

# UC Irvine

## Faculty Publications

### Title

A fast Newton-Krylov solver for seasonally varying global ocean biogeochemistry models

### Permalink

<https://escholarship.org/uc/item/4q79k44c>

### Journal

Ocean Modelling, 23(1-2)

### ISSN

14635003

### Authors

Li, Xingwen  
Primeau, Francois W

### Publication Date

2008

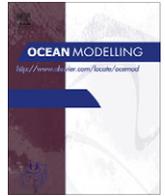
### DOI

10.1016/j.ocemod.2008.03.001

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



# A fast Newton–Krylov solver for seasonally varying global ocean biogeochemistry models

Xingwen Li, François W. Primeau \*

Department of Earth System Science, University of California, Irvine, CA 92697-3100, United States

## ARTICLE INFO

### Article history:

Received 28 September 2007  
 Received in revised form 27 February 2008  
 Accepted 3 March 2008  
 Available online 19 March 2008

## ABSTRACT

We present a computationally-efficient method for obtaining the fully spun-up state of a seasonally-varying global ocean biogeochemistry model. The solver uses a Newton–Krylov method to find the fixed points of the map that assigns to an initial state the value of the model state at the end of a one-year run. Apart from the preconditioner, which we describe in the paper, the method relies on a black-box public-domain Newton–Krylov solver that does not require the explicit construction of the model's Jacobian matrix. Applied to the  $\text{PO}_4$  plus dissolved organic phosphorus (DOP) cycle of an Ocean Carbon Model Intercomparison Project II (OCMIP-2) type model, the solver is more than two orders of magnitude faster than the traditional time-stepping method for spinning up the model. The efficiency of the solver is illustrated by using the seasonally varying globally-gridded  $\text{PO}_4$  climatology to objectively optimize the parameters that control the mean lifetime of semi-labile DOP and the fraction of new production allocated to DOP. The optimization study demonstrates that the information in the seasonal variations of  $\text{PO}_4$  do not provide a significantly stronger constraint than the annually averaged data used in previous optimization studies.

© 2008 Elsevier Ltd. All rights reserved.

## 1. Introduction

Ocean-biogeochemistry models have many uncertain parameters that cannot be directly measured in the field or laboratory. These parameters must be constrained by adjusting them so that the model output matches the available observations as well as possible. It is therefore desirable to develop objective parameter optimization methods that go beyond the trial-and-error approach.

The major impediment to applying automatic parameter optimization methods to global ocean-biogeochemistry models is that global models take several thousand years to reach a new equilibrium each time a parameter is perturbed. This timescale for transients to die out and the model to reach its equilibrium is set by the most slowly decaying eigen-mode of the model's advection-diffusion equation. For a global model the slowest decaying mode has a timescale of roughly 800 years (Primeau, 2005; Khatiwala et al., 2005) and projects directly on the transport of nutrient rich deep water back to the sunlit surface ocean. The long runs needed to reach equilibrium have made it impractical to perform the large number of runs needed to systematically optimize parameters objectively even for relatively coarse resolution models.

Here, we present a new implicit solver that greatly reduces the computational time needed for obtaining equilibrium solutions of a seasonally varying global ocean-biogeochemistry model. We

build on the recent work of Kwon and Primeau (2006, 2008) who introduced the use of an implicit iterative solver based on Newton's method to compute the steady state of an ocean-biogeochemistry model. The key new development we present here is a solver that can obtain cyclo-stationary solutions of a seasonally varying model.

The inclusion of seasonal effects is important for biogeochemistry because the upper ocean exhibits strong seasonal variations in both circulation and export production. Potentially important rectification effects cannot be captured in a steady model. It is known for example, that the  $T$ - $S$  properties of the main thermocline reflect more closely those of the end-of-winter mixed layer rather than the annual-mean state and one expects that Stommel's mixed-layer "Demon" is also at work biasing the coupling between the surface mixed layer and the main thermocline for nutrients, albeit modified by the seasonality of the sinking particle flux (Stommel, 1979; Williams et al., 1995). The first goal of this paper is to present the formulation and implementation of a new solver and demonstrate its efficiency in the context of the phosphorus cycle as implemented in phase 2 of the Ocean carbon Model Intercomparison Project (OCMIP-2, R. Najjar and J. Orr, Design of OCMIP-2 simulations of chlorofluorocarbons, the solubility pump and common biogeochemistry, 1998, <http://www.ipsl.jussieu.fr/OCMIP/phase2/simulations/design.ps>, hereinafter referred to as Najjar and Orr online document). A key advantage of our new solver is that it does not require the explicit construction of the model's Jacobian matrix and should therefore be easy to implement in any ocean

\* Corresponding author. Tel.: +1 949 824 9435; fax: +1 949 8243874.  
 E-mail address: [fprimeau@uci.edu](mailto:fprimeau@uci.edu) (F.W. Primeau).

biogeochemistry model. In fact the solver we use is a professionally written public domain Newton–Krylov solver written by Kelley (2003). All that is needed to apply the method is to formulate a preconditioner and to implement a simple interface between the model and the black-box Newton–Krylov solver. We give only a basic idea behind the Newton–Krylov algorithm because the implementation details can be found in the book by Kelley (2003). Our focus is on the formulation of the interface between the model and the solver which we describe in Section 2 and the preconditioner which we described in Section 4 after first reviewing the model formulation in Section 3. The formulation of an effective preconditioner is the key to the method’s efficiency.

In addition to describing the solver itself, we also illustrate how the fast solver can be used effectively in a parameter optimization study. The optimization study we conduct is motivated by the previous work of Kwon and Primeau (2006, hereinafter referred to as KP06). In their optimization study of the parameters controlling the cycling of phosphorus using the annual-average global PO<sub>4</sub> climatology, KP06 found that the parameters controlling the mean lifetime of dissolved organic phosphorus (DOP) and the fraction of new production allocated to DOP were not well constrained independently. They speculated that perhaps the information in the seasonal variations of the PO<sub>4</sub> climatology that was left out in the cost function for their steady-state model might help to better constrain the parameters. The second goal of the present study is to test this hypothesis. To our knowledge ours is the first such optimization study of a seasonally varying global model. The formulation of the cost function and the optimization results are presented in Section 6.

