

A computational framework for simulation of biogeochemical tracers in the ocean

Samar Khatiwala¹

Received 4 January 2007; revised 7 April 2007; accepted 23 May 2007; published 4 July 2007.

[1] A novel computational framework is introduced for the efficient simulation of chemical and biological tracers in ocean models. The framework is based on the “transport matrix” formulation, a scheme for capturing the complex three-dimensional transport of tracers in a general circulation model (GCM) as a sparse matrix, thus reducing the task of simulating tracers to a sequence of simple matrix-vector products. The principal advantages of this formulation are efficiency and convenience. It is many orders of magnitude more efficient than GCMs, allowing us to address problems that are currently either difficult or unaffordable with GCMs. The scheme also allows us to quickly “prototype” new biogeochemical parameterizations or “plug in” existing ones. This paper describes the key features and advantages of the transport matrix method, and illustrates its application to a series of realistic problems in chemical and biological oceanography. The examples range from simulation of a transient tracer (SF₆) to adjoint sensitivity of a complex coupled biogeochemical model. Finally, the paper describes an efficient, portable, and freely available implementation of this computational scheme that provides the necessary framework for simulating any biogeochemical tracer.

Citation: Khatiwala, S. (2007), A computational framework for simulation of biogeochemical tracers in the ocean, *Global Biogeochem. Cycles*, 21, GB3001, doi:10.1029/2007GB002923.

1. Introduction

[2] Coupled physical-biogeochemical models of the ocean are important tools in the quest to understand the cycling of chemical and biological tracers such as nutrients and carbon. They also allow us to quantify the ocean uptake of climatically important gases such as CO₂, and clarify the interpretation of various paleoceanographic tracers as proxies of past ocean circulation. Finally, these models place quantitative constraints on key physical processes such as mixing, and provide a basis for more rigorously evaluating the underlying physical circulation model.

[3] Global modeling of chemical and biological tracers has traditionally been performed within the framework of “box models” and three-dimensional ocean general circulation models (GCMs). While box models are simple to use, GCMs have the advantage of spatial resolution, ability to represent important physical and biogeochemical processes, and a circulation field that is dynamically consistent. However, the increasing sophistication and complexity of modern GCMs have likely also hindered their more widespread use by the biogeochemical and paleoceanographic community. GCMs remain technically challenging to use,

both in terms of incorporating new biogeochemical processes or simply running them. They are also computationally demanding, as many biogeochemical tracers require several thousand years to reach equilibrium. These factors likely discourage potential users and thus limit our ability to fully exploit the tools of numerical modeling to address fundamental problems in chemical and biological oceanography.

[4] Here I introduce a novel computational framework for simulating biogeochemical tracers in the ocean. At its core is the “transport matrix method”, or “TMM”, of *Khatiwala et al.* [2005] (KVC, hereafter) for efficient simulation of passive tracers in ocean models. The essential idea is that the discrete tracer transport operator of a GCM can be written as a sparse matrix, which may be efficiently constructed by “probing” the GCM with a passive tracer. This empirical approach ensures that the circulation embedded in the “transport matrix” accurately represents the complex three-dimensional advective-diffusive transport (including all sub-grid-scale parameterizations) of the underlying GCM. Once the matrix has been derived, the GCM can be dispensed with, as simulating a tracer is reduced to a sequence of simple matrix-vector products. Consequently, this allows us to focus on the biogeochemical processes of interest, rather than the details of transporting tracers around.

[5] The matrix method has several advantages over GCMs or conventional “offline” tracer models. First, it is up to many orders of magnitude more efficient. Second, it can directly compute steady state solutions of the tracer

¹Lamont-Doherty Earth Observatory, Columbia University, Palisades, New York, USA.

equations, thus circumventing expensive, transient integrations. Third, it allows us to perform calculations such as adjoint sensitivity analysis that are currently either difficult or unaffordable with GCMs. Finally, it is very convenient to use, allowing us to quickly “prototype” (implement and test) novel biogeochemical parameterizations, or “plug in” existing ones, and rapidly obtain results, all in the context of a realistic underlying circulation field. Moreover, the resulting biogeochemical models are independent of any particular GCM, which should encourage their adoption by a wider community of users. Given these advantages, an increasing number of scientists have expressed interest in using the transport matrix method in their own research. Problems currently being addressed with the TMM include: inverse modeling of anthropogenic carbon uptake by the ocean; optimizing biogeochemical parameterizations through model-data comparison; understanding the impact of the oceanic iron cycle on atmospheric $p\text{CO}_2$; and prognostic modeling of paleoceanographic proxies. The TMM has also been used in a graduate course in biogeochemical modeling at Lamont-Doherty.

[6] The goal of this paper is to bring the transport matrix method to a wider audience. It is directed at both observationalists, who may be interested in a potentially powerful tool for prototyping new biogeochemical models or testing novel hypotheses, and modelers, who may find the TMM an efficient alternative to GCMs and offline tracer models. In the following, I provide a brief description of the matrix formulation and summarize its primary features and advantages. The bulk of the paper is devoted to a series of examples from chemical and biological oceanography. These range from a transient tracer simulation to adjoint sensitivity analysis of a coupled biogeochemical model of carbon, phosphorus, and iron. (A brief but self contained description of the “adjoint method” is given in Appendix A.) The examples serve to illustrate the flexibility and efficiency of the matrix formulation. Indeed, many of the computations shown were performed on a desktop computer in minutes to a few days at most. Using traditional models would take far longer and require substantially greater resources. Finally, while the matrix formulation is independent of any particular implementation; that is, the only requirement is the facility to perform matrix-vector products, I describe a portable and efficient implementation of the TMM that greatly facilitates the use of this technique. The software, which is based on the freely available PETSc scientific library from Argonne National Laboratories, makes it relatively easy to add new biogeochemical models or plug in existing ones. Furthermore, the code can be run without modification on most computers, sequential or parallel.

2. Transport Matrix Method

[7] In this section, I briefly describe the ideas underlying the TMM. A more detailed description can be found in KVC. The time evolution of passive tracers in the ocean is governed by a linear advection-diffusion equation. A GCM solves a discretized version of this equation. The transport matrix approach is based on the recognition that for a passive tracer, the discretized advection-diffusion equation

can be written as a linear matrix equation. At the simplest level, this may be written as

$$\frac{d\mathbf{c}}{dt} = \mathbf{A}(t)\mathbf{c} + \mathbf{q}(t), \quad (1)$$

where \mathbf{c} is the vector of tracer concentrations at the grid points of the GCM, i.e., a vector representation of a discretized three-dimensional tracer field. \mathbf{q} represents sources and sinks, including surface and bottom (sediment) fluxes. The matrix $\mathbf{A}(t)$ is the “transport matrix” (“TM”) which results from discretization of the advection-diffusion operator and includes the effects of advection, diffusion, and various (parameterized) sub-grid-scale processes. The significance of writing the tracer equation in matrix form is that it simplifies the problem of solving the advection-diffusion equation to a sequence of matrix-vector products.

[8] While the above formulation (also considered by Primeau [2005]) has its advantages, it does not capture the full range of numerical schemes implemented in modern ocean GCMs. For instance, many OGCMs now implement more sophisticated numerical schemes based on the “direct space-time” approach wherein both the time and space derivatives are simultaneously and self consistently discretized resulting in a fully discrete equation [e.g., *Hundsdoerfer and Verwer*, 2003]. This is in contrast to “method-of-lines” schemes wherein the spatial operator is first discretized, resulting in a system of ODEs (equation (1)). The latter is then time stepped via a stable integration scheme. GCMs also typically use some form of operator splitting wherein horizontal advection and diffusion are treated explicitly in time, while vertical mixing (and, increasingly, advection) is treated implicitly in time. This is to ensure stability while allowing for a reasonably large time step. Implicit treatment of vertical mixing is also essential when dealing with deep convection and in some subgrid mixing schemes such as the K profile parameterization (KPP) of *Large et al.* [1994]. To accommodate these complexities, we write the discrete tracer equation as

$$\mathbf{c}^{n+1} = \mathbf{A}_i^n (\mathbf{A}_e^n \mathbf{c}^n + \mathbf{q}^n), \quad (2)$$

where n is the time step index, \mathbf{A}_e the TM for the explicit-in-time component of advection-diffusion, and \mathbf{A}_i the matrix for implicit transport. Note that, as written, equation (2) appears to be fully explicit in time. This is because \mathbf{A}_i is in fact the inverse of the operator arising in the implicit solution of the advection-diffusion equation. It is important to recognize that any linear transport equation, however complex, can be written in matrix form, although it is often difficult to obtain the explicit form of the matrix. Indeed, it would be impractical to directly code the advection-diffusion equation, including various parameterizations, for an arbitrary ocean geometry in the form of equation (2). Instead, KVC describe an empirical procedure that allows us to efficiently construct \mathbf{A}_e and \mathbf{A}_i by “probing” the GCM with a passive tracer. The procedure involves performing a forward integration of the GCM with a suite of passive tracers, each of which is reset after every time step to a particular spatial pattern. These patterns or “basis functions” are automatically constructed

using graph coloring methods [e.g., *Curtis et al.*, 1974; *Gebremedhin et al.*, 2005]. This empirical approach ensures that the circulation embedded in the transport matrix, by construction, satisfies the equations of motion and incorporates transport due to all parameterized sub-grid-scale processes (eddy-induced mixing, convective adjustment, etc.) represented in the GCM. It is also consistent with the geometry and numerics (advection scheme, time-stepping method, implicit mixing, etc) of the underlying GCM. The KVC algorithm can be applied to most OGCMs, requiring little or no modification to the underlying GCM code. Once the TM has been derived, the GCM can be dispensed with. Tracer simulations are performed entirely via equation (1) or equation (2). This “decoupling” of the biogeochemical process model (the “ \mathbf{q} ” term) from the tracer transport implies that we are not tied to any particular GCM. The same biogeochemistry code can be used without modification with TM’s derived from different GCMs or even with different geometries or resolutions. This should be contrasted with offline models, which are typically based on, and tied to, a particular GCM code and hence not independent of that GCM. Circulation fields obtained from one GCM cannot therefore be readily used, if at all, with an offline model derived from a different GCM.

3. Key Features of the Transport Matrix Formulation

[9] The transport matrix method has a number of important features that I summarize below.

