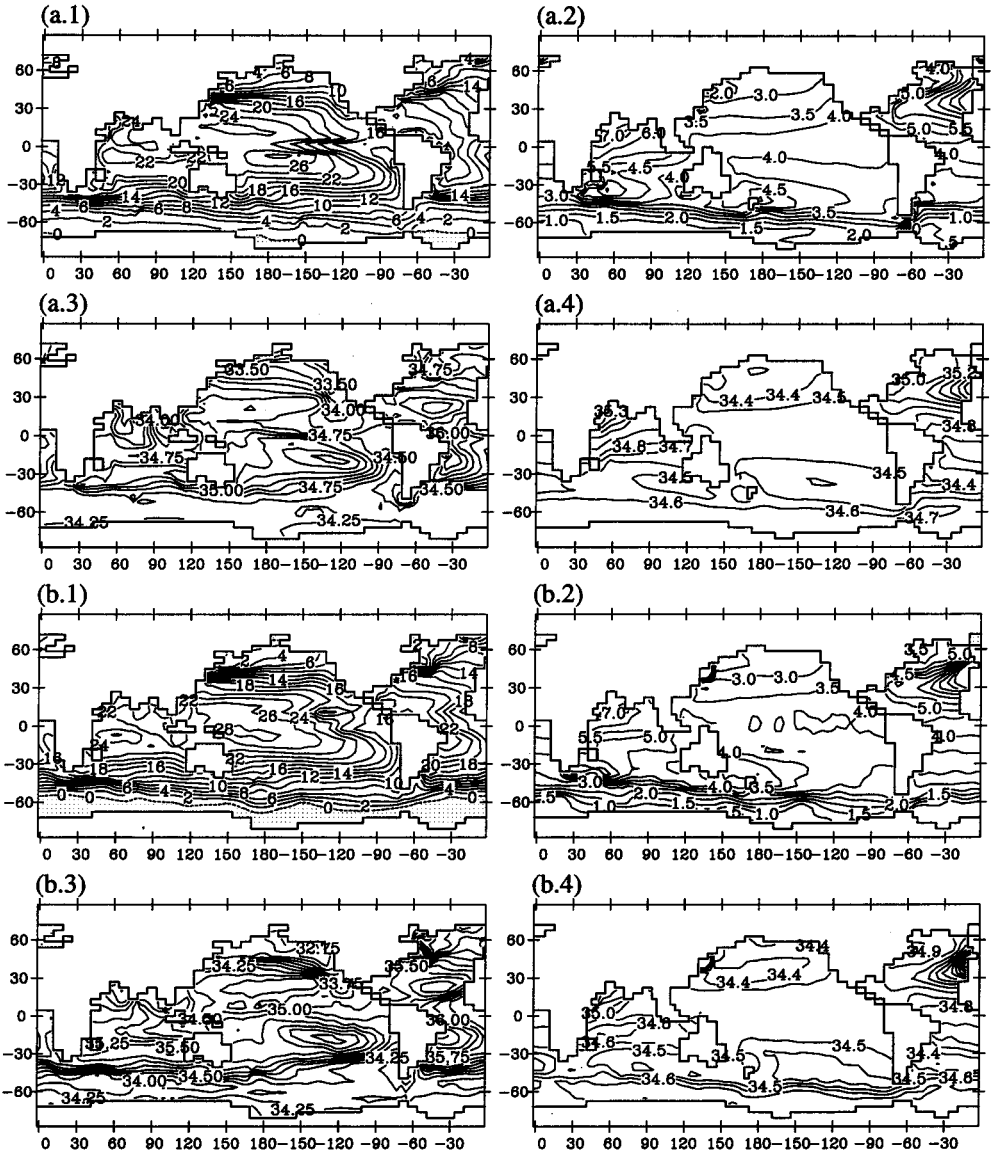

Figure 4: (a) Optimization solution. Uppermost two panels: solution for the temperature at levels 2 and 7 (depths of 85m and 1130m). Second row of panels: salinity solution at the same levels. (b) Levitus analysis. Third row of panels: Levitus temperature at model levels 2 and 7. Bottom two panels: Levitus salinity at model levels 2 and 7. Contour intervals for a.1–a.4 (and for b.1–b.4) are 2.0°C, 0.5°C, 0.25°/$_{oo}$and 0.1°/$_{oo}$, respectively. Negative areas are dotted.

