Strategies for nested and eddy-permitting state estimation

Geoffrey Gebbie,[1,2,3] Patrick Heimbach,[2] and Carl Wunsch[2]

[1] Both ocean process and prediction studies increasingly rely on state estimation (i.e., data assimilation) to provide the most complete representation of how the ocean circulates. This study applies the formalism and methodology of state estimation, recently developed for the global, coarse-resolution problem, to eddy-permitting state estimation in regional domains. Two major challenges exist for a state estimate that is nested inside a global state estimate: (1) estimation of open-boundary conditions consistent with information interior and exterior to the regional domain and (2) estimation with the higher-resolution models of regional studies. Here a least squares cost function defines the problem of minimizing the misfit between a North Atlantic regional general circulation model and actual observations, including those of the Subduction Experiment. A first experiment, using a novel ''multiscale'' method to constrain the large-scale regional circulation, shows that the use of the adjoint of both an eddy-permitting model and its coarse-resolution twin leads to a solution of the least squares problem in a computationally practical way. Therefore no fundamental obstacle exists to constraining the large-scale regional circulation nested within a global circulation. A second experiment in the North Atlantic shows that the model circulation can be constrained to the full observational signal, including eddy variability, as observed at selected point locations. Both experiments in this study produce eddy-permitting state estimates which are exactly self-consistent with the equations of motion as embodied by a general circulation model. Therefore dynamical balances can be diagnosed and easily interpreted; in particular, a companion paper uses the state estimates to determine eddy subduction rates in the North Atlantic.

1. Introduction

[2] The statistical combination of observations and a numerical model, termed state estimation (or in the borrowed meteorological nomenclature, data assimilation), provides a way to reconstruct the realistic time-evolving, three-dimensional circulation of the ocean, using both the newly available global ocean data sets and the best of modern numerical general circulation models. Recent advances, such as those of the ECCO (Estimating the Circulation and Climate of the Ocean) Consortium [*Stammer et al.*, 2002; *Fukumori*, 2002; *Stammer et al.*, 2003, 2004], have demonstrated the practicality of state estimation in global, coarse-resolution ocean models. The forbidding high-dimensionality of the problem due to the large number of degrees of freedom in fluid flows, even at 1 or 2 degree spatial resolution, did not prove to be a fundamental obstacle. Although these previous results are useful, many regional studies, such as those outlined by CLIVAR (the International Research Program on Climate Variability and Predictability), require higher model resolution than has been used globally. The methods previously used for global estimates therefore should be applied to regional problems. In shifting the focus from the coarse-resolution global scale to high-resolution regional estimation, two issues emerge: the need to deal with open boundaries and the potential numerical problems arising with higher-resolution ocean models.

[3] For state estimation in the typical oceanographic context where observations have been collected over a past interval of time, the so-called adjoint method [e.g., *LeDimet and Talagrand*, 1986; *Thacker and Long*, 1988] is a natural way to combine observations with a model such that the resulting estimate is exactly self-consistent with the model itself (i.e., no nonphysical sources to ''keep the model on track''). The adjoint method is an optimization method based upon Lagrange multipliers, and in the case of a limited area model, the investigator searches for open-boundary conditions (in addition to other uncertain model parameters and surface boundary conditions) to find a model trajectory which reproduces the observations suffi-

---

[1]Massachusetts Institute of Technology/Woods Hole Oceanographic Institution Joint Program in Oceanography, Cambridge, Massachusetts, USA.
[2]Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.
[3]Now at Department of Earth and Planetary Sciences, Harvard University, Cambridge, Massachusetts, USA.

ciently well. The goodness of fit of the model is measured by a scalar cost function, typically a sum of squared model-data misfits, and state estimation reduces to nothing more than a giant least squares problem.

[4] Regional state estimation with the adjoint method has been attempted before, but with widely differing goals. The feasibility of estimating open-boundary conditions in quasi-geostrophic (QG) models was shown by *Moore* [1991], *Seiler* [1993], and *Nechaev and Yaremchuk* [1994]. A number of investigators [e.g., *Schröter et al.*, 1993; *Gunson and Malanotte-Rizzoli*, 1996; *Cong et al.*, 1998; *Yaremchuk and Maximenko*, 2002] used these QG models with resolution varying from 1/4° [*Schröter et al.*, 1993] to roughly 1/8° [*Cong et al.*, 1998] to estimate the eddy-scale circulation over monthly timescales. In contrast, our goal here is to attempt regional state estimation with a general circulation model (GCM) and real observations over 1–2 years. We choose a GCM and real observations both because we wish to form an "optimal" estimate of ocean circulation for scientific purposes, and because consistency between model and data may be more likely in this case. Regional state estimation with GCMs at non eddy-permitting resolution (1° and 2°) has been reported by *Zhang and Marotzke* [1999], *Ferron and Marotzke* [2003], and *Ayoub* [2005]. The primary focus of this paper is to extend the methodology of regional state estimation with GCMs to a resolution similar to the quasi-geostrophic studies. An ultimate goal is state estimation at true eddy-resolving scales, but we wish to explore the methodology at eddy-permitting resolution first.

[5] An ever present challenge in regional simulation and state estimation is finding open-boundary conditions that are physically compatible with the interior circulation. In this work, an existing global state estimate is used as a first guess of the uncertain open boundary conditions. When regional state estimation is performed with the same model and state estimation codes as the global state estimate, this approach could be called "nested state estimation." By the use of similar tools in the global and regional problems, the mismatch between the boundary and interior circulation may be lessened.

[6] The organization of the paper follows. Section 2 presents methods for estimating open-boundary conditions in GCMs at high regional resolution. The methodology is applied to a region of the ocean with a wide variety of actual observations. Specifically, we estimate the ocean state during the Subduction Experiment of the eastern subtropical North Atlantic Ocean [*Brink et al.*, 1995], and the model and data are introduced in section 3. In the effort to perform nested estimation, we take an incremental approach. One, we wish to find the large-scale regional circulation over the one to two year period in which extensive observations are available (detailed in section 4). This is a longer time window than many of the previous regional state estimation studies. Two, a remaining question is whether the model can fit the full observational signal, both large-scale and small-scale, using this methodology; this is explored in section 5. Both experiments provide improved estimates of the circulation that could be used for further scientific research, but the focus here is primarily on the methodology. Analysis of subduction rates and dynamical processes is presented in a companion paper (G. Gebbie, Does eddy subduction matter

in the northeast Atlantic Ocean?, submitted to *Journal of Geophysical Research*, 2006).

2. Methods for Nested State Estimation

[7] In a mathematical sense, there is little doubt that regional state estimation methods resolve the ill posedness of the forward open-boundary problem (see *Bennett and Kloeden* [1981] and *Oliger and Sundström* [1978] for a deeper review and a definition of well posedness). In a practical sense, however, the complexity of GCMs and modern data sets may make the search for the solution of the least squares problem difficult. For example, the model may not be capable of a circulation that is consistent with real observations, and then either the model or the assumed errors in the data must be reevaluated. Another issue is the efficiency of optimization schemes; when searching over a large number of control parameters, can a solution be found with present-day computational resources?

[8] To address these practical issues, this section offers three techniques to increase the efficiency of solving the nested state estimation problem, and the likelihood that the solution is acceptable. One, we propose a multiscale method where the model is used at multiple resolutions and only the large-scale circulation is explicitly constrained to make the optimization more efficient. Two, the open-boundary velocity field has special physics which may have slowed the search process for previous investigators. We propose a decomposition of the velocity field into depth-independent and depth-dependent components, and show that this simplifies the optimization problem. Three, the resulting state estimates of previous investigators showed a number of nonphysical characteristics near the boundary. We propose a simple numerical amendment to the estimation problem, the added constraint of thermal wind balance, to improve the resulting state estimate.

2.1. A Multiscale Method

[9] In the simplest form of nested state estimation, a global (or larger domain), coarse-resolution estimate, $\widetilde{\mathbf{x}}_{global}(t)$, is used to give the necessary boundary conditions for the regional, high-resolution model, $\mathbf{q}_{REG}(t)$ (where capitals are used to emphasize that the regional model has high resolution). The mapping from the global estimate can be represented by $\mathbf{q}_{REG}(t) = f(\widetilde{\mathbf{x}}_{global}(t))$ where f is the mapping function and the boundary conditions change with time. The following discussion applies to the case of Dirichlet boundary conditions, a common situation in GCMs in which the entire state must be prescribed at the boundaries and hence $\mathbf{q}_{REG}(t)$ includes open-boundary temperature, salinity, and two components of horizontal velocity. The discussion can be applied to other boundary conditions equally well by changing the form of f and $\mathbf{q}_{REG}(t)$. The simplest form of nested state estimation could be made more efficient, however. The boundary conditions can be improved by forming a coarse-resolution, regional state estimate ($\widetilde{\mathbf{x}}_{reg}(t)$, lowercase "reg" for low resolution) and then using this new coarse resolution estimate as a first guess for the high-resolution regional model. This involves interpolating the coarse-resolution boundary conditions to higher resolution: $\mathbf{q}_{REG}(t) = g(\widetilde{\mathbf{x}}_{reg}(t))$, where g is the interpolator function. The final state estimate, a high-

resolution regional product, is labeled $\widetilde{\mathbf{x}}_{REG}(t)$. This approach, which we call a "multiscale method," is similar to the idea of *Köhl and Willebrand* [2002], who used a reduced-resolution adjoint. The multiscale method can be computationally efficient because iterations of the search procedure are done cheaply at coarse resolution and the number of costly iterations at high resolution is reduced. In many common cases (such as the situation to be discussed in section 3), the computational cost of the coarse-resolution regional model is negligible compared to the high-resolution model. In such an event, any changes to open-boundary conditions which lead to greater consistency with observations is computationally helpful (although this criterion needs to be evaluated on a case-by-case basis).

