Finding the Steady State of a General Circulation Model Through Data Assimilation: Application to the North Atlantic Ocean

JOCHEM MAROTZKE AND CARL WUNSCH

Center for Meteorology and Physical Oceanography, Department of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Cambridge

An attempt is made to find a steady state of a general circulation model consistent with error estimates of the North Atlantic hydrography during 1981–1985 and with wind and thermodynamic boundary conditions. The quadratic misfit of the model state from the data is minimized using the general circulation model together with its adjoint. A state which is steady within limits of estimated observational error is found, but no such state is simultaneously consistent with the observed hydrographic and surface flux fields. Model dynamics are able to sharpen gradients which were overly smooth in the mapped data, producing a meridional overturning cell with a maximum value of about 21 Sv. At equilibrium, the model must produce its own water masses and has a strong tendency to go toward wintertime conditions. The consequence is an estimated model surface temperature systematically lower than in the hydrography (which comes from all seasons). The western boundary current thermocline is also both colder and fresher than the mapped hydrography, and the overall meridional heat transport is low (about 0.6×10^{15} W). The results suggest that the concept of a realistic steady state North Atlantic circulation has reached the end of its utility, that models with realistic property fluxes and divergences must have much higher resolution, and that the open-ocean boundary conditions must be formulated as control variables.

1. INTRODUCTION

Almost since the beginning of oceanography, attempts have been made to use hydrographic data to form pictures of the ocean circulation. Such pictures are sought for a number of reasons. In a climate context, they are one way of making improved estimates of the crucial air-sea interactions, whose present uncertainty is unacceptably large. Any attempt to forecast climate states rests on the ability to initialize a model with accurate estimates of present-day circulation.

Because the data available have necessarily been nonsynoptic, numerous issues arise in the construction of circulation estimates. Among the questions we would list are the following. Is there an existing steady state which the data have all measured? If there is such a state, what is the nature of the inevitable aliasing, and what does it do to the resulting synthesis? Given the nonlinear nature of the hydrodynamical equations, is it best to form space-time averages of data before use, or are the data best used in unaveraged form? Are the resulting pictures of the circulation consistent with equivalent estimated climatologies of boundary conditions (wind/buoyancy)? How can we make realistic estimates of the uncertainty of the result?

The difficulty of these questions is such that after 100 years of effort, we are not entirely sure of the answer to any one of them. Quite apart from any other issues, a simple answer to the question of the degree and import of the aliasing would demand sampling the ocean circulation at a rate sufficient to avoid aliasing and to produce quantitative estimates of the full frequency/wavenumber spectrum of its variability in such a way as to permit computation of the effects of undersampling. Because there is little short-term prospect for the requisite database, considerable interest focuses on a simpler question: Is the existing hydrography

internally consistent both with itself and with the equations of motion for a steady-circulation ocean? Finding answers to this question would produce several benefits. If inconsistencies are found, we obtain clues to the validity of the steady state assumption and to the extent and effects of aliasing. Equally important, the issue of how to deal with the combination of a realistic oceanic circulation model and realistic hydrographic data (by definition, at least, noisy) raises a sequence of difficult technical questions about methodology.

The present paper is directed at determining whether a general circulation model (GCM) can indeed demonstrate quantitative consistency or the lack of it with the modern hydrography of the North Atlantic Ocean. Much of our focus is inevitably on a sequence of practical issues that do not normally arise in conventional hydrographic analysis or in conventional modeling alone. This paper is, in part, a progress report toward the ultimate goal of synthesizing hydrography with dynamics whether or not a steady state is actually indicated. We believe that much of the experience outlined here will be of use to others engaged in similar enterprises with different models and data.

As suggested by *Willebrand and Wunsch* [1990], a community effort was undertaken to use the North Atlantic as a setting for understanding the hydrographic estimation problem. Some of the results of that effort are described by *Martel and Wunsch* [1993], *Mercier et al.* [1993], and *Schiller* [1993].

The present effort is most immediately an outgrowth of the results described by *Tziperman et al.* [1992a, b]. They used a reduced form of the Geophysical Fluid Dynamics Laboratory (GFDL) general circulation model [*Bryan*, 1969; *Cox*, 1984] and produced a numerically nearly exact "adjoint" version of the model. The adjoint model is an element of what is usually called the Pontryagin principle: it corresponds to the Lagrange multipliers used to impose the model as a constraint in a minimization problem. Numerically, the Lagrange multipliers are useful as a method for finding the minimum by search. *Marotzke* [1992] reported on our early

attempts to exploit the existence of the adjoint model and his proposed modifications of the numerical scheme originally outlined by *Tziperman and Thacker* [1989]. In the present paper, we will use the model, its adjoint, the modern North Atlantic hydrography, and *Marotzke*'s [1992] proposed changes in the algorithms (*Tziperman et al.* [1992b] also discuss the implications of Marotzke's findings).

The use of Pontryagin principle methods represents a natural extension of earlier work [e.g., *Wunsch and Grant*, 1982] using linear inverse methods on this same problem. In moving from modest-sized, linear constraint problems to massive numbers of finite-difference constraints representing time-evolving, nonlinear systems, one must find methods for overcoming the resulting numerical burden. *Martel and Wunsch* [1993] invert a steady, linear finite-difference model; *Mercier et al.* [1993] invert a steady nonlinear one; *Wunsch* [1988] inverted a linear, time-dependent system; and several other examples exist in the literature. But the combination of nonlinearity with a time-evolving model raises computational issues which none of these earlier studies had to address, and addressing them is our purpose here.

Much of this paper is given over to addressing a myriad of practical problems. Although a particular model, data set, and region of the ocean are under consideration, the issues are generic, to be faced by anyone using a GCM with real (as opposed to synthetic) observations. Because we believe that the future of quantitative understanding of the ocean circulation necessarily requires use of similar optimization methods, some of the technical issues are worth describing in some detail.

This paper is organized as follows. Section 2 describes the data base, and section 3 describes the model and the "adjoint technique." The best fit with the data obtained so far is described in detail in section 4, alternative experiments are outlined in section 5, uncertainty estimates are discussed in section 6, and section 7 summarizes the situation.

2. THE DATABASE

The hydrographic data used were all obtained between 1981 and 1985 as described by *Fukumori et al.* [1991] and *Fukumori and Wunsch* [1991]. A gridded "climatology" was formed from these data, as described in Appendix A. This climatology was essentially an objective mapping onto a regular grid with 1° longitude by 2° latitude spacing that was transformed to the model vertical grid (Table 1). Potential temperature is used.

In any estimation problem the data must be accompanied by an estimate of their uncertainty. For the hydrography we use the formal errors provided by the objective analysis algorithm. Following *Olbers et al.* [1985], however, the formal error is replaced by a constant whenever the former is judged to be too small relative to various systematic errors not included in the formal error; the threshold values are 0.01°C and 0.01 practical salinity unit (psu) for temperature and salinity, respectively. Following *Fukumori et al.* [1991], an additional uncertainty was included in the two uppermost layers (0 to 100 m) to accommodate the sampling over all seasons. Table 2 lists the horizontal averages of the standard deviations in temperature and salinity that are used.

The surface heat flux climatology used was produced by *Isemer et al.* [1989], who modified work by *Bunker* [1976]. The specific estimates based upon a requirement for a

northward heat transport across 25°N of 1.2 petawatts (PW; 1 PW = 10^{15} W) were used, consistent with *Hall and Bryden* [1982] and others. (Note, however, that *Isemer et al.* [1989] preferred an estimate with northward heat transport of only 1.0 PW across 25°N). As with the hydrographic data, the surface heat fluxes are subsampled onto the model grid. The formal error ascribed to the surface heat flux data are the rms estimates produced by *Isemer et al.* [1989]. They derive only from the uncertainties in the parameters entering the air-sea interaction formulas; typical values of the formal uncertainty lie between 25 and 40 W/m^2 and are probably an underestimate. Use of the *Isemer et al.* [1989] heat fluxes represents a temporal inconsistency, because the ship observations entering *Bunker*'s [1976] data set were obtained in the years 1941–1972.

The evaporation minus precipitation (*E-P*) climatology of *Schmitt et al.* [1989] was used. They produced no estimate of the uncertainty, and it was set here arbitrarily to a uniform value of 0.3 m/yr.

Using the routine analyses of the European Centre for Medium Range Weather Forecasts, *Trenberth et al.* [1989] produced global maps of monthly mean wind stress for the years 1980 through 1986 on a 2.5° by 2.5° grid. From the North Atlantic portion of this data set, we computed a time average and interpolated it onto the model grid. The temporal standard deviation is assumed to represent the observational uncertainty and is likewise interpolated to the model grid.

3. MODEL DESCRIPTION

The model extends from 9.5° to 59.5°N and from 80° to 9°W with a resolution of 1° zonally, 2° meridionally, and 24 levels vertically. The anisotropic resolution is chosen to limit the computational burden while allowing for the preservation of some of the frontal structure in the boundary currents. Realistic bottom topography and coastlines are used (Figure 1). The northern and southern boundaries are closed walls; to include the effects of the heat and salt transports across these latitude circles, sponge layers of thickness 2° are used at the southern and northern boundaries, in which temperature and salinity are restored to the "data" with a time constant of 1 day. Since there are no hydrographic data for parts of the Caribbean Sea and the Gulf of Mexico, the Antilles form the model boundaries in the southwest (Figure 1), and open boundary conditions are employed in a very crude way. The inflow from the Straits of Florida into the open Atlantic is assumed to be known (Figures 2a and 2b), the distribution of transports with depth and temperature is inferred from *Leaman et al.* [1987, 1989], and inflow salinities are assumed to be the horizontal averages of the Atlantis 109 section Florida-Fort Pierce at 27°N [*Roemmich and Wunsch*, 1985]. The total transport of the Florida Current is set to 30 Sv (1 Sv = 10^6 m^3/s).

The outflows from the open Atlantic into the Caribbean occur through narrow and shallow passages between the islands and cannot be modeled with the coarse resolution used here. Following the results of *Roemmich* [1981], an 8-Sv flow through the Windward Passage and a 22-Sv flow through the eastern passages are imposed, with the appropriate distribution in depth (Figure 2a). The temperature/ salinity profiles of the outflows are specified from the observed hydrography.

TABLE 1. Parameters of Vertical Grid

Value at Grid Point:

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
DZ	50	50	100	100	100	100	100	100	100	100	100	200	300	300	400	400	400	500	500	400	400	400	400	400
ZDZ	50	100	200	300	400	500	600	700	800	900	1000	1200	1500	1800	2200	2600	3000	3500	4000	4400	4800	5200	5600	6000
N	25	75	150	250	350	450	550	650	750	850	950	1100	1350	1650	2000	2400	2800	3250	3750	4200	4600	5000	5400	5800

DZ, grid box thickness; ZDZ, depth of box bottom; Z, depth of grid point.

TABLE 2. Horizontally Averaged Standard Errors for Temperature and Salinity as a Function of Depth

k	$\langle \sigma_{T,k} \rangle$	$\langle \sigma_{S,k} \rangle$
1	2.90	0.424
2	2.87	0.362
3	1.28	0.185
4	1.17	0.172
5	1.17	0.167
6	1.13	0.155
7	1.05	0.135
8	0.917	0.111
9	0.765	0.0935
10	0.624	0.0852
11	0.520	0.0830
12	0.441	0.0819
13	0.335	0.0633
14	0.164	0.0303
15	0.0692	0.0131
16	0.0520	0.0100
17	0.0493	0.0100
18	0.0354	0.0100
19	0.0251	0.0100
20	0.0487	0.0101
21	0.0529	0.0102
22	0.0520	0.0101
23	0.0458	0.0100
24	0.0122	0.0100

Here, $\langle \sigma_{T,k} \rangle$ is standard error for temperature, and $\langle \sigma_{S,k} \rangle$ is standard error for salinity. No that for every grid point, the formal mapping error from the objective analysis has been replaced by 0.01°C and 0.01 psu when the formal error was smaller.

The model code derives from work by *Tziperman et al.* [1992a, b] with the "forward" computation based upon the GCM of *Bryan* [1969] and *Cox* [1984]. The GCM is, however, used with simplified momentum equations as follows.

1. Local time derivatives are neglected, thereby eliminating internal gravity and external Rossby waves, the latter in conjunction with the rigid lid approximation, which is a part of the original model. External gravity waves are also eliminated.

2. Momentum advection is also neglected, consistent with the resolution, which is still far coarser than the first internal radius of deformation.

