



Accelerated simulation of passive tracers in ocean circulation models

Samar Khatiwala ^{a,*}, Martin Visbeck ^a, Mark A. Cane ^a

^a *Division of Ocean and Climate Physics, Oceanography 201B, Rte 9W, Lamont-Doherty Earth Observatory, Columbia University, Palisades, NY 10964, USA*

Received 3 October 2003; received in revised form 10 March 2004; accepted 14 April 2004

Available online 5 May 2004

Abstract

A novel strategy is proposed for the efficient simulation of geochemical tracers in ocean models. The method captures the tracer advection and diffusion in a general circulation model (GCM) without any alteration (or even knowledge) of the GCM code. In comparison with offline tracer models, the proposed method is considerably more efficient and automatically includes all parameterizations of unresolved processes present in the most sophisticated GCMs. A comparison with a global configuration of the MIT GCM shows that the scheme can capture the complex three-dimensional transport of a state-of-the-art GCM. A key advantage of the proposed technique is the ability to directly compute steady-state solutions, a facility particularly well-suited to tracers such as natural radiocarbon. This capability is applied to develop a novel algorithm for accelerating the dynamical adjustment of ocean models.

© 2004 Elsevier Ltd. All rights reserved.

1. Introduction

Geochemical tracers such as ¹⁴C (radiocarbon) and chlorofluorocarbons (CFCs) have significantly contributed to our understanding of ocean circulation and climate. Many important aspects of the ocean's circulation, such as the poleward transport of heat, water mass transformation and ventilation, and the uptake of anthropogenic CO₂, are controlled by dynamical

* Corresponding author. Tel.: +1-854-365-8454; fax: +1-845-365-8736.
E-mail address: spk@ldeo.columbia.edu (S. Khatiwala).

processes (e.g., mixing and convection) that are difficult to directly observe. Our knowledge of these processes, and their impact on the large scale circulation must be indirectly inferred from observations of chemical tracer distributions. Tracers have proved to be particularly powerful tools when combined with ocean general circulation models (GCMs). (See England and Maier-Reimer (2001) for a review.) However, the ability of oceanographers to exploit the tools of numerical models and chemical tracers through long simulations on climatic timescales is severely limited by available computational resources. Long integrations are a necessity. Tracers such as natural radiocarbon take several thousand years to approach a steady-state, a problem exacerbated by the fact that the dynamical adjustment timescale, especially for the deep ocean, is also hundreds to thousands of years. Consequently, such simulations at sufficient resolution remain out of the reach of all but the fastest supercomputers.

Both problems have been long recognized. For passive tracers, one solution is to use an “offline” model in which tracers are advected and diffused by a flow field archived from a previous GCM run. While this allows the use of somewhat longer time steps, the computational expense is still very large. More significantly, offline models have difficulty in accurately representing vertical fluxes due to deep convection, and can also potentially distort the seasonal cycle (Ribbe and Tomczak, 1997; England and Maier-Reimer, 2001). Since convection at high latitudes is the primary mechanism by which many climatically important tracers penetrate into the ocean interior, this is a fundamental limitation of offline models. Another solution (Aumont et al., 1998), developed for steady-state rather than transient tracers, is to initialize the model with a state closer to the equilibrium solution. The initial condition is produced from integrating a coarse-resolution offline model. This method is, however, only appropriate for steady state tracers and cannot be used to study, say, CFCs.

A number of ad hoc acceleration techniques have been proposed for the dynamical adjustment problem. One such method uses unequal time steps in the tracer and momentum equations (Bryan and Lewis, 1979; Bryan, 1984). While this use of “distorted physics” can considerably speed up the approach to steady-state, it still requires long integrations, as the equilibration of temperature and salinity in the deep ocean is rate limited by the slow along and across isopycnal diffusion of tracers and can take several thousand years (Danabasoglu et al., 1996). Additionally, there is also the potential for distorting the seasonal cycle, although the issue is still under debate (e.g., Wang, 2001; Huang and Pedlosky, 2002).

Motivated by the potential range of applications to climate and tracer modeling, we have developed a novel technique, called the “matrix method”, for the accurate and efficient simulation of passive tracers. Compared with offline tracer models, this method is both significantly more efficient and automatically includes all parameterizations of physical processes such as convection and eddy-induced transport found in state-of-the-art ocean GCMs. A key advantage of the technique over more traditional approaches is the facility to directly compute steady state solutions, without simulating the transient. Furthermore, the scheme provides the adjoint tracer transport model at no additional computational cost. Here, we describe the matrix method and present results from several numerical experiments to illustrate its application to tracer problems, and to show that the technique can successfully capture the complex three-dimensional transport of a state-of-the-art ocean general circulation model. As an important application of the steady-state capability of the scheme, we also present a new algorithm for accelerated spin-up of ocean GCMs.

2. A matrix approach to tracer transport

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. Our approach is based on the recognition that for a passive tracer, the discretized advection-diffusion equation can be written as a linear matrix equation:

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

where \mathbf{c} is the vector of tracer concentrations at the grid points of the GCM. The matrix $\mathbf{A}(t)$ is the “transport matrix” which results from discretization of the advection-diffusion operator and includes the effects of advection, diffusion, subgrid scale (SGS) processes, and surface boundary conditions (BCs). Note that \mathbf{c} is merely a vector representation of a discretized three-dimensional tracer field. The influence of boundary conditions can be seen explicitly by splitting \mathbf{A} into an “interior” matrix \mathbf{A}_I and a “boundary” matrix \mathbf{B} :

$$\frac{d\mathbf{c}}{dt} = \mathbf{A}_I(t)\mathbf{c} + \mathbf{B}(t)\mathbf{c}_b(t), \quad (2)$$

where \mathbf{c} is now the vector of interior tracer concentrations. \mathbf{B} is the operator which transports time-dependent surface boundary conditions $\mathbf{c}_b(t)$ into the interior. These equations are readily generalized to include time-dependent, linear (as for radioactive decay) or nonlinear (as for biological tracers) sources and sinks, as well as flux BCs.

