

Practical Global Oceanic State Estimation

Carl Wunsch and Patrick Heimbach

Department of Earth, Atmospheric and Planetary Sciences

Massachusetts Institute of Technology

Cambridge MA 02139 USA

cwunsch@mit.edu, heimbach@mit.edu

April 26, 2006

Abstract

The problem of oceanographic state estimation is described and contrasted with the meteorological process of data assimilation. In practice, all such methods reduce, on the computer, to forms of least-squares. The global oceanographic problem is at the present time focussed primarily on smoothing, rather than forecasting, and the data types are unlike meteorological ones. As formulated in the consortium Estimating the Circulation and Climate of the Ocean (ECCO), an automatic differentiation tool is used to calculate the so-called adjoint code, and the method of Lagrange multipliers used to render the problem one of unconstrained least-squares minimization. Major problems today lie less with the numerical algorithms (least-squares problems can be solved by many means) than with the issues of data and model error. The results of ongoing calculations suggest that the solutions are now useful for scientific purposes. Both methodology and applications are developing in a number of different directions.

1 Introduction

In physical oceanography, the problem of combining observations with numerical models differs in a number of significant ways from its practice in the atmospheric sciences. It is these differences that lead us to use the terminology “state estimation” to distinguish the oceanographers’ problems and methods from those employed under the label “data assimilation” in numerical weather prediction. “Data assimilation” is an apt term, and were it not for its prior use in the meteorological forecast community, it would be the terminology of choice. But meteorologists, faced with the goal of daily weather forecasting, have developed sophisticated techniques

directed at their own particular problems, along with an opaque terminology not easily penetrable by outsiders. Because much of oceanography has goals distinct from forecasting, the direct application of meteorological methods is often not appropriate.

This paper is primarily based upon experience with calculating the oceanic state within an academic consortium “Estimating the Circulation and Climate of the Oceans” (ECCO). The consortium was established in 1998 to exploit the new global data sets that became available during the World Ocean Circulation Experiment (WOCE) as well as the much more skillful ocean general circulation models that had also been developed.

In general terms, anyone faced with trying to calculate a system state by combining a physical model with a set of observations is faced with an estimation problem. Estimation theory is an enormous subject with a wide range of methods available. Among the techniques used are maximum likelihood, minimum variance, and Bayesian ones. In a somewhat different context, these are all methods for solution of “inverse problems,” although there has been a tendency by some authors to define inverse problems as referring only to steady-state situations. (But for a more general point of view, see Bennett, 2002.) In the geophysical sciences, inverse methods were introduced by Backus and Gilbert (1965; see especially Parker, 1994) for situations in which the parameters sought represented a static earth.

Much of the theory underlying meteorological state estimation, as in the original inverse theory of Backus and Gilbert (1965), formulates the problems in terms of continuous time and space. The combination of continuous, infinite domain, fields with finite, discrete observations leads to an interesting set of mathematical issues in functional analysis (see Parker, 1994; Talagrand, 1997; Bennett, 1992; 2002). Oceanographers, coming later to this problem, enter it at a time when computational resources are widespread and far more powerful than they were in the early days of the meteorological problem. Thus, the point of view taken within the ECCO Consortium is, both for practical and pedagogical reasons, that the problems are discrete and finite, *ab initio*, and that in practice, they can all be regarded as versions of classical non-linear least-squares.

Justifying this point of view requires some effort. The fundamental assumption is that an adequate discretization of the fluid dynamical and thermodynamical equations governing the fluid flow is available. “Adequate” means that with sufficient grid refinement or increase in the number of basis functions, that one could not distinguish the numerical solution from that of the underlying continuous formulation that generated it. Whether it exists when required is a separate problem—the assumption here is that such a discretization *could*, and if important enough, *will* eventually exist. Alternatively, one is simply asserting that a finite dimensional discrete scheme, no matter how large that dimension might prove to be in practice, can always be found

that is sufficient for the purpose at hand. Is there a physically meaningful system that is not in principle susceptible to rendering in finite dimensions? We are assuming, without proof, that the answer is “no.” (The expression “physically meaningful” is intended to exclude analogues of the pathological functions used by mathematicians as counter examples to physically-based intuition. A simple representative is the curve of unbounded variation, $\sin(1/t)$, $|t| \rightarrow 0$, which probably does not occur anywhere in nature.)

There is another aspect to this problem. All atmospheric and oceanic circulation models contain semi-empirically determined sub-gridscale processes. An example is the invocation of biharmonic friction operators, ∇^4 , to mimic small-scale turbulent dissipation. Many of the problems involved in solving inverse problems in continuous formulations are generated by these high derivatives e.g., the singular behavior of an analytical Green function—see Bennett (1992). What is unclear is whether the purely discrete rendering of the biharmonic operator is any less convincing as a representation of what is basically an unknown process? If the discrete rendering of this and other terms is regarded as adequate, all issues of convergence, continuity, differentiability, existence, essentially disappear. From a pedagogical point of view, many of the most important estimation methods (Kalman filters, adjoints, etc.) can be reduced to the practice of conventional least-squares (Wunsch, 1996, 2006). That is, we have deliberately removed the most intriguing mathematical issues from the problem—in the interest both of practical utility, and as a way of making the methodologies accessible to scientists interested in phenomenology rather than mathematics. (Analogous reductions are widespread in practical problems. A particularly well-known one is Levinson’s (1947) conversion of Norbert Wiener’s mathematically challenging filter theory—based on spectral factorization by Wiener-Hopf methods—to the least-squares form widely exploited in exploration geophysics.)

The data sets employed by oceanographers differ considerably from those employed by meteorologists. As discussed below, these differences do not create any fundamental obstacles, but require changes in the details of representation.

2 The Goal

Meteorological data assimilation has been driven by the compelling need to forecast the weather. The dominant problem in oceanography at the present time is quite different—it is to understand how the system works. Observing the ocean is technically far more difficult than observing the atmosphere—fundamentally because of its opacity to electromagnetic radiation. Thus, until comparatively recently, oceanographers of necessity had to treat the ocean as an unchanging, essentially static, system. With the data sets that have accumulated since approximately the

beginning of WOCE, about 1992, marked by the launch of the high precision altimetric mission TOPEX/POSEIDON, it becomes possible to contemplate calculating the time-evolving ocean state—to begin the process of answering the questions of whether, how, and by how much and why, the ocean circulation changes? In the terminology of control theory, the meteorological problem has been primarily one of filtering and prediction (extrapolation); the oceanographic one is primarily one of “smoothing” (interpolation). Both of these emphases can ultimately be expected to change and overlap, but they describe the present situation. Note in particular, that ocean models are major components of climate models. Such models are run out into the indefinite future, and absent any general understanding of ocean model errors, it is almost impossible to interpret the result. State estimation as practiced here, becomes in part, a serious attempt to understand which elements of ocean models are likely skillful, and capable of integration with bounded errors for finite times.