3. A Rayleigh friction law is used instead of the horizontal and vertical eddy viscosity parameterizations of the original model. The governing equations used are
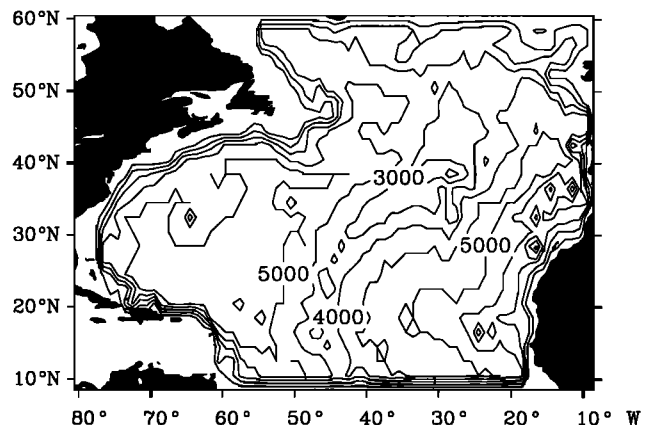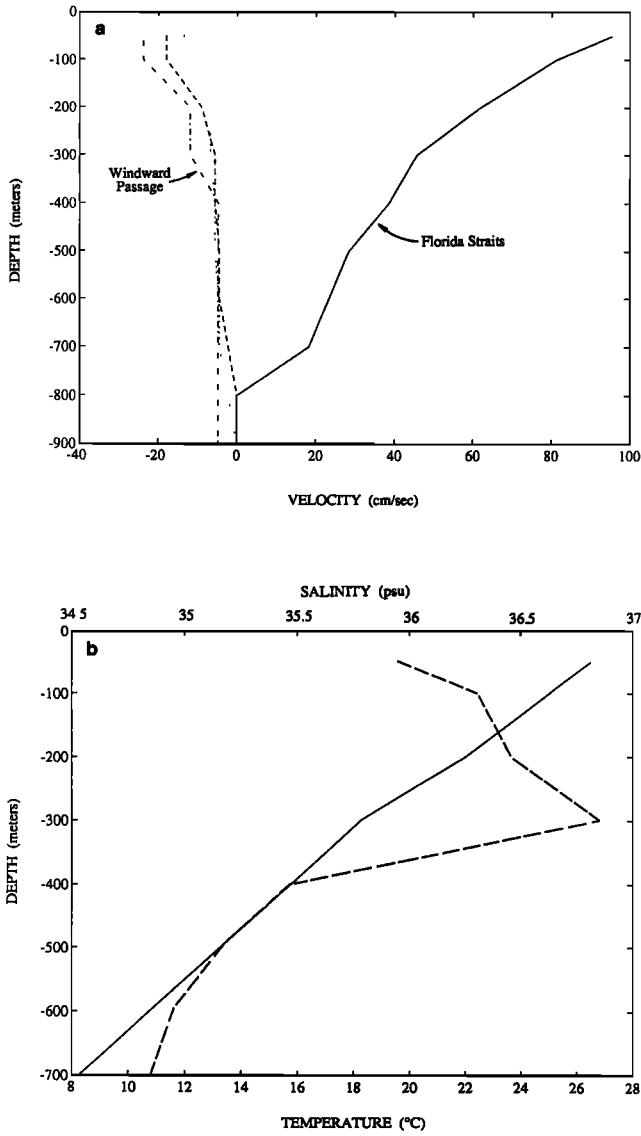

Fig. 1. Model bathymetry. Contour interval, 1000 m.

Fig. 2. (a) Prescribed velocities at open lateral boundaries as a function of depth; positive values denote inflow. Unlabeled curves are the eastern passages. (b) Prescribed profiles of inflow temperature (solid curve) and salinity (dashed curve) in the Florida Straits as a function of depth.

$$\mathbf{f} \times \mathbf{u} = -\frac{1}{\rho_0} \nabla p - r\mathbf{u} \qquad (1a)$$

$$\frac{\partial \theta}{\partial t} + (\mathbf{u} \cdot \nabla)\theta + w\frac{\partial \theta}{\partial z} = A_{HH}\nabla^2\theta + A_{HV}\frac{\partial^2\theta}{\partial z^2} + C_\theta \qquad (1b)$$

$$\frac{\partial S}{\partial t} + (\mathbf{u} \cdot \nabla)S + w\frac{\partial S}{\partial z} = A_{HH}\nabla^2 S + A_{HV}\frac{\partial^2 S}{\partial z^2} + C_S \qquad (1c)$$

$$\nabla \cdot \mathbf{u} + \frac{\partial w}{\partial z} = 0 \qquad (1d)$$

$$0 = -\frac{\partial p}{\partial z} - g\rho \qquad (1e)$$

$$\rho = \rho(S, \theta, p)$$

The notation is conventional, with $\mathbf{u}$ as the horizontal velocity, ∇ as the horizontal gradient, r as a Rayleigh damping coefficient, and θ as the potential temperature.

As in the work by Bryan and Cox [1972], the equation of state is a polynomial approximation; C_θ and C_S represent removal of static instability by a convective adjustment process. Bottom and lateral walls are insulating, and the sidewalls have no-slip boundary conditions. Finally, the rigid lid approximation is made. Apart from slight differences in the velocity boundary conditions, the only deviation from the standard Bryan and Cox model lies in the almost-geostrophic momentum equation $(1a)$.

As a consequence of these simplifications, the velocity field is everywhere geostrophic except possibly where frictional effects are strong. Thus the model does not resolve the geostrophic adjustment process. Because the seasonal cycle is absent, the neglect of the external-mode Rossby waves should be acceptable. The role of closure schemes in GCMs is poorly understood, and the consequences of using Rayleigh friction are less clear than the other simplifications; the change was introduced to make the diagnosis of the velocity field computationally easy (see the discussion by Salmon [1986]). The complete conservation equations for heat and salt, including a convective adjustment procedure to eliminate static instability, are employed. We expect these equations to represent the dominant constraints set by the model, which justifies further the simplifications in the momentum equations. The time step used is 12 hours in most of the experiments. Mixing parameter values are A_{HH} 10^3 m^2 s^{-1}, $A_{HV} = 10^{-4}$ m^2 s^{-1}, and $r = 5 \times 10^{-7}$ s^{-1}.

In its use of the conservation equations for heat and salt along with geostrophy, the present model form is similar in its physics to the GCM developed by Maier-Reimer et al. [1993]. The numerical scheme is, however, entirely different.

The adjoint to the GFDL model was originally developed by C. Thacker, R. Long, and S.-M. Hwang at the Atlantic Oceanographic and Meteorological Laboratory in Miami; the simplifications to the model as used here were made by E. Tziperman. The equations constituting the adjoint model can be derived as an application of the Pontryagin minimum principle in control theory [e.g., Luenberger, 1979; Talagrand and Courtier, 1987; Thacker and Long, 1988; Wunsch, 1988].

The basic idea can be stated as follows. Let $\mathbf{x}(t)$ be the state vector describing the ocean at time t (the state vector consists of the dependent variables temperature and salinity explicitly computed by the model in its time stepping). The discretized form of the model is written very generally and symbolically as

$$\mathbf{x}(t + 1) = \mathcal{L}[\mathbf{x}(t), \mathbf{Bq}(t), \mathbf{\Gamma u}(t)] \qquad (2)$$

Here the time step is taken as unity. The matrices $\mathbf{B}$ and $\mathbf{\Gamma}$ are used to distribute the time dependences in the known "controls" $\mathbf{q}(t)$ and the unknown ones $\mathbf{u}(t)$ over the model grid. The decomposition into $\mathbf{B}$, $\mathbf{q}$, $\mathbf{\Gamma}$, and $\mathbf{u}(t)$ is not unique; see any textbook on control methods, e.g., Bryson and Ho [1975].

Most observations are linear combinations of the elements of the state vector written in the form

$$\mathbf{y}(t) \equiv \mathbf{E}(t)\mathbf{x}(t) + \mathbf{n}(t) \qquad (3)$$

where $\mathbf{n}(t)$ represents the inevitable observation errors, and $\mathbf{E}$ is a matrix relating the state vector to whatever is observed. In general, if $\mathbf{x}(t)$ computed from the model is substituted into (3), $\mathbf{y}(t)$ thus predicted will differ from the values actually observed. A quantitative measure of the misfit can be obtained from a quadratic cost or objective function:

$$J = \sum_t [\mathbf{y}(t) - \mathbf{E}(t)\mathbf{x}(t)]^T \mathbf{W}_D(t)[\mathbf{y}(t) - \mathbf{E}(t)\mathbf{x}(t)] \quad (4)$$

where $\mathbf{W}_D$ is a weighting matrix, most commonly the inverse of the observation error covariance matrix. In the present situation, the only nonvanishing term is the last one, $t = t_f$, the observations being treated as being obtained from a terminal time with the ocean in nearly steady state. Our aim is to minimize this objective function by determining a new estimate $\bar{\mathbf{x}}(t)$, subject to the requirement that it be a solution to the model equations. By enforcing the model through Lagrange multipliers $\boldsymbol{\mu}(t)$, the problem is made one of unconstrained minimization. That is, we append the model to the cost function as

$$J' = J - 2 \sum_{t=0}^{t_f-1} \boldsymbol{\mu}(t+1)^T\{\mathbf{x}(t+1)$$

$$- \mathscr{L}[\mathbf{x}(t), \mathbf{Bq}(t), \mathbf{\Gamma u}(t)]\} \quad (5)$$

As outlined in the introduction, one of our main goals is to ask whether a steady state ocean circulation as described by the particular model is consistent with all the observations available. To find such a steady state, another term was added to the objective function (4) penalizing the quadratic difference between the model state at the end of the integration period and the initial state. The objective function now reads

$$J' = J + \sum_t \mathbf{u}(t)^T\mathbf{Q}\mathbf{u}(t) + [\mathbf{x}(t_f) - \mathbf{x}(0)]^T\mathbf{W}_s[\mathbf{x}(t_f) - \mathbf{x}(0)]$$

$$- 2 \sum_{t=0}^{t_f-1} \boldsymbol{\mu}(t+1)^T\{\mathbf{x}(t+1) - \mathscr{L}[\mathbf{x}(t), \mathbf{Bq}(t), \mathbf{\Gamma u}(t)]\} \quad (6)$$

following the discussion by *Tziperman and Thacker* [1989], *Tziperman et al.* [1992a, b], and *Marotzke* [1992]. $\mathbf{Q}$ is a weight matrix representing a priori beliefs about the unknown control terms (boundary condition uncertainty) in the model. $\mathbf{W}_s$ is a weight matrix specifying how stringently steadiness is demanded. In contrast to the original suggestion by *Tziperman and Thacker* [1989], the integration is carried over more than one time step to ensure that the algorithm finds the minimum of the objective function [*Marotzke*, 1992]. The matrices $\mathbf{W}_D$, $\mathbf{W}_s$, and $\mathbf{Q}$ are assumed diagonal; each element of $\mathbf{W}_D$ is formed by dividing the inverse observation error variance by the total number of grid points N.

By setting the derivatives of J' to zero, we find the conditions for stationary values of J' (minimum of J) in the form

$$\frac{1}{2}\frac{\partial J'}{\partial \mathbf{u}(t)} = \mathbf{Q}\mathbf{u}(t) + \left(\frac{\partial \mathscr{L}}{\partial \mathbf{u}(t)}\right)^T \boldsymbol{\mu}(t+1) = 0, \quad (7a)$$

$$0 \le t \le t_f - 1$$

$$\frac{1}{2}\frac{\partial J'}{\partial \boldsymbol{\mu}(t)} = \mathbf{x}(t+1) - \mathscr{L}[\mathbf{x}(t), \mathbf{Bq}(t), \mathbf{\Gamma u}(t)] = 0, \quad (7b)$$

$$0 \le t \le t_f - 1$$

$$\frac{1}{2}\frac{\partial J'}{\partial \mathbf{x}(t)} = -\boldsymbol{\mu}(t) + \left(\frac{\partial \mathscr{L}}{\partial \mathbf{x}(t)}\right)^T \boldsymbol{\mu}(t+1) = 0, \quad (7c)$$

$$1 \le t \le t_f - 1$$

$$\frac{1}{2}\frac{\partial J'}{\partial \mathbf{x}(0)} = -\mathbf{W}_S\{\mathbf{x}(t_f) - \mathbf{x}(0)\} - \left(\frac{\partial \mathscr{L}}{\partial \mathbf{x}(0)}\right)^T \boldsymbol{\mu}(1) = 0 \quad (7d)$$

$$\frac{1}{2}\frac{\partial J'}{\partial \mathbf{x}(t_f)} = \mathbf{E}^T(t_f)\mathbf{W}_D(t_f)\{\mathbf{E}(t_f)\mathbf{x}(t_f) - \mathbf{y}(t_f)\}$$

$$+ \mathbf{W}_S\{\mathbf{x}(t_f) - \mathbf{x}(0)\} - \boldsymbol{\mu}(t_f) = 0 \quad (7e)$$

That is, we seek the combination of initial and boundary conditions which results in the best fit of the model state to the data over the complete time history.

In principle, we need to solve the large set of simultaneous coupled nonlinear difference equations in (7). To make determining a solution practical, we start with a first estimate of the unknown controls, initial state and boundary forcing, and construct a first-guess time history of the model state by integrating (7b) forward in time. Equation (7c) is then integrated backward to produce estimates of $\boldsymbol{\mu}(t)$. Equation (7a) will not usually vanish, and the nonzero values provide the gradient of J with respect to the controls in terms of the Lagrange multipliers. A standard optimizing search algorithm (here, the conjugate gradient method) can use this information to improve the initial estimates of $\mathbf{u}$. Another forward and backward iteration is then made, the cycle continuing until convergence is obtained [e.g., *Thacker and Long*, 1988]. In practice, $\mathbf{u}(t)$ was set to constant values over the times of the forward and backward iterations. As used here, then, the main merit of the adjoint model resides in the gradient information it provides, which makes possible a very efficient minimization of the objective function, rather than in the physics of $\boldsymbol{\mu}(t)$ themselves.