[10] A subtlety of the multiscale method is that the same objective function should be used for both the coarse- and high-resolution problems. To be physically consistent, however, the objective function weights must differ with resolution because the representation error of the model changes with resolution. At coarse resolution, the weighting must be such that the expected cost function value associated with each observation is one when the large-scale circulation is consistent with the observations. For example, a term of the cost function has the form $(x - x_o)^2/\sigma^2$, where x is an individual observation, $x - x_o$ is the model-data misfit, and σ is the expected error. Because a coarse-resolution model does not resolve motions at scales less than the grid spacing, such information in the observations must be treated as noise (i.e., representation error). The observations could be prefiltered to include only the large-scale signal, but the least squares formulation naturally handles such a situation anyway. The weights are set to be inversely proportional to the expected noise, and are therefore decreased. The expected noise can be computed in the wave number band of interest via any independent model, such as the spectrum of *Zang and Wunsch* [2001]. The high-resolution problem proceeds similarly; the small-scale observational signal is considered a representation error and may again be treated as noise. In addition, the eddy-permitting model partially predicts an eddy field which cannot be expected to match the observations perfectly. Therefore the total expected error is the sum of both the small-scale observational signal and the model variability at those scales, and the weights are chosen accordingly. Using this multiscale approach, the high-resolution model is only constrained to the large-scale ocean circulation.

2.2. Estimating Open-Boundary Velocities

[11] Previous investigators have shown that the estimation of open boundary velocities in a GCM has led to complications in the optimization procedure. When *Ferron and Marotzke* [2003] sought a regional Indian Ocean state estimate, a two-step estimation process was necessary: one step to estimate the initial conditions and surface forcing, and a second step to estimate the open-boundary conditions. One possible explanation for the necessity of a two-step approach is ill conditioning of the problem by the extreme sensitivity to certain control parameters. As an example, consider a regional model and the sensitivity of the sea surface height to the open-boundary velocity. By conservation of volume, the sea surface reacts to a mean inflow by the relation: $d\overline{\eta}/dt = A_{xz}/A_{xy} \cdot \overline{V}_\perp$, where $\overline{\eta}$ is the regional

mean sea surface height, A_{xz} is the cross-sectional area of the open boundary, A_{xy} is the sea surface area, and $\overline{V}_\perp$ is the regional mean velocity normal to the open boundary. In a square region with 1000 km sides, an imbalance of 1 mm/s generates a sea surface rise of 1 m in approximately 12 days, a magnitude beyond anything ever observed. As a means of crude comparison between different sensitivities, the sensitivity due to domain-integrated sea surface height can be normalized by the expected deviation of boundary-integrated velocity to give a dimensionless quantity. All sensitivities can be normalized this way, and the preconditioning step of optimization routines does exactly this [see *Gill et al.*, 1986]. Even after this normalization, the sensitivity to the open-boundary velocity is orders of magnitude larger than any other physically based sensitivities in the ocean, and may cause numerical difficulties.

[12] To remove these difficulties, we propose that the normal component of open-boundary velocity be decomposed as:

$$V_\perp(x,z) = V_1(x,z) + V_2(x) + V_3, \tag{1}$$

where V_1 is the depth-varying component of velocity, V_2 is the depth-independent component, and V_3 is a constant. These different components of the velocity field are expected to have different magnitudes, and are individually scaled. The depth-dependent velocity is defined to have no net inflow in any column,

$$\sum_{z=1}^{Z} V_1(x,z) \cdot \Delta h(z) = 0, \tag{2}$$

where z is the vertical level number, Z is the total number of levels, Δh is the thickness of a model level, and the equation holds at all grid points along the boundary. The depth-independent velocity is set to have no net inflow into the regional domain:

$$\sum^{ob} V_2(x) \cdot \Delta x \cdot H(x) = 0, \tag{3}$$

where *ob* means that the summation holds everywhere along the open boundary, Δx is the length of a grid cell along the open boundary and H is the total depth of the ocean. V_3 represents the scaled net volume flux into the domain which should be nearly zero, although some short-lived, unbalanced volume fluxes have been documented [*Wunsch and Gill*, 1976; *Fu et al.*, 2001]. One can let V_3 vary within an expected range and estimate its value. Alternatively, one can enforce a hard constraint, $V_3 = 0$, making the optimization problem easier to solve. We take this approach throughout the remainder of this work. The tradeoff is that one quantity cannot be estimated.

2.3. Maintaining Thermal Wind Balance

[13] Even when using a nested state estimation approach with compatible global and regional models, some mismatch between the open-boundary conditions and the interior circulation is likely. Nonphysical boundary layers may form on the open boundaries. Sponge layers are one device to minimize these effects. Boundary jets are dominated by

Table 1. Parameters and Specifications for Coarse- and High-Resolution State Estimation

	2°	1/6°
Horizontal resolution	(167–218) km × 222 km	(14.2–18.2) km × 18.5 km
Grid points	20 × 16 × 23 vertical levels	192 × 168 × 23 vertical levels
Time step	3600 s = 1 hr.	900 s = 15 min.
Lap. horizontal viscosity	5×10^4 m²/s	0
Lap. horizontal diffusivity	1×10^3 m²/s	0
Biharmonic horizontal viscosity/diffusivity	0	2×10^{11} m⁴/s
Vertical viscosity	1×10^{-3} m²/s	1×10^{-3} m²/s
Vertical diffusivity	1×10^{-5} m²/s	1×10^{-5} m²/s
State vector	1.70×10^4 elements	3.14×10^6 elements
Control vector	9.11×10^4 elements	5.49×10^6 elements
Observations	1.28×10^7 elements	1.28×10^7 elements
Model input	7.68×10^5 forcing elements	7.98×10^7 forcing elements
Model output	1.50×10^8 estimated elements	1.09×10^{11} estimated elements
Processors	1 processor	24–48 processors
Computational time	2 cpu hours/iteration	400 cpu hours/iteration
Search iterations	≈40 iterations	≈120 iterations
Total computer time	≈80 hours (2.3 days)	≈50,000 hours (5.7 years)

ageostrophic currents which strongly affect the heat, buoy-ancy, and momentum budgets along the boundary.

[14] To decrease the impact of nonphysical boundary features, thermal wind balance is imposed by appending an extra term to the cost function in which deviations from balance are penalized. This is sometimes called a soft constraint. Thermal wind balance at the open boundary is: $\partial V_\perp / \partial z = -g/(\rho_0 f) \cdot \partial \rho / \partial x$, where x is the along-boundary coordinate, z is the depth coordinate, g is gravity, ρ is density with some reference value ρ_o, and f is the Coriolis parameter. Rearranging,

$$F(x,z) = \partial V_\perp / \partial z + g/(\rho_0 f) \cdot \partial \rho / \partial x = 0, \qquad (4)$$

which must be appropriately discretized [see *Gebbie*, 2004]. For the entire open boundary at all times, discretized values of $F(x, z)$ are appended in vector form, $\mathbf{F} = [F(1, 1), F(2, 1) \ldots F(x_{max}, z_{max})]^T$. The additional term to the objective function then takes the form of a weighted squared deviation: $\mathbf{F}^T \mathbf{W}_F \mathbf{F}$, where $\mathbf{W}_F$ is a weight matrix. Deviations from thermal wind are proportional to the Rossby number, thus permitting an estimate of $\mathbf{W}_F$. As model resolution increases, the Rossby number may also increase, and the weighting of the thermal wind constraint should decrease. In any case, this constraint may force the characteristics of the interior circulation and boundary to be similar.

[15] An alternate method, based upon the work of *Stevens* [1991] and *Zhang and Marotzke* [1999], is to diagnose the open-boundary velocities from the imposed density struc-ture on the boundary. This idea can be extended to the open-boundary estimation problem by defining the control vector as only the temperature and salinity, and then adding a so-called hard constraint to the model. In practice, unfortu-nately, the previous investigators found that the thermal wind mapping is noisy, and it does not give any information on the depth-independent component of the flow.

3. Model Setup

[16] The eastern subtropical North Atlantic Ocean (here-after, northeast Atlantic) is a favorable location for experi-ments with nested state estimation; the region hosted an intensive field campaign to collect oceanographic and meteorological observations known as the subduction experiment [*Brink et al.*, 1995]. No intense western boundary current is present, and the nonlinearities that earlier investigators encountered may not be present here [e.g. *Schröter et al.*, 1993; *Cong et al.*, 1998]. This section applies the open-boundary state estimation concepts to the northeast Atlantic.

3.1. Eddy-Permitting Model With Open Boundaries

[17] The model used here, as in the work by *Stammer et al.* [2002, 2003], is the ECCO version of the MITgcm [*Marshall et al.*, 1997; *Marotzke et al.*, 1999], a state-of-the-art GCM. The high-resolution regional model has a hori-zontal grid spacing of 1/6° by 1/6°, or roughly 15 km. With the Rossby radius of deformation between 25 and 45 km in this region, the model is eddy permitting. A coarse-resolution regional twin model is run at 2° resolution with some adjustments of the model parameters (Table 1 is a detailed list of parameters). The model domain contains most of the eastern subtropical gyre of the North Atlantic (see the domain of Figure 1). The eddy kinetic energy of the forward high-resolution model is typically 50–75% of TOPEX/POSEIDON observations in the northern half of the do-main, and 90% in the southern half.