The representation Eq. (1) (and Eq. (2)) of the advection-diffusion equation is not new. Nevertheless, it is useful to illustrate the idea of representing tracer transport as a matrix equation through a simple example. The one-dimensional diffusion equation,

$$\frac{\partial C}{\partial t} = \kappa \frac{\partial^2 C}{\partial x^2} \quad 0 < x < L$$

with boundary conditions

$$C(x=0, t) = a(t) \quad \text{and} \quad C(x=L, t) = b(t),$$

is sufficiently simple that the matrices \mathbf{A}_I and \mathbf{B} can be derived explicitly. Discretizing in space with a centered scheme we obtain,

$$\frac{dC_i}{dt} = \frac{\kappa}{(\Delta x)^2} (C_{i+1} - 2C_i + C_{i-1}), \quad i = 1, \dots, N,$$

$$C_0 = a(t) \quad \text{and} \quad C_{N+1} = b(t),$$

where Δx is the grid spacing, and C_i is the tracer concentration at $x = i\Delta x$, $i = 0, \dots, N + 1$. In matrix form, we then have

$$\frac{d}{dt} \begin{bmatrix} C_1 \\ C_2 \\ \vdots \\ C_{N-1} \\ C_N \end{bmatrix} = \frac{\kappa}{(\Delta x)^2} \begin{bmatrix} -2 & 1 & & & \\ 1 & -2 & 1 & & \\ & & \ddots & & \\ & & & 1 & -2 & 1 \\ & & & & 1 & -2 \end{bmatrix} \begin{bmatrix} C_1 \\ C_2 \\ \vdots \\ C_{N-1} \\ C_N \end{bmatrix} + \begin{bmatrix} \kappa/(\Delta x)^2 & 0 \\ 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \\ 0 & \kappa/(\Delta x)^2 \end{bmatrix} \begin{bmatrix} C_0 \\ C_{N+1} \end{bmatrix}$$

which, with $\mathbf{c} \equiv (C_1 \dots C_N)^T$ and $\mathbf{c}_b \equiv (C_0 C_{N+1})^T$, is in the form of Eq. (2) with the matrices \mathbf{A}_1 and \mathbf{B} explicitly written out. While this is a particularly simple example, 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.

3. An empirical scheme for obtaining \mathbf{A}

Clearly, it would be quite time consuming to directly code the advection-diffusion equation, including various parameterizations, for an arbitrary ocean geometry in the form of Eq. (1). Recognizing this, we instead take an *empirical* approach to the problem: we will estimate the elements of \mathbf{A} by utilizing a GCM. The underlying idea exploits the following fact: \mathbf{A} times a unit vector Δ_j (a vector with “1” in the j th row, and “0” in all other rows) equals the j th column of \mathbf{A} . This suggests a recipe for computing \mathbf{A} :

- In a GCM, initialize the tracer field to a “unit vector” Δ_j , i.e., “1” at a single grid point, and “0” elsewhere.
- Take one time step, and evaluate the finite difference tendency.
- Reinitialize the tracer field to Δ_j , integrate the GCM one more time step, and recompute the tendency.

This procedure is repeated for a specified number of steps, and the tendencies averaged (see below). Finally, we equate the average tendency to the j th column of \mathbf{A} . By repeating this procedure for other unit vectors, all columns of \mathbf{A} can be computed.

Note that the tracer field is reset after each time step. This reinitialization distinguishes our method from, say, that of Stammer and Wunsch (1996), who computed the response to Gaussian (rather than point) perturbations. Mathematically, their approach was directed at obtaining empirical (dynamical) Green’s functions, in essence, $\exp \mathbf{A}t$, while we are interested in \mathbf{A} itself. (Because \mathbf{A} is extremely sparse while $\exp \mathbf{A}t$ is a full matrix, there are advantages to working with \mathbf{A} .)

The averaging of the tendency requires further comment. If the GCM transport is time-independent, a single time step is sufficient to estimate \mathbf{A} . In practice, however, this is seldom the case and averaging the tendency in time is crucial, particularly for numerical stability. (The short time scales associated with convective mixing may potentially make the matrix stiff.) The appropriate averaging period will be determined by the variability of the transport and the timescale of interest. For instance, if the transport has a seasonal cycle, an averaging period of 1 month may be acceptable. It remains to be seen whether the transport due to a turbulent flow (as in a high resolution model) can be accurately captured by our approach.

The naive approach described above has two principle drawbacks, which we address immediately below. First, the method is not very practical as each column of \mathbf{A} , i.e., each grid box, requires a separate tracer simulation. Second, the unit vectors used to assemble the matrix may not be the optimum basis from a numerical standpoint.

3.1. A more efficient approach

Physically, the sparseness is due to the fact that in a single time step tracer can only spread a finite distance, i.e., 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.) To exploit this limited “connectivity”, we subdivide the domain into a nonoverlapping set of “tiles”. Each tile is composed of a grid cell and an associated “halo” of neighboring grid cells to which tracer can spread in a single time step. (The halo can extend beyond the nearest neighbor as it typically does in the vertical.) Instead of initializing the tracer field to a unit vector, we set it to a checkerboard pattern of 1’s at the center of each tile and 0’s elsewhere. In a single time step, tracer concentrations within each tile evolve independently of those within other tiles. Thus the tiles within a set can be treated as a single “tracer”, and we may compute a large number of columns of \mathbf{A} simultaneously. The process must be repeated for each point within a tile, so the number of tracers required is on the order of the (maximum) number of points within a tile (see Appendix B). The process of “tiling” an irregular 3-d domain is readily automated.

The improvement in efficiency is dramatic and increases with the size of the problem. For example, a 2.8° global model with $\approx 53,000$ grid points requires 53,000 independent tracers with the first approach. The number of independent tracers required in the second approach is roughly $10 \times$ (number of vertical levels), here, 200. This number does not increase as the horizontal resolution increases. Being independent, they may be integrated in parallel (see below). A further gain in efficiency, without loss of accuracy, can be obtained by limiting the halo to nearest neighbors in the vertical in regions where deep convection is known to be absent.