Specification of $\mathbf{W}_S$ is equivalent to an upper bound on residual time rates of change of the solution (in an rms sense). If τ marks a time larger than or equal to the adjustment timescale of the processes one wants to equilibrate, it is sensible to demand that the equilibrium to be determined be steady up to $\delta T/\tau$, where δT is a typical observation error (likewise for salinity). This concept can be translated into a recipe of how to specify $\mathbf{W}_S$ provided $\mathbf{W}_D$ is given. The changes in temperature and salinity during the model integration are extrapolated over time τ, $[\mathbf{x}(t_f) - \mathbf{x}(0)]\tau/t$, where t is the integration time. If the extrapolated drifts are smaller than the observation error, their timescale is so large that they should not be penalized heavily. The elements of $\mathbf{W}_S$ are chosen here to be horizontally uniform and are constructed according to

$$W_{s,ijk} = \langle \sigma_k \rangle^{-2} \frac{1}{N} \left(\frac{\tau}{t}\right)^2 \quad (8)$$

where $\langle \sigma_k \rangle$ is the horizontal mean of the assumed observation error at depth level k. $N = 71 \times 25 \times 24 = 42,600$ is the total number of grid points; ij denotes the horizontal grid

points; and t is the integration time. Note that if τ is much larger than t, the elements of $\mathbf{W}_s$ are much larger than those of $\mathbf{W}_D$, which makes it impossible to find the minimum of the objective function [see *Marotzke*, 1992, for details].

Because the physical model is nonlinear, the time history of each forward run has to be stored to provide the information necessary to step the Lagrange multipliers backward in time in the associated adjoint run. Ideally, one would store the roughly 10^5 variables completely describing the model state (temperature and salinity at all grid points or, additionally, velocities to avoid redundant computations) at every single time step. For 1000 time steps or more, the storage requirements become (almost) prohibitive. To alleviate this problem, the forward history is stored only at certain intervals, which are chosen to be larger (typically 10 to 25 time steps) as the model comes closer to a steady state. Special attention must be paid to the storage of the time history of convection events which may occur intermittently and thus be missed by the sampling scheme. Whenever convective adjustment has occurred during one of these intervals of 10–25 time steps, the adjoint model assumes that convection has been active throughout the entire interval, because over the order of 10 days, other processes are too slow to moderate the vertical homogenization.

Several outcomes can emerge from what is a process of constrained optimization. (1) A stationary point of (6) is found, with value consistent with the prior estimated uncertainties and with the model. (2) A stationary point of the cost function is found but with a value too large to be acceptable, there being inadmissible unsteadiness or deviations from the observations or both. (3) Several minima, all equally acceptable, can be found.

Result 1 would show that the model equations (1) and (2), are fully consistent with all of the observations, and there would be a uniquely determined best-estimated oceanic state. The implications are that no more complicated (either in terms of physics or spatial resolution) model could be justified and that the model could then provide best estimates of secondary quantities of interest, such as Gulf Stream transports and heat flux divergences. The production of a model proving wholly consistent with everything observed about the ocean might be regarded as a summary statement of the goals of physical oceanography.

Result 2 would prove that the model was inadequate in one or more ways for describing what is observed. In such a situation, the result may contain strong clues about how the model must be modified to gain consistency with the observations.

Result 3 would show not only that the model is adequate to describe all the observations but that the observational data set cannot uniquely describe the ocean circulation, there being several oceanic states all consistent with what is known.

Our method is one of nonlinear optimization by search algorithm. Like most nonlinear problems, this one entails technical difficulties. Typically, if a minimum is found, whether acceptable or not, the proof that one has a global minimum or even that there are no others equally acceptable can be made plausible only by search; it cannot usually be done rigorously.

4. EXPERIMENT 1: THE STANDARD EXPERIMENT

Finding a Minimum

The results of four model runs differing in their experimental strategy or weights or both will be presented here.
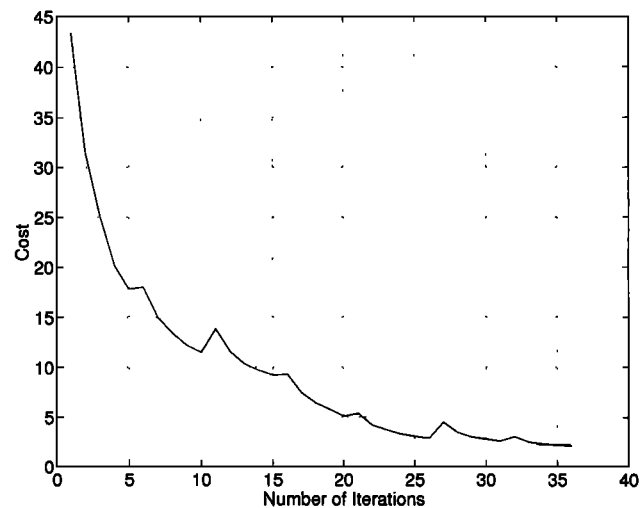

Fig. 3. History of the objective (or cost) function as the minimization in the standard experiment proceeds. The nonmonotonic behavior occurs when the minimization procedure is restarted using the final rather than the initial state of the previous iteration.

Experiment 1, the standard experiment, represents what we judge to be the best solution found so far. It starts from the observed hydrography and, for reasons to be explained below, is first run in a forward mode for 250 days, forced by the observed surface fluxes. The final state of this "spin-up" represents the first estimate of the initial conditions for the minimization; the observed surface fluxes are the first estimates of the boundary conditions. The cost function is defined according to (6), with $\mathbf{Ex}$ representing observed temperature and salinity. Only the hydrographic data misfit and the deviations from steadiness are penalized in the standard experiment; deviations of the surface fluxes from their climatologies are not. In effect, arbitrarily large deviations from the surface flux climatologies would be tolerated, reflecting their large uncertainty $(\|\mathbf{Q}\| \to 0)$.

Starting from the final state of the spin-up, the model performs 15 conjugate gradient steps, each comprising runs of the forward and adjoint model over 100 days. As the iterations become inefficient in reducing the cost function, the integration time is increased to 500 days and another 21 iteration steps are taken. Short integration times at the beginning of the experiments serve to mutually adjust the forcing and near-surface fields (see *Marotzke* [1992] for a detailed discussion) at low computational cost. After conjugate gradient steps 5, 10, 15, 20, 26, and 31, the final state of the preceding iteration step is taken as the initial condition of the subsequent step. In this way, some of the gross adjustment processes that have taken place during early iteration steps need not be repeated. Starting from a value of 43.3, the cost function decreases to 2.1 after iteration step 36; the progress of the minimization is displayed in Figure 3; at this stage a minimum has been found.

To demonstrate that a minimum has indeed been reached, the initial conditions are perturbed by a small change in sea surface temperature in two locations, and the changes in cost and gradient are computed. The two grid points represent mid-ocean regions of contrasting behavior, where convection either is absent (Figure 4a) or reaches a maximum depth (Figure 4b). Figure 4a shows the result for perturbations of
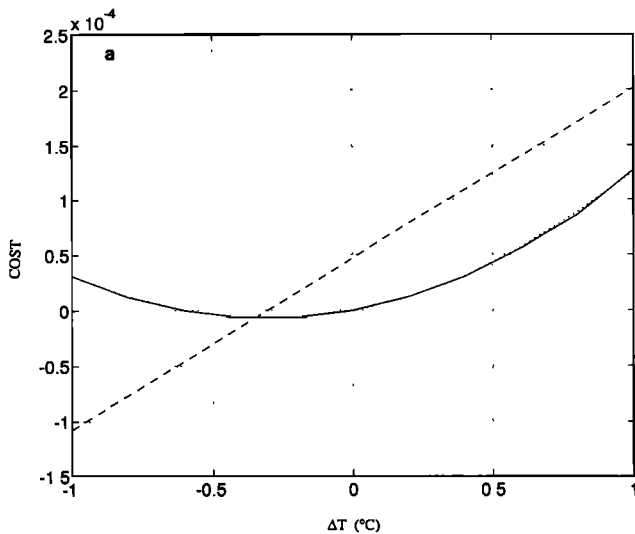
Fig. 4a. Cost function (solid curve) at an open-ocean point without convection occurring, in the vicinity of the optimal solution of the standard experiment as the sea surface temperature is perturbed. The classical parabolic shape is seen. The estimated gradient (dashed curve) is nearly linear, and its zero almost coincides with the minimum of the parabola.

sea surface temperature between $-1°$ and $+1°C$ at $16.5°N$, $31.5°W$. The value of J for zero perturbation has been subtracted from all values. The standard experiment ends in the immediate vicinity of the minimum (at least with respect to variations in this particular state vector element); the "true" minimum is smaller by only 10^{-5}. The local cost function has an almost perfect parabolic shape: the gradient as computed from the adjoint is a near-perfect straight line. Both the curvature of J and the slope of the gradient are estimates of a particular diagonal element of the Hessian in the vicinity of the minimum. If the inverse square root of the curvature (or the
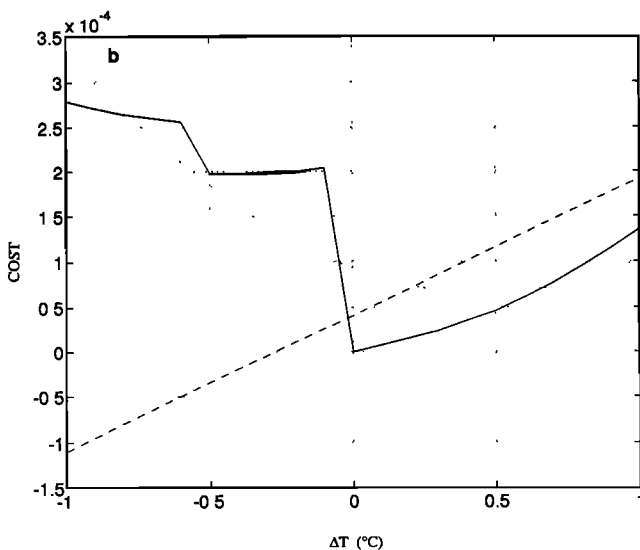
Fig. 4b. Same as Figure 4a except at a grid point where convection is present. The discontinuities introduced into the cost function by convection are plainly evident. Nonetheless, the best-fitting parabola (dotted curve) and the estimated gradient (dashed curve) do an adequate job of locating the near-optimal solution.

slope of the gradient) is taken as a measure of the uncertainty of the surface temperature estimate at that point, we obtain $0.4°C$ (note that the curvature and the slope of the gradient have to be multiplied by twice the number of grid points and the number of grid points, respectively, to obtain an inverse variance). Uncertainties are taken up further in section 6.

Figure 4b shows the corresponding experiment for perturbations of surface temperature at $46.5°N$, $31.5°W$. Here, a minimum has been found to within less than $0.1°C$. For positive temperature perturbations, the objective function again has a locally parabolic shape, but for negative perturbations, changes in convective activity lead to jumps in the cost function with smooth variations in between. The gradient computed from the adjoint is again a straight line, but it is seen not to integrate back to the true function. This result occurs because the calculation of the gradient is based on the adjoint equations to the linearized forward model; changes in convective activity (a nonlinear phenomenon) are not taken into account. The dotted curve in Figure 4b shows the best fit of a parabola to J in the vicinity of the minimum. Although the quadratic fit is poor, the estimated curvature provides a reasonable order of magnitude for the gross behavior of the objective function. Here, estimates of the uncertainty of the particular surface temperature element of the state vector are $0.3°$ and $0.4°C$ based on the quadratic fit and the slope of the gradient, respectively.

At this minimum the contributions from temperature and salinity data misfit and temperature and salinity steadiness misfit are 1.0, 0.6, 0.2, and 0.3, respectively. Consistency with the model requires each of these four numbers to be of the order of unity, since the global rms error in each term should be consistent with the prior estimates of their standard deviations. (Because the normalization entering the weight factors does not take into account that only about 60% of the model grid points are ocean points, one strictly wants each of the four numbers to be around 0.6.) Experiment 1 thus meets the most fundamental requirements for a compatible solution. The minimization finds a compromise between the data and the steadiness constraints; moreover, the residuals of the fit are compatible with the prior estimates of the errors, at least according to a global measure (the cost function). The solution found by experiment 1 is now examined to investigate whether it is oceanographically acceptable in all respects, given that determination of an acceptable minimum is a necessary but not sufficient condition for a solution. Note in particular that systematic deviations of pointwise values in the model from the observations can lead to rejection of a solution, even if the formal difference variance is acceptable.

The Overturning Cell

Many North Atlantic circulation models [e.g., *Manabe and Stouffer*, 1988; *Martel and Wunsch*, 1993] fail to be vigorous enough in the thermohaline component of the flow. The zonally integrated "overturning cell" apparently carries a mass transport notably weaker than that required by direct observations and with a consequent too-weak meridional heat flux. For the present case, Figure 5 shows the meridional overturning stream function (i.e., the zonally integrated meridional velocity, summed up starting from the surface) at the end of the optimization run. To make the definition of a stream function meaningful in the presence of zonal outflows, the implied northward flow in the Caribbean Sea and
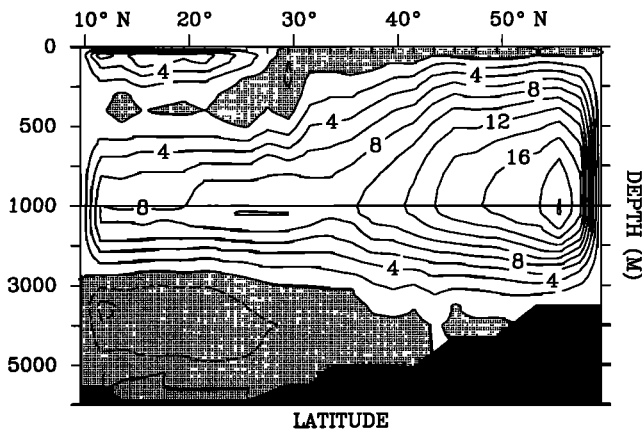
Fig. 5.   Meridional overturning stream function (in Sverdrups) at the end of the standard experiment. Positive values denote clockwise circulation; contour interval, 2 Sv.

the Gulf of Mexico is added to the midocean contributions, assuming no net vertical motion outside of the model domain. The maximum overturning strength is about 20 Sv and occurs just south of the northern sponge layer. At the southern boundary, 4.8 Sv enter below 3800-m depth, and about 13 Sv leave between 1400 and 3800 m. Eight Sverdrups enter above 1000-m depth, and although one cannot expect an exact correspondence between the water masses and the zonally averaged currents, Figure 5 reproduces the classical picture of Antarctic Bottom Water flowing below and counter to the North Atlantic Deep Water, above which the Antarctic Intermediate Water enters the North Atlantic. Note that around and south of 30°N, the overturning strength is only 10 Sv, compared to 18 Sv reported by, for example, *Hall and Bryden* [1982] and *Roemmich and Wunsch* [1985]. As a consequence, the northward heat transport is low (Figure 6); the maximum is 0.52 PW at 27.5°N. The part of the heat transport attributable to the meridional overturning pattern (i.e., the vertical correlation of the zonal averages of velocity and temperature) is 0.33 PW at 27.5°N, in contrast to 1.2 PW [e.g., *Hall and Bryden*, 1982]. The strength of the heat transport is also apparently weakened by too small a contribution (0.16 PW) leaving the southern sponge layer. Note that heat transport cannot be rigorously defined where the model circulation is not mass-conserving; the contribution of the flow through the Caribbean is crudely taken into account by adding terms of the form $Q_F U_P / U_F$ to the mid-ocean meridional temperature flow, where U_P is the volume transport through the passages into the Caribbean, and U_F and Q_F are the volume transport and temperature flow through the Florida Straits, respectively.