[18] Using the ECCO-MITgcm model, the goal is to obtain a best description of the oceanic circulation in this region through a least squares fit of the model to real observations for the year June, 1992, to June, 1993. The first guess of the time-dependent boundary values and initial conditions of the regional model are taken from the 2° ECCO global estimate. The National Centers for Environ-mental Prediction (NCEP) Reanalysis daily sensible and latent heat fluxes and twice daily surface wind stresses are used as first-guess forcing fields [*Kalnay et al.*, 1996]. What renders the method practical is the availability of an automatic differentiation (AD) tool called TAF (Tranforma-tions of Algorithms in Fortran [see *Giering and Kaminski*, 1998; *Heimbach et al.*, 2005]. This software permits a (semi) automatic production of the Fortran code for the adjoint of the MITgcm.
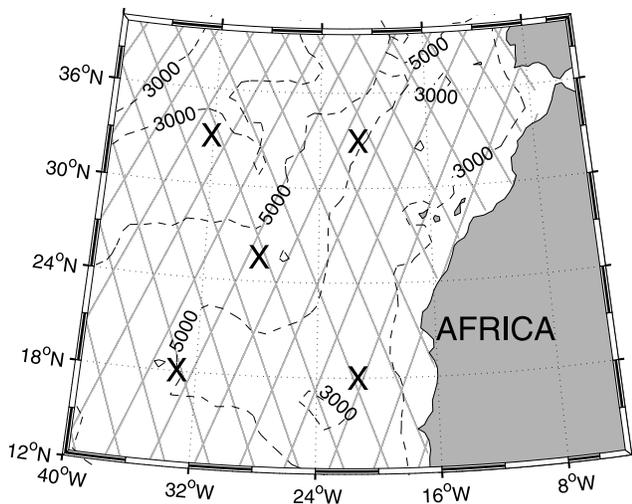
Figure 1. Domain of the regional general circulation model and the Subduction Experiment. The subduction experiment was an intensive field experiment designed to study the subduction of fluid from the mixed layer into the main thermocline. The five moorings of the subduction experiment are marked by crosses. TOPEX/POSEIDON altimetric tracks are marked by thin, diagonal lines. Depth contours at 3000 and 5000 m are marked with dashed lines.

3.2. Northeast Atlantic Least Squares Problem

[19] The cost function of the northeast Atlantic is listed in Table 2. Its many terms include both the TOPEX/POSEIDON altimeter data and mooring data from the Subduction Experiment (Figure 1). The form of the objective function terms is illustrated with the mooring temperature misfit term: $\sum^t (T(t) - T_{SubEx}(t))^T \mathbf{W}_{SubEx}(T(t) - T_{SubEx}(t))$, where $T(t)$ is the model temperature and $T_{SubEx}(t)$ is the mooring temperature. The weight matrix, $\mathbf{W}_{SubEx}$, is diagonal with values that vary as a function of horizontal location, depth, and data type. The next three terms are the climatological misfits; they constrain the estimate to the three-dimensional, monthly varying Levitus atlas [*Levitus and Boyer*, 1994], and the monthly varying climatology of sea surface temperatures [*Reynolds and Smith*, 1994].

[20] The last 11 terms in Table 2 are control penalty terms; they constrain the control parameters to lie within a certain range of their initial guess. The control penalty terms take the place of an explicit model error term in *J*. Here, the controls include (1) initial temperature and salinity, (2) surface heat flux and freshwater flux, (3) surface wind stress, and (4) open-boundary conditions (as formulated in section 2). Although mixing coefficients are highly uncertain, they are not included in the control parameters because of the computational burden. The mixing coefficients are not static, however, as the KPP boundary layer scheme computes new mixing coefficients for each model integration. For further information on the controls, see Table 2.

[21] A well-known problem with diagonal weight matrices is the emergence of small-scale noise in the estimated control fields which is unphysical because the dominant atmospheric variability is at longer wavelengths. Ideally, **W** would be the inverse of the error covariance matrix, which would suppress the small-scale noise. In the absence of estimates of the true covariance, an ad hoc method is used to impose a smoothness criterion on the solution by appending

Table 2. Terms of the Cost Function[a]

Terms	Frequency	Number	$\langle n^T n \rangle^{1/2}$ (2°)	$\langle n^T n \rangle^{1/2}$ (1/6°)	Reference	Notes
Observational Terms						
Mooring temperature	monthly average	12	0.3°C	0.6°C	*Brink et al.* [1995]	gridded on model levels
Mooring zonal velocity	monthly average	12	3.0 cm/s	6.0 cm/s	*Brink et al.* [1995]	gridded on model levels
Mooring meridional velocity	monthly average	12	3.0 cm/s	6.0 m/s	*Brink et al.* [1995]	gridded on model levels
TOPEX/POSEIDON anomaly	daily	365	8.0 cm	12.0 cm	*Fu et al.* [1994], *Tai and Kuhn* [1995]	pointwise
TOPEX/POSEIDON mean		1	10.0 cm	10.0 cm	*Lemoine et al.* [1997]	relative to geoid, gridded
Climatological terms						
Levitus temperature	monthly	12	0.3°C	0.6°C	*Levitus and Boyer* [1994]	gridded on model levels
Levitus salinity	monthly	12	0.2	0.4	*Levitus et al.* [1994]	gridded on model levels
Reynolds SST	monthly	12	0.5°C	1.0°C	*Reynolds and Smith* [1994]	
Initial conditions						
ECCO temperature	1 Jun 1992	1	0.2°C	0.4°C	*Stammer et al.* [2002]	identical grid
ECCO salinity	1 Jun 1992	1	0.1	0.2	*Stammer et al.* [2002]	identical grid
Surface forcing						
NCEP net heat flux	every 10 days	37	20. W/m²	20. W/m²	*Kalnay et al.* [1996]	NCEP Reanalysis Project
NCEP freshwater flux	every 10 days	37	8.0×10^{-7} m/s	8.0×10^{-7} m/s	*Kalnay et al.* [1996]	forward model forced daily
NCEP wind stress	every 10 days	37	0.02 N/m²	0.02 N/m²	*Kalnay et al.* [1996]	forward model forced 2x day
Open-boundary terms						
ECCO temperature	monthly	12	0.2°C	0.4°C	*Stammer et al.* [2002]	
ECCO salinity	monthly	12	0.1	0.2	*Stammer et al.* [2002]	
ECCO tangential velocity	monthly	12	3.0 cm/s	6.0 cm/s	*Stammer et al.* [2002]	
ECCO normal velocity	monthly	12	3.0 cm/s	6.0 cm/s	*Stammer et al.* [2002]	
Thermal wind deviation	monthly	12	0.1 cm/s	0.2 cm/s	*Pond and Pickard* [1983]	assume Ro = 0.1

[a]The first column introduces 19 types of terms in the cost function, which are further divided into five main categories. The second column refers to the frequency of comparison between the model and prior information. The third column gives the total number of terms in time per cost function type. The fourth column gives the average expected error in the 2° optimization problem, which is used for calculating weights. The fifth column is the same information but for the eddy-permitting problem. References are given for the data source or error estimates. The notes give some extra information regarding the cost function evaluation.
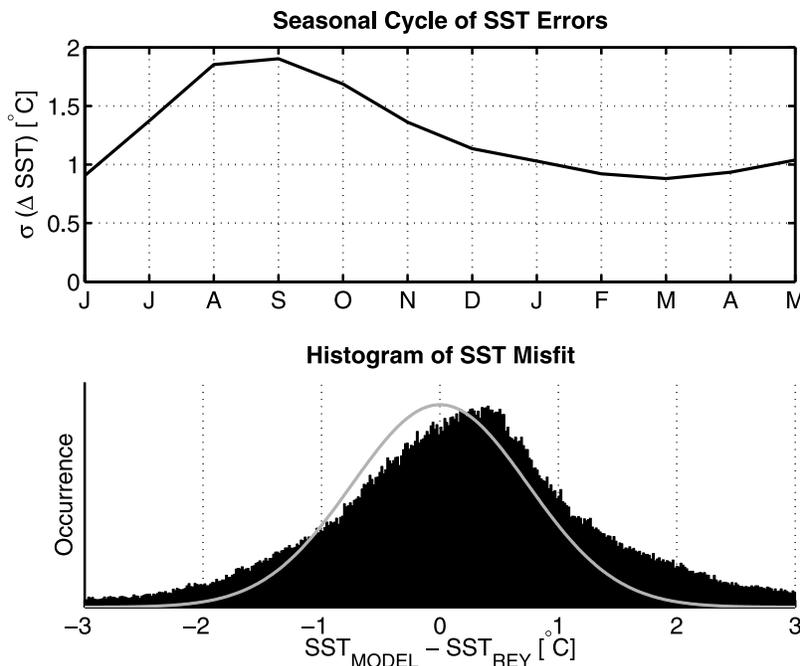
Figure 2. (top) Standard deviation of SST misfit as a function of month. The SST misfit is defined as the difference between the model prediction and the Reynolds SST. (bottom) Histogram of SST misfit and the assumed prior error statistics (gray line).

terms $(\nabla^2 \mathbf{u})^T \mathbf{W}_2 (\nabla^2 \mathbf{u})$ to the cost function, where $\mathbf{u}$ represents a control adjustment to the surface boundary conditions. *Gebbie* [2004] and *Lea et al.* [2006] give formulae to calculate $\mathbf{W}_2$ based on the covariance length scale. As is well known, the effect of such numerical derivative terms is to smooth the result, but they are only an interim substitute for the true inverse covariance [see also *Bennett*, 1992].