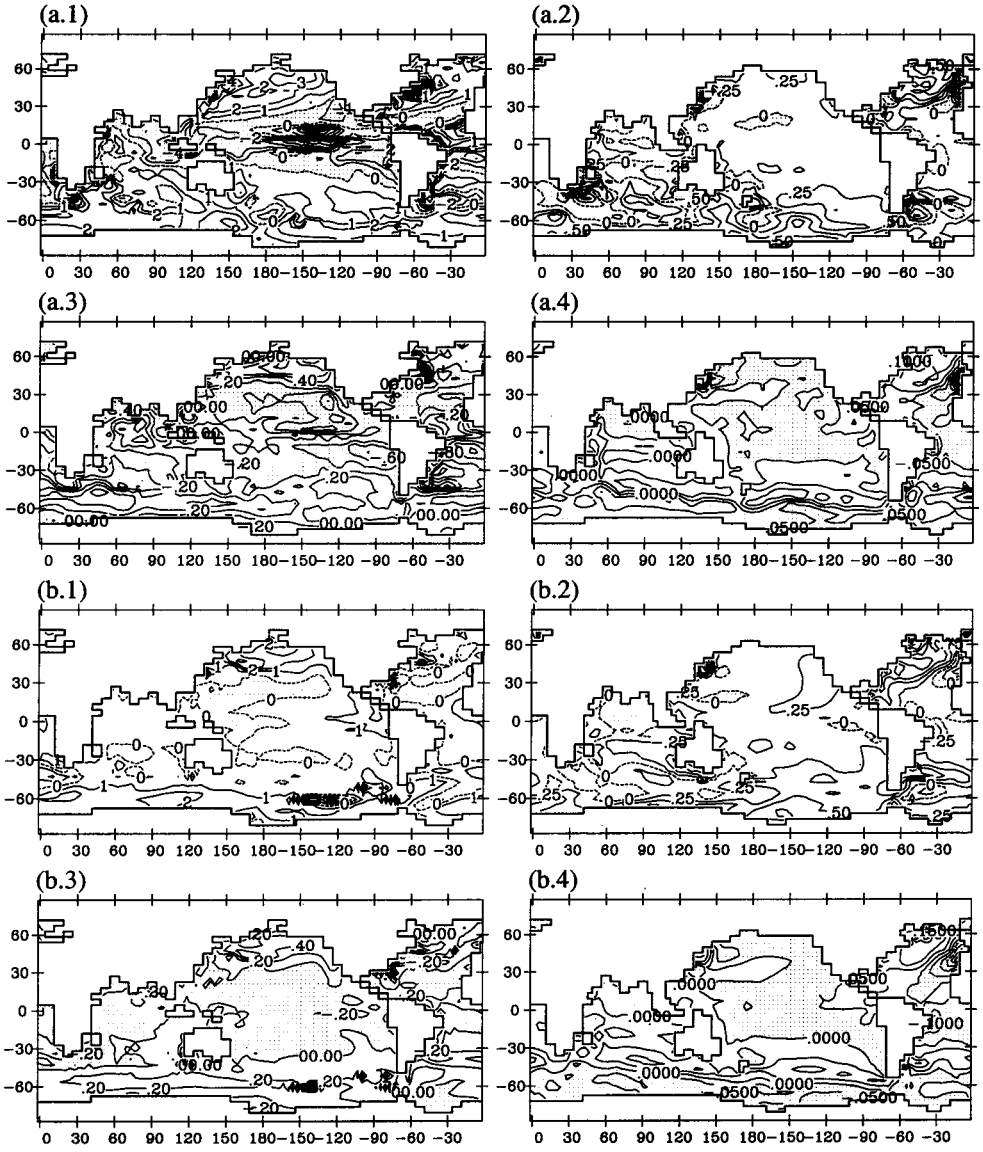

Figure 5: (a) Data residuals for the optimization solution (run (f) in Table 3). Uppermost two panels: data temperature-residuals at levels 2 and 7; second raw of panels: same, for salinity. (b) Steady residuals for the optimization (run (f) in Table 3). Third raw of panels: steady temperature-residuals at levels 2 and 7 ; bottom raw of panels: same, for salinity. Here and elsewhere in the manuscript, plots of data residuals are of optimization solution minus the corresponding data. Similarly, plots of steady residuals are of the model solution after a two-year integration time minus the initial state, extrapolated to 15 years by multiplying the difference by 7.5. Contour intervals for panels a.1–a.4 (as well as for b.1–b.4) are 1.0°C, 0.25°C, 0.2°/$_{oo}$and 0.05°/$_{oo}$, respectively. Negative areas are dotted.