Sea Surface Temperatures

Sea surface temperatures (SST) are displayed in Figure 7. The climatology (Figure 7a) shows a front, albeit blurred, along the western boundary north of 34°N. Farther north, the front marks the northward and eastward turns of the North Atlantic current (compare to the dynamic topography, 0/3000 dbars [*Fukumori et al.*, 1991, Fig. 100]). A conspicuous feature is the lack of any frontal structure north of the Florida Straits because of insufficient data coverage and grid resolution (see the discussion by *Martel and Wunsch* [1993]).

The optimized SST shows a front north of the Florida Straits

(Figure 7b); the model is thus able to correct partially for the data insufficiency there. Between 50° and 65°W, however, the front in the data has almost disappeared, largely because of a strong cooling. For example, the 20°C isotherm is shifted southward by about 1000 km. The difference map (model minus data; Figure 7c) shows that the model is up to 6°C colder in two large areas: the western half of the basin between 32° and 42°N and the entire northeastern Atlantic north of 40°N. A data misfit of 4°–6°C appears quite large; note, however, that the assumed SST data error has an average of almost ±3°C since it encompasses the annual cycle. When the data misfit is weighted by the observation standard error, it exceeds 3 standard deviations only at a few grid points (not shown). Thus the magnitude of the difference between model surface temperature and the data is within the limits set by seasonal variations. Ideally, however, Figure 7c should be white noise in space, because the weight matrix $\mathbf{W}_D$ is diagonal. Instead, there are large-scale structures in the data residuals, reflecting incompatibilities that are not properly accounted for by the assumption of spatially uncorrelated errors.

Surface Temperature Temporal Drift

Figure 7d shows the temporal drift of surface temperature over the integration cycle upon extrapolation over $\tau = 10$ years, that is, $(T_f - T_0)\tau/t$ is displayed, where T_0 is initial temperature, T_f is temperature at the final time, $\tau = 10$ years is the time entering the steadiness weights (see equation (8)), and $t = 500$ days is the integration time. The residuals of the steadiness constraint are quite noisy, but an overall warming tendency exists, with maxima of 4°C around 40°N and 10°C in a small area near the northern buffer zone. In other areas the normalized temporal drifts seldom exceed 1 standard deviation (not shown). Overall, the modeled surface temperature is consistent within error bars with the assumption of a steady state ocean.

Surface Salinity

Observed surface salinity (Figure 8a) shows a more pronounced front along the western edge than does observed surface temperature. This front is widely absent in the optimization (Figure 8b), owing to the disappearance of the low-salinity water along the American coast. In contrast, the
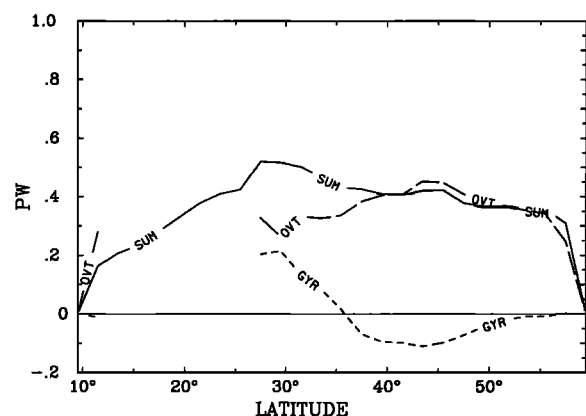

Fig. 6.   Standard experiment; northward heat transport as a function of latitude. Solid curve, total transport; long-dashed curve, overturning contribution; short-dashed curve, horizontal gyre contribution.
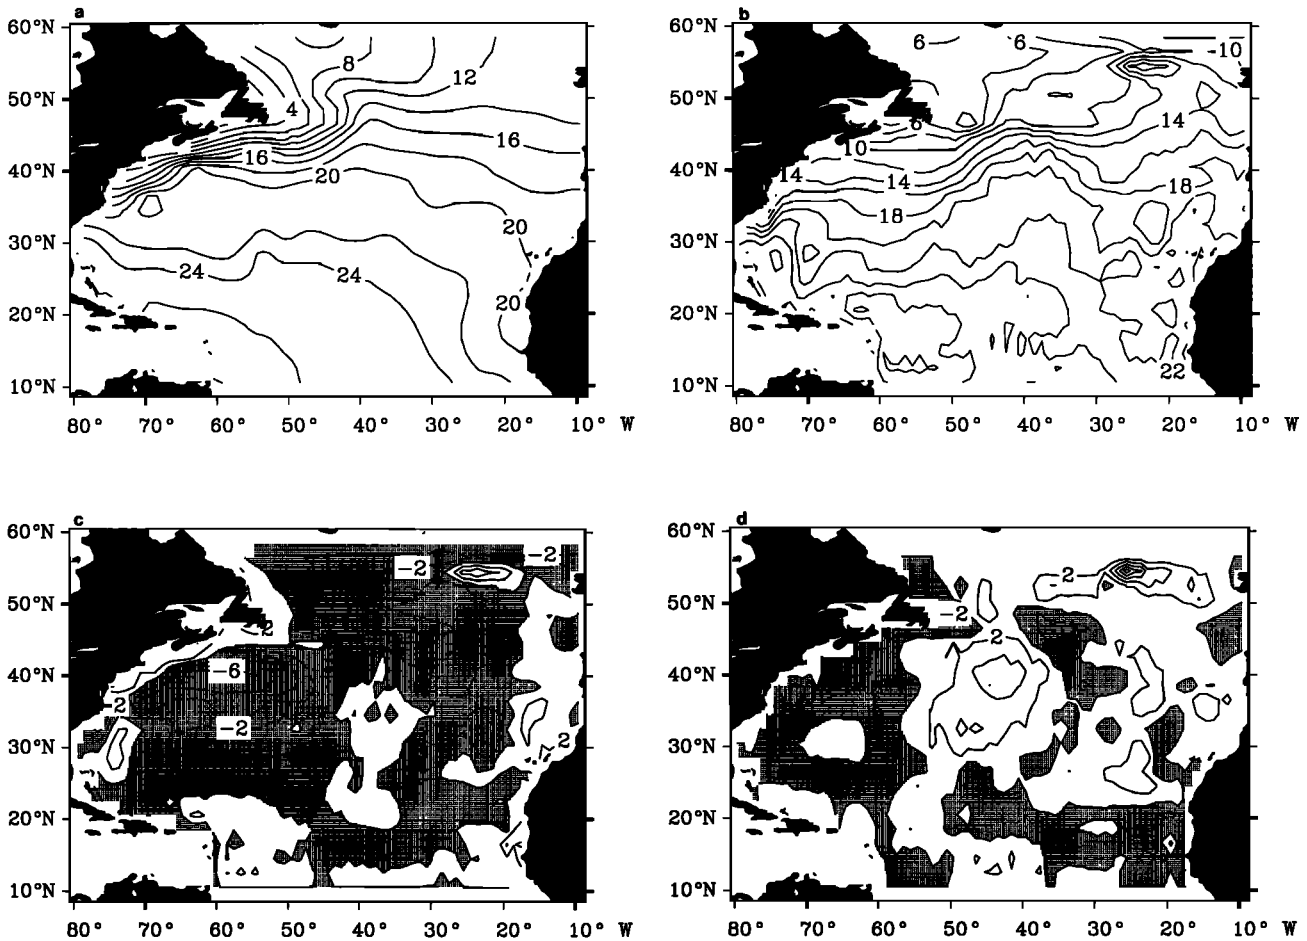
Fig. 7. Sea surface temperature, standard experiment. (a) Climatology; contour interval, 2°C. (b) Optimization; contour interval, 2°C. (c) Optimization minus climatology; contour interval, 2°C. Negative areas are shaded. (d) Temporal drift of optimized state extrapolated over $\tau = 10$ years; contour interval, 2°C. Negative areas are shaded.

two maps coincide very well in the interior (e.g., the maintenance of the subtropical salinity maximum), as confirmed by the map of normalized data residuals (not shown).

Optimized horizontal velocities in level 2 (75 m; Figure 9) form a continuous western boundary current which leaves the coast at about 35°N (the latitude of Cape Hatteras), turns northward at 48°W, and turns eastward again at 42°N to enter the northeastern Atlantic as a broad, fairly smooth current. At 34°N, the Gulf Stream carries 39 Sv to the north, as measured from the coast to the zero line of the northward velocity. The interior flow field is very noisy, but branching southeast of the Grand Banks can be identified with a weak current heading toward the Azores. An eastern boundary current off Africa is clearly visible; another notable feature is the short-circuiting of parts of the inflow from the Straits of Florida to supply the outflow through the Windward Passage. Note the total absence of a counterflow inshore of the Gulf Stream, which could explain why the low salinity of near-coastal waters cannot be maintained. South of 16°N, the flow field is extremely noisy, which might be because of upstream effects of the very crude formulation of the open boundaries, possibly in conjunction with the sponge layer.

Surface Fluxes

The estimation procedure employed here yields as part of the solution optimal estimates of the notoriously poorly

known surface fluxes of heat, freshwater, and momentum. Technically, these are the boundary conditions which, together with the initial conditions, form the set of independent variables that uniquely define the time history of a specific model run. The optimized surface fluxes should represent improved estimates, having been deduced from fairly well-known temperature/salinity and velocity fields as opposed to error-prone bulk parameterizations. But the model accuracy and precision enter here directly. In this standard model, the surface fluxes were employed with formally infinite uncertainties; hence whatever skill exists in the climatologies was not used. Our discussion of surface fluxes is thus based upon the residuals of the model budgets and does not yet contain optimal estimates.

Figures 10a and 10b display climatological and estimated surface heat fluxes, respectively. The "observed" fields [*Isemer et al.*, 1989] show heat loss over the entire North Atlantic except for the upwelling region off Africa. Heat loss over the whole subtropical gyre, while counterintuitive, is necessary in their results to match 1.2 PW of northward transport at 25°N. Three distinct maxima of heat loss occur (over the Gulf Stream Extension, east of the Grand Banks, and adjacent to the northern model boundary). The model roughly maintains strengths and locations of the heat loss maxima (Figure 10b) in the climatology. In between the maxima, however, model air-sea exchange is very different.
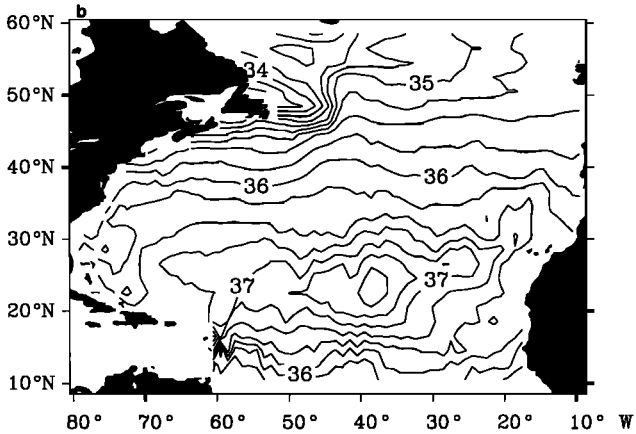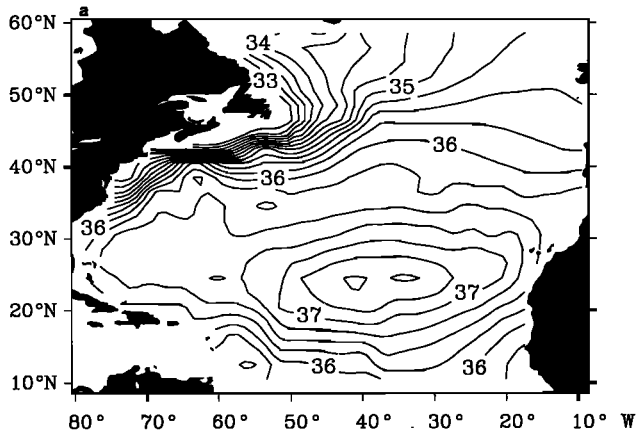
Fig. 8. Sea surface salinity, standard experiment. (*a*) Climatology; contour interval, 0.25 practical salinity unit (psu). (*b*) Optimization; contour interval, 0.25 psu.
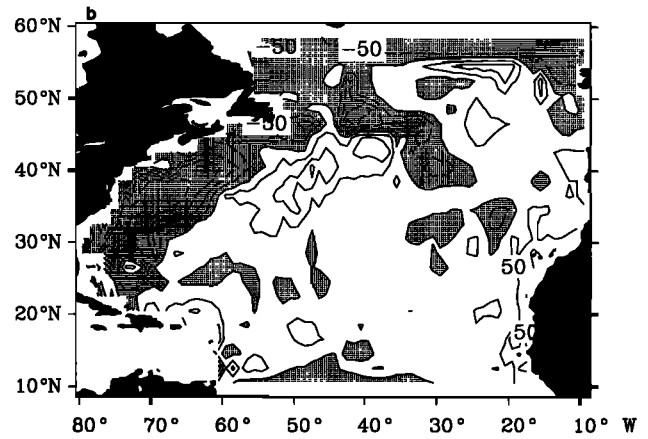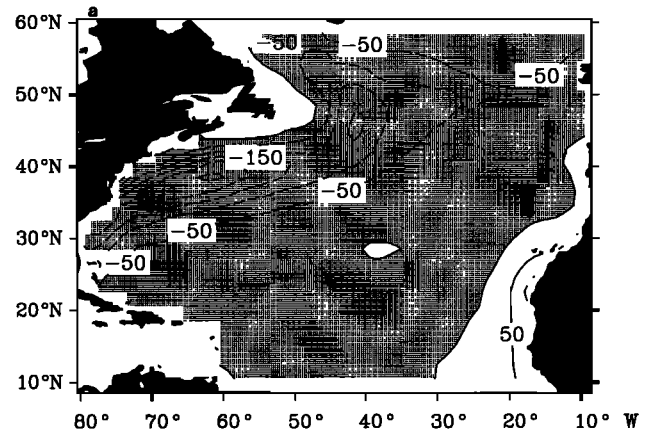