3.1. Sparsity of the Transport Matrix

[10] The TM’s \mathbf{A}_e and \mathbf{A}_i are both extremely sparse. Physically, this is due to the finite speed of advection and diffusion, i.e., in a single time step tracer can only spread a finite distance since grid cells typically communicate only with their nearest neighbors. In the vertical, convective adjustment can spread the tracer further, but this is restricted to small regions at high latitudes and is (typically) captured entirely by \mathbf{A}_i . The sparsity of the TM’s can be exploited in many ways. First, storage and computational requirements are drastically reduced. Second, programming languages such as MATLAB now provide extensive support for sparse matrices. And third, it allows us to take advantage of efficient and scalable computational kernels for sparse matrices provided by widely available numerical libraries such as PETSc and Sparse BLAS.

3.2. Seasonal Cycle

[11] Variability in the circulation can be accounted for by appropriate time averaging. For example, if the transport has a seasonal cycle, the GCM can be integrated for a full year and an average TM computed for each month, resulting in 12 different \mathbf{A}_e ’s and \mathbf{A}_i ’s. Linear interpolation in time can then be used to obtain the TM at any time instant. In numerous experiments, I have found this practice to be quite effective in capturing the seasonal cycle.

3.3. Longer Time Steps

[12] In ocean GCMs, the time step is generally limited by the dynamics to a few hours. Models that are “synchronously” integrated use this same (small) time step for both the dynamical and tracer (both active and passive) equations. (Offline tracer models can generally take longer time steps.) This can be prohibitively expensive for many problems, especially those involving biogeochemical tracers that take several thousand years to approach a steady state. A particular advantage of the TM method is that it allows us to take much longer time steps ($O(\text{days})$) than that of the underlying GCM with minimal loss of accuracy. Specifically, to increase the effective time step in equation (2) by a factor m , we modify \mathbf{A}_e and \mathbf{A}_i as follows:

$$\mathbf{A}_e \rightarrow \mathbf{I} + m(\mathbf{A}_e - \mathbf{I})$$

$$\mathbf{A}_i \rightarrow \mathbf{A}_i^m.$$

Here \mathbf{I} is the identity matrix and dt is the tracer time step in the underlying GCM. The first of these relations follow from rewriting the explicit-in-time term as a forward Euler time step, but with a time step of mdt . The second relationship follows from applying the implicit transport term m times. The modified TMs have the exact same sparsity as the original TMs.

3.4. Coarse Graining of the Transport Matrix

[13] The transport matrix can be “coarse grained” to any desired resolution via a simple linear transformation of the form $\mathbf{A} \rightarrow \mathbf{MAM}$, where \mathbf{M} and \mathbf{M} are sparse matrices that depend on the geometry of the problem. As it is the tracer fluxes that are averaged onto a coarser grid, this is more accurate than integrating the GCM at the (lower) resolution at which the matrix is required. (Of course, it is also possible to directly compute the TM at a lower resolution than that of the underlying GCM.) Thus, for instance, a TM computed from a high-resolution GCM can be averaged down to a lower resolution TM. Simulations carried out with the coarser TM can potentially be orders of magnitude faster than at the full resolution while retaining, to a good degree, many of the features of the higher-resolution model. To illustrate this idea, I have derived a zonally-averaged TM from the TM of a 2.8° resolution GCM, and subsequently used the coarse grained matrix to compute the steady state (natural) radiocarbon distribution in the model. The number of grid boxes in the full model is ≈ 53000 , but is only ≈ 2000 in the zonally-averaged one. (The zonal averaging is performed in each ocean basin.) As can be seen in Figure 1, the results compare favorably with those obtained with the full three-dimensional TM, and subsequently zonally averaged.

[14] One intriguing application of this feature is to spinning up biogeochemical tracers in eddy resolving OGCMs. The idea is to use a coarsened TM of the eddy resolving GCM to compute an equilibrium solution, that in turn can be used as an initial condition for the GCM. If the matrix solution is sufficiently close to the true equilibrium solution (defined in some average sense) of the eddy

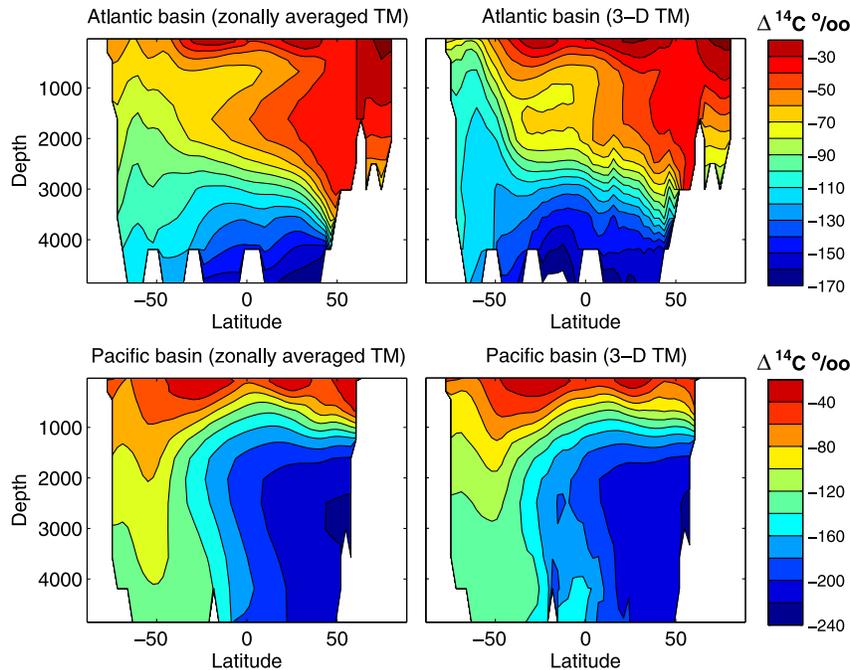


Figure 1. Plots comparing the zonally averaged distribution of $\Delta^{14}\text{C}$ in the (top) Atlantic and the (bottom) Pacific basins. Solutions shown on the left were obtained using the zonally averaged transport matrix. Those on the right were obtained by zonally averaging the solution computed with the transport matrix for the full three-dimensional model. To simulate radiocarbon, I use the idealized approach of *Toggweiler et al.* [1989] in which the tracer concentration c is restored to 100 units in the surface layer with a wind speed-dependent timescale.

resolving GCM, this approach could potentially circumvent long transient integrations of the GCM.

3.5. Tracers With Prescribed Surface Boundary Conditions

[15] In some problems, tracers are subject to a prescribed concentration boundary condition (BC) at the surface. An example is the “ideal age” tracer which is subject to a zero BC [e.g., *Thiele and Sarmiento, 1990*]. To solve such problems, we convert equation (2) into an equation for the time evolution of the interior tracer field as follows. We split \mathbf{A}_e into an “interior” matrix \mathbf{A}_e^i , and a “boundary” matrix \mathbf{B}_e (and similarly for \mathbf{A}_i). The (square) interior matrix describes the transport of tracer between interior grid points. To construct it, we simply extract the rows and columns of \mathbf{A}_e corresponding to interior grid points. The (rectangular) boundary matrix describes the exchange of tracer between the surface boundary points and the interior, and is composed of the rows and columns of \mathbf{A}_e corresponding to interior and boundary grid points, respectively. Equation (2) then becomes

$$\mathbf{c}_i^{n+1} = \mathbf{A}_i^i (\mathbf{A}_e^i \mathbf{c}_i^n + \mathbf{B}_e \mathbf{c}_B^n + \mathbf{q}_i^n) + \mathbf{B}_i \mathbf{c}_B^{n+1}, \quad (3)$$

where \mathbf{c}_i is the vector of interior tracer concentrations, \mathbf{c}_B the vector of prescribed, possibly time-dependent, surface boundary values, and \mathbf{q}_i the interior source/sink term. We will apply equation (3) below.

3.6. Steady State Solutions

[16] In many cases, we are interested in the equilibrium distribution of tracers. As mentioned above, obtaining such solutions via direct integration of the GCM is computationally very demanding. However, in those instances in which the seasonal cycle can be neglected (e.g., by using an annual mean transport matrix), equation (2) allows us to directly compute steady state distributions without the need for long transient integrations. This is because a fundamental difference between the TMM and a conventional time-stepping model is that in the former, the transport operator is an explicit matrix, while in the latter it is only implicitly defined via the time-stepping equations. (This fact is exploited when deriving the TM.) Hence in the TMM we can write the steady solution as a system of linear or nonlinear equations and bring to bear all the machinery developed to solve such equations.

[17] Specifically, for linear source/sink and boundary terms, the equilibrium distribution is the solution to a simple linear system of equations. This follows from equation (2) by setting $\mathbf{c} \equiv \mathbf{c}^{n+1} = \mathbf{c}^n$ and solving for \mathbf{c} . This is how the radiocarbon distributions displayed in Figure 1 were computed. Another example is the steady state ideal age or mean age, τ_m , a useful diagnostic of ocean ventilation. In a conventional GCM, the calculation of τ_m requires integrating a tracer with a unit source in the interior and a surface BC of zero concentration to steady state, an expensive computation even at coarse resolution. In our notation,