## 2. Equilibrium solutions expressed as the fixed points of a map

In this section, we describe the interface needed to cast the output of a biogeochemistry model in a form that can be used by the Newton solver. For the steady-state case considered by KP06, the equations governing the biogeochemistry model can be written symbolically as

$$\frac{d\mathbf{x}(t)}{dt} = F(\mathbf{x}(t)), \quad (1)$$

where  $\mathbf{x}(t)$  is the model state at time  $t$ . The usual approach to spinning-up the model to equilibrium is to integrate Eq. (1) forward in time until the transients die out so that  $d\mathbf{x}/dt \rightarrow 0$  and  $\mathbf{x}$  becomes independent of time. To find the equilibrium state using Newton’s method one sets  $d\mathbf{x}/dt = 0$  and seeks the solution to the resulting time-independent coupled system of algebraic equations,  $F(\mathbf{x}_{\text{eq}}) = 0$ . A sequence of Newton iterates that yields  $\mathbf{x}_{\text{eq}}$  when it converges is

$$\mathbf{x}_{n+1} = \mathbf{x}_n - \left[ \frac{\partial F}{\partial \mathbf{x}} \right]_{\mathbf{x}_n}^{-1} F(\mathbf{x}_n), \quad (2)$$

where  $[\partial F / \partial \mathbf{x}]_{\mathbf{x}_n}$  is the Jacobian matrix of partial derivatives of  $F$  with respect to the model state  $\mathbf{x}$  evaluated at  $\mathbf{x}_n$ . States for which the time-tendency are zero within machine precision are typically obtained after 5 or 6 iterations. The efficiency of the solver used by Kwon and Primeau (2006, 2008) derives from the fact that the Jacobian matrix for finite-difference models is very sparse. Because the matrix has mostly zero elements it can be inverted efficiently using a modern sparse-matrix factorization algorithm.

We now wish to apply a similar method to a seasonally varying model. The difficulty is that the equilibrium state is now time-dependent so we cannot simply set  $d\mathbf{x}/dt = 0$ . The governing equation for the seasonally varying model is of the form

$$\frac{d\mathbf{x}(t)}{dt} = F(t, \mathbf{x}(t)), \quad (3)$$

with an explicit time-dependence in  $F$  due to the seasonality of the circulation, temperature field, solar radiation, etc. We are interested in the case where the explicit time dependence of  $F$  is periodic with period  $T = 1$  year, i.e.  $F(t + T, \mathbf{x}) = F(t, \mathbf{x})$ . In an offline transport model it is easy to ensure that the system satisfies this condition by simply recycling the pre-saved one-year time series of velocity and eddy-diffusivity. In a biogeochemistry model that is run concurrently with ocean dynamics, one can ensure that the system is periodic by recycling a one-year time series of surface fluxes of momentum, fresh-water and heat, and by resetting the dynamical state of the model (temperature, salinity and momentum) to the same state at the beginning of each year. To find the periodic equilibrium solutions of equation (3) we want to recast the problem in a form that can be solved using Newton’s method. To this end, we define the map,  $M : \mathbf{x}_0 \rightarrow \mathbf{x}$ , that assigns to the initial model state  $\mathbf{x}(t_0) \equiv \mathbf{x}_0$ , the value of the solution one period later, i.e.  $\mathbf{x}(t_0 + T)$ . Evaluating the map,  $M$ , is equivalent to running the model forward in time for a one year period starting from the initial state  $\mathbf{x}_0$ . The equilibrium solutions of Eq. (3) correspond to fixed points of  $M$ , i.e.  $M(\mathbf{x}_{\text{eq}}) = \mathbf{x}_{\text{eq}}$ . We note in passing that the usual approach of spinning up an ocean model by time-stepping the model forward is equivalent to successive applications of the map  $M$ , in other words the model state at the end of one year is used as the initial condition for starting the next year of integration. Spinning-up the model in this way is therefore equivalent to applying Picard’s method (e.g. Pozrikidis, 1998),  $\mathbf{x}_{n+1} = M(\mathbf{x}_n)$ , to find the fixed point of  $M$ . An alternative to Picard’s method with a better convergence rate is to apply Newton’s method to find the roots of  $G(\mathbf{x}_{\text{eq}}) \equiv M(\mathbf{x}_{\text{eq}}) - \mathbf{x}_{\text{eq}} = 0$ . The use of a map as an intermediary in the application of Newton’s method has been applied successfully in the chemical engineering community to obtain periodic states of cyclically operated chemical processes, (see for example Van Noorden et al., 2003.) Here, we apply the method on a much grander scale to the global ocean viewed as a chemical reactor operated cyclically by the periodic orbit of the Earth around the Sun.

Newton’s method applied to  $G(\mathbf{x}_{\text{eq}}) = 0$  yields

$$\mathbf{x}_{n+1} = \mathbf{x}_n - \left[ \frac{\partial G}{\partial \mathbf{x}} \right]_{\mathbf{x}_n}^{-1} G(\mathbf{x}_n). \quad (4)$$

Unfortunately, the Jacobian of  $G$  unlike that of  $F$  is not sparse. Even for a low-order finite difference scheme, a tracer perturbation at one grid point can spread to many grid points during the integration of the model over a full year. Consequently, the direct inversion approach used for the steady-state case by Kwon and Primeau (2006, 2008) is no longer feasible. Instead of explicitly constructing and inverting the Jacobian matrix of  $G$ , we use an iterative solver that needs only Jacobian-vector products. These in turn are computed by the solver itself from  $G$  using a forward-difference directional derivative; all that is needed to apply the solver is a subroutine that evaluates  $G$ . Evaluating  $G(\mathbf{x})$  is as simple as running the model forward for one year starting with  $\mathbf{x}$  as initial condition and then subtracting this initial state from the model state at the end of the one-year run. The nonlinear solver then uses an iterative linear solver, to approximate the solution to Eq. (4). A Newton–Krylov solver therefore consists of two iterative solvers: an outer Newton solver which forms the iterations of Eq. (4) and an inner Krylov solver that uses an iterative approach to “invert” the Jacobian. The particular linear iterative solver we use is the GMRES algorithm (Saad and Schultz, 1986) as implemented by Kelley (2003). GMRES approximates the solution of a linear system  $\mathbf{A}\mathbf{x} = \mathbf{b}$  with a sum of the form,

$$\mathbf{x}_n = \mathbf{x}_0 + \sum_{k=0}^{n-1} a_k \mathbf{A}^k \mathbf{r}_0, \quad (5)$$

where  $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$  and where  $\mathbf{x}_0$  is the initial iterate. The coefficients  $a_k$  are determined using linear least squares to minimize  $\|\mathbf{b} - \mathbf{A}\mathbf{x}_n\|$ . The details of the implementation are taken care by the black-box solver.