Fig. 10. (*a*) Surface heat gain, climatology; contour interval, 50 W/m². (*b*) Surface heat gain, standard experiment; contour interval, 50 W/m².

Over the bulk of the subtropical gyre, the model ocean gains heat at a rate of up to 50 W/m²; around 40°N, 50°W, heat gain is largely between 50 and 100 W/m², which is a deviation from the climatology of more than 200 W/m². The much reduced heat loss over most of the basin is consistent with the low meridional transport. The heat flux residual (atmo-

sphere to ocean, model minus data; not shown) is positive almost everywhere, so although a misfit of the order of 50 W/m² locally might be within observational error, the large-scale structure marks an incompatibility between the modeled heat fluxes and the climatology and thus with the direct estimates of northward heat transport.

In contrast, the surface freshwater fluxes E-P do not deviate substantially from the climatology (Figures 11*a* and 11*b*). The model shows less freshwater loss south of about 25°N, but the difference is typically 0.3 m/yr or less, well within the (guessed) uncertainty of the observations. Modeled E-P is larger than the climatology in a latitude belt from 25°N to about 40°N by up to 0.5 m/yr, farther north the model requires a larger freshwater input than the climatology suggests; over the western part of the basin the difference is typically between 0.5 and 1 m/yr.

Fields at Depth

The horizontal velocity field at 2400-m depth (Figure 12) shows strong overflows in the northern buffer zone and a continuous southward deep western boundary current (DWBC) following the continental slope. The flow field at 3250 m (not shown) is similar to that in Figure 12; the total southward transport of the DWBC is 16 Sv at 34°N, measured from the coast to the zero line of the meridional

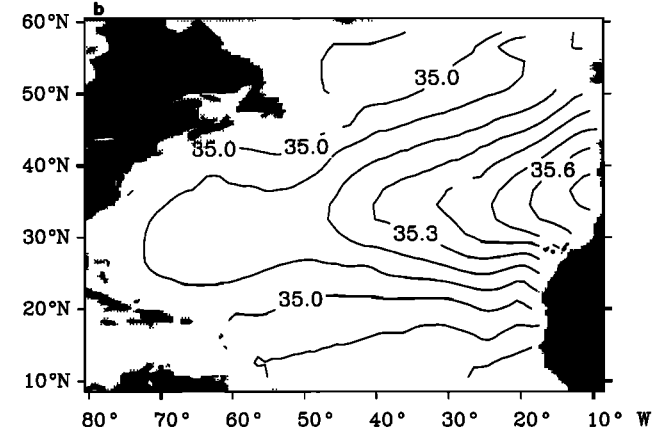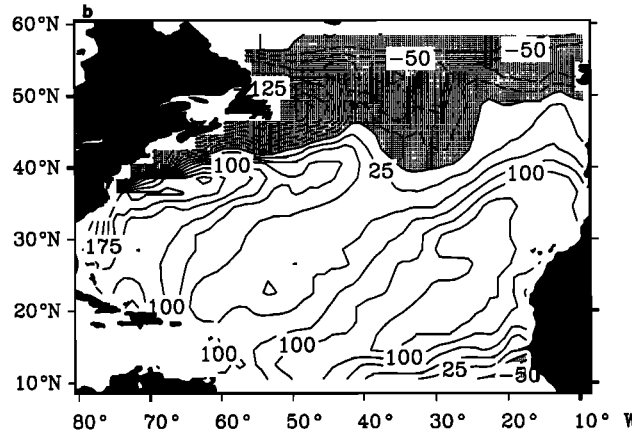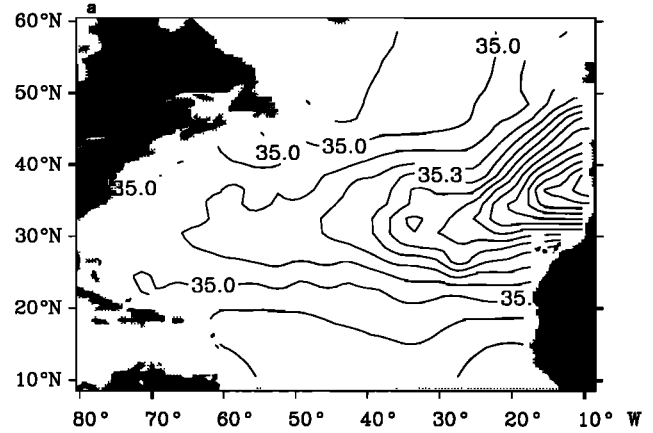**Fig. 9.** Standard experiment; velocity at 75-m depth (level 2). Reference vector denotes 50 cm/s.
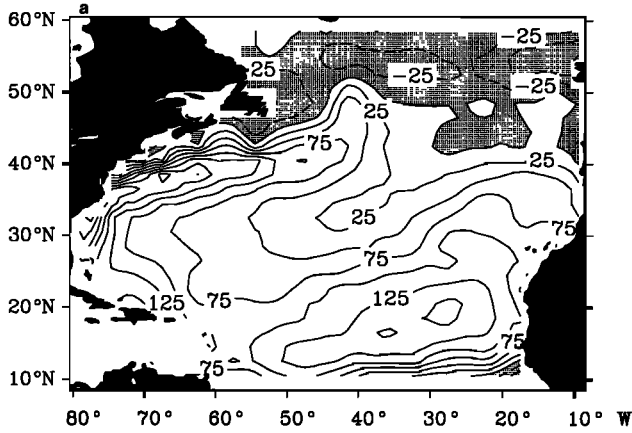
Fig. 11. Evaporation minus precipitation, standard experiment. (a) Climatology; contour interval, 25 cm/yr. (b) Optimization; contour interval, 25 cm/yr.

Fig. 13. Salinity at 1100 m, standard experiment. (a) Climatology; contour interval, 0.1. (b) Optimization; contour interval, 0.1.

velocity, which is 10° offshore. At 26°N, the DWBC is reduced to 8 Sv.

Omitting the influx of Mediterranean water leads to incompatibilities, as the salinity at 1100 m shows. The "data" (Figure 13a) display a front around 36°N at the eastern

Fig. 12. Standard experiment, velocity at 2400-m depth (level 16). Reference vector denotes 10 cm/s.

boundary, with maximum salinity of 36.2 psu. The front is considerably weaker in the optimization (Figure 13b); maximum salinity is now reduced to 35.7 psu. The difference plot (not shown) displays a strong reduction in salinity near the observed maximum. Not only is the modeled salinity tongue weaker than the observed one, but there is also a strong freshening trend in time. The temporal drift of salinity at 1100 m, extrapolated over $t = 10$ years, has a minimum of -0.5 psu. The normalized residuals of both data fit and steadiness demand are about 10 and 5 standard deviations, respectively, near the salinity maximum. The omission of the inflow of Mediterranean water is incompatible with the observed hydrography and the assumption that the model remains fairly steady over 10 years.

Zonal Sections

The model performance on a zonal section along 36°N is illustrated in Figure 14. The measured temperatures (Figure 14a) show a front near the western boundary; surface temperatures are between 20° and 22°C over most of the section. Calculated temperatures are considerably lower than observed ones west of 40°W, and the temperature front is more confined to the continental margin. The model is colder by more than 1°C to a depth of 900 m (Figure 14b) over the entire region between 50° and 70°W. The tendency toward lower temperatures over the western basin is not
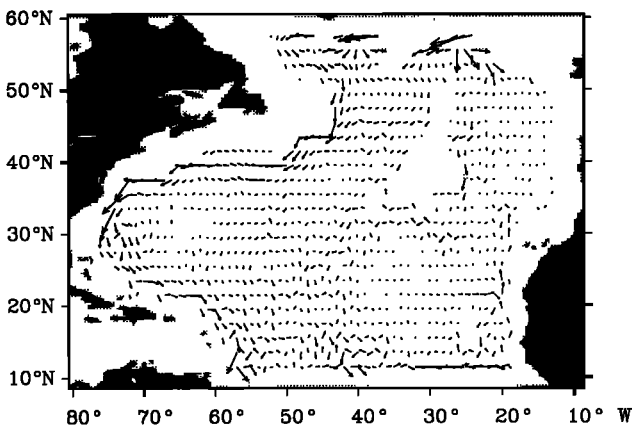
flow field is very noisy, with somewhat stronger southward patches. The DWBC is recognizable as two adjacent bands extending from 700 to 3500 m. Compared to the corresponding figure in *Olbers et al.*'s [1985] beta-spiral fit to *Levitus*'s [1982] climatology, Figure 14c shows Gulf Stream and recirculation more confined to the west and a much more structured near-surface flow field in the eastern North Atlantic. Overall, our sections of horizontal velocities resemble more the inverse solutions by, for example, *Wunsch and Grant* [1982] and *Martel and Wunsch* [1993] than the results of *Olbers et al.* [1985].

The failure of the gridded data to support a Gulf Stream north of the Florida Straits is apparent in Figure 15a, showing temperature along 30°N. In the absence of a nearby hydrographic section, no information is available at 30°N about the steep downward slope of the thermocline off the American continent. The salinity data along the same section do exhibit a front near the western boundary, which is, however, too weak. From a geostrophic fit to the data (including the Ekman flux), the zonally averaged meridional velocity between 30° and 35°N and above 800 m is southward (this very strong departure from other, independent evidence was the reason for running the 250-day spin-up, as discussed in detail below). The model is successful in creating the front necessary to support the Gulf Stream geostrophically, as Figure 15b shows. The complicated structure with the doming at 78°W and associated southward velocities is not yet understood; note, however, the funda-
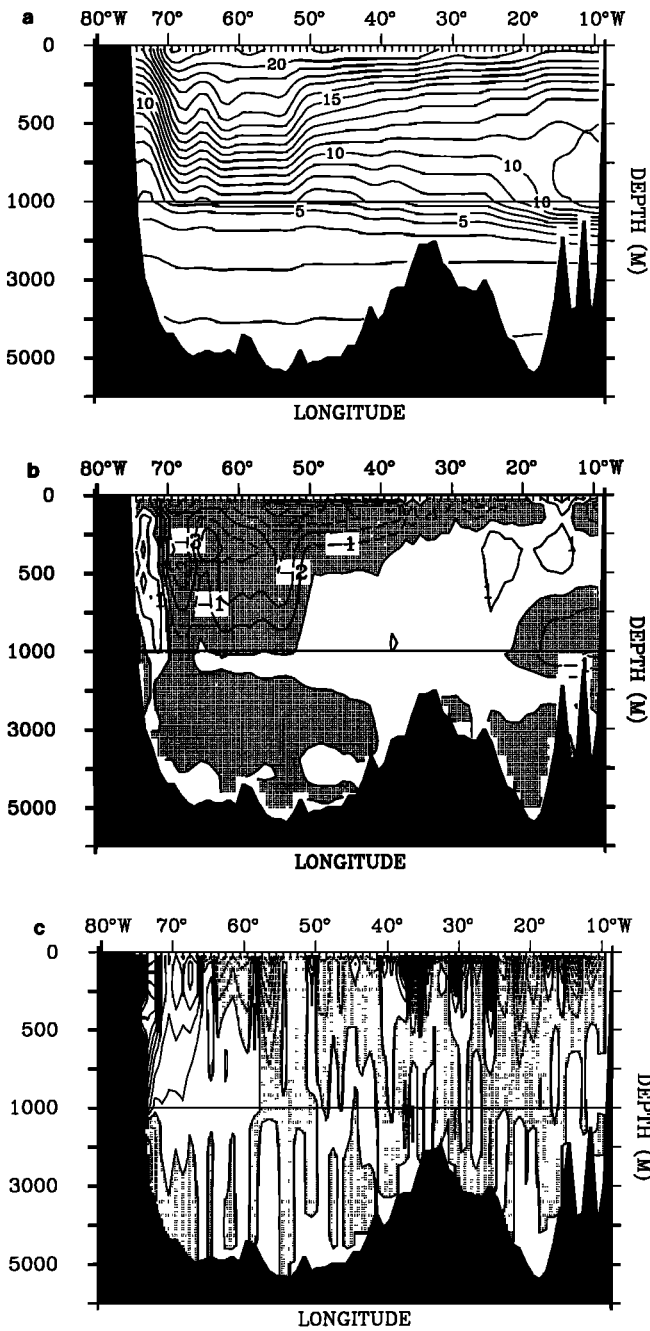
Fig. 14. Zonal section across 36°N, standard experiment. (a) Temperature, climatology; contour interval, 1°C. (b) Temperature, optimization minus climatology; contour interval, 1°C. Negative areas are shaded. (c) Meridional velocity; contour interval, 1 cm/s. Negative areas are shaded.

confined to the surface layers; the model seems generally unable to maintain the bowl structure of the western subtropical gyre. The positive anomalies at the western boundary are associated with the westward shift of the temperature front and hence the maximum surface velocities (when the model is initialized with the data, maximum Gulf Stream velocities are three or four grid points into the interior).