3.2. Smoother basis vectors

The naive method proposed above requires calculating the response (tendency) to a “delta” initial condition (IC) so that each tendency directly gives a column of \mathbf{A} . Unit vectors, however, are also highly discontinuous, which leads to numerical inaccuracies. Linear advection schemes generally result in overshoots in the face of strong tracer gradients; delta initial conditions obviously create very strong gradients. Nonlinear advection methods (“flux limiters”) are numerically more stable, but generally tend toward an “upwind” scheme for strong tracer gradients. (While the advection-diffusion operator is always linear for a passive tracer, it is often implemented using nonlinear methods.) Upwind schemes are notorious for their numerical diffusivity, and as a consequence the matrix calculated from unit vectors will be too diffusive.

A solution to this problem is to use smoother basis vectors (initial conditions). The task of assembling \mathbf{A} from GCM tendencies is no longer as straightforward since the “response” at any grid point is now a combination of the “impulse” responses due to the initial tracer concentration at many different grid points. To obtain \mathbf{A} , these responses must be deconvolved. Moreover, the basis vectors must allow the computational efficiency of the method to be maintained. As a first step in investigating the feasibility of carrying out this complex procedure, we have modified the basic algorithm to apply ICs which are smooth in the vertical only. Arguably, since transport in the vertical plays a key role in ocean ventilation, any refinement of our method must first be

directed at improving this aspect of the calculation. (The general problem of a 3-d smooth basis is rather more difficult, and one we are currently addressing.)

The modified algorithm is as follows. We subdivide the domain into nonoverlapping tiles which extend over the entire water column. At the center of each tile, we apply a nonzero IC that varies smoothly in the vertical at that horizontal position, but is zero elsewhere in the tile (see Fig. 1). Since the IC is spread out over many grid points, there is no longer a one-to-one correspondence between a tendency vector and a column of \mathbf{A} . However, the tiles are still decoupled from each other, and we can associate a “local” transport matrix $\mathbf{A}_{\text{local}}$ with each tile, characterizing the interactions between the grid cells of that tile. $\mathbf{A}_{\text{local}}$ is of size $n \times nz$, where n is the total number of grid points in the tile, and nz the number of points in the vertical. Focus now on a particular tile. Denoting the IC in the vertical by ϕ_i , the GCM is integrated as before, and we equate the tendency vector \mathbf{T}_i to $\mathbf{A}_{\text{local}}\phi_i$. We repeat these steps for nz linearly independent ICs (vectors), resulting in nz tendency vectors. Defining matrices \mathbf{T} and Φ with \mathbf{T}_i and ϕ_i as their i th columns, respectively, the above procedure may be written as: $\mathbf{T} = \mathbf{A}_{\text{local}}\Phi$. Since the ϕ_i (the columns of Φ) have been chosen to be linearly independent, Φ^{-1} is guaranteed to exist. We then have: $\mathbf{A}_{\text{local}} = \mathbf{T}\Phi^{-1}$. The elements of $\mathbf{A}_{\text{local}}$ go directly into \mathbf{A} . (Note that Φ^{-1} may be pre-computed.)

This new level of complexity in the algorithm can be handled offline. Thus, once the GCM has been modified to reinitialize the tracer field and compute tracer tendencies at each time step, no further changes to the GCM code are required. (See Appendix A for a discussion of the technical issues involved in implementing the matrix algorithm.) Note too, that the use of smoother basis functions does not increase the number of GCM integrations.

3.3. Important aspects of the matrix algorithm

- The circulation embedded in \mathbf{A} , by construction, satisfies the (finite difference) equations of motion and automatically incorporates transport due to all parameterized subgrid scale processes (eddy-induced mixing, convective adjustment, etc.) represented in the GCM.
- The resolution at which we wish to simulate tracers is flexible since there need not be one-to-one correspondence between the elements of \mathbf{c} and the GCM grid cells. The vector \mathbf{c} can be taken as a spatially averaged version of the GCM tracer field. (\mathbf{A} is then computed by taking an




Fig. 1. Schematic to illustrate matrix algorithm. Shown (left) is a vertical column of the GCM (gray) surrounded by a “halo” of grid cells. On the right is a typical basis vector (ϕ), i.e., initial condition, applied to the central column of the tile.

appropriate weighted average of the GCM tendencies.) In the extreme case in which \mathbf{c} has been reduced to a few elements, Eq. (1) would resemble a traditional box model, with the added advantage of dynamically consistent fluxes between “boxes”. This “coarse graining” considerably reduces the computational task of estimating \mathbf{A} . Note that, it is tracer fluxes that are averaged onto a coarser grid; the GCM itself is still run at full resolution. Arguably, this may be more accurate than integrating the GCM at the (lower) resolution at which the matrix is required.

- The transport matrix is extremely sparse due to the finite speed of advection and diffusion. For example, an $O(53000 \times 53000)$ transport matrix for the global ocean (at 2.8° resolution) has fewer than 0.03% of its elements nonzero. Due to this sparseness, storage and computational requirements are drastically reduced.

- The algorithm is trivially parallelizable. Since each “tracer” can be integrated independently of other tracers, they can be simulated in parallel on different processors. With the widespread availability of “Beowulf” clusters \mathbf{A} can be computed very efficiently in a reasonable amount of time. For example, on a 20 processor, commodity Beowulf cluster it takes ≈ 1 h to compute a seasonally varying transport matrix for the global ocean at 2.8° resolution.

- Variability (seasonality) in the circulation can be readily accounted for. For instance, the GCM could be integrated for a full year, and the average tendency for each month computed, resulting in 12 different \mathbf{A} 's. The optimum averaging period will be problem specific.