[22] Two experiments are carried out in northeast Atlantic. Section 4 seeks the large-scale circulation consistent with observations, and is called experiment 1. Using the multiscale method of section 2, the global estimate, $\widetilde{\mathbf{x}}_{global}(t)$, is used as a first guess for a coarse-resolution regional model, then the resulting coarse-resolution regional state estimate, $\widetilde{\mathbf{x}}_{reg}(t)$, is used to find a high-resolution regional state estimate, $\widetilde{\mathbf{x}}_{REG}(t)$. Section 5 poses the more stringent problem of finding the circulation consistent with the full observational signal, and is called experiment 2.

4. Search for the Large-Scale Circulation Consistent With Observations

[23] The first step in any state estimation procedure is a prediction of the observations by model simulation, hereafter termed the unconstrained simulation or simply simulation. Despite a visual similarity between the modeled circulation and observations, both the 2° and 1/6° simulations show large-scale hydrographic deficiencies. Sea surface temperature approaches 35°C in the northern basin (30–40°N). Overly warm sea surface temperatures are also associated with a weakened Canary Current in the simulation. Anomalously warm SST is a ubiquitous problem of numerical model products including the ECCO state esti-

mate. (Here we have used the original ECCO state estimate from the adjoint method, 1992–1997 [*Stammer et al.*, 2002]. Later estimates with improved upper ocean physics [*Stammer et al.*, 2004] do not have the same preponderance of overly warm sea surface temperatures.) Surface layers of the model are too warm in the summer (Figure 2) because the seasonal mixed layer does not deepen enough. Another major deficiency of the simulation is the meridional slope of the winter mixed layer base; the mixed layer deepens to the south, reaching a depth of 220 m, at 22°N. In contrast, observations and climatologies alike show that the mixed layer shoals equatorward, a crucial feature because horizontal flow across a sloping boundary leads to "subduction," or volume flux from the mixed layer to the main thermocline [*Woods*, 1985; *Marshall et al.*, 1993]. The abnormally cool surface layers of the eastern subtropics between 20–30°N are caused by large western boundary heat fluxes, and corresponding heat flux divergence near the western boundary. Adjustments to the control vector are needed to bring the model into consistency with observations.

4.1. Coarse-Resolution Optimization

[24] In the multiscale approach, we first search for a coarse-resolution regional state estimate ($\widetilde{\mathbf{x}}_{reg}(t)$). Using the quasi-Newton method of *Gilbert and Lemaréchal* [1989] and adjoint-computed gradients, the ocean circulation is brought within observational uncertainty in 22 iterations of the coarse forward and adjoint models (see Figure 3, top). Thirty more iterations do not significantly change the solution. Considering the control vector contains about 100,000 elements (i.e., $N_{iterations} \ll N_{controls}$), convergence to the solution is rapid. The control variables
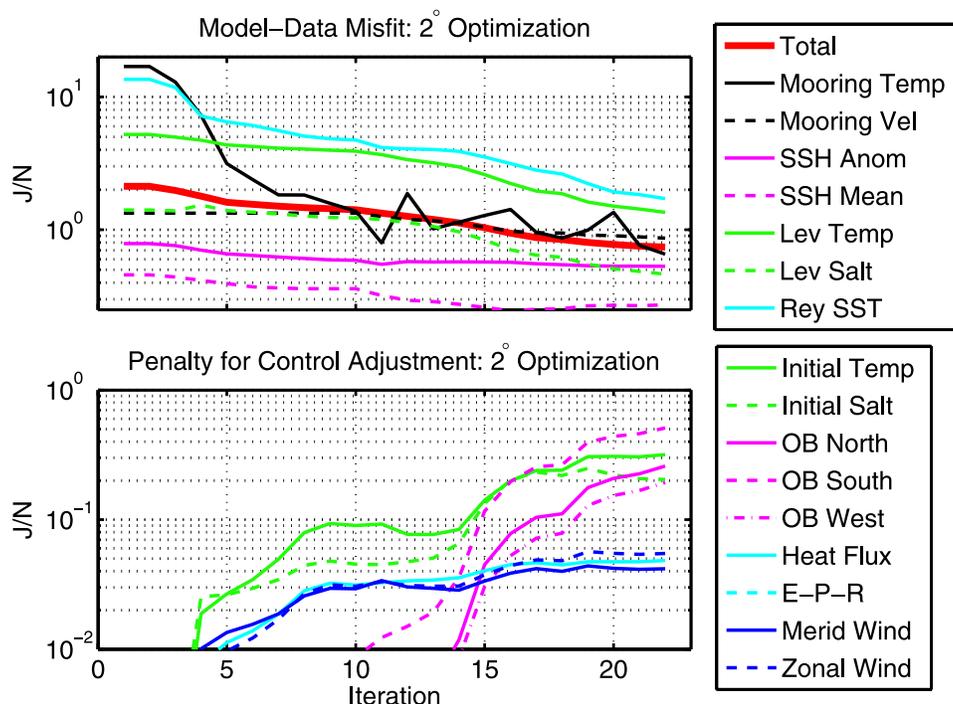
Figure 3. (top) Normalized model-data misfit as a function of iteration of the search method for the coarse-resolution optimization. The cost function value for each type of observation, J, is normalized by the number of each kind of observation, N. A value of O(1) is expected for all lines. (bottom) Normalized contribution to the cost function of the penalty for control adjustments, which reflects the size of the control adjustments, $\|u\|^2$. Note the different scale on the y axis in this plot.

quadratically converge upon the minimum of the cost function (Figure 3, bottom); this convergence rate is the theoretical one for the quasi-Newton method with a parabolically shaped cost function [*Press et al.*, 1990]. Such a topology is expected for a diffusive coarse-resolution ocean model. The adjustments of the control variables lie within the range of uncertainty in those variables. It is not surprising that the method works so well for a coarse-resolution model because it is a nearly linear system.

[25] What does the coarse-resolution model adjust in order to fit the observations? The gradient of the cost function with respect to the controls gives a simple estimate of the relative importance of each control for a particular model trajectory. Each control variable has an expected magnitude; this magnitude is used to nondimensionalize the variables in the search algorithm [e.g., *Gill et al.*, 1986]. By this simple measure, the initial temperature field and time-varying open-boundary conditions are most important over the first year of integration. The memory of initial conditions extends well beyond 1 year, as both forward model studies [*Griffies and Bryan*, 1997] and adjoint studies [*Hill et al.*, 2004] showed a memory of almost ten years in the upper ocean. The initial temperature field is warmed between 20–30°N where the model simulation is too cold. With this major adjustment, the estimated mixed layer depth shoals toward the south, and never reaches a depth greater than 170 m, in close accord with observations. Such an improvement in the mixed layer structure is crucial for surface layer process studies.

4.2. Application to the Eddy-Permitting Model

[26] The control adjustments from the regional, coarse-resolution state estimate are used as first-guess boundary conditions for the eddy-permitting model. Linear interpolation is used to change resolution of the controls. The coarse-resolution estimate is expected to have some skill in predicting the large-scale aspect of the ocean observations. To examine this hypothesis, two eddy-permitting model trajectories are compared: a run with zero control adjustments and another with coarse-resolution estimated controls. A comparison of the two cost function values (iterations 0 and 1, Figure 4) shows that the coarse-resolution controls do improve the eddy-permitting model simulation. These adjustments decrease the total observational cost function elements by 10%, primarily by bringing the model closer to the Levitus climatological temperature and Reynolds SST. Large-scale hydrographic biases in the high-resolution upper ocean are corrected by the coarse-resolution controls.

[27] Were the predictions made by the coarse-resolution model useful in the eddy-permitting case? In this model setup, 22 coarse-resolution iterations were performed at negligible cost relative to one high-resolution iteration, so any improvement in the cost function makes the coarse-resolution state estimation useful. However, the coarse-resolution control adjustments made a much bigger impact on the coarse-resolution model than the eddy-permitting one (compare Figures 3 and 4). On the basis of this result, the coarse-resolution optimization was useful here, but we
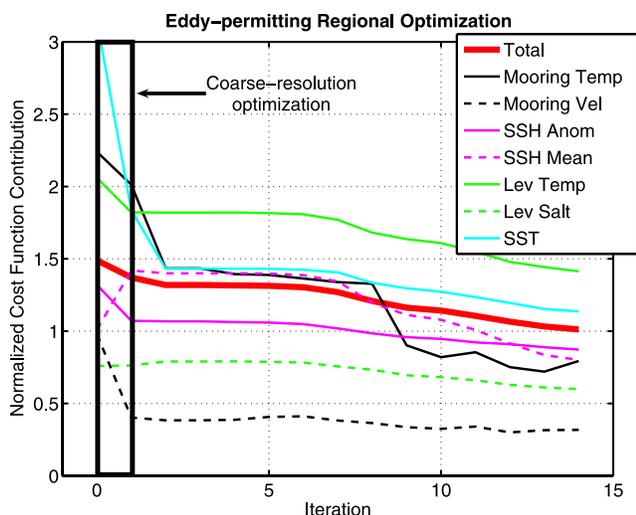
Figure 4. Normalized model-data misfit as a function of iteration of the search method for the eddy-permitting optimization. A value of O(1) is expected. Iteration 0 is the eddy-permitting model simulation (unconstrained). Iteration 1 is the eddy-permitting model with control adjustments from the coarse-resolution optimization (see box). Iterations 1–14 proceed with a quasi-Newton descent algorithm and gradients computed from the adjoint of the eddy-permitting model. The SST contribution to the cost function was everywhere divided by two for presentation purposes.

cannot prove that it will be useful in all cases. A different perspective on this result is that the high-resolution model dramatically decreased the cost function on its own merits.