Meridional velocity along 36°N (Figure 14c) shows strong northward flow (maximum 19 cm/s) between 75° and 65°W to depths between 600 and 1000 m. To the east is a broader and shallower southward return flow; still farther eastward, the
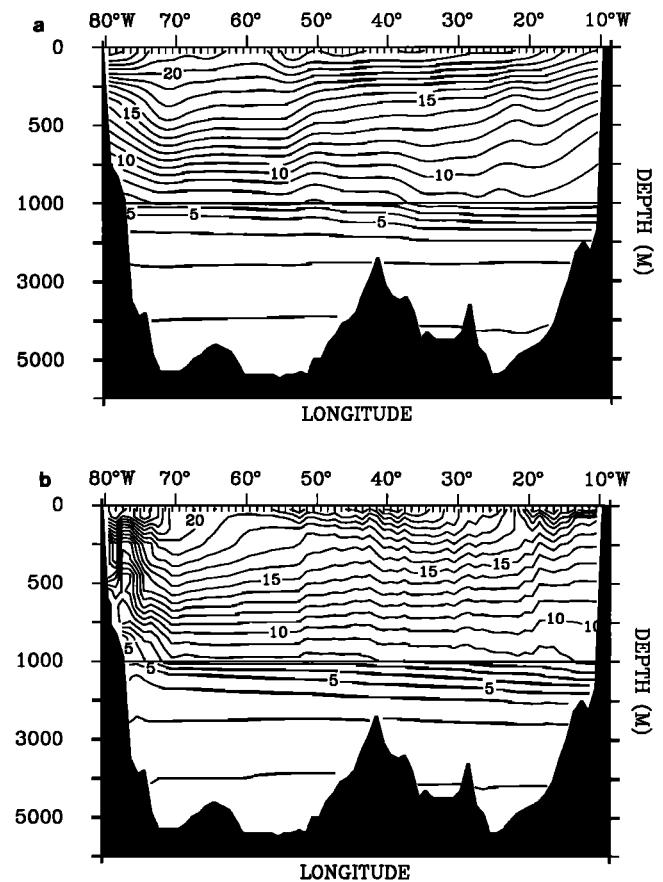
Fig. 15. Zonal section across 30°N, standard experiment. (a) Temperature, climatology; contour interval, 1°C. (b) Optimized temperatures; contour interval, 1°C.

mental advantage of employing a nonlinear inverse model. The combined information about inflow conditions and dynamics change the insufficient hydrographic data and thus provide a more consistent estimate.

The Heat Budget

Figure 16a shows a meridional temperature section along 30°W. The thermocline is deep in the subtropical gyre and shallow approaching the equator; the stratification is stable everywhere. In contrast, the optimized section (Figure 16b) shows convective activity, as identified from vertically outcropping isotherms over large parts. Around 45°N, convention extends to 600-m depth. The setting up of convective activity is associated with a much cooler sea surface; for example, the 10°C sea surface isotherm in the optimization is 800 km farther south than in the data, which locally corresponds to a cooling of up to 5°C.

As the model must represent the observed thermocline water masses, a steady state demand implies that a mechanism for the renewal of these water masses must be supplied. For the very high and very low temperatures, this renewal occurs in the sponge layers through the artificial heat and salt sources, mimicking water mass formation outside the model domain. The bulk of the thermocline water, however, has to be formed within the model domain through convective sinking if there are significant advective (into the sponge layers) and diffusive losses. It is plausible that the model tends toward wintertime conditions to form the water masses at the same rate as they are being destroyed at the boundaries and by interior diffusion. The computed SST resembles wintertime climatologies [e.g., *Shea et al.*, 1990], whereas the hydrographic data used here have a strong bias toward summer observations [*Fukumori et al.*, 1991]. The layer with temperature between 17° and 18°C is much thicker than in the data, so the optimized model forms its 18°C water through convection south of the Gulf Stream.

5. ALTERNATIVE EXPERIMENTS

The standard experiment produced a single minimum of acceptable value. If there are other such minima, we never encountered them despite variations in the search procedure. Four major differences between a model steady state and observations were identifiable: (1) systematically low SST; (2) smearing out of fronts near the western edge, around 36°N; (3) a weak northward mass and consequent heat transport; and (4) weakening Mediterranean salt tongue. The last incompatibility is well understood and could readily be rectified. The low surface temperatures are plausibly explained by the tendency to migrate toward winter conditions.

It is possible, however, that the result of our standard experiment is only a local minimum, not a global one. Advective and convective nonlinearities of the forward model and the large dimension of the state vector mean that there may well be multiple local minima of J of similar or even greater depth (consistent with experience using an idealized GCM [*Marotzke*, 1992]). The systematic difference between the modeled SST and the data might mean only that the solution being used reflects an "accidental" choice of a particular solution in which the conjugate gradient search terminated in a local minimum. We therefore seek, by
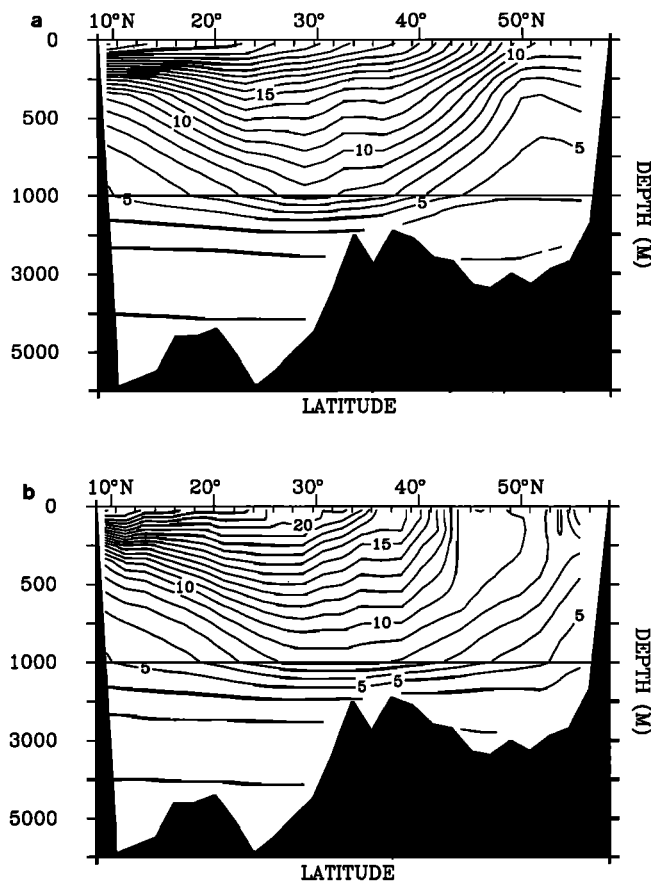

Fig. 16. Meridional section across 30°W, standard experiment. (a) Temperature, climatology; contour interval, 1°C. (b) Optimized temperatures; contour interval, 1°C.

modifying the cost function and therefore the search path, alternative acceptable solutions.

Experiment 2: Enforcing the Data

In experiment 2, tests are made to determine whether the model is fundamentally unable to produce a steady thermocline hydrography closer to the data. The model is forced toward the observations by putting very large weights on the thermocline data. Here we are "curve fitting," the weights being chosen according to taste and convenience, rather than "estimating," where the weights are inverse covariances.

The steadiness weights W_s remain the same as in the standard experiment, but the data weights were modified so that the comparatively large one previously applied to layer 14 is applied now both to layer 14 and to everywhere above it (see Table 2). The model is required to stay close to the data in the entire thermocline, including the western edge, and at the surface.

The resulting minimum of J', obtained after 53 conjugate gradient iterations, is 9.8, which is more than a fourfold increase compared to results from the standard experiment and is too large to characterize a formally compatible solution (the more so since the data misfit contributions to J are roughly an order of magnitude larger than the steady misfit contributions). SST and, to a slightly lesser degree, surface salinity remain close to the observations (not shown), but the
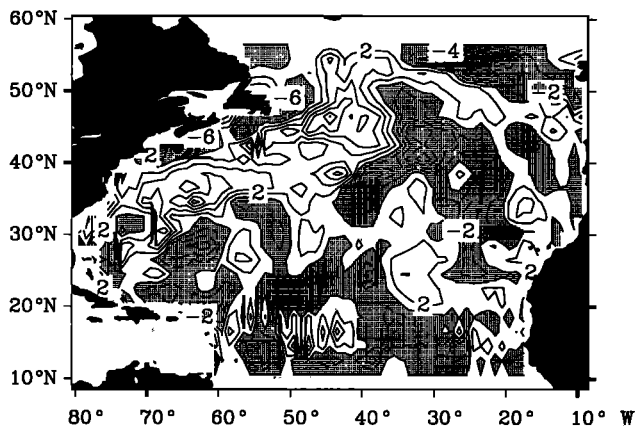
Fig. 17. Experiment 2 (enforcing the data), sea surface temperature. Temporal drift of optimized state extrapolated over $\tau = 10$ years; contour interval, 2.0°C. Negative areas are shaded.

Fig. 19. Experiment 3 (enforcing the meridional flux), sea surface temperature. Temporal drift of optimized state extrapolated over $\tau = 10$ years; contour interval, 2.0°C. Negative areas are shaded.

steady residuals in SST show large-scale drifts of 6°–8°C over 10 years over the Gulf Stream extension and east of the Grand Banks (Figure 17). The normalized steadiness residuals in SST (not shown) are above 3 standard deviations only in isolated grid points, however. Because of the very large penalties on data misfits, the normalized data residuals take very large values (up to 40 standard deviations if the large weights are interpreted as reciprocal variances). In the western parts of the 36°N section, the data residuals are still considerable (Figure 18; compare with Figure 14b), resulting in very large contributions to J (10–20 standard deviations pointwise).

Although not a rigorous statistical result, this experiment is interpreted as showing that the climatologies we are using are not consistent with a steady state circulation (as determined from the present GCM), and we reject this solution.

Experiment 3: Enforcing the Meridional Flux

Meridional overturning and northward heat transports are very similar in experiments 1 and 2; in both cases the maximum heat transport at 27.5°N is only 0.5 PW. In experiment 3 we try to enforce a stronger heat transport by also penalizing deviations of the estimated surface fluxes
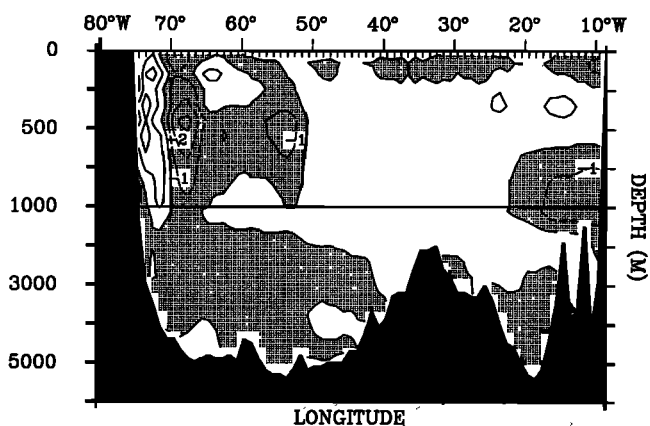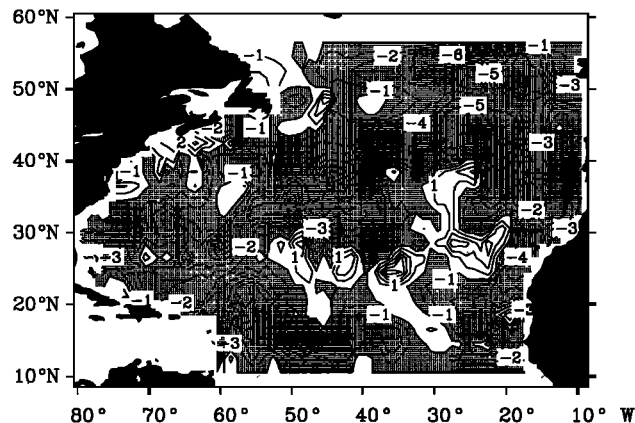

Fig. 18. Zonal temperature section across 36°N, experiment 2 (enforcing climatology). Optimization minus climatology; contour interval, 1°C. Negative areas are shaded.

from their respective climatologies (recall that the climatology was forced to consistency with 1.2 PW across 25°N).

Penalty terms for deviations from observed heat flux, E-P, and wind stress were introduced into J with diagonal weights. As in experiment 2, the weights are no longer purely statistically based but instead represent an experiment to force the model toward the surface climatology.