- There are a number of constraints on the elements of the transport matrix. First, the elements of each row of \mathbf{A} sum to zero. This follows from the fact that a uniform tracer field will not change in time. This condition implies that \mathbf{A} has a nonzero nullspace and is therefore not invertible. The interior matrix \mathbf{A}_I , however, is of full rank and its inverse does exist. Thus, with appropriate boundary conditions, the forward tracer problem is well-posed with a unique steady-state solution (if one exists, as it will if \mathbf{A} and the BC are time-independent). Second, the diagonal elements of \mathbf{A} are all < 0 (strictly, ≤ 0 , but diffusion will ensure that the diagonal elements are negative), while the off-diagonal elements are ≥ 0 . This is obvious from the way \mathbf{A} is constructed. When coupled with the Gerschgorin circle theorem, these two conditions imply that all the eigenvalues of \mathbf{A} have a negative real part. Thus \mathbf{A} is a stable matrix. Finally, the (volume) weighted sum of the elements of every column of \mathbf{A} is zero. This follows from mass conservation.

- The technique can be readily implemented on any GCM with only minor modifications to the source code. Much of the complexity of the algorithm is in the pre- and post-processing of GCM input and output data, and is independent of the GCM, so it does not require a new code for each new GCM.

- The matrix formulation reduces the problem of tracer simulation to solving a system of coupled ordinary differential equations (ODEs). (Note that once the transport matrix has been computed, the GCM can be dispensed with.) These may be solved either analytically or by numerical integration. Indeed, if \mathbf{A} is time-independent, i.e., the transport is stationary, then the exact analytical solution to Eq. (1) with time-independent source (or BC) \mathbf{q} is

$$\mathbf{c}(t) = e^{\mathbf{A}t}\mathbf{c}(0) + \int_0^t e^{\mathbf{A}(t-t')}\mathbf{q} dt'. \quad (3)$$

(Recall, that $\mathbf{G}(t; t') = e^{\mathbf{A}(t-t')}$ is the Green's function for this problem.) This solution can be very efficiently evaluated by fast matrix exponential (Krylov) methods (Sidje, 1998). These “matrix-free” algorithms, which do not require the explicit computation of $e^{\mathbf{A}t}$ (a formidable task), are not

memory intensive and are well-suited to large, sparse problems. The solutions shown below were computed using this method. (Fortran and MATLAB implementations of the fast matrix exponential algorithm are available in the ΕΡΟΚΙΤ package (Sidje, 1998).)

In the general case with time-dependent \mathbf{A} and \mathbf{c}_b , and arbitrary sources and sinks, an array of efficient solvers for initial value problems can be brought to bear on the problem (e.g., Press et al., 1992). The matrix method is especially well-suited to widely available software packages with built-in support for integration of ODEs and sparse matrices: MATLAB, for example. Thus, even large systems of equations (corresponding to high spatial resolution) can be rapidly integrated out to many thousands of years on a desktop computer.

Finally, we note that if \mathbf{A} is time-independent, or can be treated as such (say, an annual mean) then steady-state solutions can also be trivially computed by solving a system of linear equations of the form $\mathbf{A}\mathbf{c}_{\text{steady}} = \mathbf{b}$. This allows us to directly compute equilibrium fields of tracers such as radiocarbon, circumventing the need to perform very long GCM integrations.

- *Adjoint tracer transport:* The adjoint to the tracer equation is required in many applications (e.g., Holzer and Hall, 2000). 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.) However, when written in matrix form, the adjoint is simply the transpose of \mathbf{A} (i.e., \mathbf{A}^T) and thus trivial to compute. (The adjoint equation to be integrated is now $-\mathbf{dc}/\mathbf{dt} = \mathbf{A}^T\mathbf{c} + \dots$).

4. Numerical results

To demonstrate the feasibility and usefulness of the matrix method, we have implemented the scheme in the MIT GCM, a state-of-the-art primitive equation model (Marshall et al., 1997). The MIT model features a variety of parameterizations to represent unresolved processes, such as isopycnal thickness diffusion (Gent and McWilliams, 1990) (GM), convective adjustment, and the K-profile parametrization (KPP) of Large et al. (1994). In the experiments discussed below, both GM and KPP are used. Unless stated otherwise, passive tracers are advected with the “superbee” second order nonlinear flux limiting scheme (Roe, 1985). An operator splitting method is used to improve performance in multidimensions. Active tracers are advected with a 2d order centered scheme, which requires a two level time stepping algorithm, here Adams–Bashforth, for stability. Below, we show results from two sets of experiments. In both experiments, the transport matrix was computed using the smooth basis algorithm of Section 3.2. (The basis vectors have the form $\phi_i \sim \exp(-|1 - z/z_i|)$.)

4.1. Ocean sector

As a test case, we consider a simple ocean sector, 1000 m deep, and extending 40° zonally, and 60° meridionally. The resolution is 2° in the horizontal, and 100 m in the vertical. The model was forced at the surface by an idealized cosine wind stress, and by warming (cooling) at low (high) latitudes (by a relaxation BC on surface temperature). The forcing produces a gyre circulation in the horizontal and a meridional overturning cell driven by convection at high latitude. The cir-

ulation was spun-up to a steady-state, and the matrix computed from a 1 month average of the tendencies.

4.1.1. Initial value problem

We first compare the GCM and matrix solutions to an initial value problem. The tracer was initialized with a Gaussian concentration in the surface layer (Fig. 2) and zero in other layers. The GCM was then integrated for 30 years, and the corresponding matrix solution obtained by evaluating Eq. (3) with the fast matrix exponential algorithm. Since the initial condition is both localized in the horizontal and discontinuous in the vertical, the problem is numerically quite challenging, requiring a nonlinear flux limiting advection scheme in the GCM. Thus, this test of the matrix method is quite stringent.

The matrix solution is somewhat more diffuse than in the GCM as expected, since the basis vectors are discontinuous in the horizontal. Nonetheless, differences between the matrix and GCM solution are a few percent at most. The matrix method has largely reproduced the complex




Fig. 2. Passive tracer simulations to verify matrix solution. Panels on the left display the GCM solution, while those on the right show the difference between the matrix and GCM solution. Top three panels compare the solutions to an initial value problem. Shown is the surface concentration at time $t = 0, 5,$ and 30 years. Bottom panels compare the mean age at 450 m.

3-d transport of the GCM. This is confirmed by simulations of a more conventional passive tracer, the steady-state ideal age or mean age, τ_m (e.g., Thiele and Sarmiento, 1990).