The ECCO consortium thus has undertaken the smoothing problem of using a stored data set over a finite time interval, for the purpose of making best estimates of the oceanic state during that time in such a way that the final estimate would be dynamically and kinematically self-consistent, accompanied by an understanding of the structure of residual data/model misfits. Solution of the smoothing problem makes the result differ fundamentally from what is sometimes called, in the meteorological literature “reanalysis” (e.g., Kalnay et al., 1996). Such calculations maintain a fixed numerical model configuration over a finite time interval, but produce an analysis at time, t_a , using data only for $t \leq t_a$, in an approximate “filtering” calculation. The omission of data in the formal future, $t > t_a$, means that the system state estimate can undergo jumps, implying implicit non-physical sources, and rendering very difficult the physical interpretation of the time-evolving state. Methods have been employed (e.g., Bloom et al., 1996) to smooth out the discontinuities over finite times, but still leaving artificial imbalances in the solution. Reanalyses done sequentially in the forward direction alone can leave important global imbalances: for example (Stammer et al., 2004), the Kalnay et al. (1996) estimates do not conserve freshwater or enthalpy. Such artifacts were deemed unsuitable for ECCO purposes, and is one reason that conventional meteorological practice is not obviously suitable for the oceanographic problem.

Some fundamental assumptions were made, ones that side-step the underlying mathematical problems:

- (1) The ocean state could be adequately represented by the particular GCM of Marshall et al. (1997), if run at adequate resolution. This model is an approximate numerical representation of the full non-linear Navier-Stokes and thermodynamical equations on a sphere. There is no proof of the numerical convergence of any known GCM in the limit of arbitrarily fine resolution.

(Griffies, 2005, provides a concise overview of ocean modelling issues and fundamentals.)

(2) Oceanic observations and model state variables are sufficiently close to Gaussian to justify use of a minimum variance estimator (essentially that the probability densities are unimodal). Although there are fragmentary tests of the probability densities of oceanic and atmospheric variables, there is no general description or theory.

Neither assumptions (1) nor (2) is rigorously correct, but they provide both a starting point and a focus for constant vigilance.

(3) It is possible to use oceanographic data of arbitrary type. As will see in the next sections, the data types used in practice vary greatly.

Adopting the notation of Wunsch (1996), the algebraic expression of the problem is as follows: Let $\mathbf{x}(t)$ be the model state vector at discrete time t . Although an unnecessary assumption, in practice, all of our data at time t are linearly related to the state vector as,

$$\mathbf{E}(t) \mathbf{x}(t) + \mathbf{n}(t) = \mathbf{y}(t),$$

where $\mathbf{n}(t)$ is a stochastic noise vector such that,

$$\langle \mathbf{n}(t) \rangle = 0, \quad \langle \mathbf{n}(t) \mathbf{n}(t')^T \rangle = \delta_{tt'} \mathbf{R}(t).$$

The GCM is represented as,

$$\mathbf{x}(t) = \mathbf{L}(\mathbf{x}(t - \Delta t), \mathbf{B}(t - \Delta t) \mathbf{q}(t - \Delta t), \mathbf{\Gamma}(t - \Delta t) \mathbf{u}(t - \Delta t), t - \Delta t), \quad (1)$$

where \mathbf{L} is a non-linear operator (a FORTRAN code), $\mathbf{q}(t)$ are known sources/sinks/boundary conditions. $\mathbf{u}(t)$ is the unknown control vector with moments,

$$\langle \mathbf{u}(t) \rangle = 0, \quad \langle \mathbf{u}(t) \mathbf{u}(t')^T \rangle = \delta_{tt'} \mathbf{Q}(t).$$

The matrices \mathbf{B} , $\mathbf{\Gamma}$, are available to distribute the forcing and control vectors over the state vector in a flexible way (see Wunsch, 1996).

The state vector, $\mathbf{x}(t)$, contains just enough information about the flow to march the model one time step into the future, given externally prescribed boundary conditions and sources or sinks. In the present case, it includes three components of velocity, surface elevation, temperature, and salinity at each grid point. The control vector contains any model variables subject to adjustment, here initial conditions on temperature and salinity, time-dependent corrections to the externally prescribed meteorological forcing, as well as parameters such as mixing coefficients. An important point, commonly only implicitly acknowledged, is that the $\mathbf{u}(t)$ also include the model error. Little is known of model error and an important unsolved mathematical problem concerns the representation of model error as a function of time and space. In

practice, structures that are incorrectly rendered by a model are treated as data error with little understanding of the consequences. Initial conditions are represented as $\tilde{\mathbf{x}}(0) = \mathbf{x}_0$, such that $\langle (\mathbf{x}_0 - \mathbf{x}(0))(\mathbf{x}_0 - \mathbf{x}(0))^T \rangle = \mathbf{P}(0)$. (Tildes represent estimates. $\mathbf{x}(t)$ is the hypothetical true value.) That is, there is an uncertain initial condition. (From here, we assume $\Delta t = 1$.)

Adopting the point of view that we are solving a least-squares problem, we seek to minimize the objective function,

$$J = \sum_{t=1}^{t_f} [\mathbf{y}(t) - \mathbf{E}(t)\mathbf{x}(t)]^T \mathbf{R}(t)^{-1} [\mathbf{y}(t) - \mathbf{E}(t)\mathbf{x}(t)] + [\mathbf{x}_0 - \mathbf{x}(0)]^T \mathbf{P}(0)^{-1} [\mathbf{x}_0 - \mathbf{x}(0)] \quad (2)$$

$$+ \sum_{t=0}^{t_f-1} \mathbf{u}(t)^T \mathbf{Q}(t)^{-1} \mathbf{u}(t),$$

subject to Eq. (1). Under the Gaussian assumption (implying a linear dynamical model), a solution would coincide with the minimum variance estimate and would simultaneously be a maximum likelihood solution—assuming the error covariance matrices \mathbf{R}, \mathbf{Q} are properly specified. As written, however, because the \mathbf{R}, \mathbf{Q} are poorly known, this problem is best regarded as one of constrained least-squares (curve fitting). One approach to solving it would be to use a sequential method, such as an extended Kalman filter, followed by e.g., an extended form of the RTS smoother (Rauch et al., 1965). Within ECCO, such an approach in approximate form has been undertaken by Fukumori et al. (1999)

Here we proceed to reduce the problem from constrained to unconstrained least-squares by introducing vector Lagrange multipliers, $\mu(t)$ augmenting the objective function,

$$J' = J - 2 \sum_{t=1}^{t_f} \mu(t)^T [\mathbf{x}(t) - \mathbf{L}(\mathbf{x}(t-1), \mathbf{B}(t-\Delta t)\mathbf{q}(t-1), \mathbf{\Gamma}(t-1)\mathbf{u}(t-1), t-1)]. \quad (3)$$