The convergence rate of the inner iterative solver depends critically on preconditioning the matrix  $\mathbf{A}$  by pre-multiplying it by an approximate inverse  $\mathbf{P} \approx \mathbf{A}^{-1}$ . The closer  $\mathbf{P}$  is to the exact inverse of  $\mathbf{A}$  the fewer iterations will be needed. However, there is a trade-off between the computational costs associated with the construction and application of  $\mathbf{P}$  and the need to minimize the number of inner iterations. The preconditioner we have implemented is the so-called “physics-based” preconditioner in that it is inspired by the form of the governing equations. We therefore briefly describe the biogeochemistry model equations before returning to the formulation of the preconditioner.

### 3. Model description

The OCMIP-2 type phosphorus cycle model is described in detail in KP06. The circulation model is also the same as the one used by KP06. The governing equations of phosphate,  $\text{PO}_4$ , and dissolved organic phosphorus DOP are

$$\begin{aligned} \frac{\partial [\text{PO}_4]}{\partial t} &= \nabla \cdot \mathbf{J}_{\text{PO}_4} + S_{\text{PO}_4} - ([\text{PO}_4] - \overline{\text{PO}_4}) / \tau_{\text{PO}_4}, \\ \frac{\partial [\text{DOP}]}{\partial t} &= \nabla \cdot \mathbf{J}_{\text{DOP}} + S_{\text{DOP}}, \end{aligned} \quad (6)$$

where  $S_{\text{PO}_4}$ , and  $S_{\text{DOP}}$  are the source-sink terms of  $\text{PO}_4$  and DOP due to biological uptake and remineralization and where  $\mathbf{J}_{\text{PO}_4} \equiv (\mathbf{u} - \mathbf{K}\nabla)[\text{PO}_4]$  and  $\mathbf{J}_{\text{DOP}} \equiv (\mathbf{u} - \mathbf{K}\nabla)[\text{DOP}]$  are the advective-diffusive fluxes of  $\text{PO}_4$  and DOP for a fluid velocity  $\mathbf{u}$  and eddy-diffusion tensor  $\mathbf{K}$ . We have also added a slow phosphate restoring term with  $\overline{\text{PO}_4} = 2.2 \text{ mmol/m}^3$  and with  $\tau_{\text{PO}_4} = 53,000$  years set to the mean residence time of  $\text{PO}_4$  in the ocean. This extra source-sink term eliminates the null space in the Jacobian of  $G$  associated with the conservation of total phosphorus in the original model. Apart from making the Jacobian invertible the extra term ensures that the iterative solver will always find a solution whose total phosphate inventory equals the value estimated from observations. Without such a constraint the total amount of phosphorus would be free to drift during the Newton–Krylov iterations. The geological timescale associated with this term ensures that it has little effect on the spatio-temporal  $\text{PO}_4$  distribution.

The biological source-sink terms are defined as follows

$$\begin{aligned} S_{\text{PO}_4} &= \kappa [\text{DOP}] + \begin{cases} -j_{\text{PO}_4} & z > z_c \\ -\frac{\partial}{\partial z} (z/z_c)^2 (1 - \sigma) \int_0^{z_c} j_{\text{PO}_4} dz & z < z_c \end{cases} \\ S_{\text{DOP}} &= -\kappa [\text{DOP}] + \begin{cases} \sigma j_{\text{PO}_4} & z > z_c \\ 0 & z < z_c \end{cases} \end{aligned} \quad (7)$$

in which

$$j_{\text{PO}_4} = \frac{1}{\tau} ([\text{PO}_4] - [\text{PO}_4]_{\text{obs}}) \times \Theta([\text{PO}_4] - [\text{PO}_4]_{\text{obs}}). \quad (8)$$

Here,  $j_{\text{PO}_4}$  is the organic phosphorus production, also defined as the new production, following Anderson and Sarmiento (1995). It is determined by restoring the modeled  $[\text{PO}_4]$  to the observed monthly averaged values,  $[\text{PO}_4]_{\text{obs}}$ , with a time scale of  $\tau = 30$  days. The function  $\Theta(x) \equiv \frac{1}{2} [1 + \tanh(x/\lambda)]$  is a smoothed step-function that is used to shut off production when the modeled  $[\text{PO}_4]$  drops below the observed value. This shutting off of the production when nutrient concentrations drop below a threshold introduces the essential nonlinearity in the model. The parameter  $\lambda$  controls the steepness of the smoothed-step. In the limit  $\lambda \rightarrow 0$ ,  $\Theta(x)$  becomes the step-function. We chose a value of  $\lambda = 10^{-5} \text{ mmol m}^{-3}$  which is small enough to give results that are indistinguishable from solutions obtained

with the original step-function used in the OCMIP formulation. The parameter  $\kappa$  defines the inverse e-folding time scale for the remineralization of DOP.  $\sigma$  is the fraction of the production that is allocated to DOP. The remaining fraction,  $1 - \sigma$  is exported out of the euphotic zone as sinking particles (POP). The exponent  $\alpha$  controls the attenuation of the POP flux with depth as particles are remineralized and organic phosphorus is returned into inorganic  $[\text{PO}_4]$ .

The model is solved using an offline model with a one day time-step. The advection-diffusion transport operator uses monthly averaged velocity and eddy-diffusion tensor fields that were obtained from a dynamical OGCM that was spun up for more than 6000 years. Our particular offline model is integrated forward in time using a second order Crank Nicholson scheme for the advection-diffusion terms and a third order Adams Bashforth scheme for the biogeochemical source-sink terms. The off-line model periodically cycles through the 12  $\mathbf{u}$  and  $\mathbf{K}$  fields. It is important to emphasize, that the implementation and efficiency of the Newton–Krylov solver does not depend in any way on the choice of time-integration scheme used in the offline model. We chose this particular time-integration scheme because it was easy for us to implement.

The dynamical OGCM used to generate the circulation fields is a version of the climate model of the Canadian Center for Climate Modeling and Analysis, based on the NCAR CSM Ocean Model (Pacanowski et al., 1993; NCAR CSM Ocean Model Technical Note (NCAR/TN-423+STR) [http://www.cesm.ucar.edu/models/ocn-ncm/Doc1\\_3.html](http://www.cesm.ucar.edu/models/ocn-ncm/Doc1_3.html)). The transport model has 29 vertical levels with thickness of 50 m near the surface to 300 m in the deepest level and a horizontal resolution about  $3.75^\circ \times 3.75^\circ$ . The water mass ventilation properties of the annually averaged circulation are described in detail in Primeau (2005), Primeau and Holzer (2006), Holzer and Primeau (2006, 2008).