[28] The multiscale method is presented here with two models: one at coarse resolution, one at high resolution. Conceptually, this need not be the case and the twin models could be made into a whole family. In this application, for example, it would be interesting to have three models: $2°$, $1/2°$, and $1/6°$. With smaller jumps in resolution, the adjustments at one resolution may better apply to the next higher-resolution model. The tradeoff, of course, is the human power necessary to set up a family of models.

[29] Now, the multiscale method proceeds to search for a high-resolution regional state estimate ($\widetilde{\mathbf{x}}_{REG}(t)$). First, one must determine if any solution exists to the least squares problem, because no solution will be found if model is inadequate, the a priori error assumptions are incorrect, or the optimization method fails. The search for a solution utilizes the quasi-Newton descent method with gradients calculated from the adjoint of the eddy-permitting model. Improvement of the model trajectory comes at a slower pace due to the increased search space dimension. Nevertheless, fifteen iterations bring the large-scale state estimate within the expected error of the observations (Figure 4). The total cost function value reaches the normalized value of $J/N \approx 1$, corresponding to a root-mean-square error that is equal to the prior expected error. (The sum of N squared, independent Gaussian deviates with standard deviation 1 is N, with an expected spread given by a χ^2 distribution.) A solution exists and is found.

[30] Although we have looked at the value of the cost function and the distribution of control adjustments, it is also advantageous to look at the spatial patterns of the control adjustments to see if they look physically reasonable. Generally speaking, the control adjustments from the coarse-resolution optimization were larger-scale and had a larger magnitude than the high-resolution adjustments. One implication of this result is that the first-guess controls for the high-resolution problem are important because they do not change much in the optimization procedure. Therefore the multiscale method may be important for more than computational reasons; the method can provide a better first guess with which to make further small-scale adjustments. Without the improved first guess, the high-resolution optimization scheme would have searched for fine-scale adjustments around a deficient large-scale circulation.

[31] The result is a time-evolving, three-dimensional estimate of the ocean circulation ($\widetilde{\mathbf{x}}_{REG}(t)$) which reasonably fits a wide variety of available information and exactly follows the dynamics of the MITgcm (Figure 5; an animated version of Figure 5 is available in the auxiliary material[1]). In addition, we now have improved estimates of the initial eddy field, open-boundary conditions, wind stresses, and air-sea fluxes. The state estimate is ideal for the study of the role of eddies in subduction because it is dynamically consistent and it permits eddy-scale motions (completed in the companion paper).

4.3. Posterior Checks

[32] A stringent posterior test is to compare the state estimate with observations that were withheld from the optimization. In the northeast Atlantic, withheld data include hydrographic sections. The WOCE AR11 section along $33°W$ was completed in November 1992 [*Pallant et al.*, 1995]. The transect passes the western moorings at $19°N$ and $33°N$, but no observations were used in the intervening 1500 km. In general, the upper layer hydrographic structure is significantly improved in the state estimate relative to the withheld WOCE hydrography; data-model misfits are no larger than $1-2°C$ (see Figure 6). Although the state estimate is an improvement, systematic errors do remain. Estimated surface temperature is as much as $1°$ warmer than observed, yet is erroneously cold at the base of the mixed layer (50–100 m depth). Mixed layer model deficiencies lead to this discrepancy. Overall, we conclude that the general success of the state estimate in reproducing withheld data lends confidence to the state estimate even away from sites of observations, although mixed layer dynamics could still be improved further.

4.4. Open-Boundary Adjustments

[33] Section 2.3 postulated that appending a soft constraint to the cost function, namely that the open boundary is in thermal wind balance, would lessen the typical nonphysical features seen along open boundaries in numerical simulations. Figure 7 shows the character of the western boundary in the unconstrained model simulation and in the constrained model. At $1°$ longitude away from the western
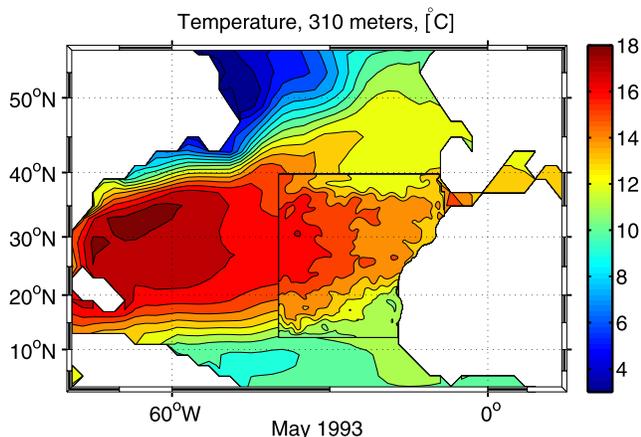
Figure 5. Nested view of the 1/6° regional state estimate inside the 2° ECCO state estimate. Potential temperature at 310 m depth, with a contour interval of 1°C, is shown. The boundary between the two estimates (black line) is discontinuous in temperature because of the open-boundary control adjustments. An animated version of this figure is available with the auxiliary material.

boundary, the average kinetic energy is elevated relative to the interior; we will refer to this as a "nonphysical boundary jet" even though such a feature does not necessarily appear in all velocity snapshots. The boundary jet has increased shear on both sides and therefore an increased local Rossby

number. The imposed soft constraint of thermal wind balance decreases the mismatch in the interior and boundary Rossby Number, and consequently decreases the strength of the boundary jet. The thermal wind constraint does not eliminate nonphysical features, but it does lessen the negative impact of the open boundaries.

[34] The open-boundary adjustments from state estimation greatly affect the estimated Azores Current, a zonal current at $33-36°N$ that extends from the mid-Atlantic to the Gulf of Cadiz near the Mediterranean Sea outflow. The ECCO 2° global state estimate ($\widetilde{\mathbf{x}}_{global}(t)$) includes an Azores Current with reasonable transport but zonal velocities that are too weak. In the regional simulations, the Azores Current meanders southward and does not follow the tight, zonal trajectory of the true current. After adjustment of the western boundary to have a more narrow and rapid Azores Current inflow, the state estimate has a zonal jet at 36°N in accordance with the known characteristics of the current.