The mean age is a frequently computed 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. Note that τ_m is the solution of

$$\mathbf{A}_I \tau_m = -\mathbf{1}, \quad (4)$$

where $\mathbf{1}$ is a vector of “1”s. With our method it may be computed directly. As seen in Fig. 2 (bottom panels) the mean age distributions in the GCM and matrix are in good agreement (the fields here are smoother than the initial value problem). A more detailed discussion of the errors appears below. For now, we simply note that with a linear advection scheme the matrix and GCM give identical results (not shown), since the matrix algorithm captures exactly the GCMs transport.

4.1.2. Steady-state circulation and temperature

The matrix also allows us to directly compute steady-state solutions for active tracers such as T and S . In addition, the symmetric and anti-symmetric parts of \mathbf{A} contain information about the “advective” and “diffusive” components, respectively, of the transport. Thus, given \mathbf{A} , we can deduce velocity and diffusivity fields. Fig. 3 compares the velocity and steady-state temperature at two levels in the GCM and from the matrix. Note that for the GCM, what is shown is the *Eulerian* velocity, while the matrix solution provides an *effective* velocity. (The effective velocity is the sum of the Eulerian mean and the “eddy-induced” components. In our coarse resolution configuration, the GM eddy parameterization will dominate the later.) The GCM and matrix solutions are again in reasonable agreement. (We do not expect exact agreement, since the GCM uses a different advection scheme and time-stepping procedure for active tracers.)




Fig. 3. Steady-state temperature (contours) and velocity at 450 m (top) and 950 m (bottom). Left column: GCM solution. Right column: difference between matrix and GCM solutions.

4.2. Global ocean

As a more realistic test, we present results from a global configuration of the MIT model. The version we use was set up to simulate biogeochemical tracers as part of the Ocean Carbon Model Intercomparison Project (M. Follows, personal communication). The model, with 2.8° resolution and 15 vertical levels, was forced with seasonally varying fluxes of heat, salt, and momentum at the surface. In addition, surface temperature and salinity were restored to the Levitus climatology (Levitus et al., 1998) with a 1 month timescale. The model was integrated for 5000 years until it reached equilibrium. We used the equilibrium state of the model to derive the transport matrix at monthly mean resolution (see Appendix B for further details).

We asserted above that for stationary transport, any linear advection scheme will lead to identical results with the matrix and GCM. To illustrate, we compare the matrix and GCM solutions for the transient ideal age computed using three different advection schemes: a non-linear, flux limiting scheme (“superbee”), and two linear schemes (1st order upwind and 3-d order upwind biased). The 3-d order scheme (“DST3”) is one of a family of so called “direct space–time” methods which perform well even in the presence of sharp gradients. Results are shown in Fig. 4.

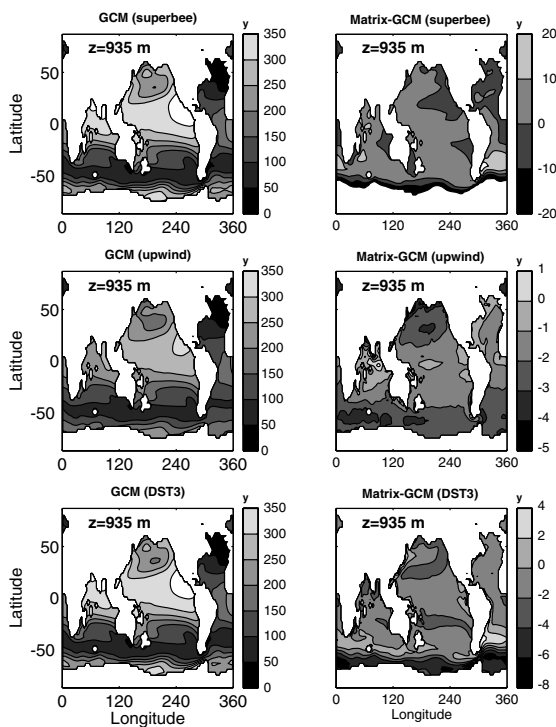


Fig. 4. Transient ideal age (after 500 years) in the GCM (left) simulated with different advection schemes (from top to bottom: superbee, upwind, and DST3). The tracer field at 935 m is shown. Panels on the right display the difference between the corresponding matrix solution and the GCM solution.

In the case of a nonlinear scheme, ages in the matrix solution are on average younger than in the GCM. Because the basis vectors are not smooth in the horizontal, the matrix scheme mimics the diffusive limit of the nonlinear advection scheme, yielding a more diffusive solution. These errors are greatly diminished when a linear scheme is used. (The errors are not zero, however, as we have neglected the seasonal cycle when computing the matrix solution.) While the upwind method is intrinsically more diffusive (middle panels), the DST3 scheme (bottom) does remarkably well, and the solution is practically identical to the superbee solution. (For these relatively “smooth” tracer fields, it is not clear which of these should be considered more “accurate”).

We next compare the steady-state temperature fields in the GCM and the matrix model. The temperature equation in the GCM has the form:

$$\frac{\partial T}{\partial t} = \mathcal{A}_{\text{GCM}}(T) - \lambda(T - T_{\text{Levitus}}) + Q/\rho_0 c_p \Delta z_{\text{sfc}}, \quad (5)$$

where \mathcal{A}_{GCM} is the transport operator of the GCM, Q the surface heat flux, and λ an inverse timescale ($\neq 0$ only in the surface layer). To obtain the corresponding matrix solution, $\mathbf{T}_{\text{steady}}$, we solve

$$\overline{\mathbf{A}}\mathbf{T}_{\text{steady}} - \lambda(\mathbf{T}_{\text{steady}} - \mathbf{T}_{\text{Levitus}}) + \mathbf{Q}/\rho_0 c_p \Delta z_{\text{sfc}} = 0, \quad (6)$$

where $\overline{\mathbf{A}}$ is the annually averaged transport matrix. (The DST3 advection scheme is used to compute the matrix.) Fig. 5 compares the annual mean temperature in the GCM at 455 and 935 m, with the matrix solution. The maximum error in the matrix solution is ≈ 1 °C. The difference between the two solutions may be due to several factors. First, in computing $\mathbf{T}_{\text{steady}}$, we have ignored the seasonal cycle in boundary conditions and transport. Second, as mentioned above, temperature in the GCM is advected by different advection and time-stepping schemes (2d order centered and Adams–Bashforth, respectively). In the matrix computation we use DST3 and a forward Euler step. The transport operators that appear in the GCM and matrix temperature equations are therefore not identical. It is also possible that the restoring terms in Eqs. (5) and (6)




Fig. 5. Steady-state temperature at 455 m (top) and 935 m (bottom) in the GCM (left) and that based on the transport matrix (right). (The matrix-GCM difference is shown.)

may amplify these differences. For many applications, the residual difference may not be insignificant. However, if only an approximate solution is required (as, for example, in the problem of dynamical spin-up discussed below), a difference of this magnitude may not be problematic.