### 4. Preconditioner

Our starting point for formulating the preconditioner is based on the insight gained from previous water mass ventilation studies (Primeau, 2005; Primeau and Holzer, 2006) that the long spin-up runs needed for ocean biogeochemistry models to reach equilibrium are due to the slow ventilation timescale for the deep ocean. Even though the preconditioner we use and describe below is “physics based”, it does not have to strictly satisfy any physical or biogeochemical principles. We are free to construct the preconditioner by neglecting most of the biological sources and sinks. We kept only the slow  $\text{PO}_4$  phosphate restoring term and the first-order DOP decay rate. These two terms were kept because they eliminate the transport operator’s null space. In general, retaining one sink term for each tracer in the model is sufficient to assure that the null space will be eliminated. Because these source sink terms are linear in the model state they do not introduce any state dependence in the preconditioner. In more complex models that include tracers with only a nonlinear sink term one would have to linearize the retained nonlinear sink terms about an appropriate time-independent reference state. Furthermore, since the seasonal variations of the circulation in the deep ocean are not very pronounced, we treated the circulation as if it was steady and used the annually averaged circulation to construct an approximate inverse to  $\partial G / \partial \mathbf{x}$ .

The resulting preconditioner we used is given by  $\mathbf{P} = \mathbf{Q}^{-1} - \mathbf{I}$  with

$$\mathbf{Q} = \begin{bmatrix} \langle \mathbf{A} \rangle - \mathbf{I} / \tau_{\text{PO}_4} & \kappa \mathbf{I} \\ \mathbf{0} & \langle \mathbf{A} \rangle - \kappa \mathbf{I} \end{bmatrix}, \quad (9)$$

where  $\mathbf{I}$  is an  $n \times n$  identity matrix with  $n$  equal to the total number of grid-boxes in the model and where

$$\langle \mathbf{A} \rangle = -\nabla \cdot [(\mathbf{u}) - \langle \mathbf{K} \rangle \cdot \nabla], \quad (10)$$

is the advection-diffusion transport operator constructed from the annually averaged velocity and eddy-diffusivity tensor fields. The advection-diffusion equation when discretized using a finite-difference scheme naturally yields a sparse matrix operator because the instantaneous tracer-flux divergence at any given grid point depends only on the tracer concentration in a small neighborhood of grid-cells close to the point at which the flux divergence is being computed. The sparsity of  $\langle \mathbf{A} \rangle$  makes it possible to efficiently compute and store the LU factorization of the  $\mathbf{Q}$  from which the effect of pre-multiplying by  $\mathbf{P} = \mathbf{Q}^{-1} - \mathbf{I}$  can be computed accurately and efficiently (Golub and Van Loan, 1989). Some other important properties of  $\mathbf{P}$  are that: (1) it is independent of the biogeochemistry model state as well as time independent and so needs to be constructed only once, (2) the numerical scheme used to construct  $\langle \mathbf{A} \rangle$  need not match exactly the numerical scheme used for the time dependent tracer transport model so that the time-dependent scheme could in principle use a high-order upwind scheme with flux limiters, and (3) the matrix  $\mathbf{Q}$  has a block triangular structure which means that only the diagonal blocks of the form  $\langle \mathbf{A} \rangle - \mathbf{D}$ , where  $\mathbf{D}$  is a diagonal matrix whose coefficients are the linear Taylor series coefficients of the retained sink terms, need to be factored in order to apply  $\mathbf{P}$ . This last point suggest that when constructing  $\mathbf{Q}$  for biogeochemical models that carry a large number of tracers it should be especially advantageous to neglect source-sink terms so as to retain a block-triangular structure.

The above preconditioner is applied by pre-multiplying  $G$  by  $\mathbf{P}$ . This is done by solving  $\mathbf{Q}\tilde{G} = G$  using the precomputed and stored LU decomposition of  $\mathbf{Q}$  and passing  $\tilde{G} = G$  instead of  $G$  to the black-box Newton-Krylov solver. Left-multiplying by  $\mathbf{P}$  does not change the fixed points of  $M$  but greatly accelerates the convergence rate of the solver as we will see in the next section.

## 5. Convergence rate to equilibrium

The convergence rate of the new solver is shown in Fig. 1 where the root mean square (rms) drift of the model solution is plotted as a function of the number of evaluations of the map  $M$ , i.e. as a function of the number of one-year runs used. The figure compares the convergence rate of the traditional time-stepping method (Picard iterations) which takes in excess of 3000 years to reduce the rms drift below  $10^{-9} \text{ mmol m}^{-3} \text{ yr}^{-1}$ , the Newton-Krylov solver without the application of the preconditioner which requires approxi-

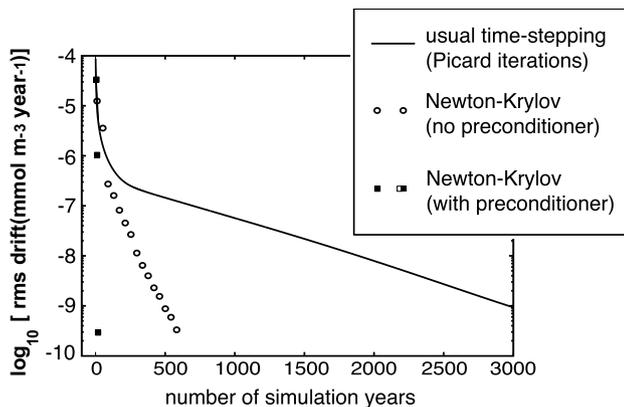


Fig. 1. Root mean square differences of  $\mathbf{x}(t+T) - \mathbf{x}(t)$ , where  $T = 1$  yr, for the usual time-stepping spin up (solid line) and for the Newton-Krylov solver with no a preconditioner (open circles) and for the Newton-Krylov solver with the preconditioner (solid squares). Although it is difficult to determine from the figure, a total number of 24 simulation years were needed for the pre-conditioned Newton-Krylov solver to converge.