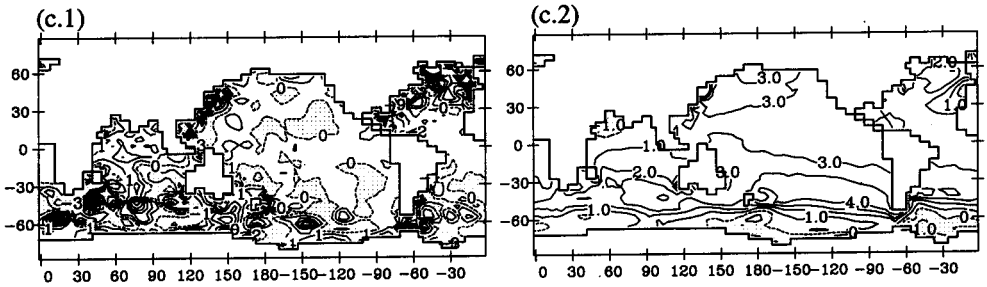

Figure 5: (c) Left panel: Steady temperature-residuals at level 7 estimated for the Levitus analysis; right panel: data temperature-residuals at level 7 estimated for the steady state model solution of Fig. 1. Contour intervals are 1.0°C. Negative areas are dotted.

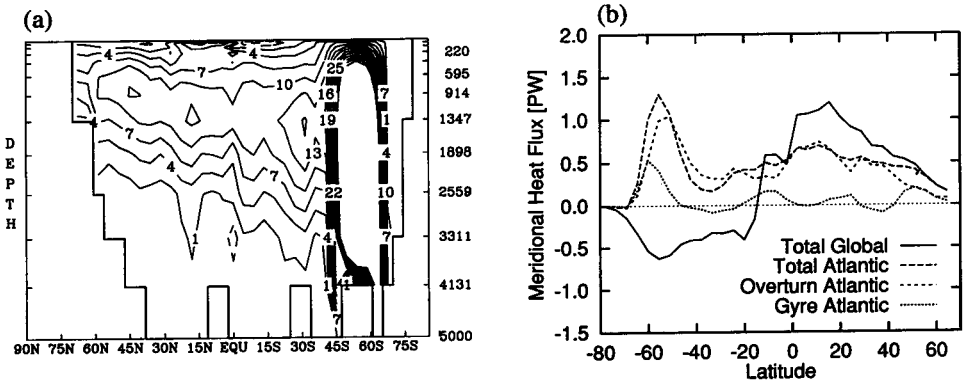

Figure 6: (a) North Atlantic meridional stream function for optimization solution (run (f) in Table 3). (b) Meridional heat flux for optimization solution; see caption of Fig. 1b.

This is seen from Table 3 which shows that the cost values for the Levitus data and the steady state solution are significantly larger than for the optimal solution. Fig. 5c shows the steady temperature residuals at level 7 estimated for the Levitus analysis as well as the temperature data residuals at level 7 estimated at the steady state model solution. Clearly both the data and the steady state are not optimal in the sense that they minimize one type of cost terms (data or steady penalties), but on the expense of a large increase in the other cost terms.

The North Atlantic overturning circulation for the optimal solution is shown in Fig. 6a. The overturning circulation at 30N is only 10Sv instead of the expected 16-20Sv. This feature of the solution cannot be considered an improvement over the prognostic run of Fig. 1. Fig. 6b shows the meridional heat flux for the optimal solution. Again, no significant improvement is obtained over the prognostic model solution of Fig. 1, and the northward heat flux carried by the North Atlantic ocean at 25N is still significantly less than the expected 1PW (10^{15} watts). These limitations of the meridional circulation

and meridional heat flux for the optimal solution are not surprising, considering the model performance in these areas. It seems that the only appropriate solution is to improve the prognostic model, perhaps by using isopycnal mixing or another eddy mixing parameterization [19].