As an additional example, Fig. 6 shows the equilibrium distribution of (natural) $\Delta^{14}\text{C}$ predicted by the matrix model. The equation solved was

$$\overline{\mathbf{A}}\mathbf{c} - k(\mathbf{c} - 100) - \lambda\mathbf{c} = 0, \quad (7)$$

where following Toggweiler et al. (1989), \mathbf{c} is related to conventional $\Delta^{14}\text{C}$ units, k is an exchange coefficient with a simple wind speed dependence, and λ is the decay constant for ^{14}C . For comparison, we also show a reconstruction of the $\Delta^{14}\text{C}$ field based on *Levitus T/S* (S. Peacock, personal communication).

We conclude by discussing the eigen-decomposition of the transport matrix. The eigenvalues and eigenvectors of \mathbf{A} , which would be difficult to obtain from a GCM, can provide great insight into the timescales of ocean transport (Haine and Hall, 2002). The spectrum of eigenvalues shown in Fig. 7 indicate that the longest “e-folding” time scale in this configuration of the MIT model is of $\text{O}(900)$ years). The spatial structure of the eigenvector associated with this smallest eigenvalue is displayed in Fig. 8. The largest amplitude is in the north Pacific, with the amplitude in the Atlantic and Southern Ocean close to zero. The interpretation is that anomalies in the north Pacific take the longest to damp out. In fact, the distribution looks quite similar to the ideal age distribution (Fig. 4). (The eigen-decomposition of the boundary propagator shows that the spatial structure of the mean age is heavily weighted toward that of the gravest mode. We thank one of the referee’s for pointing this out.) In the vertical, the maximum amplitude is always at mid-depth.

4.3. Accelerated dynamical adjustment

A fundamental computational problem encountered in climate modeling is the slow dynamical adjustment of ocean GCMs. The rate limiting step is the very slow diffusive adjustment of abyssal temperature and salinity fields. The fact that the matrix formulation allows us to directly compute steady-state solutions for the active tracers, suggests a solution to the problem. While Eq. (6) is




Fig. 6. Plots of predicted (left) and observed (right) $\Delta^{14}\text{C}$ at 935 and 1000 m, respectively.

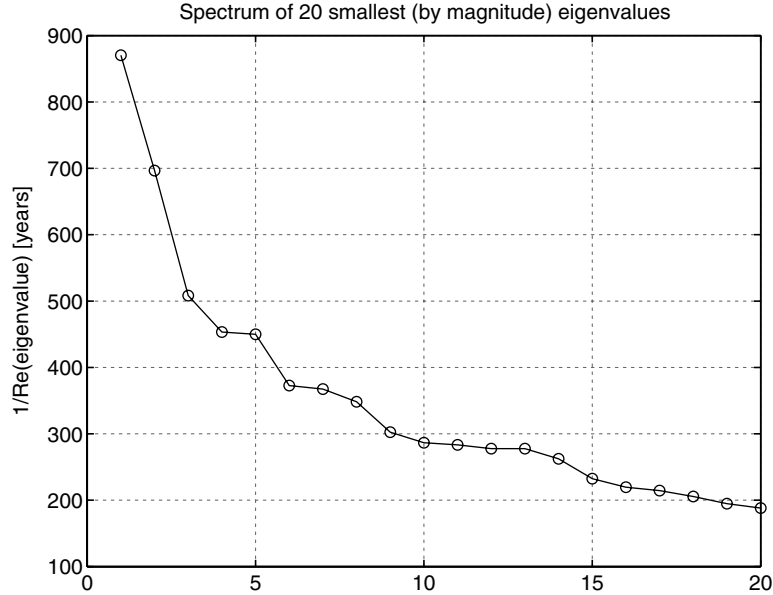


Fig. 7. Spectrum of 20 smallest (by magnitude) eigenvalues of a transport matrix for the global ocean. Plotted are the reciprocal of the real part of the eigenvalues.

strictly valid only if \mathbf{A} is computed from an equilibrium circulation, it can still be used to predict an approximate steady-state solution if \mathbf{A} is not stationary. Since the underlying problem is nonlinear, an iterative scheme is called for. A possible algorithm is as follows. Given the instantaneous circulation of the GCM,

- (1) Compute the transport matrix, $\mathbf{A}(t)$.
- (2) Compute approximate steady-state solutions ($\mathbf{T}_{\text{steady}}, \mathbf{S}_{\text{steady}}$) from Eq. (6) (and a similar equation for salinity).
- (3) Initialize the GCM with new (T, S) fields which are a combination of the existing GCM fields ($\mathbf{T}_{\text{GCM}}, \mathbf{S}_{\text{GCM}}$) and the approximate steady-state solutions.
- (4) Integrate the GCM for a short period. In preliminary experiments we have found that the GCM dynamics geostrophically adjust quite rapidly to the newly imposed density field, as is expected from geostrophic adjustment theory.
- (5) Check for convergence. If the solution has not converged, repeat steps 1–5.