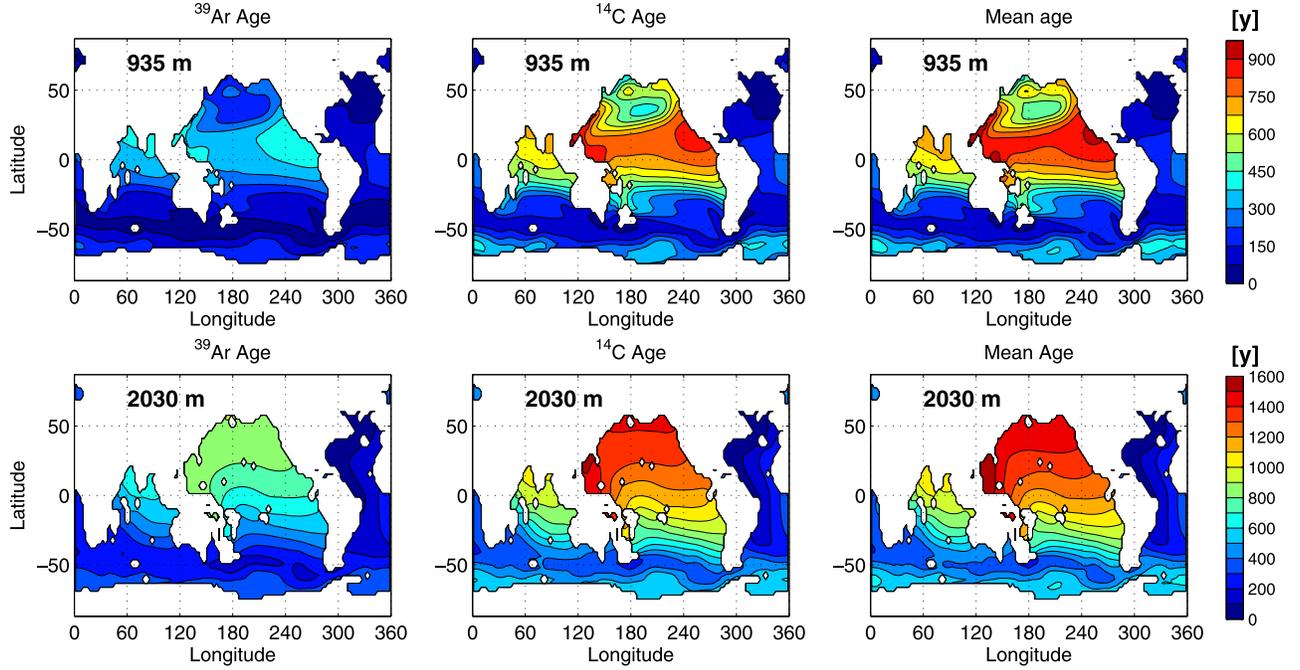


Figure 2. Distribution of (top) ^{39}Ar age, (middle) ^{14}C age, and (right) mean age at (top) 935 m and (bottom) 2030 m obtained with the annual mean TM for the 2.8° model.

$\mathbf{c}_B = \mathbf{0}$, and $\mathbf{q}_I = dt\mathbf{1}$, where $\mathbf{1}$ is a vector of 1's and dt is the time step. Substituting these in equation (3) we have

$$(\mathbf{A}_i^1 \mathbf{A}_e^1 - \mathbf{I})\tau_m = -dt\mathbf{A}_i^1 \mathbf{1},$$

which can be readily solved for τ_m . Figure 2 shows the solution obtained using the annual mean TM for a 2.8° resolution global model (see section 4 for details). For comparison, Figure 2 also shows ages based on two radioactive tracers, ^{39}Ar (half-life of 269 years) and ^{14}C (half-life of 5370 years) obtained by solving an analogous equation. It is interesting to note that $\tau_{^{39}\text{Ar}} < \tau_{^{14}\text{C}} \leq \tau_m$. This ordering is consistent with the well known effect of mixing on tracer ages, i.e., tracer ages are biased toward younger values (relative to the mean age) with the bias increasing with decreasing decay time [Khatiwala *et al.*, 2001; Waugh *et al.*, 2003].

[18] For reasonably sized problems (up to 2° resolution, global), the linear system can be solved with sparse direct solvers such as MATLAB's "\" operator or SuperLU [Demmel *et al.*, 1999; Li and Demmel, 2003]. (A comprehensive list of sparse direct solvers is maintained at: <http://www.cise.ufl.edu/research/sparse/codes/>.) At higher resolution, it becomes more difficult to store and factor the coefficient matrix. The reason is that while \mathbf{A}_e and \mathbf{A}_i are extremely sparse, their product $\mathbf{A}_i \mathbf{A}_e$ is significantly denser (by a factor of $O(nz/5)$, where nz is the number of grid points in the vertical, and the "5" is related to the stencil of the spatial discretization). Such problems can be addressed using iterative solvers based on Krylov subspace methods [Saad, 2003].

[19] In the general case of nonlinear sources and sinks, the steady solution satisfies a nonlinear system of equations of the form

$$\mathbf{F}(\mathbf{c}) = (\mathbf{A}_i \mathbf{A}_e - \mathbf{I})\mathbf{c} + \mathbf{A}_i \mathbf{q}(\mathbf{c}) = 0. \quad (4)$$

This equation is readily extended to multiple tracers. One of the most common methods for solving such a system is Newton's method [Kelley, 1995]. To apply Newton's method, we require the Jacobian $\partial \mathbf{F} / \partial \mathbf{c}$. With the matrix formulation, this reduces to the problem of computing $\partial \mathbf{q} / \partial \mathbf{c}$, since the Jacobian of the transport term is simply $\mathbf{A}_i \mathbf{A}_e - \mathbf{I}$. For relatively simple problems, it may be possible to construct $\partial \mathbf{q} / \partial \mathbf{c}$ (or an approximation) by hand [e.g., Kwon and Primeau, 2006]. For complex biogeochemical source/sink terms involving multiple tracers, the use of automatic differentiation tools (see Appendix A) makes this task quite feasible.

[20] As a typical example of a nonlinear problem, we consider a simple model of ocean carbon based on the OCMIP-2 protocols (Abiotic model) [Orr *et al.*, 1999b]. This model simulates total dissolved inorganic carbon (DIC), and has a parameterization for air-sea exchange of CO_2 at the surface (the so-called solubility pump). No biological effects are included. The source/sink term at the i -th surface grid point is

$$q_i^n = dt \left\{ -\frac{V_{\text{gas}i}}{dz} (1 - F_{\text{ice}i}) \left([\text{CO}_2^{\text{aq}}]_i^n - \alpha_i p \text{CO}_2^{\text{atm}n} \right) + [\text{DIC}]_g \frac{Q^{\text{freshwater}}}{dz} \right\}. \quad (5)$$

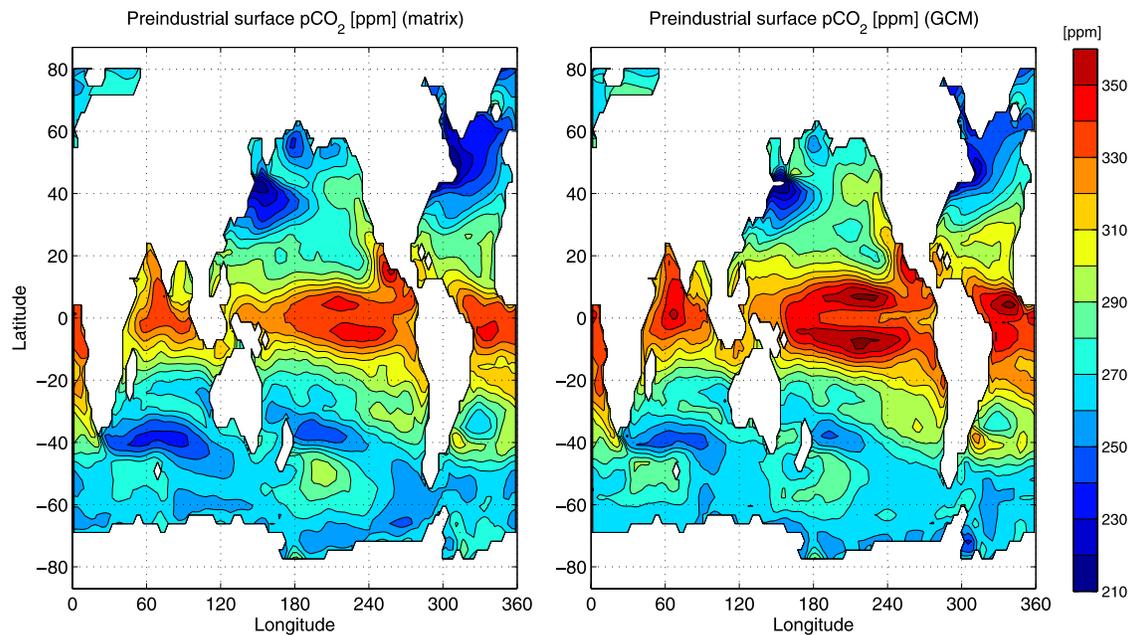


Figure 3. Map of preindustrial surface $p\text{CO}_2$ simulated using (left) the annual mean transport matrix and (right) a GCM with full seasonal cycle and annually averaged. The GCM results were computed by M. Follows and S. Dutkiewicz at MIT as part of OCMIP-2. They are available online at <http://dods.ipsl.jussieu.fr/ocmip/phase2/MIT/>.

Here V_{gas} is the gas transfer velocity parameterized according to Wanninkhof [1992] (see Orr *et al.* [1999a] for details), F_{ice} the local sea ice fraction, and α the solubility (a function of local temperature (T) and salinity (S)). The surface aqueous CO_2 concentration depends on DIC, alkalinity (an empirical function of S), T , and S through the nonlinear equilibrium carbonate chemistry. The last term on the RHS represents a “virtual” flux of DIC due to the effective freshwater flux ($E - P + \text{salinity restoring}$), and is proportional to the global surface mean DIC concentration ($[\text{DIC}]_{\text{g}}$). (T , S , and $Q^{\text{freshwater}}$ are available from the underlying GCM.) To compute the steady state solution for the preindustrial DIC, I apply Newton’s method, holding the atmospheric $p\text{CO}_2$ fixed at 278 ppm in equation (5). To apply this technique, we require the Jacobian of \mathbf{q} , $\partial\mathbf{q}/\partial[\text{DIC}]$, or a close approximation to it. The problem can be greatly simplified by ignoring the virtual flux term, thus making the Jacobian diagonal. (The virtual flux term is only ignored during Jacobian construction. The final steady state solution incorporates the full effect of this term.) A good approximation to the required derivative $\partial[\text{CO}_2^{\text{aq}}]/\partial[\text{DIC}]$ is provided by the carbonate chemistry solver (at no additional cost, since this derivative is also required to compute the aqueous CO_2 concentration). The method, implemented in a dozen lines of MATLAB code, converges very rapidly to a steady state. Figure 3 (left) shows the preindustrial surface $p\text{CO}_2$ predicted by the model. Also shown (right) is the annual mean preindustrial surface $p\text{CO}_2$ computed “online” in a similarly (but not identically) configured GCM with full seasonal cycle and integrated for several thousand years to equilibrium.

Considering the absence of a seasonal cycle in the matrix computation, the close similarity between the matrix and GCM solutions is quite encouraging. I should note that if seasonality is a particular concern, spinning up the model with a seasonally varying TM is a viable and inexpensive alternative. A 5000 year spinup with the TMM takes less than 1 day on a fast workstation.