As is quite clear from Table 3, most of the cost reduction as compared to the Levitus analysis or steady state solution has been obtained during the robust diagnostics initialization run [entry (c) in Table 3. Still, the cost reduction during the optimization itself is not negligible (Fig. 3), in particular for the steady penalties. Fig. 7 shows the steady residuals at the end of the robust diagnostic solution. The general picture is of fairly significantly reduced steady residuals in the optimization as compared to the robust diagnostics (compare Fig. 5b.1 and Fig. 7a). The reduction is spread over the entire domain, showing again the effectiveness of the optimization. A similar comparison of the salinity steady residuals (not shown) shows a similar reduction. A comparison of the distribution of the data residuals does not show a significant difference between the robust diagnostics solution and the optimization, as may be expected from the results in Table 3.

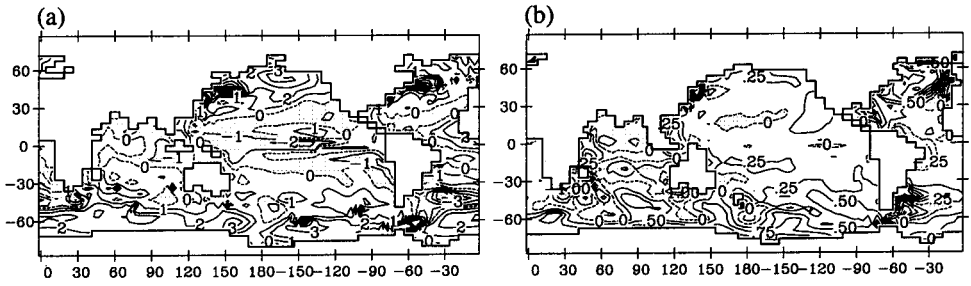

Figure 7: (a) Steady temperature residuals for robust diagnostic solution (run (c) in Table 3) at level 2; (b) same, at level 7. Contour intervals are 1.0°C and 0.25°C, respectively. Negative areas are dotted.

Note that during the optimization, the steady residuals were reduced significantly more than the data residuals [compare entries (c) and (f) in Table 3]. This indicates that a relatively small change in the temperature and salinity fields can induce a larger change in the steady penalties. This asymmetry between the steady and data penalties is again a possible indication that the cost function is not sufficiently well conditioned. Perhaps a better cost formulation may be able to better balance steady and data penalties. Interestingly, the asymmetry observed here between the data and steady penalties is of an opposite nature to that seen in the cost section of Fig. 2, where a small change to the steady penalties (between $r = 0$ and $r = 1$) involved a large change in the data penalties. Clearly the highly nonlinear structure of the steady penalties, reflecting the nonlinear model equations, accounts for these complex behaviors of the steady penalties.

4.2 Restoring vs fixed-flux surface boundary conditions

Let us now consider the issue of fixed-flux vs restoring surface boundary conditions in inverse problems based on an ocean GCM. Tziperman et al. [5] have shown that when

using fixed flux boundary conditions with the flux as a control variable, small errors may be amplified by the optimization in areas of deep convection resulting in huge errors in the calculated heat flux. Furthermore, Tziperman et al. [6] found a very large discrepancy between their optimization solution for the SST and the data, and suggested that this is due to the use of flux rather than restoring conditions. Marotzke and Wunsch [8] encountered a similar large discrepancy in SST which they interpreted as a drift towards winter conditions and felt that this is the result of the absence of seasonal cycle in their model.

We would like to suggest here that these large SST discrepancies may be eliminated by the use of restoring boundary conditions. We further argue that such a boundary condition formulation is more physically motivated as well as more successful from a practical point of view.

Based on the success of the robust diagnostic approach in obtaining a near-optimal solution to the least square optimization problems, we shall base our discussion on the two robust diagnostic solutions represented by entries (d) and (e) in Table 3. Run (d) uses flux boundary conditions with the surface fluxes of heat and fresh water specified to be the climatological data sets described in section 2, while run (e) uses restoring boundary conditions and combines the climatological flux data and the restoring to the observed SST using the extended robust diagnostics approach [18] described in section 3.2.