[35] Sea surface height variance in the TOPEX/POSEIDON altimeter is elevated in a zonal patch near the Azores Current, where some of the eddy variability is due to meandering of the zonal jet. In the model simulation, which does not have a realistic Azores Current, the estimated SSH variance is too weak (Figure 8). In the high-resolution state estimate ($\widetilde{\mathbf{x}}_{REG}(t)$), the region of relatively high SSH variance extends zonally across more of the domain, as seen in the observations. Although we have not explicitly constrained the small-scale circulation, the state estimate with the improved large-scale circulation produces a better small-scale circula-
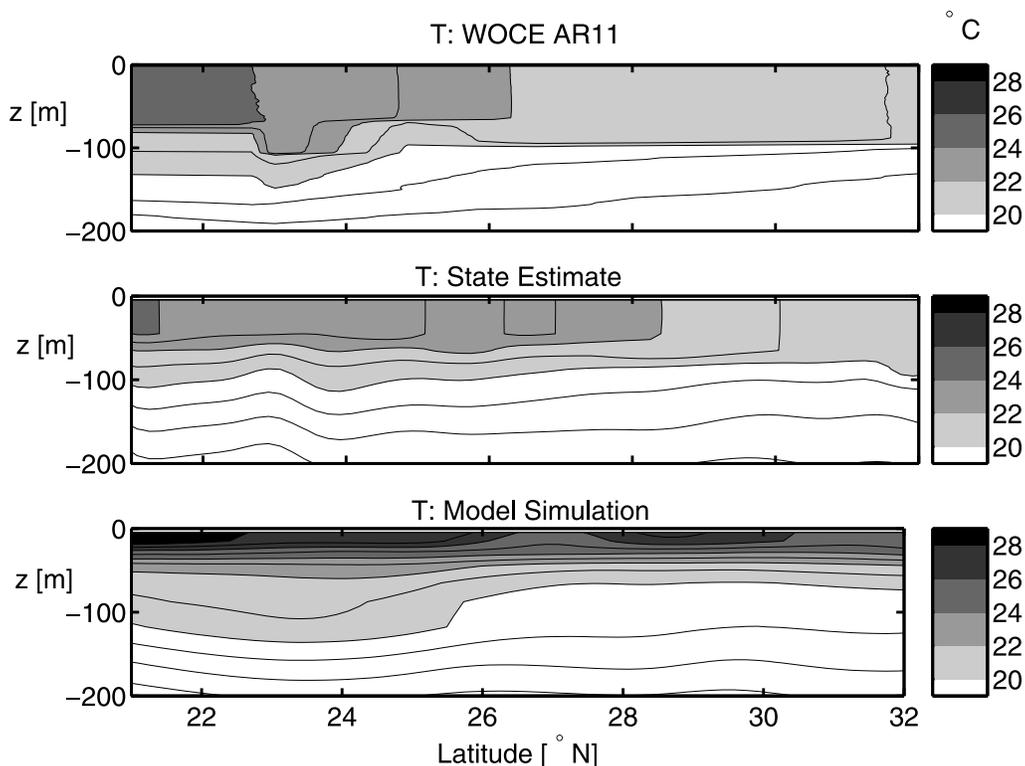
Figure 6. Meridional sections of potential temperature along the WOCE AR11 line (33°W) in November 1992. (top) Observations from WOCE [*Pallant et al.*, 1995]. (middle) Temperature in the state estimate. (bottom) Temperature in the unconstrained model simulation. The contour interval is 1°C.
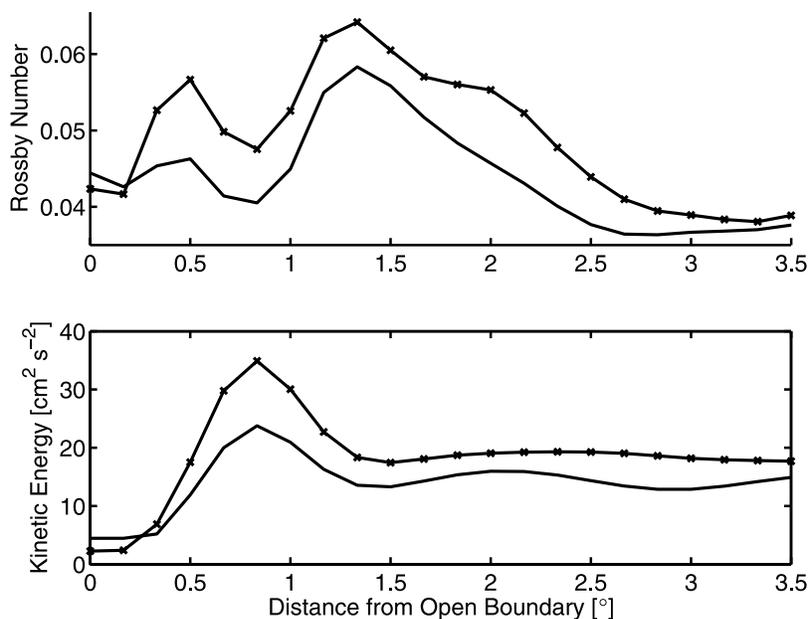
Figure 7. Composite of (top) Rossby Number and (bottom) kinetic energy in the unconstrained model simulation (solid line with crosses) and the state estimate (solid line) as a function of degrees longitude away from the western boundary. The composite is formed by averaging over depth, time, and meridional extent of the domain. The state estimate differs from the unconstrained simulation by the fact that thermal wind balance has been imposed as a soft constraint in the open-boundary layer.
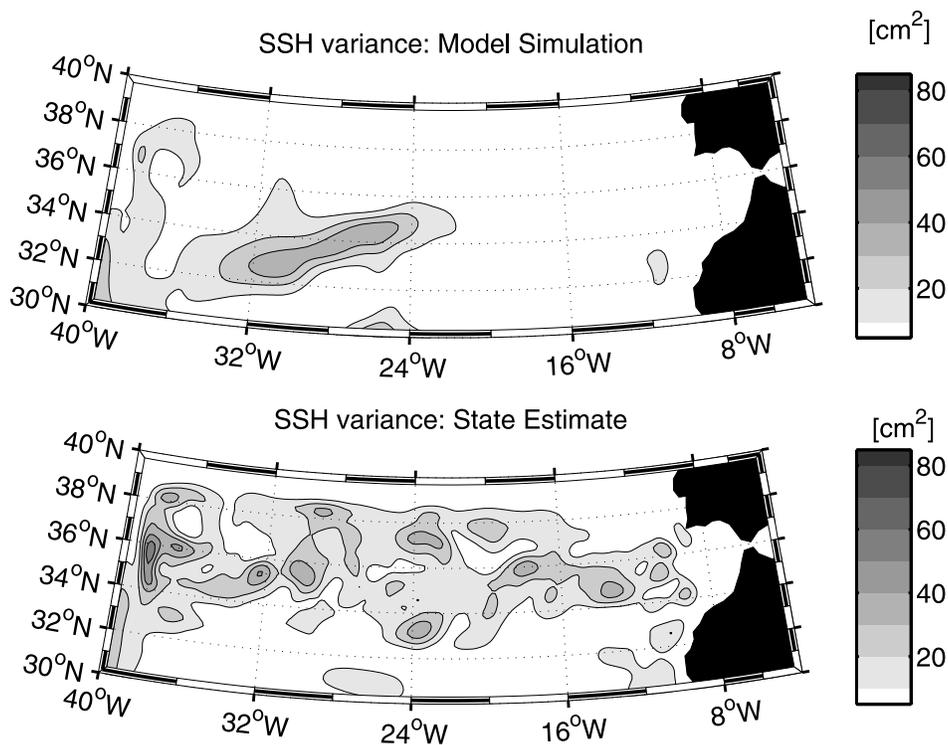
Figure 8. Sea surface height variance in the Azores current subregion as seen in (top) the model simulation and (bottom) the state estimate. Figure 8 (bottom) has the same general spatial structure of TOPEX/POSEIDON observations and roughly 60% of the energy. The contour interval is 10 cm^2.

Table 3. Terms of the Cost Function for the Optimization of the "Full Observational Signal" at the Moorings[a]

Terms	Frequency	Number	$\langle n^T n \rangle^{1/2}$	Reference	Notes
Observational Terms					
Mooring temperature	monthly average	12	0.1°C	*Brink et al.* [1995]	gridded on model levels
Mooring zonal velocity	monthly average	12	0.5 cm/s	*Brink et al.* [1995]	gridded on model levels
Mooring meridional velocity	monthly average	12	0.5 cm/s	*Brink et al.* [1995]	gridded on model levels
Initial conditions					
ECCO temperature	1 Jun 1992	1	0.4°C	*Stammer et al.* [2002]	identical grid
ECCO salinity	1 Jun 1992	1	0.2	*Stammer et al.* [2002]	identical grid
Surface forcing					
NCEP net heat flux	every 10 days	37	20. W/m^2	*Kalnay et al.* [1996]	NCEP Reanalysis Project
NCEP freshwater flux	every 10 days	37	8.0×10^{-7} m/s	*Kalnay et al.* [1996]	forward model forced daily
NCEP wind stress	every 10 days	37	0.02 N/m^2	*Kalnay et al.* [1996]	forward model forced 2x day
Open-boundary terms					
ECCO temperature	monthly	12	0.4°C	*Stammer et al.* [2002]	
ECCO salinity	monthly	12	0.2	*Stammer et al.* [2002]	
ECCO tangential velocity	monthly	12	6.0 cm/s	*Stammer et al.* [2002]	
ECCO normal velocity	monthly	12	6.0 cm/s	*Stammer et al.* [2002]	
Thermal wind deviation	monthly	12	0.2 cm/s	*Pond and Pickard* [1983]	assume Ro = 0.1

tion. This implies that the large-scale state estimate may also estimate the statistics of the small-scale circulation in an automatic way through the information inherent in the dynamics of the GCM.

5. Search for Circulation Consistent With the Full Observational Signal

[36] After finding a circulation consistent with the large-scale observational signal (experiment 1), the next question is whether a model can be simultaneously constrained to both the large and small-scale data signal (experiment 2). The technical implementation of this new problem is similar to the previous experiment: only the observational weights must be increased in order to correspond to the decrease in the expected errors (see Table 3). Although the mathematical transformation between the two problems is straightforward, the new least squares problem poses a more stringent test than the original. In order to increase the prospects of success, the TOPEX/POSEIDON altimetry terms of the cost function are omitted. This test is a somewhat simpler one: fit the full observational signal of the moorings.

[37] The optimization is started with our best prior information about the circulation; in this case, we can use the results of section 4, the large-scale regional state estimate, as our first guess. With 28 iterations of the forward adjoint model, the average mooring data-model misfit decreases from 7.6 σ to 1.8 σ where σ is the expected error. The cost function does not monotonically decrease with iteration. We also found it necessary to change the weights of the background climatologies to restart the stalled optimization. Probably 30–50 iterations of the forward adjoint model are needed for complete convergence, but are not computationally feasible here. There seems to be no fundamental obstacle to produce the full observational signal at the mooring sites.

[38] As an example of the state estimate reproducing higher-frequency variability of the mooring observations, consider the time series of temperature anomaly at the southwest mooring (Figure 9). This time series was chosen because of the cold event passing through the thermocline

for one month in January. Although the unconstrained model simulation did not capture such an event, it is present in the state estimate. From the observations alone, the spatial extent of this feature cannot be determined. How did the state estimate "create" this feature? The adjoint-calculated gradient (not shown) has two bands of increased sensitivity to the cost function: the Azores Current and the North Equatorial Current. Previous studies, including *Gill et al.* [1974], have shown the basic state North Equatorial Current, the site of the southwest mooring, to be baroclinically unstable. *Galanti and Tziperman* [2003] convincingly showed that baroclinic instability increases the sensitivity of these regions because eddies can grow and transport information away from their formation site. In the optimization here, a small perturbation in the initial eddy field leads to large changes in the eddy field at later times. The southwest mooring is near an eddy-scale adjustment in the initial conditions, and the state estimate predicts that the January cold event was caused by a feature with length scale of approximately 100 km. There is no reason to believe that this actually happened at the southwest mooring, but this is because the observing system could not capture the feature more fully; it is not a limitation of the model or the state estimation method.

