Evolution of Algorithms for the Ocean Free Surface

V. Balaji
SGI/GFDL

ECMWF TeraComputing Workshop
19 November 1998

Overview

- The role of the ocean circulation in the climate system
- MOM: the GFDL Modular Ocean Model
- Algorithms for the barotropic mode
- Conclusions

The Ocean in the Climate System

Water has a high heat capacity and density and interacts with many chemical constituents of the atmosphere, including CO_2. The ocean therefore has a key role to play in many aspects of the variability of the climate system. Its persistent circulations redistribute heat and fresh water in ways that act as important controls on the climate of adjacent land masses. Exchanges at the ocean surface are a key component in the global cycle of CO_2 and other gases, and are a function of the dynamics. This plot thickens further when we consider feedbacks due to ocean ecosystems. Further, slow modes of the ocean involving deep water can retain climatic memory over a timescale of $\sim$ 1000 years.

Equations

Ocean models integrate the Navier-Stokes equations for a hydrostatic Boussinesq fluid on a sphere with complex bottom topography, and a free upper surface exchanging momentum, heat, moisture and other constituents with the atmosphere.

The Boussinesq approximation disregards variations in density in all terms other than the buoyancy, thus excluding sound and shock waves. The hydrostatic approximation treats the scale of horizontal motions as being much larger than vertical motions, thus excluding convection. (Non-hydrostatic ocean models have recently become available).

Equations

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot (\mathbf{u}\mathbf{u}) = -2\mathbf{\Omega} \times \mathbf{u} - \frac{\nabla p}{\rho} + \nabla \cdot (K_m \nabla \mathbf{u}) + \mathbf{F} \tag{1}$$

$$\frac{\partial p}{\partial z} = -g\rho \tag{2}$$

$$\frac{\partial w}{\partial z} = -\nabla \cdot \mathbf{u} \tag{3}$$

$$\frac{\partial \theta}{\partial t} + \nabla \cdot (\mathbf{u}\theta) = F_\theta \tag{4}$$

$$\frac{\partial s}{\partial t} + \nabla \cdot (\mathbf{u}s) = F_s \tag{5}$$

$$\rho = \rho(\theta, s, z) \tag{6}$$

The Hydrostatic Approximation

The vertical acceleration equation would give:

$$\frac{\partial \rho w}{\partial t} + \nabla \cdot (\rho \mathbf{u} w) = -\frac{\partial p}{\partial z} - g\rho \tag{7}$$

At large enough scales, the LHS term is a small difference of much larger terms on the RHS. This equation may no longer be the best one for extracting w. We instead set $w = 0$ here and extract w from the continuity equation:

$$\frac{\partial w}{\partial z} = -\nabla \cdot \mathbf{u} \tag{8}$$

Balance models

The hydrostatic treatment is the simplest example of the use of *balance models* to reduce the number of prognostic equations. Provided we have a theory that posits a certain balance between the model unknowns, we can use balance models to extract quantities when the number of unknowns exceeds the number of prognostic equations.

Examples are the use of the quasi-geostrophic approximation, geostrophic momentum approximation, or gradient wind balance, with the potential vorticity as the carrier of dynamic information.