[21] An alternative to explicitly constructing the Jacobian is to use a variant of Newton’s method known as “Jacobian-free Newton-Krylov” (JFNK) [Kelley, 2003; Knoll and Keyes, 2004]. This technique has the advantage of requiring only function evaluations (i.e., $\mathbf{F}(\mathbf{c})$) and avoiding explicit computation and storage of the Jacobian matrix (impractical for large problems). Moreover, it has recently been shown that JFNK can also be used for seasonally forced problems (T. Merlis and S. Khatiwala, Matrix-free Newton-Krylov methods for fast dynamical spin up of ocean models, submitted to *Ocean Modelling*, 2007).

3.7. Multiple Tracers

[22] Many biogeochemical process models involve multiple interacting tracers. Equation (2) then applies to each tracer, with the source/sink term \mathbf{q} being, most generally, a nonlinear function of all the tracer species. Typically, at any grid point, this dependence is restricted to tracer concentrations in the same vertical column. As discussed later, this is an important simplification that can be exploited to great advantage.

3.8. Adjoint Tracer Transport

[23] The adjoint to the tracer equation is essential in many applications, including studies of ocean ventilation [e.g.,

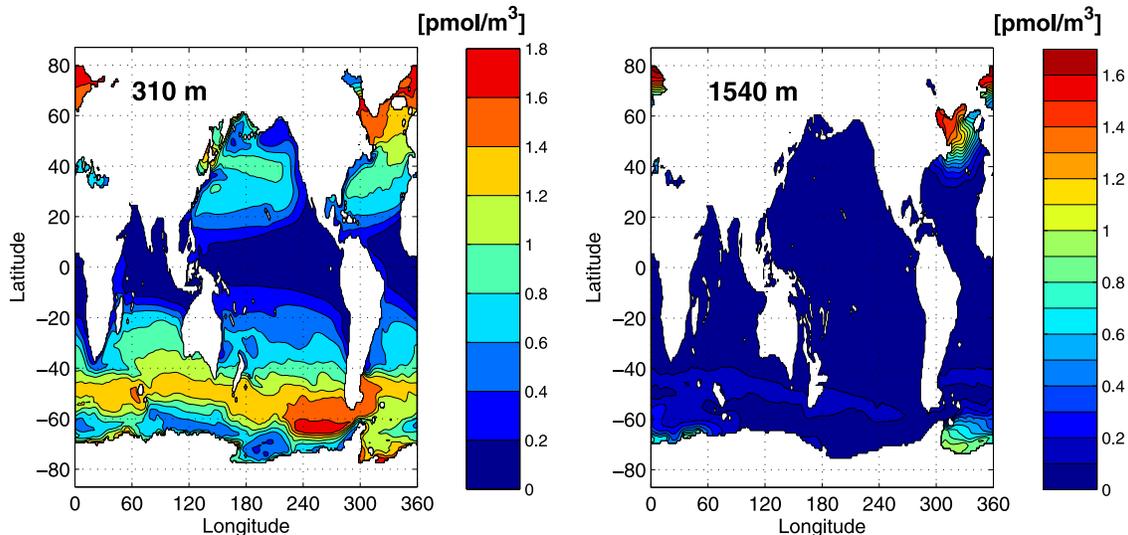


Figure 4. Distribution of SF_6 in 2002 at (left) 310 m and (right) 1540 m in the ECCO-based transport matrix simulation.

Holzer and Hall, 2000; Fukumori et al., 2004; Holzer and Primeau, 2006], optimization of biogeochemical model parameters [e.g., Fennel et al. 2001], inverse modeling [e.g., Schlitzer, 1999; Kaminski and Heimann, 2001], and parameter sensitivity analysis [e.g., Li and Wunsch, 2004; Dutkiewicz et al., 2006]. For a conventional GCM, even an offline one, construction of an adjoint, which must typically be integrated backward in time, remains a nontrivial task. The recent availability of “adjoint compilers” [Giering and Kaminski, 1998] has made the task more manageable, but far from routine. When written in matrix form, however, construction of the adjoint becomes considerably easier, as the adjoint of the advection-diffusion term involves a simple transpose of the transport matrices (section 4.2).

4. Illustrative Examples

[24] I now present a series of examples to illustrate how the matrix method can be applied to simulate chemical and biological tracers. To perform the simulations, I use two different transport matrices based on different configurations of the MIT GCM, a state-of-the-art primitive equation model [Marshall et al., 1997]. The first is a TM derived from a 2.8° global configuration of the MIT model with 15 vertical levels, and forced with monthly mean climatological fluxes of momentum, heat, and freshwater. In addition, surface temperature and salinity are weakly restored to the Levitus climatology [Levitus et al., 1998]. I use the equilibrium state of the model, obtained after a 5000 year integration, to derive \mathbf{A}_e and \mathbf{A}_i at monthly mean resolution. While the size of the TMs is 52749×52749 , only 0.19% (0.02%) of the elements of \mathbf{A}_e (\mathbf{A}_i) are nonzero.

[25] The second configuration is the data assimilated model from the ECCO (“Estimating the Circulation and Climate of the Ocean”) consortium [Stammer et al., 2004]. The ECCO strategy involves adjusting air-sea fluxes to

bring the model into consistency (within error limits) with the data. As discussed by Wunsch and Heimbach [2007], this is essentially a least-squares fit to the data constrained, exactly, by the models’ dynamics. This optimization, has been performed using a 1° , 23 vertical level, global configuration of the MIT GCM for the 13 year period, 1992–2004. The result is a dynamically consistent estimate of the ocean’s circulation and hydrography over that period [Stammer et al., 2004]. Data used in this procedure include drifter velocity, WOCE hydrography measurements, global XBT data, PALACE and ARGO temperature and salinity profiles, and satellite sea surface height. The monthly mean TM derived from this model represents a climatology over the 13 year assimilation period. The size of the ECCO TM is 682604×682604 , with 0.016% (0.0029%) of the elements of \mathbf{A}_e (\mathbf{A}_i) being nonzero. Both configurations use a third-order direct space time advection scheme with operator splitting for tracers, and a variety of parameterizations to represent unresolved processes, including isopycnal thickness diffusion [Gent and McWilliams, 1990] and KPP for vertical mixing. All of these features are automatically incorporated into the TM. The simulations described below were performed using the PETSc-based implementation of the TMM mentioned in the introduction (see Appendix B).

4.1. Transient Tracer Simulation

[26] As a first example, I simulate SF_6 , an anthropogenic tracer, using the 1° ECCO-derived TM. Unlike the CFCs, SF_6 concentrations in the atmosphere are still rising, and there is considerable interest in using it as a ventilation tracer. The time-dependent integration shown here is based on the OCMIP-2 protocols for simulating CFCs [Orr et al., 1999a], but modified appropriately for SF_6 . The source/sink term arising from air-sea fluxes is similar to that for carbon (see section 3.6): $q^n = -atk ([\text{SF}_6]^n - [\text{SF}_6]_{\text{sat}}^n)$. Here k is the gas transfer coefficient, specifically the transfer velocity

divided by the thickness of the surface grid box. As in the abiotic carbon model, the gas transfer velocity is parameterized according to *Wanninkhof* [1992]. The surface saturation concentration is based on the atmospheric history of SF_6 tabulated by *Maiss and Brenninkmeijer* [1998]. I use monthly mean TM's, and the monthly mean fields of gas transfer velocity and sea ice fraction provided by OCMIP-2. All variables are linearly interpolated in time. As a practical matter, because of the size of each TM ($O(1 \text{ GB})$), the calculation was performed on 8 processors of a commodity Beowulf cluster. The entire 50-year integration (1953–2002) took ≈ 4 hours. Figure 4 shows the simulated SF_6 distribution in 2002 at shallow and intermediate depths. The spatial patterns are very similar to those for CFC-12 (not shown).

4.2. Adjoint Transit Time Distribution

[27] Transit time distributions (TTDs) provide a conceptual framework for understanding the transport of tracers in geophysical fluids such as the ocean and atmosphere [*Holzer and Hall*, 2000]. TTDs formalize the heuristic notion that in a turbulent fluid there is a multiplicity of pathways connecting two points, each associated with a different “transit time.” They are a useful diagnostic of tracer transport and ocean ventilation, and have been applied to many problems, including the interpretation of transient tracers [*Khatiwala et al.*, 2001; *Waugh et al.*, 2003, 2004], ocean uptake of anthropogenic carbon [*Haine and Hall*, 2002; *Hall et al.*, 2004], and interpretation of paleoceanographic proxies [*Rutberg and Peacock*, 2006].

[28] In a numerical model, it is straightforward to simulate TTDs. They satisfy the advection-diffusion equation with no sources/sinks (i.e., they are passive tracers), have a zero initial condition, and are subject to an impulse (“delta”) boundary condition in time at the surface. We can also simulate “marginal” TTDs by partitioning the surface into patches and applying the delta BC on any given patch [e.g., *Primeau*, 2005]. The time-dependent response (tracer concentration) at any interior point can be interpreted as a distribution of transit times relative to that patch. Its time integral is the fraction of water at that grid point that was last in contact with (i.e., “ventilated” from) the surface patch over which the delta was applied.

[29] As a specific application, consider the problem of quantifying source regions for a particular interior water-mass defined, for example, by its T/S properties. With a forward model, to compute the fraction of water that originated at every surface patch or grid point requires one forward integration of the model with a unit step BC at the corresponding patch. With this BC, the solution at any interior location at time t is, in effect, the time integral from 0 to t of the marginal TTD with respect to the surface patch. The steady state solution, i.e., $t \rightarrow \infty$, averaged over the interior volume of interest, is the fraction of water in that volume that was last ventilated at the surface patch. If the number of surface grid points/patches is large, this procedure can be prohibitively expensive. Instead, we can use the “adjoint” tracer model to perform the calculation far more efficiently as follows. Consider simulating a tracer with an initial concentration of zero units, and a step BC at the

j th surface grid point. We denote this BC by the vector Δ_j , i.e., a unit vector with a “1” in the j th row. Repeated application of equation (3), with $\mathbf{c}_1^0 = \mathbf{0}$, $\mathbf{q}_1 = \mathbf{0}$, and $\mathbf{c}_B = \Delta_j$, yields the following solution for the interior tracer concentration at time step n (ignoring seasonality for clarity of presentation):

$$\mathbf{c}_j^n = \left\{ \sum_{k=0}^{n-2} \mathbf{M}_b \mathbf{M}_1^{n-2-k} \right\} \Delta_j \equiv \mathbf{G}^n \Delta_j, \quad (6)$$

with $\mathbf{M}_1 \equiv \mathbf{A}_1^T \mathbf{A}_1$ and $\mathbf{M}_b \equiv \mathbf{A}_1^T \mathbf{B}_e + \mathbf{B}_i$.