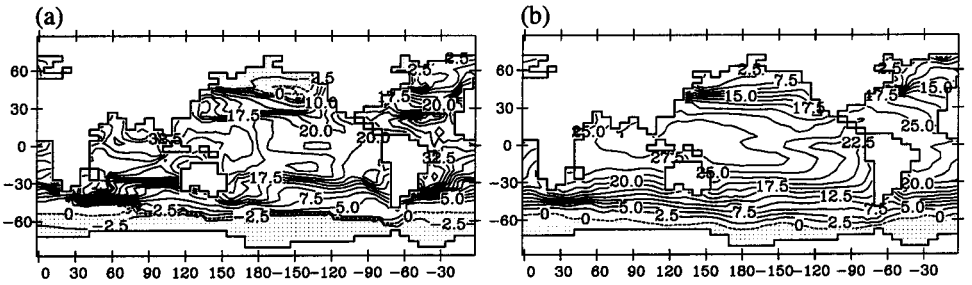

Figure 8: SST for (a) robust diagnostics run using flux conditions with the climatological heat and fresh water flux data [entry (d) in Table 3), (b) an extended robust run using restoring surface boundary conditions [entry (e) in the Table 3]. Contour intervals are 2.5°C. Negative areas are dotted.

Fig. 8 shows the SST for both runs. The surface temperature field for run (d), using flux boundary conditions with climatological flux data, is very far from the observed field. Note that the temperature and salinity at all levels are still restored in this run to the Levitus data by the robust diagnostic term in the model equations. The restoring time, however, is 15 years, rather than 30 to 120 days normally used for the surface fields under restoring conditions. The structure of the temperature field is consistent with a contraction of the large scale shape of the thermocline in the north-south direction, as seen in a much more pronounced form in ocean model runs under flux conditions without restoring at the interior. The mid-latitude regions and poleward are colder than the Levitus data, while the tropical regions are warmer. The large discrepancy in SST is reminiscent of the results of Tziperman et al. [6] and MW93 [8]. In our run (d), the

entire North Pacific ocean north of about 20N is significantly colder than the data, giving an impression that it tends towards a winter temperature distribution.

We note, however, that the restoring conditions run of entry (e) produces a very reasonable fit to the Levitus SST, while also being able to reduce the distance to the observed climatological fluxes (see heat-flux penalty terms for this run in Table 3). Moreover, both the data and steady penalties for the temperature and salinity under the flux conditions are significantly larger. It seems, therefore, that inverse models should use restoring conditions even when trying to estimate the optimal air-sea flux. The enforcement of the flux data can be done by including it in the cost function as in (6). Such a formulation seems capable of producing a reasonable compromise of heat-flux data, SST and interior temperature.

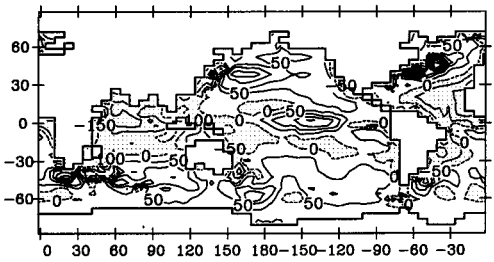

Figure 9: Heat-flux residuals for the extended robust diagnostics run of Fig. 8b. Contour intervals are 50Watts/m^2. Negative areas are dotted.

Fig. 9 shows the heat-flux data residuals for run (e), that is, the optimal heat flux of run (e) minus the climatological data of Esbensen and Kushnir [20]. There are clearly large systematic deviations from the heat-flux data in many areas such as the North Atlantic, equatorial Pacific and Indian Ocean. Large systematic heat-flux residuals in MW93 have lead the authors to suggest that the optimization's solution tends towards winter conditions with strong cooling over their entire basin. It seems to us that such large heat-flux residuals may, in fact, be related to the inability of the model to correctly simulate the North Atlantic meridional circulation [19], and therefore the meridional heat flux. Such a poor simulation of the meridional heat flux is directly linked to poor simulation of the air-sea fluxes [18], and hence the large heat-flux residuals seen in Fig. 9, and possibly also in MW93. The meridional heat flux for runs (d) and (e) is shown in Fig. 10. The run under flux conditions has a somewhat enhanced northward flux both in the northern hemisphere of the global ocean and in the North Atlantic ocean. But the price paid for this enhancement in terms of deviation from the temperature data is clearly too large. The large SST discrepancy indicate that the model cannot be forced to simulate the correct air-sea fluxes, possibly because of its inability to produce the correct overturning circulation.