J can be regarded as a regularizing requirement. One can write down explicit normal equations, which with nonlinear \mathbf{L} , will be a set of nonlinear simultaneous equations, generally

well-conditioned. They have the following structure:

$$\frac{1}{2} \frac{\partial J'}{\partial \mathbf{u}(t)} = \mathbf{Q}(t)^{-1} \mathbf{u}(t) + \left(\frac{\partial \mathbf{L}(\mathbf{x}(t), \mathbf{B}\mathbf{q}(t), \mathbf{\Gamma}\mathbf{u}(t))}{\partial \mathbf{u}(t)} \right)^T \mu(t+1) = \mathbf{0}, \quad (4)$$

$$0 \leq t \leq t_f - 1$$

$$\frac{1}{2} \frac{\partial J'}{\partial \mu(t)} = \mathbf{x}(t) - \mathbf{L}[\mathbf{x}(t-1), \mathbf{B}\mathbf{q}(t-1), \mathbf{\Gamma}\mathbf{u}(t-1)] = \mathbf{0}, \quad 1 \leq t \leq t_f \quad (5)$$

$$\frac{1}{2} \frac{\partial J'}{\partial \mathbf{x}(0)} = \mathbf{P}(0)^{-1} [\mathbf{x}(0) - \mathbf{x}_0] + \left(\frac{\partial \mathbf{L}(\mathbf{x}(0), \mathbf{B}\mathbf{q}(0), \mathbf{\Gamma}\mathbf{u}(0))}{\partial \mathbf{x}(0)} \right)^T \mu(1) = \mathbf{0}, \quad (6)$$

$$\frac{1}{2} \frac{\partial J'}{\partial \mathbf{x}(t)} = \mathbf{E}(t)^T \mathbf{R}(t)^{-1} [\mathbf{E}(t)\mathbf{x}(t) - \mathbf{y}(t)] - \mu(t) \quad (7)$$

$$+ \left(\frac{\partial \mathbf{L}(\mathbf{x}(t), \mathbf{B}\mathbf{q}(t), \mathbf{\Gamma}\mathbf{u}(t))}{\partial \mathbf{x}(t)} \right)^T \mu(t+1) = \mathbf{0}, \quad 1 \leq t \leq t_f - 1$$

$$\frac{1}{2} \frac{\partial J'}{\partial \mathbf{x}(t_f)} = \mathbf{E}(t_f)^T \mathbf{R}(t_f)^{-1} [\mathbf{E}(t_f)\mathbf{x}(t_f) - \mathbf{y}(t_f)] - \mu(t_f) = \mathbf{0}. \quad (8)$$

The use of Lagrange multipliers has become known as the ‘‘adjoint method’’, and, opaquely, as ‘‘4DVAR’’, in meteorology. The special case in control theory of a terminal constraint problem (see e.g., Wunsch, 1988) is widely known as the ‘‘Pontryagin Principle.’’ In the interests of making clear to the widest possible community what is going on, we will call it the ‘‘method of Lagrange multipliers,’’ MLM, which has meaning to anyone who has encountered classical mechanics, the calculus of variations, or constrained least-squares.

(A useful way to gain some insight into these equations is to make the connection to classical mechanics and differential geometry. Adjoint operators can be defined independent of any optimization problem (e.g., Morse and Feshbach, 1953; Lanczos, 1961) and are intimately connected to the theory of Green functions; for matrix operators, they are just the transpose. Consider too (e.g., Thiring, 1997), that the partial differential operators appearing in Eqs. (4-8), before transposition, define a tangent linear operator (or model) δL acting on a tangent bundle, TM , and mapping control vector perturbations, i.e. elements of the tangent space $T_{\mathbf{x}(0)}(M)$ at $(\mathbf{x}(0))$ into the tangent space of the model state $T_{\mathbf{x}(t)}(M)$ at time t via the model Jacobian. Their transposes, as they appear in Eqs. (4-8) act on the corresponding cotangent bundle, T^*M . They define a mapping from the cotangent space $T_{\mathbf{x}(t)}^*(M)$ into $T_{\mathbf{x}(0)}^*(M)$. (This description applies to the special case where the control vector consists of the initial conditions only. The concept is readily extended to time varying boundary controls.) The spaces T_x^* are dual spaces to T_x , and their elements are called ‘‘co-vectors’’ (note that the gradient is, strictly speaking, a co-vector rather than a vector). It can readily be shown that the transpose δ^*L is the adjoint operator of the tangent linear operator, and thus Eq. (7) defines an adjoint model operating on co-vectors along $\mathbf{x}(t)$, and which are often called ‘‘adjoint variables.’’ Following along this route

leads into differential geometric optimal control theory, a subject we will not pursue here.)

Conceptually, the oceanic problem is perhaps best regarded as a form of control problem—one seeks to find those controls (typically the wind, freshwater, and heat exchanges with the atmosphere) that drive the ocean (actually the model) through a trajectory within error bars of *all* of the observations. The problem is somewhat more complex than a standard discrete control problem in that the estimated initial conditions, and often internal parameters of the model, are subject to change as part of the computation. In analogy with the classical problem of robotic control, a robot arm is required to pass within observational error of a series of known positions, velocities, accelerations, etc. on its way to a terminal state. But the numerical description of the robotic arm (the model) is believed to contain inaccuracies that are to be removed as data are employed. Similarly, the starting conditions for the robotic arm are believed erroneous in part as well and are to be re-estimated. The goal is to deduce as best possible what the state of the robotic arm was over the duration of observation and what the actual controls were, while simultaneously improving the estimates of the model and of the initial conditions. A useful general survey of many aspects of control theory, written for the physics community, can be found in Bechhoefer (2005).

If \mathbf{L} were a linear operator, one might contemplate simply solving the normal equations by Gaussian elimination or other algorithm. If non-linear, a variety of search/descent algorithms are available. The main issue for the oceanographic problem is one of dimension. In the existing ECCO configuration, with the model written at one-degree spatial resolution and 23 vertical layers, the state vector, $\mathbf{x}(t)$ is a 5.3 million dimension vector of three components of velocity, temperature and salinity and seasurface elevation (u, v, w, T, S, η) . The model is time-stepped every hour for 13 years (1992-2004) producing a total state vector of 6.1×10^{11} elements. The control vector, $\mathbf{u}(t)$, has 310 million elements over 13 years, and the data dimension is 410 million, also over 13 years. (To restart the optimization requires a numerical state vector at time t , of about 21 million values, including tendency terms.) Thus the number of equations (albeit they are sparse) is much too large to be solved directly even with a linear model. In a sequential system such as the Kalman filter plus smoother, one is faced with the computation and storage of covariance matrices square of the dimension of $[\mathbf{x}(t), \mathbf{u}(t)]$. It is for this reason that Fukumori et al. (1999) and others have resorted to a series of approximations to the filter/smoothing equations so as to reduce the effective state and control vector dimensions.

The MLM method, if a solution can be found, does not require the use of covariance matrices for the state vector—rather it is a whole domain method in which no averaging of interim solutions takes place—in the present context, that is its great attraction. Consider how an iterative solution might work: Set, initially, $\mathbf{u}(t) = 0$, as the first guess at the controls. Integrate

Eq. (5) forward in time to produce a first estimate of $\mathbf{x}(t)$. Then integrate Eq. (7) backwards in time to produce a first estimate of $\mu(t)$. To do these calculations, the partial derivatives,

$$\left(\frac{\partial \mathbf{L}(\mathbf{x}(t), \mathbf{Bq}(t), \Gamma \mathbf{u}(t))}{\partial \mathbf{u}(t)}\right)^T, \quad \left(\frac{\partial \mathbf{L}(\mathbf{x}(t), \mathbf{Bq}(t), \Gamma \mathbf{u}(t))}{\partial \mathbf{x}(t)}\right)^T, \quad (9)$$

must be known. But these define the partial derivatives of J, J' with respect to the problem parameters. Thus we are in a position to invoke a quasi-Newton method to reduce the value of J . One then modifies the problem parameters, integrates forward and backward again (an iteration) and re-evaluates the derivatives along the way. This iterative procedure was apparently first laid out in an oceanographic context by Thacker and Long (1988) and has been used many times since then. The favored descent algorithm is that of Gilbert and Lemaréchal (1989).