[30] The term in “{ }” is the Green function \mathbf{G}^n for this problem, and since Δ_j is a unit vector, \mathbf{c}_j^n is simply the j th column of \mathbf{G} . The forward approach described above is equivalent to computing \mathbf{G} one column at a time. However, recall that we are only interested in a specific interior volume, i.e., the average of \mathbf{c}_j over some volume. This average may be written as $\bar{\mathbf{c}}_j^n \equiv \alpha^T \mathbf{c}_j$, where α is a vector that picks out the elements of \mathbf{c} that belong to the watermass of interest and performs a spatial average. To compute this average for all j (surface grid points), we form $\alpha^T \mathbf{G}^n$. This is a row vector. Its transpose is a column vector whose j th element is (for large n , i.e., in steady state) the fraction of water in the water mass that was last in contact with the j th surface grid point,

$$\mathbf{f} \equiv \lim_{n \rightarrow \infty} \bar{\mathbf{c}}^n = \lim_{n \rightarrow \infty} \mathbf{M}_b^T \left\{ \sum_{k=0}^{n-2} \mathbf{M}_1^{T^{n-2-k}} \right\} \alpha.$$

It is important to note that to evaluate the above expression we never actually compute \mathbf{G} (a large dense matrix). Instead, we use the following procedure, which is equivalent to time stepping the adjoint equations backward in time:

[31] Initialize: $\mathbf{c}_1^\dagger = \alpha$,
 $\mathbf{c}_2^\dagger = \alpha$
 for $n = 1$ to convergence do
 $\mathbf{c}_2^\dagger = \mathbf{M}_1^T \mathbf{c}_2^\dagger$
 $\mathbf{c}_1^\dagger = \mathbf{c}_1^\dagger + \mathbf{c}_2^\dagger$
 end for
 $\mathbf{f} = \mathbf{M}_b^T \mathbf{c}_1^\dagger$.

[32] An interesting application of this procedure is to the problem of identifying surface source regions that ventilate thermocline waters of the equatorial Pacific [see also *Fukumori et al.*, 2004]. There is considerable interest in this issue, both from a theoretical point of view and because of its potential role in long-term climate variability. For example, *Gu and Philander* [1997] have suggested that tropical-subtropical oceanic exchange may provide a means by which a change in surface conditions in the subtropical Pacific can affect the tropical thermocline, in turn perturbing the coupled ocean-atmosphere system, thus causing climate variability on decadal timescales. Using the annual mean ECCO TM, I have computed the fraction of water in the Niño-3 region ($5^\circ \text{S} - 5^\circ \text{N}$, $150^\circ - 90^\circ \text{W}$) between 80 and 180 m (black rectangle in Figure 5) that originates at any surface grid point. The results, displayed in Figure 5, show the extent to which subtropical surface waters ventilate the

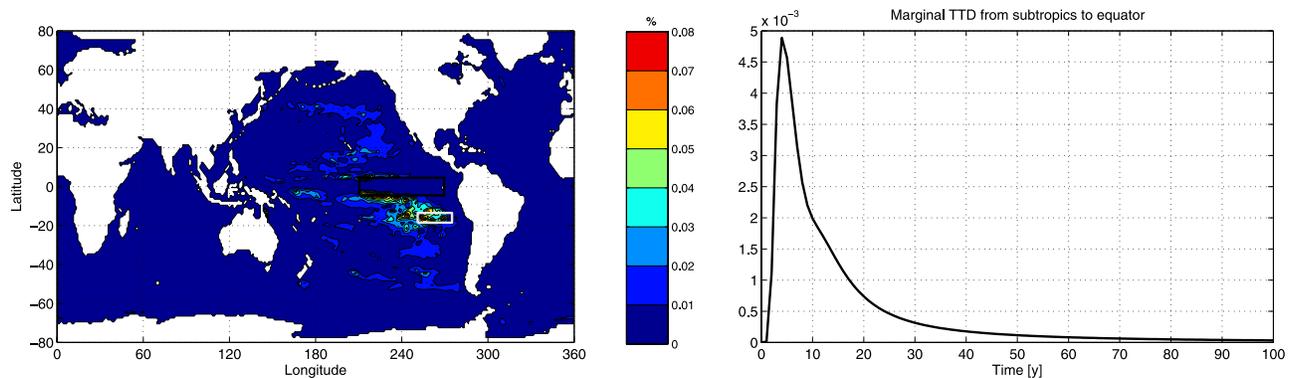


Figure 5. (left) Map showing fraction of water (in %) at any surface grid point contributing to thermocline waters (80–180 m) in the Niño-3 region of the equatorial Pacific (black rectangle). The map was computed by time stepping the ECCO-derived TM adjoint model backward for 1500 years. (right) Marginal TTD, averaged over the Niño-3 region between 80–180 m, relative to a surface patch in the Southern Hemisphere subtropics (gray rectangle on map). The area under the curve is ≈ 0.06 , indicating that roughly 6% of the water in the Niño-3 region originates from that surface patch.

equatorial thermocline. The highest fraction values occur along the southern edge of the Niño-3 box and in the southeastern subtropical Pacific. Presumably, thermocline waters are transported to the equator by the tropical and subtropical cells. Note that f is normalized such that its elements sum to 1. Summing f_j over each hemisphere we find that, consistent with previous work [Tsuchiya, 1968; Goodman *et al.*, 2005], $\approx 2/3$ of the water in the equatorial thermocline was last ventilated in the southern hemisphere, with the balance originating north of the equator. Interestingly, a small fraction also originates in the Labrador Sea, eventually upwelling at the equator. Also shown in Figure 5 (right) is the marginal TTD for the Niño-3 region relative to the southern hemisphere subtropics (gray rectangle). The mean transit time (first moment of the TTD) to the equator from this region is ≈ 1.5 years. However, owing to the skewness of the TTD, 50% of the fluid elements in the volume have a transit greater than 36 years (the median age).

4.3. Adjoint Sensitivity of Anthropogenic Carbon Uptake to Gas Exchange

[33] It is often useful to compute the sensitivity (gradient) of some model output variable (the “cost function”) to various model parameters (“control variables”). Examples of the former include the export production of biological matter, atmospheric $p\text{CO}_2$, and the misfit between the model and observations. Examples of control variables include parameters such as the rate and length scale for remineralization of organic matter and the flux of iron via aeolian dust input. Gradients can be invaluable if, for instance, we wish to optimize the parameters in a biogeochemical model by minimizing the misfit between simulation and observations. Sensitivity patterns are also of intrinsic interest in helping us better understand the factors controlling the behavior of the system.

[34] If the number of control variables is small, we can compute the sensitivity via finite differences, i.e., perturbing one parameter at a time and computing the difference between the perturbed and unperturbed cost function. For

each parameter, this requires a forward integration of the model (often to equilibrium), an approach that rapidly becomes unaffordable with the number of control variables, as would be the case if the control variable is a two- or three-dimensional field. The “adjoint method” allows us to solve this difficult problem by replacing many forward runs with a single forward integration of the model followed by a backward integration of the “adjoint” model. While a detailed description is beyond the scope of this paper, I provide a brief summary in Appendix A, along with a practical algorithm that can be applied to most biogeochemical problems. I also show how the matrix formulation greatly simplifies construction of the adjoint equations. Here I apply the adjoint method to two typical problems.

[35] As a first example, I apply it to the OCMIP abiotic carbon model of section 3.6. In particular, I compute the sensitivity of anthropogenic carbon uptake in the model to the gas transfer velocity. To simulate anthropogenic carbon, I first compute the preindustrial DIC distribution as described in section 3.6 with a fixed atmospheric $p\text{CO}_2$ of 278 ppm. Subsequently, I switch on the anthropogenic perturbation, using the historical atmospheric $p\text{CO}_2$ time series. Figure 6 shows the column inventory of anthropogenic carbon in 1995 as predicted by the model. The total simulated inventory (in 1995) is 137.8 PgC. For comparison, using transient tracer data, Waugh *et al.* [2006] estimate a total inventory (in 1994) of 134 PgC. The cost function is defined as the total amount of anthropogenic carbon in the ocean in 1995, i.e., the change in DIC inventory between 1780 (preindustrial) and 1995,

$$J = \sum_i dv_i ([\text{DIC}]_i(1995) - [\text{DIC}]_i(1780)).$$

[36] Here dv_i is the volume of the i th grid box, and the summation is over all grid points. The control vector x is a vector of length equal to the number of surface grid points. Its elements are the gas transfer velocity at the corresponding surface grid points (see equation (5)). We wish to compute the gradient of J w.r.t. x . It is important to note that

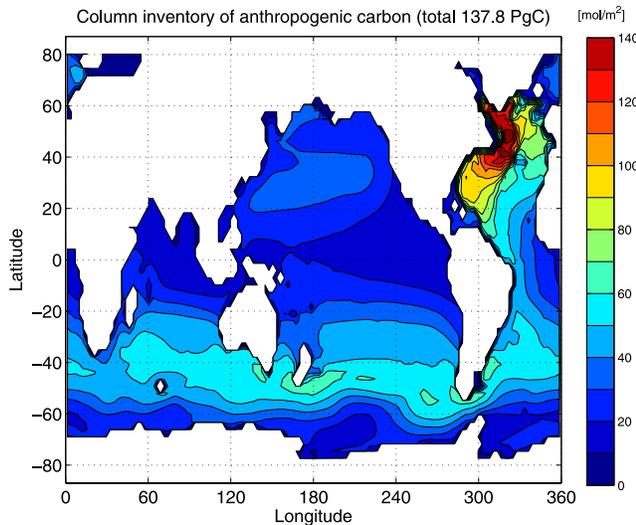


Figure 6. Column inventory of anthropogenic carbon in 1995 in the OCMIP-2 abiotic carbon model.