Runs (d) and (e) are, of course, not optimizations but solutions of a robust diagnostic model which was previously shown to closely simulate the optimal solution of a corresponding optimization. We have recently repeated the above analysis for two optimizations using restoring and flux boundary conditions correspondingly. The results

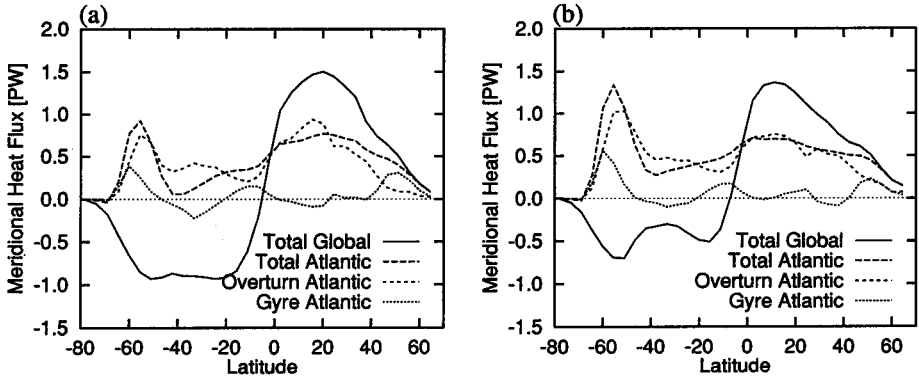

Figure 10: Meridional heat flux for both runs of Fig. 8: (a) run using fixed surface-flux boundary conditions (run (d) in Table 3); (b) run using restoring boundary conditions (run (e) in Table 3). See caption of Fig. 1b.

fully support the above conclusions reached for runs (d) and (e) and will be published elsewhere with a fuller analysis of the boundary condition formulation problem for inverse studies.

To estimate the effect of the missing seasonal cycle in our model, a comparison can be made between our results and those of Bryan and Lewis [17] who have used a very similar model, under seasonal forcing. Comparing, for example, the annually averaged meridional heat flux in their model, the to the steady state meridional heat flux in ours (Fig. 1), we see that there is not much of a difference. The addition of seasonality does not necessarily improve the simulation of the meridional heat flux (and therefore of the implied surface flux).

Let us summarize the issue of boundary condition formulation, both from the point of view of the physics, and from a practical point of view. First, from the point of view of the physics of air-sea interaction, we note that flux boundary conditions imply modeling the atmosphere by assuming that the air-sea heat flux does not depend on the SST, and ignoring the obviously important feedback between SST and air-sea heat flux. In restoring conditions, this feedback is crudely included, as the restoring to the SST observations is somewhat reminiscent of the restoring of SST to the lower atmospheric temperature which occurs in the actual coupled system. Ultimately, one may want to use more elaborate parameterizations of the air-sea heat flux as function of the SST, and perhaps use as control variables the restoring times, or the atmospheric temperature which may appear in these parameterizations. This seems to make much more physical sense than formulations in which the flux is calculated directly, ignoring the SST-flux feedback.

Second, from a more practical point of view, we note that ocean GCMs poorly simulate the observed ocean temperature when driven by specified surface heat fluxes. Similarly, when run under restoring conditions to observed SST and surface salinity, ocean models produce very poor estimate of the surface fluxes of heat and fresh water, and therefore of the meridional fluxes of heat and fresh water [18]. It seems, therefore, that ocean models

can presently produce either surface heat fluxes that are consistent with observations, or an interior solution that is consistent with observations, but not both. Now, in coupled model studies, it is presently more crucial for the ocean model to get the SST right than the heat flux, as the latter is corrected for using the artificially added flux correction. This dictates a choice of surface boundary condition formulation that is different from what was used in previous applications of the adjoint method to similar models, namely a restoring boundary condition rather than flux boundary condition. Restoring boundary conditions formulation is also consistent with our wish to use a model that can be run in a simulation mode, as this requirement cannot be met using fixed-flux surface boundary conditions.