3 Making It Work

There are a number of major issues. In some ways, the most interesting is the problem of finding the partial derivatives (9). The transposed partial derivatives define the “adjoint model.” In the discrete formulation we are using, \mathbf{L} is defined by a computer code (tens of thousands of lines of FORTRAN). and there is no analytical expression available. (Unless one prefers, as has been done, to code the underlying continuous partial differential equations as the forward model, and to manually code the analytical form of the adjoint partial differential equations. See Thacker and Long, 1988; Sirkes and Tziperman, 1997.) Our approach has been to define the adjoint of the FORTRAN-coded model. What has made this a practical possibility is the existence of automatic (or algorithmic) differentiation (AD) tools (e.g., Rall, 1981; Griewank, 2000), beginning with software such as ADIFOR. But the most important advance was that of Giering and Kaminski (1998) who, taking account of the special structures of oceanic GCMs, provided a semi-automatic AD tool for generating code for the tangent linear model (so-called forward mode—the non-transposed partial derivatives in Eq. (9)) and the adjoint model (the so-called reverse mode) corresponding to the transposed partial derivatives. (In its present configuration, the particular AD tool used is called (TAF—Transformation of Algorithms in FORTRAN).) The tools can generate higher-order derivative code, e.g., for the Hessians. The latter are of use in assessing posterior uncertainties; Hessian eigenvalues determine the principal curvature structures of the objective function, with large curvatures indicating small posterior uncertainties in the control estimates (e.g. Thacker, 1989).

The availability of the adjoint model greatly reduces the computational load of the minimization iterations (see Marotzke et al., 1999; Heimbach et al, 2005). The GCM that is used in ECCO was constructed from the beginning with its use with an AD tool in mind, as part of

the ATOC (1999) acoustic tomography project. Certain FORTRAN77 and 90 structures remain problematic in TAF, but are avoidable. AD tools permit the adjoint code to be updated and maintained in a practical way in the presence of continuing vigorous model development and improvement.

The fundamental ECCO principle, borne out in practice, is that adequate numerical methods exist for reducing the misfit between a model and data: numerical methods can be found that solve the least-squares problem. That there are challenging numerical and storage problems goes without saying, but the principle that least-squares “works” is not in doubt. (Much of the literature in this subject, done using synthetic “data,” really reduces to the uninteresting conclusion that least-squares solutions can be obtained.)

We postulate that numerical means are available for finding the stationary point(s) of Eq. (3), be it by direct solution of the normal equations, by deterministic or stochastic search, or by pure, skillful guess. Such a solution will have a physical validity directly dependent upon the weight matrices $\mathbf{Q}(t)$, $\mathbf{R}(t)$ appearing in J . The latter are conventionally supposed to be the error covariances of the data (because in the linear limit, the least-squares solution would be the maximum likelihood or minimum variance solutions as well—ones which have a ready physical interpretation). Part of $\mathbf{Q}(t)$ represents the covariances of the control vector, describing e.g., the extent to which the prior windfield is subject to modification during the calculation. ($\mathbf{Q}(t)$ also formally includes model errors.) The solution, however it is obtained, will be no better than the choice of these matrices.

The number of elements in the combined,

$$[\mathbf{y}(t) - \mathbf{E}(t)\mathbf{x}(t)]\mathbf{R}(t)^{-1}[\mathbf{y}(t) - \mathbf{E}(t)\mathbf{x}(t)], \quad (\mathbf{x}_0 - \mathbf{x}(0))\mathbf{P}(0)^{-1}(\mathbf{x}_0 - \mathbf{x}(0))^T, \quad \mathbf{u}(t)\mathbf{Q}(t)^{-1}\mathbf{u}(t),$$

is approximately 4.1×10^8 as of this writing. Every one of these terms requires a weight! Although of minor interest to mathematicians, the choice of the weights dominates the current effort. At this stage, we have a problem primarily of oceanography and meteorology rather than one of mathematics. Table 1 provides a summary of the different data types now being fit.

With some minor exceptions, the only estimates of $\mathbf{R}(t)$, $\mathbf{Q}(t)$ refer to the diagonal values alone—that is, almost nothing is known quantitatively (or more precisely, usefully) of error covariances in the observed fields. Obtaining useful, truly independent, estimates of the non-diagonal elements of these matrices is an important, ongoing, but barely underway, effort. An interim practice is to use, where regarded as reasonable, smoothing requirements on some fields, which are equivalent to specifying off-diagonal covariances. Currently, however, even the diagonal elements represent a great deal of subjective judgement and interpretation of the vast oceanographic literature referring to all of the observations used. A number of papers (Ponte et

al., 2006; Forget et al., 2006; Stammer et al., 2006) begin to describe the estimated data errors.

As noted above, because the relationship of the weight matrices to the true error covariances remains uncertain, the solutions described below should be regarded as least-squares solutions, rather than as maximum likelihood or minimum variance ones—much of the judgement as to acceptability lies with aesthetics, rather than with rigorous statistics.

4 Sample Results

A number of papers have appeared (e.g., Li and Wunsch, 2004, Hill et al., 2005; Dutkiewicz et al., 2006; Losch and Heimbach, 2006) exploiting the adjoint model interpretation as the sensitivity of the objective function J to perturbations in data and parameters. We will focus instead on some examples of the results, as yet unpublished, of the most recent optimization effort in ECCO. The main points to be made here are that the numerical system “works” with real data, that and there is much more to be done.

The results shown here are the product of hundreds of iterations of the combined forward and adjoint models over the duration of the data sets. A truly quantitative statement concerning the optimization pathway is impossible as the current configuration builds upon some years of calculation with earlier, shorter data durations, different model configurations (e.g., changed resolution and sub-components such as mixed-layer models), new data types (e.g., the ARGO floats), changed weights in J , modifications to the line-search algorithm, changes in machine architecture, etc. Earlier results have been described at length by Stammer et al. (2002a,b, 2003), Köhl et al., (2005) among others. The calculations do strain the largest available high performance computers.

Fig. 1 is intended to show the extent to which a recent solution (which is still evolving) over the interval 1992-2004 is capable of reproducing one of the major constraining data sets—in this case the sea surface height anomaly (relative to a long-term mean) during July 1997. (An animation is available at <http://ocean.mit.edu/king/mpegs/12iter125.mpeg>.) The TOPEX/-POSEIDON and related altimetric missions provide remarkably accurate measurements of the absolute global ocean surface elevation every 10 days (absolute accuracies in the range of 2-3cm; see Wunsch and Stammer, 1998; Fu and Cazenave, 2001). The misfit terms in J corresponding to the time-mean seasurface height over the 13-year model run are displayed in Fig. 2. Ideally, all these values would be order 1 with outliers consistent with a χ^2 distribution. There is a reasonable fit over the bulk of the open ocean. Serious misfits persist in coastal regions generally, in much of the Southern Ocean, and in some regions of strong variability. The scientific issue is now whether these misfits represent (A) a mis-weighting of the data; (B) model error; or

(C) simply that the iterative minimization has not had time to reduce these terms, or all of these things. As noted above, the issue of data weighting takes one into the intricacies of the underlying observations. In this particular case, calculation of the data errors involves one in the determination of Earth gravity (providing the reference gravitational surface called the “geoid”; see Wunsch and Stammer, 1998), the myriad error terms in altimetric determination of seasurface topography (see e.g., Chelton et al., 2000), and the expected temporal variability as a function of position and duration in the ocean (Ponte et al., 2006). Model error quantification involves determining the extent to which a 1° horizontal resolution, 23 vertical-layer model would be expected to reproduce the mean seasurface topography—were it perfectly known.