J depends on x in several ways. The first is through its explicit dependence on the preindustrial DIC which, in turn, is a function of the gas transfer velocity. The second is through its dependence on the DIC in 1995, which depends on the transfer velocity explicitly via the source/sink term and implicitly via the time stepping equation. Finally, the time stepping equation (equation (2)) uses the preindustrial DIC as the initial condition for the anthropogenic perturbation integration. The adjoint “propagates” these complicated nonlinear dependencies exactly and efficiently through, in essence, application of the chain rule.

[37] According to the recipe in Appendix A, we require the product of $(\partial \mathbf{q}^n / \partial [\text{DIC}]^n)^T$, $(\partial \mathbf{q}^n / \partial \mathbf{x})^T$, and $(\partial [\text{DIC}]^{1780} / \partial \mathbf{x})^T$ with a vector. The first two of these operations are readily constructed for this problem. The last one is defined implicitly via the nonlinear equation for the preindustrial steady state: $\mathbf{F}([\text{DIC}], \mathbf{x}) = \mathbf{0}$. While there are techniques for performing this operation directly via the steady state equation (see Appendix A), they are not easily generalizable to more complex situations, for example, with a seasonal cycle. Thus I avoid working with the steady state equation, and treat the problem as fully time-dependent. Starting with a uniform DIC concentration, I integrate the model forward for 4000 years to equilibrium, then apply the anthropogenic perturbation from year 1780 to 1995. The adjoint model is then run backward over this same (4000 + 215 year) period. Since the initial condition is independent of the gas exchange coefficient, it does not enter into the sensitivity calculation. The result of this procedure is shown in Figure 7. We see large positive sensitivities in the Labrador Sea and parts of the Southern Ocean. These are regions of the ocean where CO_2 is taken up by the ocean; increasing the local gas exchange coefficient leads to an increase in anthropogenic carbon uptake. In contrast, negative values are found over the tropics where CO_2 is outgassing into the atmosphere; a local increase in the exchange coefficient leads to a decrease in the carbon uptake.

4.4. Adjoint Sensitivity of a Coupled Ocean Biogeochemical Model

[38] As a final example, we consider a coupled biogeochemical model of carbon, phosphorus, and iron, based on that of *Dutkiewicz et al.* [2005]. The model includes 6 prognostic tracers: inorganic and organic forms of phosphorus, carbon, oxygen, alkalinity, and iron. As documented by *Dutkiewicz et al.* [2005, 2006], the model includes a simplified parameterization for net community production, B , which is regulated by the availability of light (photosynthetically active radiation), phosphate, and iron. Scavenging and complexation of iron is represented according to *Parekh et al.* [2005]. In addition, iron is supplied to the ocean via an aeolian dust source prescribed according to *Mahowald et al.* [2003].

[39] Some comments regarding the practical aspects of coupling a biogeochemistry module to the TMM may be useful here. The biogeochemistry model used here is available as a “package” within the MIT GCM. To couple it to the transport matrix code, I have written a “wrapper” routine to exchange data between the matrix driver code and the biogeochemical model. As mentioned previously, we only require a subroutine for computing the biogeochemical source/sink term for a single vertical column. To facilitate this, I set the number of grid points in the horizontal to 1. This is done by modifying the header file where the size of various three-dimensional tracer arrays are set (an arrangement typical of most GCMs). Thus, while the arrays are declared to be three-dimensional, they are effectively only one-dimensional (varying in z). This simple procedure, which should be applicable to other models, allows us to use the MIT biogeochemistry code without modification.

[40] In this final example, I illustrate the use of the matrix method by computing the sensitivity of net community production B to the input of iron from dust at the surface. This is an interesting question because in many parts of the ocean, particularly the so-called high nitrate, low chlorophyll regions, biological productivity is believed to

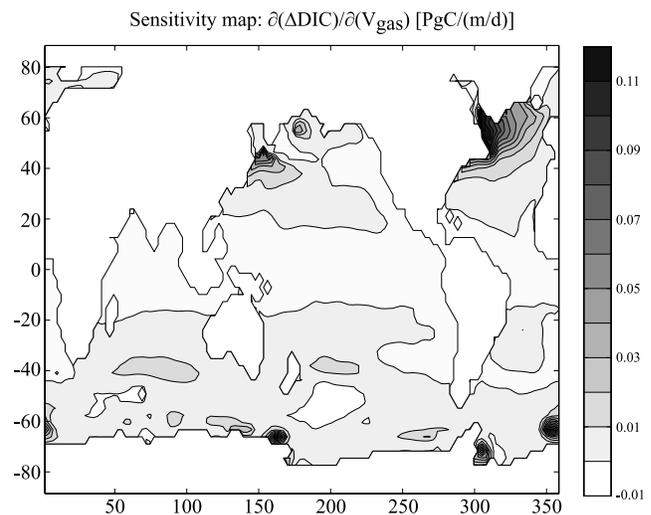


Figure 7. Map showing sensitivity of anthropogenic carbon inventory in 1995 to the gas transfer velocity.

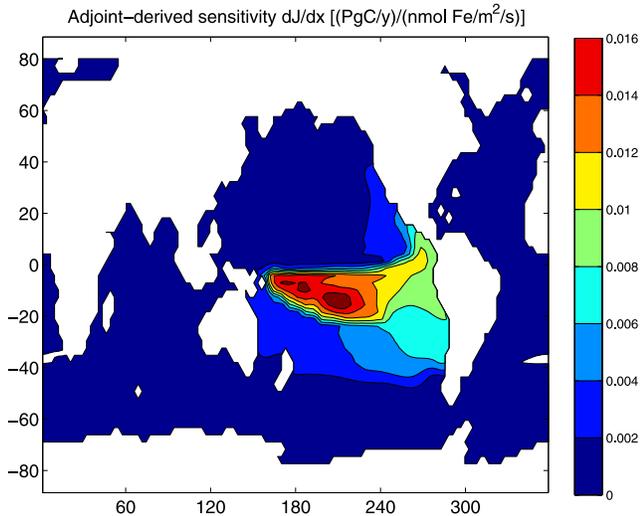


Figure 8. Map showing sensitivity of annually and globally integrated biological production to the surface flux of iron in January. Sensitivities are expressed in $[(\text{PgC yr}^{-1})/(\text{nmol Fe m}^{-2}\text{s}^{-1})]$ using a C:P Redfield ratio of 117.

be strongly modulated by the availability of iron. Additionally, iron fertilization has been hypothesized as a factor controlling atmospheric CO_2 on glacial-interglacial timescales. I use the adjoint method in conjunction with the seasonally varying 2.8° transport matrix. To generate the adjoint code, I use the publicly accessible TAMC software. An alternative is “OpenAD”, available from <http://www.openad.org>. The cost function is defined as the annually and globally integrated value of B in equilibrium:

$$J = \sum_n \sum_i B_i^n dv_i \Delta t,$$

where, n is the time index and i is the spatial index. The summation is over one year and all grid points. The control vector \mathbf{x} is the monthly mean aeolian flux of iron at the surface, i.e., \mathbf{x} is a vector of length $12 \times$ (number of surface grid points). The surface iron flux at any instant is computed by linear interpolation between the prescribed monthly mean values. This dependency is “visible” to the adjoint machinery, allowing us to compute the sensitivity of J w.r.t. the iron flux in each month. To compute the sensitivity vector $\partial J/\partial \mathbf{x}$, the forward model was integrated for 4000 years until it reached equilibrium. The cost function was evaluated in this equilibrium state. Subsequently, the adjoint model was integrated backward in time for 4000 years, resulting in the desired sensitivity. It should be noted that *Dutkiewicz et al.* [2006] have carried out a similar calculation with the full MIT GCM, but because of the large computational cost involved, the forward and backward “sweeps” were only 10 years long. Their results therefore do not represent equilibrium sensitivities. Figure 8 shows the sensitivity map for January. Sensitivity maps for the other months are very similar. The map shows that adding iron to the tropical Pacific produces the largest response in global biological production. In contrast, the

Southern Ocean shows very limited sensitivity because productivity there is primarily light limited [*Dutkiewicz et al.*, 2006]. Interestingly, the sensitivity patterns converge after ~ 100 years of backward integration. The results shown here are part of an ongoing study to improve the representation of iron in ocean biogeochemical models, and to address fundamental questions regarding the role of iron in the global carbon cycle.

5. Summary

[41] In this paper I have described the transport matrix method, a novel computational framework for simulation of chemical and biological tracers in the ocean. The method is based on representing the complex, three-dimensional transport of tracers in ocean general circulation models as a sparse matrix, reducing the problem of simulating tracers to a sequence of matrix-vector products. The efficiency and ease-of-use of the transport matrix method were illustrated by applying it to several different biogeochemical tracers. The method is particularly well suited to computations involving the adjoint of the tracer equations. It is hoped that the transport matrix framework will encourage and enable a wider community of researchers to exploit the tools of numerical modeling in studies of ocean chemistry and biology.

Appendix A: Adjoint Method

[42] In this section I provide a brief but self contained description of the adjoint method for efficient computation of sensitivities. To begin, I note that this technique is essentially an application of the chain rule. To see this, consider a fairly general and typical time stepping model for a variable \mathbf{u} ,

$$\mathbf{u}^0 = \mathbf{u}^{\text{initial}}(\mathbf{x})$$

for $n = 1$ to N do

$$\mathbf{u}^n = \mathbf{M}^{n-1}(\mathbf{u}^{n-1}, \mathbf{x})$$

end for

$$J = F(\mathbf{u}^N, \mathbf{u}^{N-1}, \mathbf{u}^{N-2}, \dots).$$

[43] Here J is the (scalar) cost function that depends on the models’ state at the final time step N , and (possibly) at previous time steps. For example, J could be the annual mean of \mathbf{u} averaged over some spatial domain. \mathbf{x} is a vector of parameters (the “control vector”) with respect to which we wish to find the sensitivity of J , i.e., $\nabla J \equiv \partial J/\partial \mathbf{x}$. \mathbf{M} is a general nonlinear function, for example, the RHS of equation (2), typically implemented in the form of computer code. Note that \mathbf{u} defines the complete state of the system. In the present context, it is the vector formed by concatenating all the prognostic tracers in the model. Finally, if the initial condition \mathbf{u}^0 represents an equilibrium solution of the biogeochemical model, then it too may depend on \mathbf{x} (see section 4.3). For clarity of presentation we take $J = F(\mathbf{u}^N)$.