5 Conclusions

This study attempted to combine a global primitive equation ocean model with climatological data for temperature, salinity and surface fluxes, using the adjoint method of data assimilation. This is a step towards using the adjoint method with the ocean component of coupled ocean-atmosphere models in the hope of achieving two ultimate goals. The first goal is the improvement of ocean climate models by estimating, for example, internal model eddy parameters and parameterizations so that the model simulations are closer to the observed ocean even when they are run without data assimilation. The second ultimate goal is the calculation of an ocean state based on the available data and the model equations to be used as initial conditions for coupled ocean-atmosphere climate simulations. While we have not achieved these goals as yet, we believe an important progress was made. Let us briefly summarize the main lessons we have learned here.

Because our goal is to work with ocean models that can also be run without data assimilation, we have taken the approach that the model used for the inverse calculation must be formulated such that it can run independently in a simulation mode. A particular consequence of this approach has been the use of restoring rather than fixed-flux surface boundary conditions. We argued that the restoring boundary conditions are both better physically motivated and more successful from a practical point of view in producing a good inverse solution.

Large discrepancies between the optimization solution for the SST and the data have been encountered in previous inversions using the adjoint method [6, 8]. These inversions used a fixed-flux surface boundary condition formulation. Marotzke and Wunsch have suggested that the SST discrepancy in their model was due to a drift towards colder surface temperatures ("winter conditions") which results from the lack of a seasonal cycle in their steady model. We have shown that the SST discrepancy can be eliminated in our model by using restoring surface boundary conditions, while still incorporating the available climatological flux data into the optimization. A comparison of our model results with a seasonal version of the present model run by Bryan and Lewis [17] seems to indicate that the annually averaged results of the seasonal model are quite close to the results of the model when driven with annually averaged steady forcing. We suspect, therefore, that at least for the present model, the use of flux conditions, rather than the lack of a seasonal cycle may be the cause of the SST discrepancies.

We have run our model in a simulation mode and have tuned the model as well as we can to produce the best possible results. Still, it is evident that according to our measure of success (the value of the cost function), the steady results with no data assimilation were grossly inconsistent with the Levitus data. Similarly, the Levitus analysis was found to be just as inconsistent with the requirement that it satisfies the steady state model equations. However, our optimization approach provided a better solution than both the steady state model solution obtained with no data assimilation and the original climatological data sets. This solution was much more consistent with both the data and steady constraints, and therefore significantly more optimal in the least square sense. The ultimate test of this optimality would be, of course, to use such a state as initial conditions for a climate simulation.

Several important lessons have been learned here concerning the cost function formulation for PE models. We have demonstrated that there are still conditioning difficulties for the dynamical constraints with the presently used cost formulations, and that there seems to be room for improvement in this area. This issue requires much further research. A second conclusion we have come to concerning the cost formulation for PE models is that it does not seem necessary to include explicit steady penalties for the baroclinic velocities and barotropic stream function. Because the velocity field very rapidly adjusts to the stratification in rotating fluids, it seems sufficient to penalize the deviations of the temperature and salinity from a steady state. Similarly, we expect that future studies using seasonal models with an optimization approach, should enforce the dynamical constraints requiring the model solution to be seasonal on the temperature and salinity fields and not on the velocity field. The fast adjustment of the velocity field to the stratification has also led us to suggest that one can do well by using only the temperature and salinity as control variables to be calculated by the optimization, allowing the model to adjust the velocity field. Such a procedure, outlined in more detailed in the previous sections, may improve the conditioning of optimization problems based on PE ocean models.

A most successful part of this study has been the use of simple assimilation method to obtain good approximations to the optimization problem. These approximations are then used to initialize the optimization, significantly reducing the minimization effort in the optimization itself. More importantly, they reduce the possibility of encountering local minima that will prevent the c-g optimization from finding the global minimum representing the desired optimal solution. We have demonstrated how dynamical constraints can be combined with data constraints using the simple robust diagnostic approach [27] to obtain a near-optimal solution. We have also shown how surface flux data may be combined with the dynamical constraints and the surface temperature and salinity data using the extended robust diagnostic approach of Tziperman and Bryan [18], again resulting in a near-optimal solution. Such simple assimilation approaches can help but not replace the optimization approach of the adjoint method, because they cannot be used to estimate parameters such as eddy coefficients etc, a goal for which the adjoint method itself is well suited.

We feel that the technical aspects of inverting complex PE ocean models treated here, as well as the more general issues we dealt with, should be useful to future studies directed at using data assimilation with ocean climate models. There is a clear and urgent necessity of improving ocean models used for climate studies, and of using these models to estimate

the ocean state as well as is allowed by the available data. We have argued here that the adjoint method is a most appropriate tool for obtaining these goals, and we feel that they should and can be achieved in the near future.