[39] To show the characteristics of the small-scale circulation, we compare the sea surface height anomalies observed from TOPEX/POSEIDON to that predicted by the state estimate. A scatterplot with one mark for each observation is shown in Figure 10. The TOPEX/POSEIDON fields were not explicitly used to constrain the model in this experiment. Over the entire domain, the average misfit in SSH is more than 10 cm, and the correlation coefficient between model and observations is $r = 0.18$. This skill is not significantly different from zero. However, if we look at locations which are within 100 km of a mooring site, the average misfit is 6.6 cm, close to the instrumental error in TOPEX/POSEIDON of 4.3 cm. The skill at predicting SSH variations near the moorings is $r = 0.72$ and is significant. On the basis of this analysis, the small-scale variability is faithfully reproduced only very near the mooring sites.

[40] When the TOPEX/POSEIDON terms of the cost function are reintroduced, the optimization frequently stalls
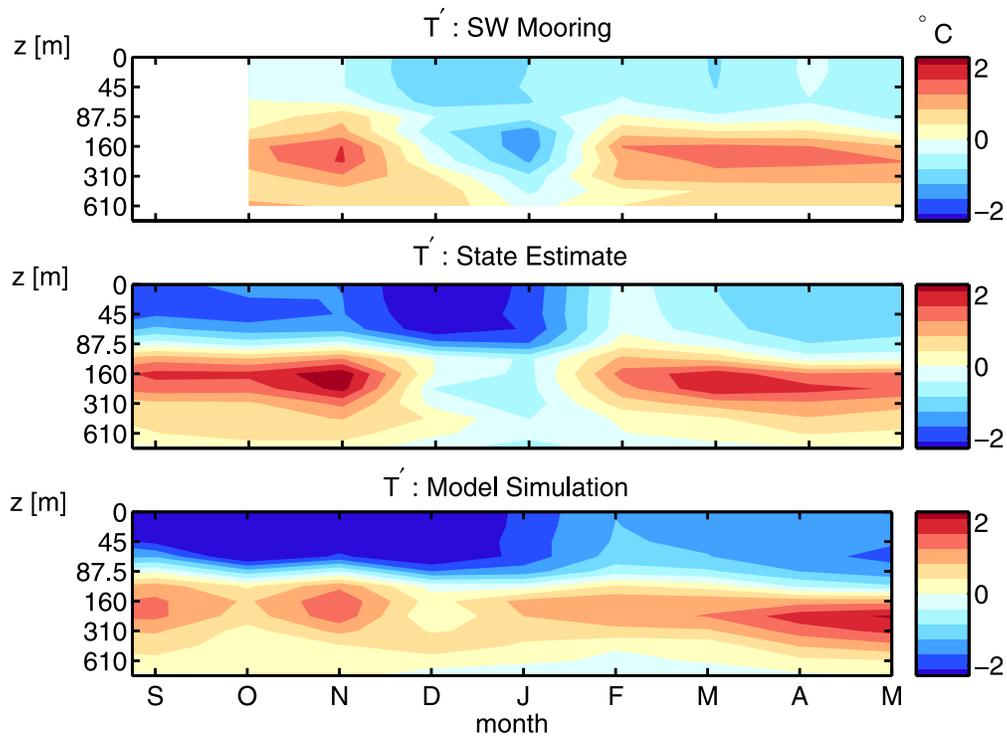
Figure 9.   Three depth-time diagrams of potential temperature anomaly at the southwest mooring site from 1 September 1992 to 1 May 1993. The anomaly is defined relative to the seasonal cycle of the Levitus climatology. (top) Observed temperature anomaly. (middle) Temperature anomaly in the state estimate. (bottom) Temperature anomaly in the unconstrained model simulation.
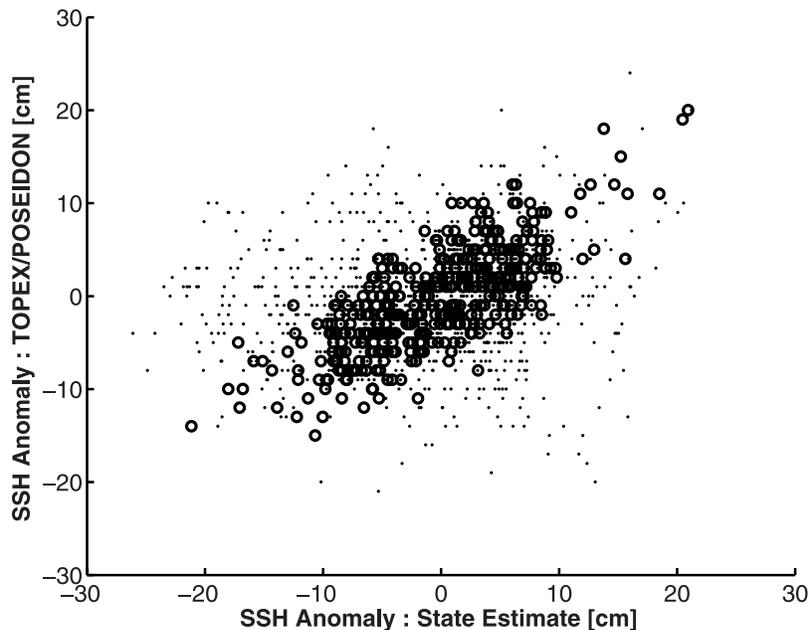
Figure 10.   Scatterplot of TOPEX/POSEIDON sea surface height anomaly versus state estimate sea surface height anomaly. Dots show comparison of SSH at distances greater than 100 km away from a Subduction Experiment mooring. Circles show SSH comparison at locations within 100 km of Subduction Experiment moorings. For increased clarity, a random selection of points are plotted here; the actual number of observations is a factor of 100 greater than what is plotted.

in control space (a commonplace occurrence with complicated nonlinear optimization problems). Many possible causes for the stalled optimization exist. For example, the eddy-scale signal in the moorings and the satellite altimeter may be inconsistent. Or the observed eddies may be incompatible with the dynamics of the model; hence the model would need improvement such as increased resolution. At this point, it is an open question.

6. Conclusions

[41] The main point of this paper is that the recent methodological advances in global state estimation, together with an improved treatment of the open-boundary conditions, may be applied to the regional, high-resolution problem in complex, state-of-the-art GCMs. The large-scale circulation of an eddy-permitting model may be constrained to actual observations over yearly time periods while remaining consistent to boundary conditions specified by a global state estimate. By deemphasizing the details of the small-scale circulation, state estimation with the adjoint method is not restricted to time windows of a few months. A model constrained to the large-scale circulation is free to produce its own small-scale circulation, and in this case, the statistics of the small-scale circulation compare reasonably with independent observations.

[42] Why was eddy-permitting state estimation successful? The search for the solution of the state estimation problem was helped by the dynamics of the eastern subtropical gyre, a more quiescent region than the western boundary or the Circumpolar Current. Adjoint-computed gradients were physically reasonable and useful in the optimization process. Furthermore, the Subduction Experiment provided a wealth of observations with which we could keep the model "on track." Although GCMs are necessarily complicated, our hypothesis is that they increase the practical likelihood of finding a state estimate because they are our most complete information source about the dynamics of the ocean, and they have the best chance of reproducing real observations.

[43] To test the limits of the methodology, we searched for a model circulation that fit the full observational signal of the five moorings of the Subduction Experiment for a year. In this case, the model appears capable of producing the small-scale circulation in the immediate proximity of the mooring sites. When adding the eddy signal of the TOPEX/POSEIDON altimeter, the optimization procedure proved problematic. The reason for this problem remains unknown, but at least two candidates exist. One, the eddy-permitting model may be unable to produce a full domain of realistic eddies, and hence the model must be rejected. Two, fitting the full observational signal of the observations may lead to a more nonlinear problem for which a different optimization scheme is necessary.

[44] This study presented three techniques to apply the adjoint method to a regional model nested inside a global state estimate: (1) a multiscale approach, (2) decomposition of the open-boundary velocities, and (3) explicit requirement of thermal wind balance at the boundary. One, the multiscale approach suggests using the global resolution state estimate as first-guess boundary conditions for a regional, coarse-resolution model, adjusting the open-

boundary conditions to solve a coarse-resolution least squares problem, then using this result as first-guess boundary conditions for a regional, high-resolution state estimation problem. For problems where a coarse-resolution model is computationally cheaper than the high-resolution model, the multiscale approach is likely more efficient. In this study, the coarse-resolution control adjustments skillfully eliminated major biases in the eddy-permitting model. Two, the interior circulation was shown to be extremely sensitive to open-boundary velocities. By separating the portion of the velocity that leads to domain-integrated SSH changes, we were able to optimize all the control variables together. Three, imposing thermal wind balance as a soft constraint to the problem was found to decrease the occurrence of nonphysical boundary jets, but they do not eliminate the problem.

[45] State estimation with true eddy-resolving models will provide a more stringent test of the methods presented here. Western boundary currents, open ocean deep convection, and sea ice formation are nonlinear processes which are not highly active in the eastern subtropical gyre, so this study cannot prove that this formulation of the nested state estimation problem will be successful everywhere. However, the techniques used here, such as the multiscale form of the cost function with a coarse-resolution twin model, suggest that there are still simple ideas that may be helpful in more nonlinear regions.

References

Ayoub, N. (2005), Estimation of boundary values in a North Atlantic circulation model using an adjoint method, *Ocean Modell.*, *12*(3–4), 319–347.

Bennett, A., and P. E. Kloeden (1981), The ill-posedness of open ocean models, *J. Phys. Oceanogr.*, *12*, 1004–1018.

Bennett, A. F. (1992), *Inverse Methods in Physical Oceanography*, Cambridge Monogr. Mech., 346 pp., Cambridge Univ. Press, New York.