[44] By the chain rule,

$$\left(\frac{\partial J}{\partial \mathbf{x}}\right)^T = \left(\frac{\partial F}{\partial \mathbf{u}^N}\right)^T \frac{\partial \mathbf{u}^N}{\partial \mathbf{x}}. \quad (\text{A1})$$

However, from the model time stepping equation we also have

$$\frac{\partial \mathbf{u}^N}{\partial \mathbf{x}} = \frac{\partial \mathbf{M}^{N-1}}{\partial \mathbf{u}^{N-1}} \frac{\partial \mathbf{u}^{N-1}}{\partial \mathbf{x}} + \frac{\partial \mathbf{M}^{N-1}}{\partial \mathbf{x}}. \quad (\text{A2})$$

Taking the transpose of the above equation and combining with equation (A1), we get

$$\begin{aligned} \frac{\partial J}{\partial \mathbf{x}} &= \left(\frac{\partial \mathbf{u}^N}{\partial \mathbf{x}} \right)^T \frac{\partial F}{\partial \mathbf{u}^N} \\ &= \left[\left(\frac{\partial \mathbf{u}^{N-1}}{\partial \mathbf{x}} \right)^T \left(\frac{\partial \mathbf{M}^{N-1}}{\partial \mathbf{u}^{N-1}} \right)^T + \left(\frac{\partial \mathbf{M}^{N-1}}{\partial \mathbf{x}} \right)^T \right] \frac{\partial F}{\partial \mathbf{u}^N}. \end{aligned} \quad (\text{A3})$$

Variables at time step $N - 1$ are in turn related to those at $N - 2$ by the time stepping equation, and so on. On the basis of this, a practical algorithm for computing ∇J is as follows:

[45] Initialize: $\mathbf{u}^\dagger = \frac{\partial F}{\partial \mathbf{u}^N}$, $\mathbf{x}^\dagger = \mathbf{0}$

for $n = N - 1$ to 0 do

$$\mathbf{x}^\dagger = \mathbf{x}^\dagger + \left(\frac{\partial \mathbf{M}^n}{\partial \mathbf{x}} \right)^T \mathbf{u}^\dagger$$

$$\mathbf{u}^\dagger = \left(\frac{\partial \mathbf{M}^n}{\partial \mathbf{u}^n} \right)^T \mathbf{u}^\dagger$$

end for

$$\nabla J = \mathbf{x}^\dagger + \left(\frac{\partial \mathbf{u}^0}{\partial \mathbf{x}} \right)^T \mathbf{u}^\dagger.$$

[46] The key feature of the adjoint method is that it replaces the difficult problem of computing large, dense Jacobian matrices with a sequence of Jacobian-transpose-vector products. As is evident from the above algorithm, this sequence of matrix-vector products is performed backward in time. Note that the various derivatives involved (“tangent adjoints”) must be evaluated about the time evolving state of the forward model. It is therefore necessary to store the entire trajectory of the forward model (known as “checkpointing”) before carrying out the adjoint run. To summarize, the complete algorithm involves an integration of the forward model up to the final time step, followed by a backward integration of the adjoint model.

[47] Each step of the above adjoint algorithm requires the operations $(\partial \mathbf{M} / \partial \mathbf{x})^T \mathbf{u}^\dagger$ and $(\partial \mathbf{M} / \partial \mathbf{u})^T \mathbf{u}^\dagger$. Automatic differentiation (AD) software (colloquially known as “adjoint compilers”) such as TAMC and TAF [Giering and Kaminski, 1998] take a description of $\mathbf{M}(\mathbf{u}, \mathbf{x})$ in the form of computer code and generate code that, given an input vector \mathbf{u}^\dagger , perform precisely these operations. (This is known as “reverse-mode” AD. “Forward-mode” AD produces code that computes the product of the Jacobian (rather than its transpose) with a vector.) In principle, AD software can be applied to the full GCM or an offline tracer model, as for instance has been done for the MIT GCM [Heimbach et al., 2002, 2005]. This is at best nontrivial. The transport matrix formulation, however, allows us to very easily construct the appropriate adjoint model. For a single time step and multiple tracers (represented by the vector \mathbf{u}), we have

$$\mathbf{M}^n(\mathbf{u}^n, \mathbf{x}) = \text{diag}(\mathbf{A}_i^n) [\text{diag}(\mathbf{A}_e^n) \mathbf{u}^n + \mathbf{q}^n(\mathbf{u}^n, \mathbf{x})]. \quad (\text{A4})$$

Here $\text{diag}(\mathbf{A}_e)$ is a block diagonal matrix with \mathbf{A}_e on the diagonal, and similarly for \mathbf{A}_i . (It is unnecessary to ever

form this matrix explicitly.) Differentiating and taking the transpose,

$$\left(\frac{\partial \mathbf{M}^n}{\partial \mathbf{u}^n} \right)^T = \text{diag}(\mathbf{A}_e^{nT} \mathbf{A}_i^{nT}) + \left(\frac{\partial \mathbf{q}^n}{\partial \mathbf{u}^n} \right)^T \text{diag}(\mathbf{A}_i^{nT})$$

$$\left(\frac{\partial \mathbf{M}^n}{\partial \mathbf{x}} \right)^T = \left(\frac{\partial \mathbf{q}^n}{\partial \mathbf{x}} \right)^T \text{diag}(\mathbf{A}_i^{nT}).$$

The key point to note is that the term involving tracer transport is easily dealt with. The only apparent complication is the biogeochemical source/sink term. However, recall that in most biogeochemical models, the “action” as it were takes place only in the vertical. That is, the source/sink term q_i at grid point i only depends on the tracer concentrations in the same vertical column. Consequently, $\partial \mathbf{q} / \partial \mathbf{u}$ has a sparse block structure, with each block corresponding to a vertical profile. Thus, to complete the solution, all we need is a subroutine that computes \mathbf{q} for a single (arbitrary) vertical profile, and the corresponding adjoint derivative code. The required adjoint-vector products can be assembled from these pieces. (This is done automatically by the PETSc-based implementation of the TMM (see Appendix B).)

[48] Finally, I briefly address the problem of computing (equilibrium) sensitivities for models in a steady state (section 4.4) or situations in which the initial condition of the time evolving system depends on the control variables (section 4.3). In either case, we require the product of $(\partial \mathbf{u}^0 / \partial \mathbf{x})^T$ with a known vector. I focus on seasonally varying problems as the purely steady problem can be solved in a straightforward manner by implicit differentiation of equation (4) [e.g., Giles, 2001; Kwon and Primeau, 2006]. As mentioned in section 4.3, the easiest (but not necessarily most efficient) way to deal with this problem is to treat the equilibrium problem in a time-dependent manner, i.e., include the spinup phase as part of the adjoint computation. This is the approach taken in the adjoint calculations shown above. A promising alternative scheme suggested by Christianson [1994] [see also Kaminski et al., 2005] is based on the recognition that the steady state problem, with or without seasonal cycle, can be posed as a fixed point iteration. Applying AD to it results in a potentially efficient adjoint which only uses the trajectory from the last iteration. It remains to be seen, however, whether this method is practicable for large-scale problems.

Appendix B: A PETSc-Based Implementation of the Transport Matrix Method

[49] As noted in section 1, the only requirement for using the TMM is the facility to perform matrix-vector products. However, as the range and size of problems to which the TMM is applied has grown, efficiency and ease of use have become a concern. To address this, I have developed a suite of computational routines that provide the necessary framework for simulating most chemical and biological tracers. The software is based on PETSc (“Portable, Extensible Toolkit for Scientific Computation”) [Balay et al., 2003], a freely available open source suite of data structures and routines for solution of large-scale linear and nonlinear

problems. PETSc is actively developed and maintained by computational scientists at Argonne National Laboratories. PETSc is also the underlying software for the NSF-funded software project, Computational Infrastructure for Geophysics (CIG).

[50] One of the key advantages of using PETSc as a building block for implementing the TMM is its support for state-of-the-art, distributed-memory routines for operating on vectors and sparse matrices, including implementations of many popular Krylov subspace solvers and preconditioners. PETSc makes it practical to prototype new mathematical and computational ideas relatively straightforwardly, while achieving high performance. Another useful aspect of PETSc is its portability. PETSc-based code will run without modification on most computers, sequential or parallel. The fact that PETSc was designed for use on parallel computers is another advantage, as it allows us to address large-scale problems. And by hiding the details of the parallelism from the end user (for example, users don't need to know anything about parallel programming), it enables us to write clean, compact code that automatically inherits the parallelism and is independent of the underlying hardware.

[51] The software I have developed is essentially a set of driver routines that implement equation (2) and its adjoint for any generic tracer problem. Its principal features are as follows: (1) Driver code provides “hooks” for user-written routines that compute biogeochemical source/sink terms. (2) Numerous utility routines, for example, for interpolating large vectors and matrices in time, are available. (3) No restrictions on problem size or number of tracers. These and most other parameters can be set at run time. (4) Facility (via data structures) for coupling the ocean model to other components (e.g., atmospheric model, sediment model). (5) Adjoint driver code with multilevel checkpointing of forward trajectory, and backward integration of adjoint model. (6) Routines for computing steady state solutions. (7) Extensive set of example code that can be modified to suit particular problems. The software, along with various transport matrices, can be obtained from the author.

[52] **Acknowledgments.** I am grateful to Stephanie Dutkiewicz and Mick Follows for sharing their expertise of the MIT biogeochemical model, and Patrick Heimbach for computing the ECCO matrix. I would like to thank to Tim Merlis, Carl Wunsch, and two anonymous referees for their comments and suggestions to improve the text. This work was funded by NSF grants OCE-0449703 and ATM-0233853. LDEO contribution 7026.