6 Acknowledgments

This study would not be possible without the efforts of Long, Huang and Thacker who have developed the adjoint code for the GFDL PE model. We are grateful to them for allowing us the use of this code. We wish to thank Jochem Marotzke for sharing with us his code improvements allowing the use of multiple time steps for the cost evaluation. Thanks also to Carl Wunsch and to two anonymous reviewers for useful comments on an earlier draft. Computer time for this study was partially provided by a grant from NCSA Illinois. Partial support for ET was provided by grant 89–00408 from the United States - Israel Binational Science Foundation. Partial support for ZS and computer time were provided by grant N00014-93-1-0831 from the Office of Naval Research - Navy Ocean Modeling Program.

References

[1] F. Le Dimet and O. Talagrand, Tellus, 38A (1986) 97.

[2] W. C. Thacker and R. B. Long, J. Geophys. Res., 93 C2 (1988) 1227.

[3] C. Wunsch, J. Geophys. Res., 93 C7 (1988) 8099.

[4] E. Tziperman and W. C. Thacker, J. Phys. Oceanogr., 19 (1989) 1471.

[5] E. Tziperman, W. C. Thacker, R. B. Long and Show-Ming Huang, J. Phys. Oceanogr., 22 (1992) 1434.

[6] E. Tziperman, W. C. Thacker, R. B. Long, Show-Ming Huang and S. Rintoul, J. Phys. Oceanogr., 22 (1992) 1458.

[7] J. Marotzke, J. Phys. Oceanogr., 22 (1992) 1556.

[8] J. Marotzke and C. Wunsch, J. Geophys. Res., 98 (1993) 20149.

[9] A. Bergamasco, P. Malanotte-Rizzoli, W. C. Thacker and R. B. Long, Deep-Sea Res.. II, 40 (1993) 1269.

[10] W. C. Thacker and R. Raghunath, J. Geophys. Res., 99 (1993) 10131.

[11] R. B. Long, S.-M. Huang 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

[12] Bennett, A. F., and P. C. McIntosh, 1982. Open ocean modeling as an inverse problem: tidal theory. J. Phys. Ocean., 12, 1004–1018.

[13] Bennett, A. F. Inverse methods in physical oceanography, Cambridge Monographs, Cambridge University Press, 346 pp, 1992.

[14] K. Bryan, J. Computat. Phys., 4 (1969) 347.

[15] A. J. Semtner, An oceanic general circulation model with bottom topography. UCLA dept of Meteorology Tech. Rep. No. 9, 1974.

[16] M. D. Cox, A primitive equation 3 dimensional model of the ocean, GFDL ocean group technical report No 1. Princeton University.

[17] K. Bryan and L. J. Lewis, J. Geophys. Res., 84 (1979) 2503.

[18] E. Tziperman and K. Bryan, J. Geophys. Res., 98 C12 (1993) 22629.

[19] C. W. Boening, F. Bryan, W. R. Holland and J. C. McWilliams, An overlooked problem in model simulations of the thermohaline circulation and heat transport in the Atlantic ocean, Manuscript, 1994.

[20] S. K. Esbensen and Y. Kushnir, The heat budget of the global ocean: an atlas based on estimates from surface marine observations, Oregon State Univ., Climate Research Institute Report No. 29.

[21] A. C. Hirst and W. Cai, J. Phys. Oceanogr., 24 (1994) 1256.

[22] S. Levitus, Climatological atlas of the world ocean, NOAA Tech. pap., 3, 1982.

[23] A. Baumgartner and E. Reichel, The world water balance, Elsevier, NY, 1975.

[24] Hellerman and Rosenstein, J. Phys. Oceanogr.13 (1983) 1093.

[25] P. E. Gill, W. Murray and M. H. Wright, Practical Optimization, Springer-Verlag, Heidelberg, 1981.

[26] Tziperman, E., W. C. Thacker and K. Bryan, 1992: Computing the steady state oceanic circulation using an optimization approach. *Dynamics of Atmospheres and Oceans.* Vol 16 No 5 pp. 379-404.

[27] J. L. Sarmiento and K. Bryan, J. Geophys. Res., 87 (1982).