The key step is in combining ($\mathbf{T}_{\text{GCM}}, \mathbf{S}_{\text{GCM}}$) with the approximate solution ($\mathbf{T}_{\text{steady}}, \mathbf{S}_{\text{steady}}$) and this can be done in several ways, such as taking a simple weighted average of the two, or relaxing the GCM (T, S) fields toward ($\mathbf{T}_{\text{steady}}, \mathbf{S}_{\text{steady}}$). Both the weighting and relaxation timescale could themselves be made state dependent, i.e., adaptive. A third possibility is to use an “overrelaxation” method (widely used in the iterative solution of linear equations) to reduce the residual (r_T, r_S) $\equiv (\mathbf{T}_{\text{GCM}} - \mathbf{T}_{\text{steady}}, \mathbf{S}_{\text{GCM}} - \mathbf{S}_{\text{steady}})$. Thus, for example, the GCM temperature at iteration $n + 1$ is predicted by $\mathbf{T}_{\text{GCM}}^{n+1} = \mathbf{T}_{\text{GCM}}^n + \omega r_T$, where ω is the overrelaxation factor. (In simple cases it




Fig. 8. Structure of “gravest” eigenvector of the global transport matrix. Shown is the horizontal structure at 5 depths, and the vertical distribution (bottom right panel) at 4 locations (marked in the bottom left panel).

is possible to calculate an optimal value of ω . It is not clear whether the same can be done here, but the theory should provide some guidance in choosing ω .) The literature on iterative solution of nonlinear equations is of course vast, and more sophisticated schemes can be developed (e.g., Kelly, 1995). The explicit availability of the linearized operator (the transport matrix \mathbf{A}) is advantageous, since we may then use \mathbf{A} or its (incomplete) LU factorization as a “pre-conditioner” to speed up convergence.

To illustrate our ideas, we have applied the overrelaxation method to the ocean sector configuration discussed previously. Fig. 9 shows the temperature at an interior grid point in the GCM. In this simple example, it takes $O(100)$ years by direct integration to reach equilibrium. By contrast, with an unoptimized implementation of the proposed iterative scheme, it takes less than 5 years to reach the same steady-state. (The matrix was computed 3 times.) The slow, diffusive relaxation of the GCM temperature is greatly accelerated by the matrix method, thus addressing one of the key problems encountered in spinning up global models. Further experimentation is under way to determine the optimum frequency and averaging period for computing the transport matrix, and the applicability of the scheme to seasonally forced, and eddy-resolving, models.

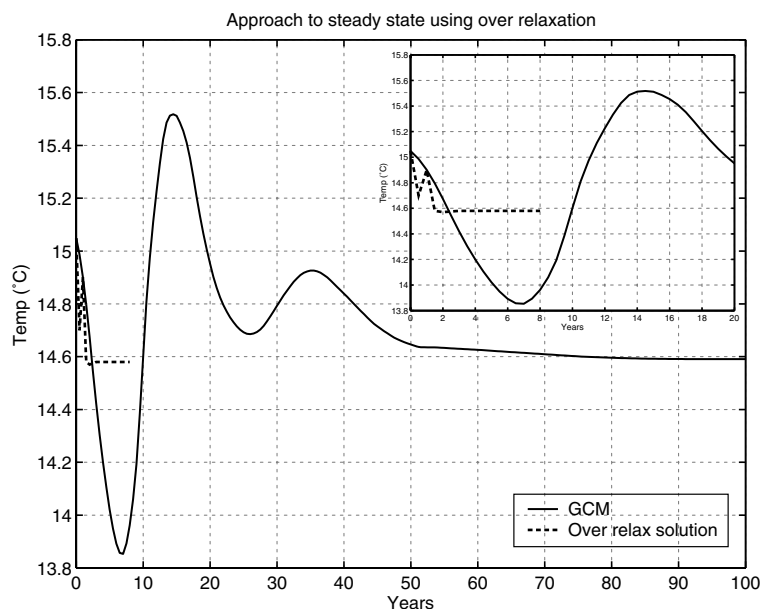


Fig. 9. Example of accelerated spin-up using the matrix approach. Shown is the mid-depth temperature in the GCM and the overrelaxation solution. Inset: expanded view of the initial 20 years.

5. Conclusions

We have developed a novel technique, the “matrix method”, for the accurate and efficient simulation of passive tracers in ocean models. The method is based on capturing the tracer transport of an ocean general circulation model in explicit matrix form. Our motivation is to make simulation of chemical and biological tracers, particularly for paleoclimate research, both convenient and routine. Initial experiments with a global, seasonally forced configuration of the MIT ocean model show that the technique is highly promising, and while some questions of accuracy remain when the method is used with nonlinear advection schemes, it is sufficiently accurate so that simulations performed with it can be compared to data. In addition to efficiently simulating tracers, we have also shown that the method can be used to accelerate the dynamical spin-up of GCMs. These applications of the matrix method should facilitate more effective use of geochemical tracers in addressing a range of oceanographic problems.

Acknowledgements

We thank Mick Follows and Stephanie Dutkiewicz for providing us with a spun-up version of the MIT GCM, Monika Kopacz for assistance with implementation of the matrix method in the MIT model, and Synte Peacock for the reconstructed $\Delta^{14}\text{C}$ data shown in Fig. 6. We also thank two anonymous referees for providing useful suggestions to improve the text. SK would like to acknowledge Francois Primeau for information regarding the fast matrix exponential method and the `EXPOKIT` package. Lamont-Doherty Earth Observatory contribution 6601.

Appendix A. Implementing the matrix algorithm

Here, we discuss some of the practical issues involved in implementing the matrix algorithm described in Section 3.2.

A.1. Alterations to the GCM code

The primary changes to the GCM involve adding code to:

- Reinitialize the tracer field at each time step.
- Compute and store tracer tendencies.
- Archive the time-averaged tendency.

We note that in some GCMs, the tendency is first computed and the tracer field subsequently updated to the next time step. In those models, the update step could be simply skipped, thus avoiding the overhead associated with reinitialization of the tracer field. However, if an implicit mixing scheme is used, or if a time step is composed of smaller steps (e.g., the Lorenz cycle in the Lamont Ocean Model) then the update/reinitialization steps are unavoidable. Predictor–corrector time-stepping schemes, for example, the Adams–Bashforth scheme available as an option in the MIT model, introduce additional complications because of their use of tendencies at multiple time steps to step forward the tracer field.