A representative misfit to an entirely different data type is shown in Fig. 3. These data are obtained by so-called ARGO floats (see Roemmich and Owens, 2000) drifting freely in the ocean, and which profile temperature and salinity on a vertical ascent to the seasurface where the data are transmitted ashore by satellites. Similar questions apply here as to the altimetry—are the data being properly weighted, and is the model being forced toward structures which are physically impossible for it?

Numerous other applications related to ECCO have appeared. Examples include Ayoub (2006) and Gebbie et al. (2006), which are regional estimates, the latter using open boundary conditions and high resolution. Stammer (2005) and Ferreira et al. (2005) discuss estimation of internal eddy-stress and mixing parameters as part of the optimization. Many others will be forthcoming.

5 Concluding Discussion

A practical system exists for solving a very large, global-scale, oceanographic least-squares problem using the method of Lagrange multipliers. The methodology relies heavily on automatic differentiation (AD) tools now specially constructed for use with oceanic general circulation models (GCMs). When written in the finite, discrete form required by computers, most of the interesting and challenging mathematical issues are side-stepped. Many of the issues in practice are oceanographic rather than mathematical, and focus on the error covariances of a multiplicity of data types.

Where is this effort going? ECCO and related projects are developing in several directions. Understanding of the data errors and acquisition of new data types is ongoing, and likely will never be completed; it is an asymptotic process. AD tools (<http://www-unix.mcs.anl.gov/OpenAD>) are under development in a number of places, directed at greater efficiency, an open-source, and new computer languages; see Naumann et al. (2006). Model errors of certain types

can be eliminated by augmenting the control vector by uncertain model elements, e.g., as Losch and Heimbach (2006) did for bottom topography. The search for greater model resolution (see Menemenlis et al., 2005) continues; this too, is likely an asymptotic process, and the search for rules for parameterizing the indefinitely present sub-gridscale motions continues. At bottom, of course, the purpose of all of this effort is to understand the ocean and the practical implications of that understanding. As one example of that use, Fig. 4 displays the estimated mass, heat and salt fluxes in the North Atlantic across 26°N as a function of time from the least-squares fit (Wunsch and Heimbach, 2006). This result is part of the ongoing debate over whether the North Atlantic circulation is undergoing some kind of important shift. Another example of the application of the results is in picking apart the heat and freshwater contributors to global sealevel rise (Wunsch, Ponte and Heimbach, 2006).

Extension of the primarily fluid physical model used here is being undertaken to include a more adequate ice dynamics and thermodynamics model. Dutkiewicz et al. (2005) have employed the same AD tools, and the background physical state for parallel study of oceanic biogeochemical cycles. Further extension of the models and methods to the study of past climate states is now getting underway.

Inadequate model resolution is likely to be a major problem for the foreseeable future. Some fluid phenomena in some regions (e.g., boundary jets, fronts, etc.) cannot be parameterized in any easy or known way, and the estimation of their properties demands a model resolution that is impossible on a global scale. To this end, developments are underway permitting embedding of high resolution models into lower resolution global estimates such as the one described here. The major issue is that the sub-model boundaries must be taken as “open,” so as to permit exchange of fluid and information between the two models. Although there are some mathematical questions about the posedness of such approaches (e.g., Bennett, 2002), no particular problems are encountered in practice (Ayoub, 2006; Gebbie et al., 2006) when the numerics are once again viewed as least-squares problems in which regularization is straightforward. The Gebbie et al. (2006) embedded model is depicted in Fig. 5. It has been argued (e.g., Lea et al., 2000) that the system may become so chaotic that tangent linear gradients fail to be useful. Real systems demonstrating a lack of useful numerical differentiability would direct one toward Monte Carlo methods, but no such situation has so far been encountered in practice.

Acknowledgements.

Supported by the National Ocean Partnership Program, the National Science Foundation ITR program, and the National Aeronautics and Space Administration through the Jason and GRACE missions. The bulk of the ECCO computations reported here were carried out at the Geophysical Fluid Dynamics Laboratory, NOAA, and at the National Center for Atmospheric

Research. Thanks are due our many ECCO partners. Comments from two anonymous reviewers were helpful.

References

- ATOC Consortium, The, 1998. Ocean climate change: comparison of acoustic tomography, satellite altimetry, and modeling. *Science*, 281, 1327-1332.
- Ayoub, N. 2006. Estimation of boundary values in a North Atlantic circulation model using an adjoint method. *Ocean Modeling*, 12(3-4), 319-347, doi:10.1016/j.ocemod.2005.06.003.
- Backus, G. E. and J. F. Gilbert, 1967. Numerical applications of a formalism for geophysical inverse theory. *Geophys. J. Roy. Astron. Soc.*, 13, 247-276.
- Bechhoefer, J., 2005. Feedback for physicists: A tutorial essay on control. *Revs. Mod. Phys.*, 77, 783-836.h
- Bennett, A. F., 1992. *Inverse Methods in Physical Oceanography*. Cambridge Un. Press, Cambridge, 346 pp.
- Bennett, A. F., 2002. *Inverse Modeling of the Ocean and Atmosphere*. Cambridge Un. Press, Cambridge, 234 pp.
- Chelton, D. B., J. C. Ries, B. J. Haines, L.-L. Fu, and P. S. Callahan, 2000. Satellite altimetry, in, *Satellite Altimetry and Earth Sciences: A Handbook of Techniques and Applications*, L.-L. Fu and A. Cazenave Eds., Academic, 1-132.
- Dutkiewicz, S., M. Follows, P. Heimbach and J. Marshall, 2005. Controls on ocean productivity and air-sea carbon flux: An adjoint model sensitivity study. *Geophys. Res. Lett.*, 33, L02603, doi:10.1029/2005GL024987.
- Ferreira, D. J. Marshall and P. Heimbach, 2005. Estimating eddy stresses by fitting dynamics to observations using a residual-mean ocean circulation model and its adjoint. *J. Phys. Oc.*, 35(10), pp. 1891-1910, doi:10.1175/JPO2785.
- Fu, L.-L. and A. Cazenave, Eds., 2001. *Satellite Altimetry and Earth Sciences. A Handbook of Techniques and Applications*. Academic, San Diego, 463 pp.
- Fukumori, I. R. Raghunath, L.-L. Fu, Y. Chao, 1999. Assimilation of TOPEX/Poseidon altimeter data into a global ocean circulation model: how good are the results? *J. Geophys. Res.* 104, 25647-25665.
- Gebbie, G., P. Heimbach and C. Wunsch, 2006. Strategies for nested and eddy-resolving state estimation. *J. Geophys. Res.*, in press.
- Giering, R. and T. Kaminski, 1998. Recipes for adjoint code construction. *ACM Trans. Math. Software*, 24, 437-474.
- Gilbert, J. C. and C. Lemaréchal, 1989. Some numerical experiments with variable-storage quasi-Newton algorithms. *Math. Prog.*, 45(B), 407-435.