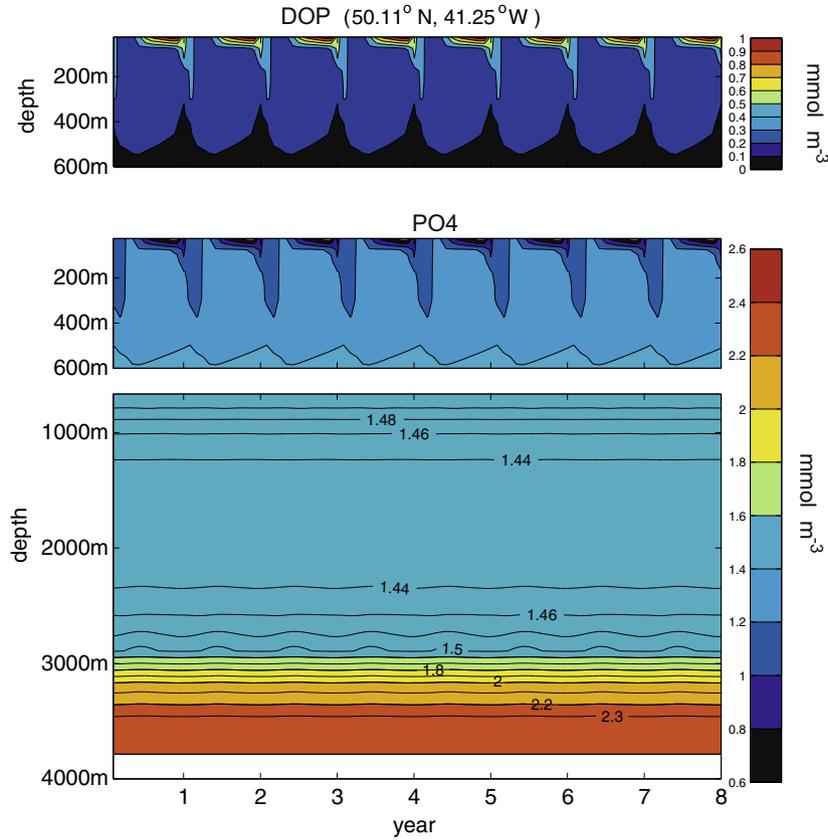
mately 500 years of integration to reduce the drift below  $10^{-9} \text{ mmol/yr}$  and the Newton-Krylov solver with the application of the preconditioner described in Section 4 which reduces the rms drift to below  $10^{-9}$  after only 24 years of integration. Since the computational cost associated with the application of the preconditioner and the overhead associated with the Newton-Krylov solver is negligible compared to a one-year run of the model the preconditioned Newton-Krylov solver is more than two orders of magnitude faster than the traditional time-stepping method for spinning up the biogeochemistry model. All three cases shown in Fig. 1 were initialized with the same observed climatological  $\text{PO}_4$  distribution and a uniform DOP distribution. If however one has a good initial iterate, the average number of years of integration needed for the preconditioned Newton-Krylov solver to converge can be substantially less. Good initial iterates are typically available when one uses the solver with a parameter optimization routine. This will be illustrated in the next section.

Depth versus time plots of the fully equilibrated DOP and  $\text{PO}_4$  concentrations for the model with parameter values ( $\sigma = 0.74$ ,  $1/\kappa = 1 \text{ yr}$ ,  $\alpha = -1.0$ ) are shown for the N Atlantic and Eastern Equatorial Pacific in Figs. 2 and 3. In the Eastern Equatorial Pacific the solution shows only a weak seasonal cycle and surface waters have elevated DOP levels throughout the year. The new-production for the model solution is very high and does not vary much with the seasons in the Eastern Equatorial Pacific (Fig. 4a and b). The over estimated production in the upwelling region of the eastern equatorial pacific is a common problem for coarse resolution models and is believed to be associated with errors in the model circulation (Najjar et al., 2007). In the N Atlantic, the solution has a strong non-sinusoidal seasonal cycle with a mixed layer that penetrates down to approximately 400 m. At depth, the solution shows only a weak seasonal cycle due to the strong attenuation of the particle remineralization profile. In the surface the DOP concentration is anti-correlated with the  $\text{PO}_4$  concentration as expected from the periodic mixing of deep waters that have high  $\text{PO}_4$  and low DOP levels and from the fact that DOP the production of DOP is associated with the consumption of  $\text{PO}_4$ . We note that the model solution is fully equilibrated even in the deep ocean and shows no drift associated with the slow adjustment from the initial condition.

## 6. Optimization of the parameters controlling the cycling of DOP

In this section, we illustrate the usefulness of the Newton-Krylov solver in an optimization study of the parameters that control the cycling of DOP. The parameters in question are  $\sigma$  the fraction of new production allocated to DOP and  $\kappa$  the first-order e-folding remineralization rate of DOP. When these parameters were optimized by minimizing a cost function formed from the steady model and the annually averaged  $\text{PO}_4$  observations, KP06 found that these two parameters were strongly correlated and could not be constrained independently. Our goal here is to test the hypothesis that the information in the seasonal variations in the  $\text{PO}_4$  data – information that was ignored in the steady-state cost function used by KP06 – can help to better resolve the parameters.

To test this hypothesis we form a cost function that measures the distance of the monthly averaged model state from the monthly averaged  $\text{PO}_4$  climatology from the World Ocean Atlas 05, (Garcia et al., 2006). The shape of the cost function near its minimum determines how well the parameters are constrained. If the information in the seasonal variations helps to better constrain the parameters, including seasonal variations in the cost function should either increase the curvature of the cost function near its minimum, or decrease the amount of parameter interactions by



**Fig. 2.** Depth versus time plot showing the fully equilibrated seasonally varying tracer fields with no visible drift. The uppermost panel shows the concentration of DOP and the lower panels show the concentration of PO<sub>4</sub> in mmol m<sup>-3</sup> in the N Atlantic (50°N, 41°W) The solution is for model parameters  $\sigma = 0.74$ ,  $1/\kappa = 1$  yr, and  $\alpha = -1.0$ .

aligning the axes of the elliptical bowl around the minimum with the parameter axes.

The cost we use is

$$C(\sigma, \kappa) = \sum_{k=1}^{12} \sum_{i=1}^n \frac{\Delta V_i ([\text{PO}_4]_{i,k}(\sigma, \kappa) - [\text{PO}_4^{\text{obs}}]_{i,k})^2}{\text{Var}[\text{PO}_4^{\text{obs}}]}, \quad (11)$$

where  $\Delta V_i$  is the volume of the  $i$ th model grid box,  $[\text{PO}_4^{\text{obs}}]_{i,k}$  is the observed PO<sub>4</sub> concentration for the  $k$ th month interpolated to the  $i$ th model grid point and  $[\text{PO}_4]_{i,k}(\sigma, \kappa)$  is the simulated equilibrium PO<sub>4</sub> concentration in the  $i$ th grid box averaged over the  $k$ th month (the model uses a one day time-step). The cost function is scaled by the spatio-temporal variance of the observed field

$$\text{Var}[\text{PO}_4^{\text{obs}}] = \sum_{k=1}^{12} \sum_{i=1}^n \Delta V_i \left( [\text{PO}_4^{\text{obs}}]_{i,k} - \langle [\text{PO}_4^{\text{obs}}] \rangle \right)^2 \quad (12)$$

where

$$\langle [\text{PO}_4^{\text{obs}}] \rangle = \frac{1}{12} \sum_{k=1}^{12} \sum_{i=1}^n \Delta V_i [\text{PO}_4^{\text{obs}}]_{i,k} = 2.2 \text{ mmol m}^{-3} \quad (13)$$

is the spatially and annually averaged PO<sub>4</sub> concentration.