A.2. Pre-processing steps

The matrix algorithm is facilitated by a sequence of preparatory steps:

- Specification of the model (GCM) geometry, including the coordinates and volumes of grid cells, the grid spacing, and the bathymetry. Care must be taken in computing cell volumes to account for partially filled cells. All information is stored in a consistent, GCM-independent format.
- Specification of the matrix model as a mapping from the GCM grid to the matrix grid. The matrix model can be at a coarser resolution than the GCM. Thus, a tracer cell of the matrix model may be composed of several GCM grid cells. (In this case, to correctly map GCM tendencies into matrix elements, the volume fraction of every GCM grid cell which contributes to a matrix cell is also required, and is pre-computed here.)
- Generation of a “linkage map” (implemented as a linked list) showing how every grid cell of the matrix model is connected to every other cell. The map takes into account both bathymetry and periodicity.
- Subdivision of the computational domain into nonoverlapping tiles based on the linkage map and the size of the halo.
- Specification of the basis vectors (ϕ_i). The same basis is used at horizontal locations with the same water depth. For efficiency, the bases and the matrices Φ^{-1} are pre-computed.

We have found it convenient to automate the entire procedure through a set of MATLAB scripts which may be transparently used with any GCM. (None of the above steps are specific to a

particular model code.) As a practical matter, we make extensive use of the various data structures available in MATLAB, including sparse matrices and cell arrays.

A.3. Matrix calculation

Calculation of the matrix requires the following steps:

- Generating initial conditions (basis vectors) to be read by the GCM.
- Integrating the GCM for specified period and archiving tendency fields.
- Mapping the GCM tendency into matrix elements.

These steps are repeated until all elements of the matrix have been computed. In practice, the entire procedure is scripted. Given a description of the GCM (maximum number of tracers per integration, etc.) and hardware (number of processors/nodes, etc.) configurations, the script automatically distributes the work onto the available processors.

Appendix B. Calculation of the global transport matrix

Computation of the transport matrix for the global configuration of the MIT model (Section 4.2) required integrating roughly 200 “tracers”. The domain can be decomposed in the horizontal into 14 sets of nonoverlapping tiles. This, times the number of vertical levels (15), gives a rough count of the number of independent tracers. Each “tracer”, with initial conditions derived from the basis vectors, was integrated for 1 year, and monthly mean tracers tendencies were archived. (Each run of the GCM carried 15 tracers.) The archived tendency fields were used to derive the transport matrix at monthly mean resolution. The calculation, which took ≈ 1 h, was performed on a 20 processor commodity Beowulf cluster.

References

- Aumont, O., Orr, J.C., Jamous, D., Monfray, P., Marti, O., Madec, G., 1998. A degradation approach to accelerate simulations to steady state in a 3-D tracer transport model of the global ocean. *Clim. Dyn.* 14, 101–116.
- Bryan, K., 1984. Accelerating the convergence to equilibrium of ocean-climate models. *J. Phys. Oceanogr.* 14, 666–873.
- Bryan, K., Lewis, L.J., 1979. A water mass model of the World Ocean. *J. Geophys. Res.* 84, 2503–2517.
- Danabasoglu, G., McWilliams, J.C., Large, W.G., 1996. Approach to equilibrium in accelerated global oceanic models. *J. Clim.* 9, 1092–1110.
- England, M.H., Maier-Reimer, E., 2001. Using chemical tracers to assess ocean models. *Rev. Geophys.* 39, 29–70.
- Gent, P.R., McWilliams, J.C., 1990. Isopycnal mixing in ocean circulation models. *J. Phys. Oceanogr.* 20, 150–155.
- Giering, R., Kaminski, T., 1998. Recipes for adjoint code generation. *ACM Trans. Math. Software* 24, 437–474.
- Haine, T.W.M., Hall, T.M., 2002. A generalized transport theory: water–mass composition and age. *J. Phys. Oceanogr.* 32, 1932–1946.
- Holzer, M., Hall, T.M., 2000. Transit-time and tracer-age distributions in geophysical flows. *J. Atmos. Sci.* 57, 3539–3558.
- Huang, R., Pedlosky, J., 2002. On aliasing Rossby waves induced by asynchronous time stepping. *Ocean Model.* 5, 65–76.

- Kelly, C.T., 1995. Iterative methods for linear and nonlinear equations. In: *Frontiers in Applied Mathematics*, vol. 16. Society for Industrial and Applied Mathematics.
- Large, W.G., McWilliams, J.C., Doney, S.C., 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. US Government Printing Office, Washington, DC.
- Marshall, J., Adcroft, A., Hill, C., Perelman, L., Heisey, C., 1997. A finite-volume, incompressible navier-stokes model for studies of the ocean on parallel computers. *J. Geophys. Res.* 102, 5733–5752.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 1992. *Numerical Recipes in FORTRAN: The Art of Scientific Programming*. Cambridge University Press, New York.
- Ribbe, J., Tomczak, M., 1997. On convection and the formation of Subantarctic Mode Water in the Fine Resolution Antarctic Model (FRAM). *J. Mar. Syst.* 13, 137–154.
- Roe, P.L., 1985. Some contributions to the modeling of discontinuous flows. *Lect. Notes Appl. Math.* 22, 163–193.
- Sidje, R.B., 1998. EXPOKIT. software package for computing matrix exponentials. *ACM Trans. Math. Software* 24 (1), 130–156.
- Stammer, D., Wunsch, C., 1996. The determination of the large-scale circulation of the Pacific Ocean from satellite altimetry using model Green's functions. *J. Geophys. Res.* 101, 18409–18432.
- Thiele, G., Sarmiento, J.L., 1990. Tracer dating and ocean ventilation. *J. Geophys. Res.* 95, 9377–9391.
- Toggweiler, J.R., Dixon, K., Bryan, K., 1989. Simulations of radiocarbon in a coarse resolution world ocean model. 1. Steady state prebomb distributions. *J. Geophys. Res.* 94, 8217–8242.
- Wang, D., 2001. A note on using the accelerated convergence method in climate models. *Tellus* 53A, 27–34.