- Griewank, A., 2000. *Evaluating Derivatives. Principles and Techniques of Algorithmic Differentiation*. SIAM, Philadelphia, 369pp.
- Griffies, S. M., 2005. Some ocean model fundamentals. In: *Ocean Weather Forecasting: An Integrated View of Oceanography*. E. P. Chassignet & J. Verron, Eds., Berlin, Germany: Springer, 19-74.
- Heimbach, P., C. Hill and R. Giering, 2005. An efficient exact adjoint of the parallel MIT general circulation model, generated via automatic differentiation. *Future Generation Computer Systems*, 21(8), c1356-1371, doi:10.1016/j.future.2004.11.010.
- Hill, C. V. Bugnion, M. Follows and J. Marshall, 2004. Evaluating carbon sequestration efficiency in an ocean circulation model by adjoint sensitivity. *J. Geophys. Res.*, 109(C11005).
- Kalnay, E. and 21 others, 1996. The NCEP/NCAR 40-year reanalysis project. *Bull. Am. Met. Soc.*, 77, 437-471.
- Köhl, A., D. Stammer, and B. Cornuelle, 2005. Interannual to decadal changes in the ECCO global synthesis. Submitted for publication.
- Lanczos, C., 1961. *Linear Differential Operators*. Van Nostrand, Princeton, 564 pp.
- Lea, D. J. M. R. Allen, and T. W. N. Haine, 2000. Sensitivity analysis of the climate of a chaotic system. *Tellus*, 52A, 523-532.
- Levinson, N., 1947. The Wiener RMS (root mean square) error criterion in filter design and prediction. *J. Math. Phys.*, 25, 261-278.
- Li, X. and C. Wunsch, 2004. An adjoint sensitivity study of chlorofluorocarbons in the North Atlantic. *J. Geophys. Res.*, 109(C1), 10.1029/2003JC002014.
- Losch, M. and P. Heimbach, 2006. Adjoint sensitivity of an ocean general circulation model to bottom topography, *J. Phys. Oceanogr.*, in press.
- Marchuk, G. I., 1995. *Adjoint Equations and Analysis of Complex Systems*. Kluwer, Dordrecht, 466 pp.
- Marotzke, J., R. Giering, K. Q. Zhang, D. Stammer, C. N. 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, 29,529-29,547.
- Marshall, J. A. Adcroft, C. Hill, L. Perelman, and C. Helsey, 1997. A finite-volume, incompressible Navier-Stokes model for studies of the ocean on parallel computers. *J. Geophys. Res.*, 102, 5753-5766.
- Menemenlis, D. C., and 13 others, 2005. NASA supercomputer improves prospects for ocean climate research. *EOS, Trans. Am Geophys. Un.*, 86, 89, 95-96.
- Morse, P. M. and H. Feshbach, 1953. *Methods of Theoretical Physics*, 2 vols. McGraw-Hill, New York, 1978pp.

Naumann, U., J. Utke, P. Heimbach, C. Hill, D. Ozyurt, C. Wunsch, M. Fagan, N. Tallent, and M. Strout, 2006. Adjoint code by source transformation with OpenAD/F. in, P. Wesseling, E. Oñate and J. Périaux, Eds. *European Conference on Computational Fluid Dynamics, ECCOMAS CFD 2006*, TU Delft, The Netherlands, 2006.

Parker, R. L., 1994. *Geophysical Inverse Theory*. Princeton U. Press, Princeton, 377 pp.

Rall, L. B., 1981. *Automatic Differentiation: Techniques and Applications*. Springer-Verlag, Berlin, 165 pp.

Rauch, H. E., F. Tung, and C. T. Striebel, 1965. Maximum likelihood estimates of linear dynamic systems. *AIAA J.*, 3, 1445-1450.

Roemmich, D. and Owens, W. B., 2000. The Argo Project: Global Ocean Observations for Understanding and Prediction of Climate Variability. *Oceanography*, 13(2), 5-50.

Sirkes, Z. and E. Tziperman, 1997. Finite difference of adjoint or adjoint of finite difference. *Mon. Wea Rev.* 125, 3373-3378.

Stammer, D., C. Wunsch, R. Giering, C. Eckert, P. Heimbach, C. Hill, J. Marotzke, J. Marshall, 2002. The global ocean state during 1992-1997, estimated from ocean observations and a general circulation model. Part I. Methodology and estimate state. *J. Geophys. Res.*, DOI: 10.1029/2001JC000888.

Stammer, D., C. Wunsch, I. Fukumori, and J. Marshall, 2002. State estimation in modern oceanographic research. *EOS*, 83(27), 289, 294-5.

Stammer, D., C. Wunsch, R. Giering, C. Eckert, P. Heimbach, J. Marotzke, A. Adcroft, C. Hill, J. Marshall, 2003. Volume, heat and freshwater transports of the global ocean circulation 1992-1997, estimated from a general circulation model constrained by WOCE data. *J. Geophys. Res.*, DOI: 10.1029/2001JC001115.

Stammer, D., 2005. Adjusting internal model errors through ocean state estimation. *J. Phys. Oc.*, 35, 1143-1153.

Talagrand, O., 1997. Assimilation of observations, an introduction. *J. Meteor. Soc. Japan*, 75, 191-209.

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

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

Thirring, W., 1997. *Classical Mathematical Physics*, 3rd Ed.. Springer-Verlag, New York, 543pp.

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

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

Wunsch, C., 1996. *The Ocean Circulation Inverse Problem*. Cambridge University Press, Cambridge, 437 pp.

Wunsch, C., 2006. *Discrete Inverse and State Estimation Problems. With Geophysical Fluid Applications*. Cambridge Un. Press, Cambridge, in press.

Wunsch, C., and P. Heimbach, 2006. Decadal changes in the North Atlantic meridional overturning and heat flux. *J. Phys. Oc.*, in press

Wunsch, C., R. Ponte, P. Heimbach, 2006. Decadal trends in global sealevel patterns and implications for global mean errors. To be published.

Wunsch, C. and D. Stammer, 1998. Satellite altimetry, the marine geoid and the oceanic general circulation. *Ann. Revs. Earth Plan. Scis.*, 26, 219-254.

Figure Captions

1. Upper panel shows the seasurface elevation as inferred from 10-days of altimetric data during July 1997. Lower panel shows the estimate made from the model constrained both to the upper panel (with data used daily) and to many other data types. Main visual signal is the large El Niño event that occurred during this time.

2. Misfit of the time mean over 13 years of the constrained model to the time-average altimetric seasurface height. An ideal solution would be everywhere of order 1, but distributed in χ^2 with the expected random outliers. Here there are strong regional misfits whose explanation is in part misweighted data, model error, and failure thus far to have achieved complete convergence.

3. Misfit over 13 years to the so-called ARGO temperature data, vertically averaged. (These data do not become available until about half-way through the calculation and are dominated in numbers by the last two years, 2003, 2004.) Again, a nominal value of one with random χ^2 outliers would denote an acceptable solution.