Experiment 3 yields a minimum J of 3.8 after 54 iterations, so J is well within the range characterizing overall compatibility. Only the deviation of model temperature from the data has an rms value of more than 1 standard deviation (1.5 σ). However, the main goal has not been achieved: the maximum northward heat transport is increased only to 0.65 PW at 30°N, with a corresponding mass transport of only 11 Sv. The difference between the model and the surface heat flux data, however, is much smaller than in experiment 1; the maximum misfit is now 80 W/m^2 compared to more than 200 W/m^2 in the standard experiment. Larger excursions from surface flux data are prohibited in experiment 3. The model reconciles weak northward transport and relatively strong heat loss to the atmosphere by producing a large temporal drift. Figure 19 shows that surface temperature decreases with time over almost the entire basin; extrapolated over 10 years, maximum cooling reaches 8°C. The basinwide coherent cooling pattern indicates a systematic bias rather than a random statistical error; the data misfit map (not shown) shows the model to be colder than the data by up to 8°C everywhere except in narrow bands off the American and African coasts. The Isemer et al. [1989] heat fluxes and the implied meridional transports are not compatible with the model in steady state.

Experiment 4: Initializing From the Data

All experiments discussed so far produce a thermocline west of 50°W near 36°N that is too cold and too fresh, and all produce a meridional transport of mass (and consequently of heat) that is too small relative to the climatologies. A purely geostrophic fit to the gridded data together with the inflow from the Florida Straits does produce a heat transport of about 1.1 PW at 27.5°N (such a fit does not address the physical consistency requirements that a dynamical model

places on different model regions). At the end of the spin-up, however, the maximum heat transport has dropped to 0.8 PW. Lower temperatures at the sea surface and in the western-edge thermocline have developed.

It is possible that in the spin-up phase, the model was driven into a "cold" region in phase space, meaning that a different first estimate of the initial conditions might lead to a better result. Why a spin-up was nevertheless run prior to the minimization becomes clear from experiment 4, in which the observed hydrography serves as the first estimate of the initial conditions. Integration time is 4.1 years, and the cost function reaches 5.8 after 36 iterations. Since the Gulf Stream is missing from the data around 30°N, there is no basinwide overturning pattern in the first estimate of the initial conditions; the flow from the Florida Straits is forced to move downward. The optimization produces only a partial remedy: the optimized initial conditions show reduced downwelling north of the Florida Straits but still no single overturning cell. After 250 days, a single overturning cell of only 12 Sv has developed. At the end of the integration, overturning strength has dropped dramatically, with a maximum of only 6 Sv. The optimization is essentially acting on the adjustment process toward maintaining a continuous Gulf Stream (and consequently a single overturning cell). Why the meridional circulation becomes so weak after 4 years is not clear, but it could simply be that the model is moving toward a particular best fit that has a weak circulation.

6. UNCERTAINTY ESTIMATES

A widely perceived drawback of the adjoint or Pontryagin principle method in data assimilation is the lack of an uncertainty estimate of the solution; in contrast, sequential estimation actually requires computation of the covariance at every time step, but the resulting large computational load renders that method infeasible as an algorithm to solve the estimation problem considered here. Estimates of the ocean circulation without some understanding of their reliability are, however, of only limited utility (they show only that a solution is possible, not what has really been determined by other than first-guess or ancillary requirements like smoothness). If a model consistent with the observations is found, one must obtain some indication of what is well and what is poorly determined. The present model is not actually consistent, and thus we could ignore the issue, but the problem is so important as a future element in the use of optimization methods that we briefly summarize the situation.

The calculations displayed in Figure 4 for two elements of the state vector were used to produce rough estimates of the uncertainties due to the finite curvature near the minimum, i.e., a few tenths of a degree for temperature. One can carry out such computations (they are expensive) for each state vector element in turn, but knowledge of the covariances is still missing.

Complete versions of these computations, leading to formal uncertainty estimates for nonlinear estimation problems, are available [e.g., *Thacker*, 1989; *Seber and Wild*, 1989]. Generally, they involve finding the Hessian of the objective function J subject to the model constraints. For problems of the dimensionality of the present model, the computational load has so far deterred anyone from actually carrying out the recipe.

Because of the great importance of the state vector uncertainty to the use of adjoint methods, an attempt was made to find an efficient algorithm for estimates of this quantity (the details of the computation are confined to Appendix B). As described in Appendix B, the result was not wholly satisfactory, and the best summary statement we can make is that we believe there are uncertainties in our results for the standard model of the order of 0.1°-0.4°C and equivalent uncertainties for salinity, but we have little or no insight into the correlations of the values. Finding methods for full determination of the uncertainties of the adjoint solutions remains on the list of very high priorities for future work.

7. DISCUSSION AND OUTLOOK

The combination of the GFDL GCM and its adjoint proves to be a practical method for combining model dynamics with a variety of observations. In a formal sense, we have shown that the model in steady state and the data are inconsistent. The methodology shows the inconsistencies quantitatively and permits exploration of various hypotheses. Integration of a GCM over decades and its comparison with observations is an extremely stringent test of both model and data: small systematic errors in either can accumulate and produce qualitative effects which are invisible in "instantaneous" inversions.

Not surprisingly, the inconsistencies confront one with fundamental questions about the model and the data. On the data side, one is driven to ask whether the uncertainty estimates are realistic, whether the (usually undiscussed) spatial covariances of the data error are significantly different from white noise, and whether the spatial and temporal aliasing of the climatologies have been adequately modeled. On the model side, a host of issues arises. Inconsistency could result from lack of spatial resolution, simplification of the boundary conditions, missing physics (e.g., linearization everywhere or in specific regions), inadequate parameterization (e.g., the convective adjustment), or general failure of a hypothesis (e.g., the ocean is not in steady state). The combination of model and data could fail for purely technical reasons, for example, if the optimization algorithm has produced a local rather than a global minimum.

These problems are difficult ones, but if the experience of the present trials is any guide, there are no fundamental obstacles to overcoming any of them. Because of the realities of oceanic observations, we believe that real progress in future understanding of the general circulation will come only by methods such as the one used here, although not necessarily exactly this one. The present result is reasonably successful, focusing quantitative machinery on inconsistencies between observations and models (theory). Although it has a long and honored tradition, the lumping together of nonsynoptic observations to produce pictures of "the" ocean circulation has long been suspected to be untenable. The present results show that we have reached a stage in which the limitations of that approach to the general circulation can no longer be defended except in a qualitative fashion.

What has been accomplished? Hydrography has been shown to be "assimilable" into a GCM. The multiple time step demand for a steady state proves practical in a way that the original single time step approach of *Tziperman and Thacker* [1989] did not.

A number of elements of the model and data require special scrutiny in any attempt to go beyond the present results. In experiment 1, the model tends toward wintertime conditions with a very deep mixed layer at high latitudes. Enforcing the observed hydrography in the thermocline (experiment 2) makes the best estimate considerably more unsteady, so experiments 1 and 2 suggest that fitting a climatological annual cycle would make consistency between model and data much easier to achieve.

One might wonder whether a very different result would have been achieved had the model been run over a much longer interval. *Schiller* [1993] has run a similar computation for as long as 10 years; he concludes that no significant changes occur when the longer durations are used, and his conclusions are very similar to ours.

In no experiment has the model been able to maintain the observed hydrography in the boundary current region. A serious candidate for an explanation is the lack of resolution: in the climatological temperature/salinity fields, the front associated with the instantaneous Gulf Stream has been smeared out through the mapping, leading to smaller geostrophic velocities. It is unlikely that a circulation conserving heat and salt can be constructed from an averaged density field in regions of strong currents where strong covariances of property and velocity extremes must occur. The usefulness of mapped data for direct model initialization and testing has probably reached its limits of utility (proper model averages might be tested against the climatologies, however).

Taking the direct estimates of meridional heat fluxes as close to the correct value, the northward heat transport has been persistently too low in all experiments (0.5 to 0.65 PW), a shortcoming in almost all published model simulations of the North Atlantic. The attempt to enforce higher transports through large weights on the surface flux observations has not been very successful (experiment 3).

Recently, the model was extended to permit meridional transport constraints to be enforced directly. Preliminary results indicate that northward mass and heat transports of 18 Sv and 1.2 PW, respectively, can be enforced at 29.5° and 35.5°N, consistent with work by *Rintoul and Wunsch* [1991]. The Gulf Stream sea surface temperatures are then 8°C higher than those observed at 34°N, and the steadiness penalties are very large, all suggesting that the model is fundamentally not compatible simultaneously with observed hydrography, transports, and steadiness. The low resolution is likely the chief culprit here. The model does not capture the downstream increase in strength of the Gulf Stream; hence the high-temperature, high-speed core and thus a considerable contribution to the temperature flux are missing. Similarly, the deep western boundary current is diffuse and lacking the observed low temperatures of section data.

Another possible candidate for the too-weak meridional circulation is the formulation of the southern boundary condition; only 8 Sv of overturning and 0.16 PW of heat transport emerge out of the southern buffer zone. Both the sponge layers and the prescribed inflow/outflow points cause numerical problems, visible especially in the vertical velocity field. Estimating the fluxes in and out of open lateral boundaries, analogous to estimating the surface fluxes, must have high priority for future work, although it is possible that one would still observe undesired upstream effects at outflow points, as are visible in the velocity field in Figure 9.

A conspicuous property of this model is a generally very noisy vertical velocity field after the optimization. A forward model run typically has noisy w along the boundary, as is common in GCMs. The optimized state shows sign changes in w typically over 1–5 grid points. A. Schiller and J. Willebrand (personal communication, 1992) performed related GCM inversions for the Atlantic with the GFDL model but with an approximate adjoint. They found very smooth hydrographic and velocity fields. The most likely explanation for the different results is the use in our calculations of a Rayleigh friction which, although tuned to give the same amount of friction on a grid-size scale as would a reasonable choice of eddy viscosity, does not couple horizontal velocities at neighboring grid points and does not have the same smoothing effect as Laplacian friction. Another possibility is the lack of gravity waves, which might adjust the small-scale structures of the updates of the density field to a smooth velocity field; in a strictly geostrophic model, noisy density creates noisy geostrophic flow which, through advection, would tend to enhance the noise in temperature and salinity.

With the possible exception of the uncertainty estimation, none of these problems is intractable. Optimization methods remain the only methods known for combining dynamics with real observations, and we expect rapid progress as the present model is modified in the directions suggested.

APPENDIX A: PREPARATION OF HYDROGRAPHIC DATA

The climatologies by *Fukumori et al.* [1991] and *Fukumori and Wunsch* [1991] are both based on the same data, which are first interpolated onto a set of standard depths; the subsequent processing differs, however. *Fukumori et al.* [1991] perform a purely horizontal objective analysis of the station data for some selected depths; the horizontal grid spacing is 1° longitude by 1° latitude. In contrast, *Fukumori and Wunsch* [1991] first project the station data onto a set of vertical empirical orthogonal functions (EOFs); the expansion coefficients up to degree 6 are then objectively analyzed on the same grid. From the gridded modal expansion coefficients and the EOFs, *Fukumori and Wunsch* [1991] construct their climatology, which should be superior to the purely horizontal objective analysis because information about vertical correlations is employed.

As *Fukumori and Wunsch* [1991] use only stations deeper than 1100 dbar to construct the EOFs, thereby excluding in particular the Atlantis 109 sections Florida-Fort Pierce at 26° and 27°N, the information about the high temperatures in the surface layers of these sections does not enter the gridded fields. Consequently, surface temperatures north of the Florida Straits are too low in the *Fukumori and Wunsch* [1991] climatology. As a remedy, it was decided to include information from the *Fukumori et al.* [1991] atlas in the following simple way. At depths of 300 dbar and deeper, the modal expansion data [*Fukumori and Wunsch*, 1991] are used. At standard depths of 0 and 50 dbar, the original station data are objectively analyzed, using exactly the same parameters as used by *Fukumori et al.* [1991]. The gridded 0- and 50-dbar values are then averaged vertically to remove static instability; one could view this as defining a mixed layer of 50-dbar depth everywhere. At each horizontal grid point, we then interpolate linearly between 50 dbar and the 300-dbar fields from the *Fukumori and Wunsch* [1991] modal expansion data. Some of the information content of the data

is removed through this rather crude merging procedure; attempts with more elaborate schemes, however, resulted in static instability in large parts of the subsurface layers and were therefore abandoned. As a final step in the preparation of the hydrographic data, they are subsampled on the model grid of 1° longitude by 2° latitude and transformed to the model vertical grid (Table 1); the temperatures are transformed to potential temperatures.

APPENDIX B: ESTIMATING THE HESSIAN

A method is presented to calculate approximately the diagonal elements of the Hessian for a steady state problem. Assume that through some minimization a steady state of a GCM has been found and that one can linearize the model around that steady state. A generic linear or linearizable model is defined from (2) by the time-stepping rule

$$\mathbf{x}(t + 1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{q}(t) + \mathbf{\Gamma}\mathbf{u}(t) \qquad (B1)$$

where $\mathbf{A}$ is the state transition matrix, and the other variables are as defined previously. In steady state, (B1) becomes

$$\mathbf{x} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{q} + \mathbf{\Gamma}\mathbf{u} \qquad (B2)$$

or

$$(\mathbf{A} - \mathbf{I})\mathbf{x} + \mathbf{B}\mathbf{q} + \mathbf{\Gamma}\mathbf{u} = 0 \qquad (B3)$$

where $\mathbf{I}$ is the identity. To apply (B3) to the problem discussed in this paper, note that although a strictly steady state has not been reached, one can interpret the residual temporal drifts of the solutions as the stochastic forcing components in (B3). The identification of stochastic interior forcing (equivalent to artificial interior heat and salt sources) with residual drifts in time is similar to the interpretation sometimes made of the "robust diagnostic" technique [see *Tziperman et al.*, 1992b].