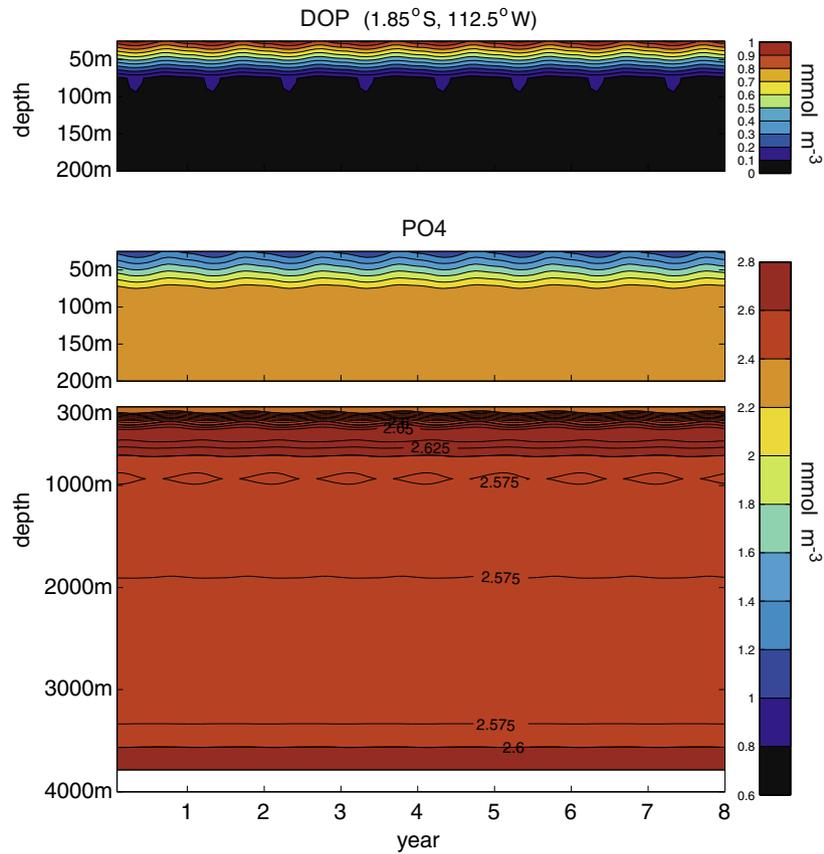
To evaluate the cost function we computed equilibrium solutions on a grid of  $\sigma$ - $\kappa$  parameter values using a time-step of  $dt = 1/72$  years. For this computation the equilibrium solution at one set of parameter values was used as the initial iterate for a nearby set of parameter values. The same preconditioner was used to compute all the solutions. Fig.5a shows the number of simulation years needed by the solver to reach equilibrium as a function of  $\sigma$  and  $\kappa$ . The plot shows that with a good initial iterate the typical number of years the model must be run to reach equilibrium is between 40 and 15 years. If one applied an automatic minimization

routine to find the optimal parameter values most of the needed equilibrium solutions would be computed in the neighborhood of the minimum of the cost function where for our example the number of years of simulation needed is around 20, making the Newton-Krylov solver more than 150 times faster than the straight time-stepping approach.

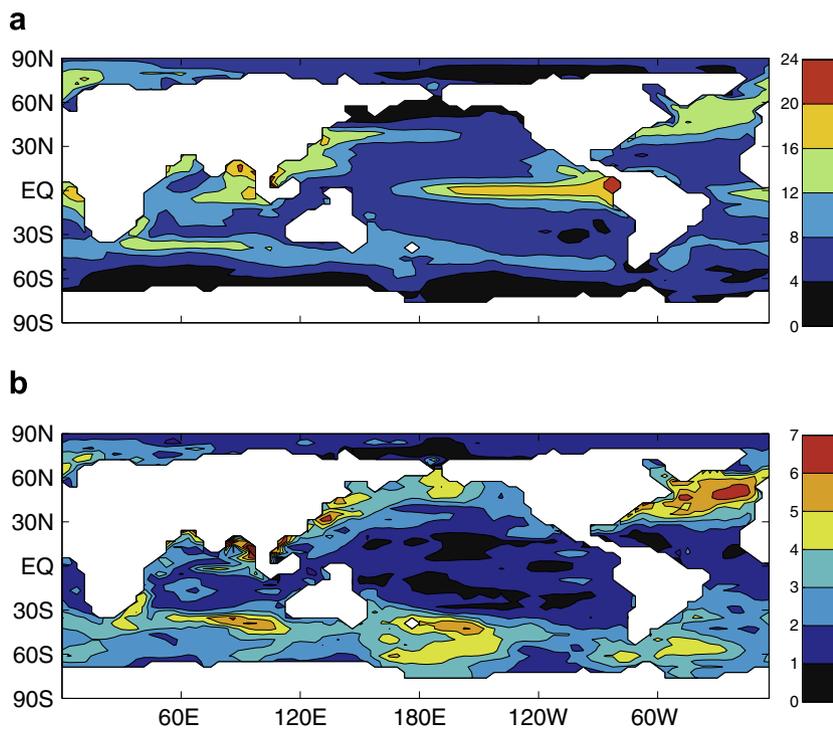
The resulting cost function is plotted as a function of  $\sigma$  and  $\kappa$  in Fig. 5b. The contours indicate the fraction of the spatio-temporal variance that was not captured by the model. The plot shows that the optimal value for  $\sigma$  is between 0.4 and 0.75, and that  $1/\kappa$  is not well constrained. The model captures slightly more than 64% of the spatio-temporal variance. This fraction is slightly less than the fraction (70%) of the spatial variance alone captured by the steady model (KP06). The shape of the cost function is also very similar for the steady and seasonally varying model. The cost function for the steady-state model is shown in Fig. 4 of KP06. The conclusion must therefore be that in the context of the OCMIP-2 type model, the information in the seasonal cycle of the PO<sub>4</sub> data does not provide a stronger constraint than the annually averaged data. This result emphasizes the importance of obtaining actual DOP or DOM data so that  $\kappa$  and  $\sigma$  can be properly constrained.