4. Zonal and vertical integrals across the constrained model at 26°N in the North Atlantic as a function of time, showing (upper panel) the total meridional mass flux (in Sverdrups - $10^6\text{m}^3/\text{s}$), the meridional heat flux in W (middle panel), and the salt flux (in kg/s) in the lowest panel. (From Wunsch and Heimbach, 2006). The absence of any clear trends supports the inference that no major climate shift is taking place today in this part of the North Atlantic Ocean.

5. Temperature field (Gebbie et al., 2006) showing nesting of an open boundary high resolution model into the global one. Open boundary properties, including flows, become part of the control vector. The temperature ranges from 3.5°C (blue) to 18°C (red).

Table 1, showing the main data sources, coverage, and some information about the originator.

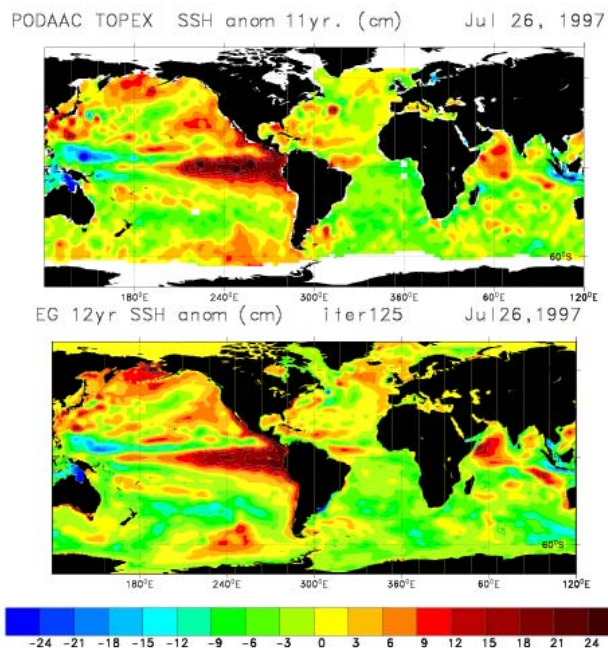


Figure 1: Upper panel shows the seasurface elevation as inferred from 10-days of altimetric data during July 1997. Lower panel shows the estimate made from the model constrained both to the upper panel (with data used daily) and to many other data types. Main visual signal is the large El Niño event that occurred during this time.

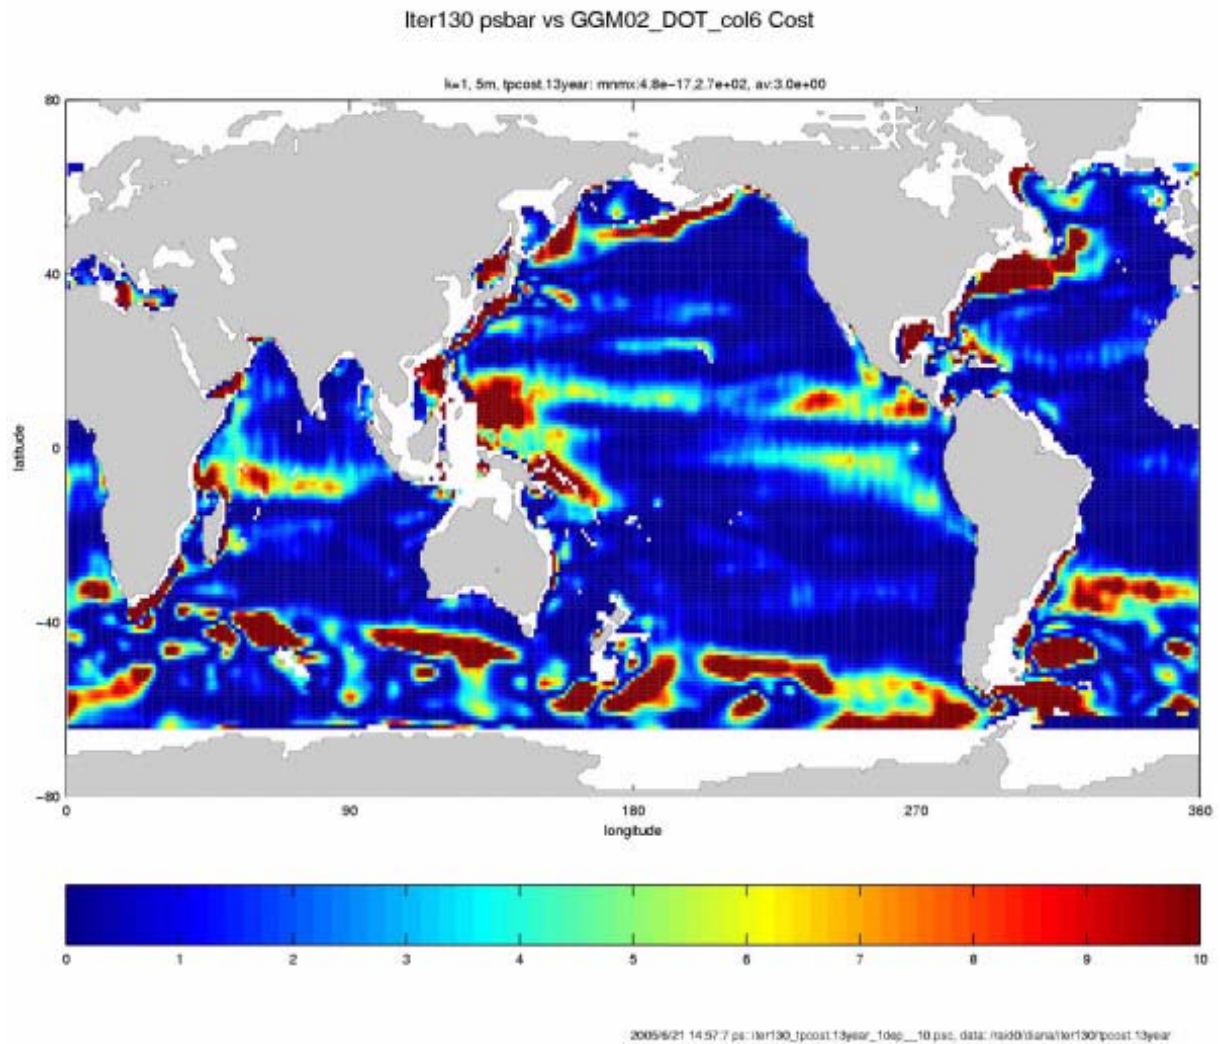


Figure 2: Misfit of the time mean over 13 years of the constrained model to the time-average altimetric seasurface height. An ideal solution would be everywhere of order 1. Here there are strong regional misfits whose explanation is in part misweighted data, model error, and failure thus far to have achieved complete convergence.