References

- Balay, S., W. D. Gropp, L. C. McInnes, and B. F. Smith (2003), *Petsc users manual*, *Tech. Rep. ANL-95/11-Rev.2.1.5*, Argonne Natl. Lab., Argonne, Ill.
- Christianson, B. (1994), Reverse accumulation and attractive fixed points, *Optim. Methods Software*, 3, 311–326.
- Curtis, A. R., M. J. D. Powell, and J. K. Reid (1974), On the estimation of sparse Jacobian matrices, *J. Inst. Math. Appl.*, 13, 117–119.
- Demmel, J. W., S. C. Eisenstat, J. R. Gilbert, X. S. Li, and J. W. H. Liu (1999), A supernodal approach to sparse partial pivoting, *SIAM J. Matrix Anal. Appl.*, 20(3), 720–755.
- Dutkiewicz, S., M. Follows, and P. Parekh (2005), Interactions of the iron and phosphorus cycles: A three-dimensional model study, *Global Biogeochem. Cycles*, 19, GB1021, doi:10.1029/2004GB002342.
- Dutkiewicz, S., M. J. Follows, P. Heimbach, and J. Marshall (2006), Controls on ocean productivity and air-sea carbon flux: An adjoint model sensitivity study, *Geophys. Res. Lett.*, 33, L02603, doi:10.1029/2005GL024987.
- Fennel, K., M. Losch, J. Schröter, and M. Wenzel (2001), Testing a marine ecosystem model: Sensitivity analysis and parameter optimization, *J. Mar. Syst.*, 28, 45–63.
- Fukumori, I., T. Lee, B. Cheng, and D. Menemenlis (2004), The origin, pathway, and destination of Nino-3 water estimated by a simulated passive tracer and its adjoint, *J. Phys. Oceanogr.*, 34, 582–604.
- Gebremedhin, A., F. Manne, and A. Pothen (2005), What color is your Jacobian? Graph coloring for computing derivatives, *SIAM Rev.*, 47, 629–705.
- Gent, P. R., and J. C. McWilliams (1990), Isopycnal mixing in ocean circulation models, *J. Phys. Oceanogr.*, 20, 150–155.
- Giering, R., and T. Kaminski (1998), Recipes for adjoint code generation, *Trans. Math. Software*, 24, 437–474.
- Giles, M. B. (2001), On the iterative solution of adjoint equations, in *Automatic Differentiation: From Simulation to Optimization*, edited by G. Corliss et al., pp. 145–152, Springer, New York.
- Goodman, P. J., W. Hazeleger, P. de Vries, and M. Cane (2005), Pathways into the Pacific Equatorial Undercurrent: A trajectory analysis, *J. Phys. Oceanogr.*, 35, 2134–2151.
- Gu, D., and S. Philander (1997), Interdecadal climate fluctuations that depend on exchanges between the tropics and extratropics, *Science*, 275, 805–807.
- Haine, T. W. M., and T. M. Hall (2002), A generalized transport theory: Water-mass composition and age, *J. Phys. Oceanogr.*, 32, 1932–1946.
- Hall, T. M., D. W. Waugh, T. W. N. Haine, P. E. Robbins, and S. Khatiwala (2004), Estimates of anthropogenic carbon in the Indian Ocean with allowance for mixing and time-varying air-sea CO₂ disequilibrium, *Global Biogeochem. Cycles*, 18, GB1031, doi:10.1029/2003GB002120.
- Heimbach, P., C. Hill, and R. Giering (2002), Automatic generation of efficient adjoint code for a parallel Navier-Stokes solver, in *Proceedings of the International Conference on Computational Science—ICCS 2002*, edited by J. J. Dongarra, P. M. A. Sloot, and C. J. K. Tan, pp. 1019–1028, Springer, New York.
- 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. Comput. Syst.*, 21, 1356–1371.
- Holzer, M., and T. M. Hall (2000), Transit-time and tracer-age distributions in geophysical flows, *J. Atmos. Sci.*, 57, 3539–3558.
- Holzer, M., and F. W. Primeau (2006), The diffusive ocean conveyor, *Geophys. Res. Lett.*, 33, L14618, doi:10.1029/2006GL026232.
- Hundsdoerfer, W., and J. G. Verwer (2003), *Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations*, 500 pp., Springer, New York.
- Kaminski, T., and M. Heimann (2001), Inverse modeling of atmospheric carbon dioxide fluxes, *Science*, 294, 259.
- Kaminski, T., R. Giering, and M. Voßbeck (2005), Efficient sensitivities for the spin-up phase, in *Automatic Differentiation: Applications, Theory, and Implementations, Lecture Notes Comput. Sci. Eng.*, vol. 50, edited by H. M. Bücker et al., pp. 283–291, Springer, New York.
- Kelley, C. T. (1995), *Iterative Methods for Linear and Nonlinear Equations*, *Frontiers Appl. Math.*, vol. 16, 166 pp., Soc. for Ind. and Appl. Math., Philadelphia, Pa.
- Kelley, C. T. (2003), *Solving Nonlinear Equations with Newton's Method*, *Fund. Algorithms*, vol. 1, 116 pp., Soc. for Ind. and Appl. Math., Philadelphia, Pa.
- Khatiwala, S., M. Visbeck, and P. Schlosser (2001), Age tracers in an ocean GCM, *Deep Sea Res., Part I*, 48, 1423–1441.
- Khatiwala, S., M. Visbeck, and M. Cane (2005), Accelerated simulation of passive tracers in ocean circulation models, *Ocean Modell.* 9, pp. 51–69, Hooke Inst. Oxford Univ., Oxford, U.K.
- Knoll, D. A., and D. E. Keyes (2004), Jacobian-free Newton-Krylov methods: A survey of approaches and applications, *J. Comput. Phys.*, 193, 357–397.
- Kwon, E. Y., and F. Primeau (2006), Sensitivity and optimization study of a biogeochemistry ocean model using an implicit solver and in-situ phosphate data, *Global Biogeochem. Cycles*, 20, GB4009, doi:10.1029/2005GB002631.
- Large, W. G., J. C. McWilliams, and S. C. Doney (1994), Oceanic vertical mixing: A review and a model with a nonlocal boundary layer parameterization, *Rev. Geophys.*, 32, 363–403.
- Levitus, S., et al. (1998), *World Ocean Database 1998*, *NOAA Atlas NESDIS 18*, NOAA, Silver Spring, Md.
- Li, X. S., and J. W. Demmel (2003), SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems, *Trans. Math. Software*, 29, 110–140.

- Li, X., and C. Wunsch (2004), An adjoint sensitivity study of chlorofluorocarbons in the North Atlantic, *J. Geophys. Res.*, *109*, C01007, doi:10.1029/2003JC002014.
- Mahowald, N., C. Lou, J. del Corral, and C. Zender (2003), Interannual variability in atmospheric mineral aerosols from a 22-year model simulation and observational data, *J. Geophys. Res.*, *108*(D12), 4352, doi:10.1029/2002JD002821.
- Maiss, M., and C. A. M. Brenninkmeijer (1998), Atmospheric SF₆: Trends, sources, and prospects, *Environ. Sci. Technol.*, *32*, 3077–3086.
- Marshall, J., A. Adcroft, C. Hill, L. Perelman, and C. Heisey (1997), A finite-volume, incompressible navier-stokes model for studies of the ocean on parallel computers, *J. Geophys. Res.*, *102*, 5733–5752.
- Orr, J. C., J. Dutay, R. G. Najjar, J. Bullister, and P. Brockmann (1999a), CFC-HOWTO: Internal OCMIP Report, technical report, Laboratoire des Sciences du Climat et l'Environnement, CEA Saclay, Gif-sur-Yvette, France.
- Orr, J. C., R. G. Najjar, C. L. Sabine, and F. Joos (1999b), Abiotic-HOWTO: Internal OCMIP Report, technical report, Laboratoire des Sciences du Climat et l'Environnement, CEA Saclay, Gif-sur-Yvette, France.
- Parekh, P., M. J. Follows, and E. A. Boyle (2005), Decoupling of iron and phosphate in the global ocean, *Global Biogeochem. Cycles*, *19*, GB2020, doi:10.1029/2004GB002280.
- Primeau, F. (2005), Characterizing transport between the surface mixed layer and the ocean interior with a forward and adjoint global ocean transport model, *J. Phys. Oceanogr.*, *35*, doi:10.1175/JPO2699.1.
- Rutberg, R. L., and S. L. Peacock (2006), High-latitude forcing of interior $\delta^{13}\text{C}$, *Paleoceanography*, *21*, PA2012, doi:10.1029/2005PA001226.
- Saad, Y. (2003), *Iterative Methods for Sparse Linear Systems*, *Class. Appl. Math.*, vol. 16, 528 pp., Soc. for Ind. and Appl. Math., Philadelphia, Pa.
- Schlitzer, R. (1999), Applying the adjoint method for biogeochemical modeling: Export of particulate organic matter in the world ocean, in *Inverse Methods in Global Biogeochemical Cycles*, *Geophys. Monogr. Ser.*, vol. 114, edited by P. Kasibhatla et al., pp. 107–124, AGU, Washington, D. C.
- Stammer, D., K. Ueyoshi, A. Köhl, 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.
- Thiele, G., and J. L. Sarmiento (1990), Tracer dating and ocean ventilation, *J. Geophys. Res.*, *95*, 9377–9391.
- Toggweiler, J. R., K. Dixon, and K. Bryan (1989), Simulations of radiocarbon in a coarse resolution world ocean model: 1. Steady state prebomb distributions, *J. Geophys. Res.*, *94*, 8217–8242.
- Tsuchiya, M. (1968), Upper waters of the intertropical Pacific Ocean, technical report, Johns Hopkins Univ., Baltimore, Md.
- Wanninkhof, R. (1992), Relationship between wind speed and gas exchange over the ocean, *J. Geophys. Res.*, *97*, 7373–7382.
- Waugh, D. W., T. M. Hall, and T. W. N. Haine (2003), Relationships among tracer ages, *J. Geophys. Res.*, *108*(C5), 3138, doi:10.1029/2002JC001325.
- Waugh, D. W., T. W. N. Haine, and T. M. Hall (2004), Transport times and anthropogenic carbon in the subpolar North Atlantic Ocean, *Deep Sea Res., Part I*, *15*, 1475–1491, doi:10.1016/j.dsr.2004.06.011.
- Waugh, D. W., T. M. Hall, B. I. McNeill, R. Key, and R. J. Matear (2006), Anthropogenic CO₂ in the oceans estimated using transit time distributions, *Tellus, Ser. B*, *58*, 376–389.
- Wunsch, C., and P. Heimbach (2007), Practical global oceanic state estimation, *Physica D*, *230*, 197–208.

S. Khatiwala, Lamont-Doherty Earth Observatory, Columbia University, Palisades, NY 10964, USA. (spk@ldeo.columbia.edu)