## 7. Discussion and summary

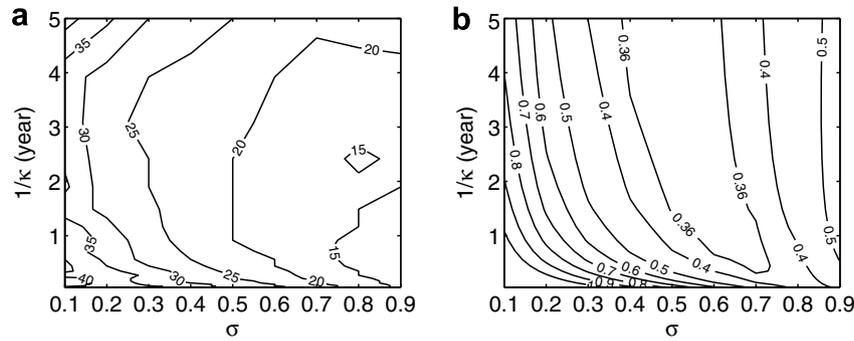
Global ocean biogeochemistry models typically need long spin-up runs of more than 3000 years before the deep ocean reaches equilibrium with the surface. This long spin-up time is due to the slowly decaying modes of the circulation model's advective-diffusive operator. In order to decrease the prohibitive computational costs associated with the need to repeatedly re-equilibrate the model each time a parameter is perturbed in the process of



**Fig. 3.** Depth versus time plot showing the fully equilibrated seasonally varying fields with no visible drift. The uppermost panel shows the concentration of DOP and the lower panels show the concentration of  $\text{PO}_4$  in  $\text{mmol m}^{-3}$  in the Eastern Equatorial Pacific at  $(2^{\circ}\text{S}, 113^{\circ}\text{W})$ . The solution is for model parameters  $\sigma = 0.74$ ,  $1/\kappa = 1$  yr, and  $\alpha = -1.0$ .



**Fig. 4.** (a) Annual mean new production in  $\text{mol C m}^{-2} \text{yr}^{-1}$ . (b) Standard deviation of the seasonal cycle of new production in  $\text{mol C m}^{-2} \text{C yr}^{-1}$ . The solution is for model parameters  $\sigma = 0.74$ ,  $1/\kappa = 1$  yr, and  $\alpha = -1.0$ .



**Fig. 5.** (a) The number of simulation years needed for the pre-conditioned Newton–Krylov solver to reach equilibrium as a function of  $\sigma$  and  $\kappa$ . The same preconditioner was used to compute all the solutions. (b) Cost function expressed as the fraction of unexplained spatio-temporal variance in the observed  $\text{PO}_4$  field as a function of  $\sigma$  and  $1/\kappa$  the parameters that control respectively the mean life-time of semi-labile DOP and the fraction of new production allocated to DOP as opposed to sinking particles.

optimizing parameters, we have developed an efficient Newton–Krylov solver that is up to two orders of magnitude faster than straight time-stepping. The solver we have developed uses a professionally written public-domain Newton–Krylov solver and does not require any modification to the existing ocean biogeochemistry model. In order to apply the method all that is needed is an interface that forms the difference between the model state at the beginning and the end of a one year run and then solves a linear system of equations in order to apply a preconditioner. Solving the linear system of equations effectively inverts the circulation model’s annually averaged transport operator and thus provides an implicit treatment of the model’s slow advective-diffusive transport modes. The applicability of implicit solvers for reducing the computational costs of spinning up global biogeochemistry models initiated by Kwon and Primeau (2006, 2008) is greatly increased by our new solver because it yields equilibrium solutions to seasonally varying models.

In this study, we have explicitly constructed a preconditioner for the phosphorus cycle of an OCMIP-2 type ocean biogeochemistry model with two active tracers. When we started this project, our search in the literature on Krylov methods yielded little in terms of general guidance or specific examples for constructing “physics based” preconditioners for advective-diffusive systems. We hope that by reporting a specific example of an effective preconditioner for an ocean transport model we might stimulate further research on general methods for constructing efficient preconditioners. Our results here suggest that the preconditioner does not necessarily need to depend on any of the complex biogeochemical parameterizations in order to be effective. Furthermore, our proposed preconditioner is based on the inverse of a block-triangular matrix. This means that only the diagonal blocks of the form  $(\mathbf{A}) - \mathbf{D}$ , where  $\mathbf{D}$  is a diagonal matrix, need to be factored. This fact means that the computer memory and cpu time needed to form, factor, store and apply the preconditioner will scale linearly with the number of tracers in the biogeochemical model. We therefore expect that preconditioners such as ours, that are based on the annually averaged advection-diffusion transport operator, will continue to be effective in more complex biogeochemistry models as long as the biogeochemical dynamics of these models does not involve any intrinsic timescales comparable to or longer than the longest timescales associated with the circulation, i.e.  $\sim 1000$  years. Whether or not it becomes necessary to modify the preconditioner should the intrinsic dynamics of the biogeochemistry model include such long timescales is difficult to predict, but we expect that the Newton–Krylov solver will perform at least as well as the un-preconditioned case which achieves a non-trivial factor of 10 speedup compared to straight time-stepping.

We have also tested that the LU decomposition of blocks of the form  $(\mathbf{A}) - \mathbf{D}$  with the same sparsity pattern as would be obtained from our model with a  $1^\circ \times 1^\circ$  resolution with 24 levels can be computed using less than 32 GB of memory in under 1.5 h using one Intel(R) Xeon(R) 2.33 GHz cpu. Furthermore, this factorization needs to be computed only once. The backsolves needed when applying the preconditioner at the end of each year of simulation took less than 10 s. The cost of computing and applying the preconditioner is therefore substantially less than running the model for one year. These results suggest that our solution method is expected to substantially speed up models with a resolution comparable to what is currently used for global climate models.