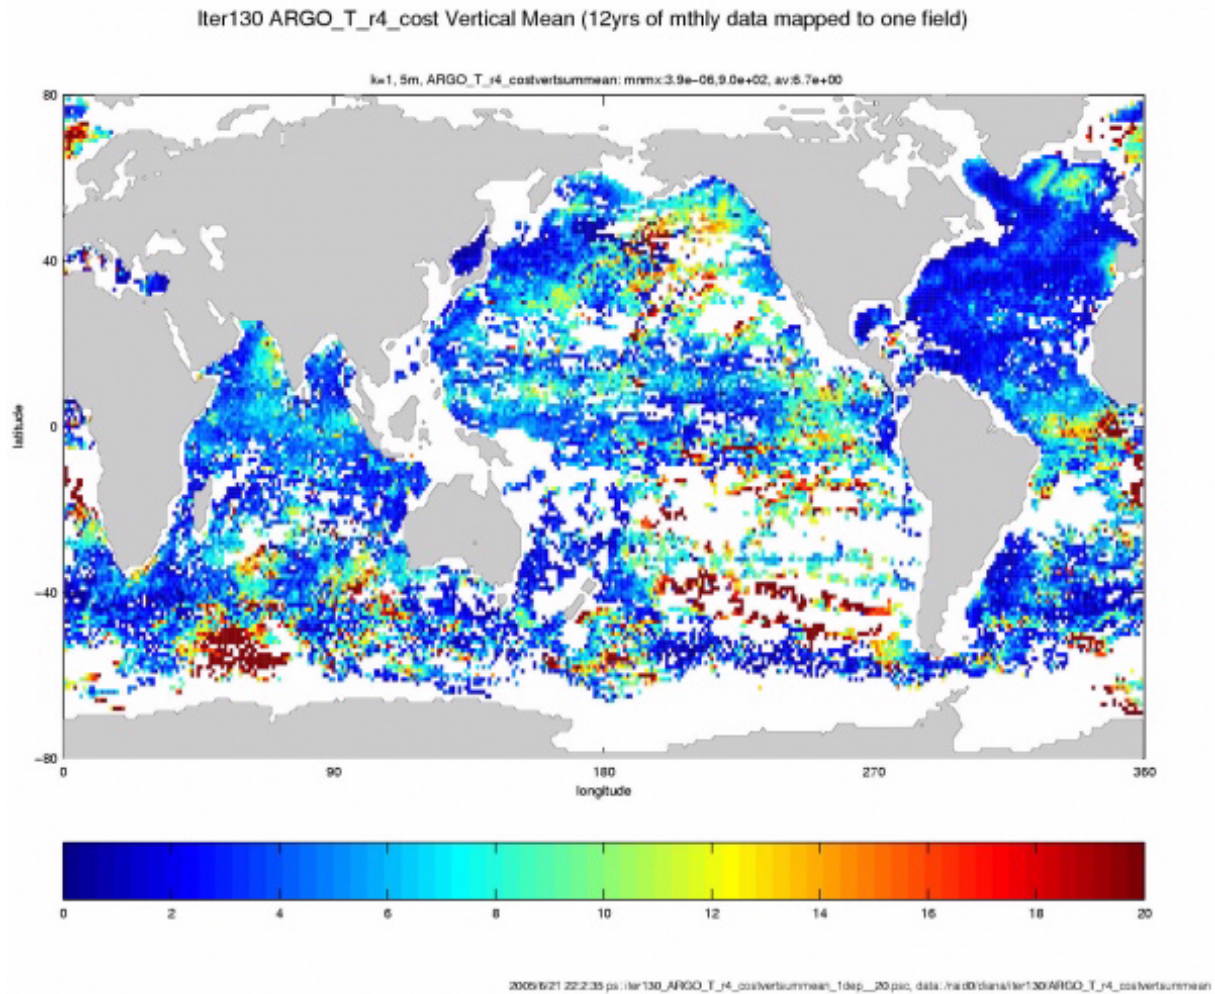


Figure 3: Misfit over 13 years to the so-called ARGO temperature data, vertically averaged. (These data do not become available until about half-way through the calculation and are dominated in numbers by the last two years, 2003, 2004.) Again, a nominal value of one would denote an acceptable solution.

iter121 30-May-2005 19:50:15 CW-home




Figure 4: Zonal and vertical integrals across the constrained model at 26°N in the North Atlantic as a function of time, showing (upper panel) the total meridional mass flux (in Sverdrups - $10^6\text{m}^3/\text{s}$), the meridional heat flux in W (middle panel), and the salt flux (in kg/s) in the lowest panel.




Figure 5: Temperature field (Gebbie et al., 2005) showing nesting of an open boundary high resolution model into the global one. Open boundary properties, including flows, become part of the control vector. The temperature ranges from 3.5°C (blue) to 18°C (red).

DATA TYPE	Source	Spatial Extent	Variable(s)	Duration	Number of values
Altimetry: TOPEX/POSEIDON	PODAAC	Global, equatorward of 65 degrees	height anomaly, temporal average	1993-2004	4500/day
Altimetry: Jason	PODAAC	Global equatorward of 65 degrees	height anomaly, temporal average	2002-2004	4500/day
Altimetry: Geosat-followon	US Navy, NOAA	Global, equatorward of 65 degrees	height anomaly	2001-2004	4300/day (gridded)
Altimetry: ERS-1/2, ENVISAT	AVISO	Global, equatorward of 81.5 degrees	height anomaly	1995-2004	3800/day
Hydrographic climatology	Gouretski and Koltermann (2004)	global, 300m to seafloor	temperature, salinity	1950-2002 inhomogeneous average	16 million
Hydrographic climatology	World Ocean Atlas (2001), Conkright et al. (2002)	global to 300m	temperature, salinity	multidecadal average seasonal cycle	included above
CTD synoptic section data	Various, including WOCE Hydro. Prog.	global, all seasons, to 3000m	temperature, salinity	NA	17 thousand
XBTs	D. Behringer (NCEP)	global, but little So. Ocean	temperature	1992-2004	1.4 million
ARGO Float profiles	IFREMER	global, above 2500m	temperature, salinity	1997-2004	2.1 million
Sea Surface Temperature	Reynolds and Smith (1999)	global	temperature	1992-2004	5.3 million
Sea Surface Salinity	(ECOP)	tropical Pacific	salinity	1992-1999	24,238
TMI	NASA/NOAA	global	temperature	1998-2003	1.5 million
Geoid (GRACE mission)	GRACE Project. (GGM02C_DOT_cel6. Tapley et al. (2003)	global	mean dynamic topography	NA	1 degree resol.
Bottom Topography	Smith&Sandwell(1997)+ETOP05	Smith/Sandwell to 72.006, ETOP05 to 79.5	water depth	NA	1 degree resol.
FORCING:					
Windstress-scatterometer	PODAAC	global	stress	1992-99, 7/1999-2004	9.4 million
Windstress	NCEP/NCAR reanalysis Kalnay et al. (1996)	global	stress	1992-2004	192x94 Gaussian grid (approx. 1.875 deg)
Heat Flux	NCEP/NCAR reanalysis	global	sensible + latent heat	1992-2004	"
Freshwater Flux	NCEP/NCAR reanalysis	global	evap/2.4x10^9-precip	1992-2004	"
Long Wave Radiation (experiment)	NCEP/NCAR reanalysis	global		1992-2004	"
WITHHELD (as of April 2006):					
tide gauges		global, sparse	sealevel		
TOGA-TAO array		equatorial oceans	velocity/temperature		
Tomographic integrals		N. Pacific	heat content		
Florida Current transport		Florida Straits	mass flux		
Float and Drifter Velocities		Global	velocity		

Figure 6: Table 1, showing the main data sources, coverage, and some information about the originator.