Brink, N. J., K. A. Moyer, R. P. Trask, and R. A. Weller (1995), The Subduction Experiment: Mooring field program and data summary, technical report, Woods Hole Oceanogr. Inst., Woods Hole, Mass.

Cong, L. Z., M. Ikeda, and R. M. Hendry (1998), Variational assimilation of Geosat altimeter data into a two-layer quasi-geostrophic model over the Newfoundland ridge and basin, *J. Geophys. Res.*, *103*(C4), 7719–7734.

Ferron, B., and J. Marotzke (2003), Impact of 4D-variational assimilation of WOCE hydrography on the meridional circulation of the Indian Ocean, *Deep Sea Res., Part II*, *50*, 2005–2021.

Fu, L.-L., E. Christensen, C. A. Yamarone, M. Lefebvre, Y. Menard, M. Dorrer, and P. Escudier (1994), TOPEX/Poseidon mission overview, *J. Geophys. Res.*, *99*(C12), 24,369–24,381.

Fu, L.-L., B. Cheng, and B. Qiu (2001), 25-day period large-scale oscillations in the Argentine Basin revealed by the TOPEX/Poseidon altimeter, *J. Phys. Oceanogr.*, *31*, 506–516.

Fukumori, I. (2002), A partitioned Kalman filter and smoother, *Mon. Weather Rev.*, *130*(5), 1370–1383.

Galanti, E., and E. Tziperman (2003), A mid-latitude ENSO teleconnection mechanism via baroclinically unstable Long Rossby Waves, *J. Phys. Oceanogr.*, 33, 1877–1888.

Gebbie, G. (2004), Subduction in an eddy-resolving state estimate of the Northeast Atlantic Ocean, Ph.D. thesis, Mass. Inst. of Technol./Woods Hole Oceanogr. Inst. Joint Program in Oceanogr., Cambridge, Mass.

Giering, R., and T. Kaminski (1998), Recipes for adjoint code construction, *Trans. Math. Software*, 24(4), 437–474.

Gilbert, J. C., and C. Lemaréchal (1989), Some numerical experiments with variable-storage quasi-Newton algorithms, *Math. Program.*, 45, 407–435.

Gill, A. E., J. S. A. Green, and A. J. Simmons (1974), Energy partition in the large-scale ocean circulation and the production of mid-ocean eddies, *Deep Sea Res.*, 21, 499–528.

Gill, P. E., W. Murray, and M. H. Wright (1986), *Practical Optimization*, 401 pp., Elsevier, New York.

Griffies, S. M., and K. Bryan (1997), Predictability of North Atlantic inter-decadal variability, *Science*, 275, 181–184.

Gunson, J. R., and P. Malanotte-Rizzoli (1996), Assimilation studies of open-ocean flows 1. Estimation of initial and boundary conditions, *J. Geophys. Res.*, 101(C12), 28,457–28,472.

Heimbach, P., C. Hill, and R. Giering (2005), An efficient exact adjoint of the parallel MIT general circulation model, generated via automatic differentiation, *Future Gen. Computer Syst.*, 21(8), 1356–1371.

Hill, C., V. Bugnion, M. Follows, and J. Marshall (2004), Evaluating carbon sequestration efficiency in an ocean circulation model by adjoint sensitivity analysis, *J. Geophys. Res.*, 109, C11005, doi:10.1029/2002JC001598.

Kalnay, E., et al. (1996), The NCEP/NCAR 40-year reanalysis project, *Bull. Am. Meteorol. Soc.*, 77(3), 432–470.

Köhl, A., and J. Willebrand (2002), An adjoint method for the assimilation of statistical characteristics into eddy-resolving ocean models, *Tellus, Ser. A*, 54(4), 406–425.

Lea, D. J., T. W. N. Haine, and R. F. Gasparovic (2006), Observability of the Irminger Sea circulation using variational data assimilation, *Q. J. R. Meteorol. Soc.*, in press.

LeDimet, F.-X., and O. Talagrand (1986), Variational algorithms for analysis and assimilation of meteorological observations: Theoretical aspects, *Tellus, Ser. A*, 38, 97–110.

Lemoine, F., et al. (1997), The development of the NASA GSFC and NIMA joint geopotential model, in *International Symposium on Gravity, Geoid and Marine Geodesy*, pp. 461–469, Springer, New York.

Levitus, S., and T. Boyer (1994), World ocean atlas 1994, volume 4: Temperature, *NOAA Atlas NESDIS 4*, U.S. Dep. of Comm., Washington, D. C.

Levitus, S., R. Burgett, and T. Boyer (1994), World ocean atlas 1994, volume 3: Salinity, *NOAA Atlas NESDIS 3*, U.S. Dep. of Comm., Washington, D. C.

Marotzke, J., R. Giering, K. Q. Zhang, D. Stammer, C. Hill, and T. Lee (1999), Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity, *J. Geophys. Res.*, 104, 529–547.

Marshall, J., A. Adcroft, C. Hill, L. Perelman, and C. Heisey (1997), Hydrostatic, quasi-hydrostatic and nonhydrostatic ocean modeling, *J. Geophys. Res.*, 102, 5753–5766.

Marshall, J. C., A. J. G. Nurser, and R. G. Williams (1993), Inferring the subduction rate and period over the North Atlantic, *J. Phys. Oceanogr.*, 23, 1315–1329.

Moore, A. M. (1991), Data assimilation in a quasi-geostrophic open-ocean model of the Gulf-Stream region using the adjoint method, *J. Phys. Oceanogr.*, 21(3), 398–427.

Nechaev, D. A., and M. I. Yaremchuk (1994), Conductivity temperature depth data assimilation into a 3-dimensional quasi-geostrophic open-ocean model, *Dyn. Atmos. Oceans*, 21(2–3), 137–165.

Oliger, J., and A. Sundström (1978), Theoretical and practical aspects of some initial value boundary in fluid dynamics, *SIAM J. Appl. Math.*, 35, 419–446.

Pallant, J. S., F. B. Bahr, T. M. Joyce, J. P. Dean, and J. R. Luyten (1995), Subduction in the subtropical gyre: Seasoar cruises data report, technical report, Woods Hole Oceanogr. Inst., Woods Hole, Mass.

Pond, S., and G. L. Pickard (1983), *Introductory Dynamical Oceanography*, 329 pp., Elsevier, New York.

Press, W. H., B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling (1990), *Numerical Recipies*, Cambridge Univ. Press, New York.

Reynolds, R. W., and T. M. Smith (1994), Improved global sea surface temperature analysis using optimum interpolation, *J. Clim.*, 7, 929–948.

Schröter, J., U. Seiler, and M. Wenzel (1993), Variational assimilation of GEOSAT data into an eddy-resolving model of the Gulf Stream Extension area, *J. Phys. Oceanogr.*, 23, 925–953.

Seiler, U. (1993), Estimation of open boundary conditions with the adjoint method, *J. Geophys. Res.*, 98, 22,855–22,870.

Stammer, D., C. Wunsch, R. Giering, C. Eckert, P. Heimbach, J. Marotzke, A. Adcroft, C. N. Hill, and J. Marshall (2002), The global ocean circulation during 1992–1997, estimated from ocean observations and a general circulation model, *J. Geophys. Res.*, 107(C9), 3118, doi:10.1029/2001JC000888.

Stammer, D., C. Wunsch, R. Giering, C. Eckert, P. Heimbach, J. Marotzke, A. Adcroft, C. Hill, and J. Marshall (2003), Volume, heat and freshwater transports of the global ocean circulation 1993–2000, estimated from a general circulation model constrained by WOCE data, *J. Geophys. Res.*, 108(C1), 3007, doi:10.1029/2001JC001115.

Stammer, D., K. Ueyoshi, A. Kohl, W. G. Large, S. A. Josey, and C. Wunsch (2004), Estimating air-sea fluxes of heat, freshwater, and momentum through global ocean data assimilation, *J. Geophys. Res.*, 109, C05023, doi:10.1029/2003JC002082.

Stevens, D. P. (1991), The open boundary condition in the United Kingdom fine-resolution Antarctic model, *J. Phys. Oceanogr.*, 21, 1494–1499.

Tai, C. K., and J. Kuhn (1995), Orbit and tide error reduction for the first 2 years of TOPEX/Poseidon, *J. Geophys. Res.*, 100, 25,353–25,363.

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

Woods, J. D. (1985), The physics of thermocline ventilation, in *Coupled Ocean-Atmosphere Models*, edited by J. C. J. Nihoul, pp. 543–590, Elsevier, New York.

Wunsch, C., and A. E. Gill (1976), Observations of equatorially trapped waves in Pacific sea level variations, *Deep Sea Res.*, 23, 371–390.

Yaremchuk, M. I., and N. A. Maximenko (2002), A dynamically consistent analysis of the mesoscale eddy field at the western North Pacific subarctic front, *J. Geophys. Res.*, 107(C12), 3223, doi:10.1029/2002JC001379.

Zang, X., and C. Wunsch (2001), Spectral description of low-frequency oceanic variability, *J. Phys. Oceanogr.*, 31, 3073–3095.

Zhang, K. Q., and J. Marotzke (1999), The importance of open-boundary estimation for an Indian Ocean GCM-Data synthesis, *J. Mar. Res.*, 57, 305–334.

G. Gebbie, Department of Earth and Planetary Sciences, Harvard University, 24 Oxford Street, Cambridge, MA 02138, USA. (gebbie@eps.harvard.edu)

P. Heimbach and C. Wunsch, Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139-4307, USA.