Our solver should also be effective for finding solutions to models that include periods other than the seasonal cycle. In particular, our solver should be applicable with little or no modification to a model that includes a diurnal cycle in addition to the annual cycle. If leap years are not taken into account, the solver should apply without modification. If leap years are included then one would have to use the solver to find the fixed point of the map that propagates the model state forward for four years instead of just one. The effects of longer period phenomena that repeat after an integral number of years could also be included by applying the solver to the map that propagates the model state forward for  $n$  years where  $n$  is the number of years in the long period cycle. This would make it possible to include some representation of the effects of low-frequency variability on biogeochemical cycling. One could, for example, include a periodic representation of the four year component of El Niño Southern Oscillation (ENSO).

To illustrate the new solver’s efficiency we have used it in a parameter optimization study of the parameters controlling the cycling of DOP. Previous studies aimed at constraining  $\sigma$ , the fraction of new production allocated to DOP and  $1/\kappa$  and the mean lifetime of semi-labile DOP, have been conducted in the context of steady models without a seasonal cycle (Yamanaka and Tajika, 1997; Kwon and Primeau, 2006, 2008). In the present study, we have explicitly computed a cost-function that measures the difference between the model simulated  $\text{PO}_4$  and the monthly  $\text{PO}_4$  climatology from the World Ocean Atlas (2005). By doing so, we were able to demonstrate that the information in the seasonal variations of the  $\text{PO}_4$  observations do not significantly improve the resolution of  $\sigma$  and  $\kappa$ . This highlights the need for direct DOP and DOM observations in order to constrain these parameters.

In this study, we did not make use of an automatic minimization algorithm to optimize the model parameters, but there are no barriers to doing so. Automatic minimization algorithms can be used synergistically to gain additional computational savings. By using a lower tolerance on the Newton–Krylov solver in the early stages of the minimization and increasing the tolerance only

in the final stages the overall number of simulation years needed to find the optimal parameters can be further reduced. In the future we plan on using the new solver in combination with an automatic minimization routine to objectively optimize the large and increasing number of parameters in state-of-the-art ocean biogeochemistry models.

### Acknowledgement

Numerous conversations with Eun Young Kwon and Keith Moore are gratefully acknowledged. Funding for this research is provided by the National Science Foundation Grants OCE 02215156 and OCE 0623647.

### References

- Anderson, L.A., Sarmiento, J.L., 1995. Global ocean phosphate and oxygen simulations. *Global Biogeochemical Cycles* 9 (4), 621–636.
- Garcia, H.E., Locarnini, R.A., Boyer, T.P., Antonov, J.I., 2006. *World Ocean Atlas 2005*. In: Levitus, S. (Ed.), *Nutrients (phosphate, nitrate, silicate)*, vol. 4. NOAA Atlas NESDIS 64, U.S. Government Printing Office, Washington, D.C., 396pp.
- Golub, G.H., Van Loan, C.F., 1989. *Matrix Computations*, second ed. The Johns Hopkins University Press, Baltimore and London. 642pp.
- Holzer, M., Primeau, F., 2006. The diffusive ocean conveyor. *Geophysical Research Letters* 33, L14618. doi:10.1029/2006GL026232.
- Holzer, M., Primeau, F.W., 2008. The path-density distribution of oceanic surface-to-surface transport. *Journal of Geophysical Research* 113, C01018. doi:10.1029/2006JC003976.
- Kelley, C.T., 2003. Solving nonlinear equations with Newton's method. *SIAM Fundamentals of Algorithms*, 104.
- Khatiwala, S., Visbeck, M., Cane, M.A., 2005. Accelerated simulation of passive tracers in ocean circulation models. *Ocean Modelling* 9, 51–69.
- Kwon, E.Y., Primeau, F., 2006. Sensitivity and optimization study of a biogeochemistry ocean model using an implicit solver and in situ phosphate data. *Global Biogeochemical Cycles* 20, GB4009.
- Kwon, E.Y., Primeau, F., 2008. Optimization and Sensitivity of a Global Biogeochemistry Ocean Model using Combined In Situ DIC, Alkalinity and Phosphate Data. *Journal of Geophysical Research, Oceans*, in press.
- Najjar, R., Orr, J., 1998. Design of OCMIP-2 simulations of chlorofluorocarbons, the solubility pump and common biogeochemistry. <<http://www.ipsl.jussieu.fr/OCMIP/phase2/simulations/design.ps>>.
- Najjar, R.G., Jin, X., Louanchi, F., Aumont, O., Caldeira, K., Doney, S.C., Dutay, J.-C., Follows, M., Gruber, N., Joos, F., Lindsay, K., Maier-Reimer, E., Matear, R.J., Matsumoto, K., Monfray, P., Mouchet, A., Orr, J.C., Plattner, G.-K., Sarmiento, J.L., Schlitzer, R., Slater, R.D., Weirig, M.-F., Yamanaka, Y., Yool, A., 2007. Impact of circulation on export production, dissolved organic matter, and dissolved oxygen in the ocean: results from phase II of the Ocean Carbon-cycle Model Intercomparison Project (OCMIP-2). *Global Biogeochemical Cycles* 21, GB3007. doi:10.1029/2006GB002857.
- Pacanowski, R.C., Dixon, K., Rosati, A., 1993. The GFDL modular ocean model users guide. Tech. Rep. 2, GFDL Ocean Group, Princeton, NJ.
- Primeau, F., 2005. Characterizing transport between the surface mixed layer and the ocean interior with a forward and adjoint global ocean transport model. *Journal of Physical Oceanography* 35 (4), 545–564. doi:10.1175/JPO2699.1.
- Primeau, F., Holzer, M., 2006. The ocean's memory of the atmosphere: residence-time and ventilation-rate distributions of water masses. *Journal of Physical Oceanography* 36, 1439–1456.
- Pozrikidis, C., 1998. *Numerical Computation in Science and Engineering*. Oxford University Press. 627pp.
- Saad, Y., Schultz, M.H., 1986. GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal of Scientific and Statistical Computing* 7, 856–869.
- Stommel, H., 1979. Determination of watermass properties of water pumped down from the Ekman layer to the geostrophic flow below. *Proceedings of the National Academy Science US* 76, 3051–3055.
- Van Noorden, T.L., Verduyn Lunel, S.M., Bliëk, A., 2003. The efficient computation of periodic states of cyclically operated chemical processes. *IMA Journal of Applied Mathematics* 68, 149–166.
- Williams, R.G., Marshall, J.C., Spall, M.A., 1995. Does Stommel's mixed-layer 'demon' work? *Journal of Physical Oceanography* 25 (12), 3089–3102.
- Yamanaka, Y., Tajika, E., 1997. The role of dissolved organic matter in the marine biogeochemical cycle: Studies using an ocean biogeochemical general circulation model. *Global Biogeochemical Cycles* 11, 599–612.