Equation (B3) is a set of simultaneous linear equations that can be viewed to determine $\mathbf{x}$ as a function of $\mathbf{q}, \mathbf{u}$. This equation has already been solved through the adjoint method, but we now interpret $\mathbf{u}$ in (B3) as "data" with zero expected value and known variance $\mathbf{Q}^{-1}$. The expected uncertainty of $\mathbf{x}$ is then

$$\langle(\hat{\mathbf{x}} - \mathbf{x})(\hat{\mathbf{x}} - \mathbf{x})^T\rangle = (\mathbf{A} - \mathbf{I})^{-1}\mathbf{\Gamma}\mathbf{Q}^{-1}\mathbf{\Gamma}^T(\mathbf{A} - \mathbf{I})^{-T} \qquad (B4)$$

Note that $\mathbf{x}$ now is constrained not only by the steady state demand (B3) but also by the data, with weight $\mathbf{W}_D$, and is determined from these two different sets of constraints as a weighted mean. The Hessian, that is, the inverse of the covariance matrix of the solution, is

$$\langle(\hat{\mathbf{x}} - \mathbf{x})(\hat{\mathbf{x}} - \mathbf{x})^T\rangle^{-1} = \mathbf{W}_D + (\mathbf{A}^T - \mathbf{I})\mathbf{\Gamma}^{-T}\mathbf{Q}\mathbf{\Gamma}^{-1}(\mathbf{A} - \mathbf{I})$$

$$(B5)$$

(In the present special case, $\mathbf{\Gamma}$ is square and diagonal.)

Equation (B5) is analogous to the variance σ^2 of the weighted mean of two measurements, with variances σ_1^2 and σ_2^2. One has then

$$\sigma^{-2} = \sigma_1^{-2} + \sigma_2^{-2} \qquad (B6)$$

Equation (B5) is readily seen as a generalization of (B6). It has been derived here somewhat heuristically, but a more rigorous derivation, starting from the second variation of the augmented cost function J' [see *Bryson and Ho*, 1975, for details], yields the same result.

Two limiting cases of (B5) are readily identified. First, let the elements of $\mathbf{Q}$ go to infinity; that is, the forcing is known exactly. This is the classic well-posed boundary value problem; the observations become unimportant, and the model response is perfectly known. Second, assume that the uncertainty on the stochastic forcing is infinite; that is, the elements of $\mathbf{Q}$ are zero, which is equivalent to not employing any dynamical constraint at all. In that limit,

$$\langle(\hat{\mathbf{x}} - \mathbf{x})(\hat{\mathbf{x}} - \mathbf{x})^T\rangle = \mathbf{W}_D^{-1} \qquad (B7)$$

and the model has not added any information.

To evaluate (B5), the state transition matrix, $\mathbf{A}$ of the numerical model must be found. For a linear model, *Fukumori et al.* [1993] obtain $\mathbf{A}$ by time stepping initial states successively constructed from unit vectors specifying the model state (i.e., different initial states in which only one element of $\mathbf{x}(0)$ is taken to be nonzero). The situation is more complicated here; the state vector is an order of magnitude larger (with about 50,000 elements), and the model is nonlinear, but the adjoint model is linear in its state vector $\mathbf{\mu}(t)$. One can thus use a numerical procedure similar to that of *Fukumori et al.* [1993] but applied to the adjoint instead of the forward model. Using successively all unit vectors of the adjoint model as "initial" states and stepping the adjoint backward one time step produces the columns of $\mathbf{A}^T$ and hence $\mathbf{A}$.

This procedure would require as many time steps of the adjoint as the state vector has elements (about 50,000), making the computation of $\mathbf{A}$ as time-consuming as the entire minimization. Furthermore, the storage requirements would be excessive ($\mathbf{A}$ has 2.5×10^9 elements). Considerable reduction in load can be obtained by observing that to a very good approximation, a unit source in the adjoint model is propagated only over a very limited volume in one time step. Note that the response of the adjoint model to a unit source, which for the sake of argument is taken to be the Lagrange multiplier for the temperature equation at grid point k, marks which elements of the state vector influence temperature at point k during the last time step of the forward model. The only immediate, long-distance propagation of information in the forward model is through the calculation of the external-mode stream function; these contributions are extremely small, however, and can be neglected. Advection and diffusion produce purely nearest-neighbor interactions over one time step in both the horizontal and the vertical. The geostrophic constraint is nonlocal in the vertical: velocity at each depth is influenced by all temperature and salinity points in the water column under consideration and by all temperature and salinity points in every neighboring column. Perturbations in the heat budget at point k are thus caused either by perturbations of the temperatures in the neighboring six boxes that are advected by the mean flow, or by perturbations of the geostrophic flow that advect mean temperatures. Velocity perturbations are caused by perturbations in both temperature and salinity over an entire 3 × 3 array of water columns.

A unit source is thus spread by the adjoint model in two ways: it is advected and diffused over neighboring boxes, but it is also distributed over the block of 3 × 3 water columns. The spreading over the entire basin through the external-mode computation is neglected, and to this approximation

every column of $\mathbf{A}^T$ has only $9 \times 24 \times 2 = 432$ nonzero elements (nine water columns; 24 layers; two tracers, i.e., temperature and salinity), which is about 1% of the size of the state vector. The banded structure of $\mathbf{A}^T$ can be used to speed up its computation by 2 orders of magnitude. Instead of integrating only one unit source backward, one "seeds" the model grid with unit sources such that the responses to the individual sources do not overlap. $\mathbf{A}^T$ is obtained in a highly compressed form, and the bookkeeping to reassign the correct column indices is very tedious. Also, the compressed $\mathbf{A}^T$ is still 300 Mbytes, and the evaluation of the triple matrix product (B4) poses some coding challenges. In the first attempt reported here, it was thus decided to compute only the diagonal elements of the Hessian, which is straightforward because Γ and $\mathbf{Q}$ are diagonal matrices. For a generic product $\mathbf{D} = \mathbf{B}^T\mathbf{C}\mathbf{B}$, where $\mathbf{C}$ is diagonal with elements C_i,

$$D_{ii} = \sum_j C_j (B_{ij}^T)^2 \qquad \text{(B8)}$$

The symbol j indexes the unit sources; inspection shows that the diagonal elements of the Hessian can be calculated "on line" without obtaining the complete $\mathbf{A}$ explicitly. As a final, possibly drastic approximation, it is assumed that the variances of the solution are the inverses of the diagonal elements of the Hessian rather than the diagonals of the inverse.

Apart from the latitude strip adjacent to the northern sponge layer, the formal error of the model solution is estimated as of the order of 0.01°C, a value too low to be a reasonable estimate of the real uncertainty of the model solution. Given the effort that went into computing it, the result is disappointing, and we must look elsewhere for an appropriate estimate.

REFERENCES

Bryan, K., A numerical method for the study of the circulation of the world ocean, *J. Comput. Phys.*, *3*, 347–376, 1969.

Bryan, K., and M. D. Cox, An approximate equation of state for numerical models of the ocean circulation, *J. Phys. Oceanogr.*, *2*, 510–514, 1972.

Bryson, A. E., Jr., and Y.-C. Ho, *Applied Optimal Control*, 481 pp., revised printing, Hemisphere, New York, 1975.

Bunker, A. F., Computations of surface energy flux and annual air-sea interaction cycles of the North Atlantic Ocean, *Mon. Weather Rev.*, *104*, 1122–1140, 1976.

Cox, M. D., A primitive equation, three-dimensional model of the ocean, *GFDL Ocean Group Tech. Rep. 1*, Geophys. Fluids Dyn. Lab., Princeton, N. J., 1984.

Fukumori, I., and C. Wunsch, Efficient representation of the North Atlantic hydrographic and chemical distributions, *Prog. Oceanogr.*, *27*, 111–195, 1991.

Fukumori, I., F. Martel, and C. Wunsch, The hydrography of the North Atlantic in the early 1980s: An atlas, *Prog. Oceanogr.*, *27*, 1–110, 1991.

Fukumori, I., J. Benveniste, C. Wunsch, and D. B. Haidvogel, Assimilation of sea surface topography into an ocean circulation model using a steady-state smoother, *J. Phys. Oceanogr.*, *23*, 1831–1855, 1993.

Hall, M. M., and H. L. Bryden, Direct estimates and mechanisms of ocean heat transport, *Deep Sea Res.*, *29*, 339–359, 1982.

Isemer, H. J., J. Willebrand, and L. Hasse, Fine adjustment of large scale air-sea energy flux parameterization by a direct estimate of ocean heat transport, *J. Clim.*, *2*, 1173–1184, 1989.

Leaman, K. D., R. L. Molinari, and P. S. Vertes, Structure and variability of the Florida Current at 27°N April 1982–July 1984, *J. Phys. Oceanogr.*, *17*, 565–583, 1987.

Leaman, K. D., E. Johns, and T. Rossby, The average distribution of volume transport and potential vorticity with temperature at three sections across the Gulf Stream, *J. Phys. Oceanogr.*, *19*, 36–51, 1989.

Levitus, S., Climatological atlas of the world ocean, *Prof. Pap. 13*, 173 pp., Natl. Oceanic and Atmos. Admin., Washington, D.C., 1982.

Luenberger, D. G., *Introduction to Dynamic Systems: Theory, Models and Applications*, 446 pp., John Wiley, New York, 1979.

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.

Manabe, S., and R. J. Stouffer, Two stable equilibria of a coupled ocean-atmosphere model, *J. Clim.*, *1*, 841–866, 1988.

Marotzke, J., The role of integration time in determining a steady state through data assimilation, *J. Phys. Oceanogr.*, *22*, 1556–1567, 1992.

Martel, F., and C. Wunsch, The North Atlantic circulation in the early 1980s—An estimate from inversion of a finite difference model, *J. Phys. Oceanogr.*, *23*, 898–924, 1993.

Mercier, H., M. Ollitrault, and P. Y. Le Traon, The contribution of Lagrangian floats to the knowledge of the North Atlantic general circulation, *J. Phys. Oceanogr.*, *23*, 689–715, 1993.

Olbers, D. J., M. Wenzel, and J. Willebrand, The inference of North Atlantic circulation patterns from climatological hydrographic data, *Rev. Geophys.*, *23*, 313–356, 1985.

Rintoul, S. R., and C. Wunsch, Mass, heat oxygen and nutrient fluxes in the North Atlantic Ocean, *Deep Sea Res., Part A*, *38*, suppl. S355–S377, 1991.

Roemmich, D., Circulation of the Caribbean Sea: A well-resolved inverse problem, *J. Geophys. Res.*, *86*, 7993–8005, 1981.

Roemmich, D., and C. Wunsch, Two transatlantic sections: Meridional circulation and heat flux in the subtropical North Atlantic Ocean, *Deep Sea Res.*, *32*, 619–664, 1985.

Salmon, R., A simplified linear ocean circulation theory, *J. Mar. Res.*, *44*, 695–711, 1986.

Schiller, A., Assimilation hydrographischer Daten in ein Zirkulationsmodell des Atlantiks auf der Grundlage des adjungierten Verfahrens, Ph.D. thesis, 157 pp., Univ. Kiel, Kiel, Germany, 1993.

Schmitt, R. W., P. S. Bogden, and C. E. Dorman, Evaporation minus precipitation and density fluxes for the North Atlantic, *J. Phys. Oceanogr.*, *19*, 1208–1221, 1989.

Seber, G. A. F., and C. J. Wild, *Nonlinear Regression*, 768 pp., John Wiley, New York, 1989.

Shea, D. J., K. E. Trenberth, and R. W. Reynolds, A global monthly sea surface temperature climatology, *NCAR Tech. Note, NCAR TN345+STR*, 167 pp., Natl. Cent. for Atmos. Res., Boulder, Colo., 1990.

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., The role of the Hessian matrix in fitting models to measurements, *J. Geophys. Res.*, *94*, 6177–6196, 1989.

Thacker, W. C., and R. B. Long, Fitting dynamics to data, *J. Geophys. Res.*, *93*, 1227–1240, 1988.

Trenberth, K. E., J. G. Olson, and W. G. Large, A global ocean wind stress climatology based on ECMWF analyses, *NCAR Tech. Note, NCAR TN-338+STR*, 93 pp., Natl. Cent. for Atmos. Res., Boulder, Colo., 1989.

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

Tziperman, E., W. C. Thacker, R. B. Long, and S.-M. Hwang, Oceanic data analysis using a general circulation model, I, Simulations, *J. Phys. Oceanogr.*, *22*, 1434–1457, 1992*a*.

Tziperman, E., W. C. Thacker, R. B. Long, S. Hwang, and S. R. Rintoul, Oceanic data analysis using a general circulation model, II, A North Atlantic model, *J. Phys. Oceanogr.*, *22*, 1458–1485, 1992*b*.

Willebrand, J., and C. Wunsch, Inversion of ocean circulation models, *Eos Trans. AGU*, *71*, 2–3, 5, 1990.

Wunsch, C., Transient tracers as a problem in control theory, *J. Geophys. Res.*, *93*, 8099–8110, 1988.

Wunsch, C., and B. Grant, Towards the general circulation of the North Atlantic ocean, *Prog. Oceanogr.*, *11*, 1–59, 1982.

J. Marotzke and C. Wunsch, Center for Meteorology and Physical Oceanography, Department of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139.