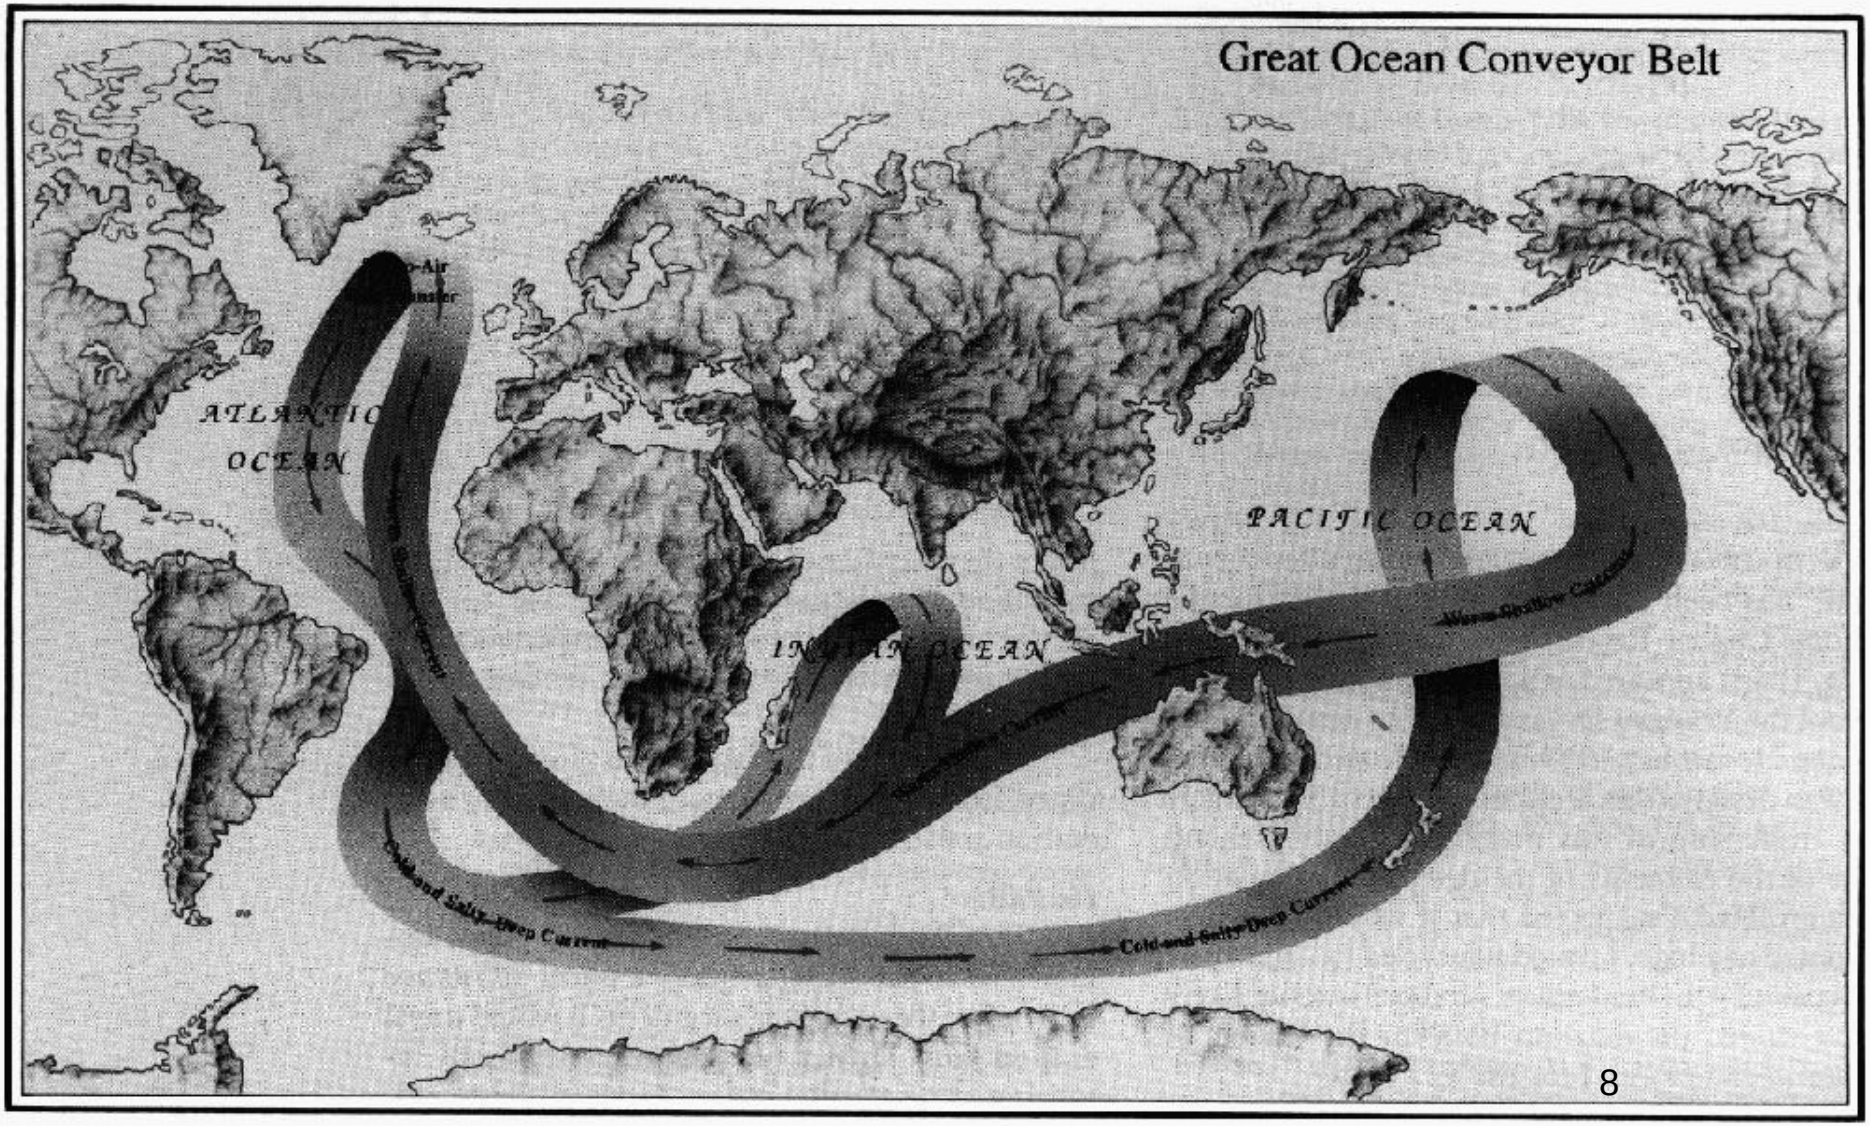

Fig. 1: The great ocean conveyor logo (Broecker, 1987). (Illustration by Joe Le Monnier, Natural History Magazine.)

Large scale motions are those governed by a rough balance between the horizontal pressure gradients and rotational inertia; motions where this balance is violated are called *eddies*. A rough scale for this separation is given by the *first baroclinic deformation radius*, on the order of 10 km in high latitudes rising toward the tropics.

Leading-order dynamics can be eddy-dominated in certain regimes, such as the Antarctic Circumpolar Current. Therefore we would like to model ocean circulations such as the ACC at eddy-resolving scales. An eddy-resolving model of the southern hemisphere would require $1080 \times 373 \times 46 \approx 18 \times 10^6$ points.

The GFDL Modular Ocean Model

MOM is a finite-difference ocean model using the equations described above. It is intended to be a community model, and has been used in studies spanning a wide range of time and space scales, from short-term studies of channel flows to coupled climate integrations on 1000-year scales. A wide range of physics packages is available for use at various resolutions.

`http://www.gfdl.gov/~kd/MOMwebpages/MOMWWW.html`

Internal and external modes

Solutions to the primitive equations can be written as a superposition over vertical eigenmodes. The 0th order, or depth-averaged, mode is referred to as the *barotropic mode* and the rest as *baroclinic modes*. For a weakly compressible fluid, the barotropic mode corresponds to elevations of the sea surface, and is also called the *external mode*.

The phase speed of the external mode is of the order of 200 m/sec for typical ocean depths, and is 2 orders of magnitude faster than the fastest internal mode. There is thus a strong incentive to separate internal and external modes in a numerical formulation.

Separate the momentum equations into a depth-averaged part and a deviation therefrom:

$$\mathbf{u} = \overline{\mathbf{u}} + \mathbf{u}' \tag{9}$$

$$\overline{\mathbf{u}} = \frac{1}{H+\eta}\int_{-H}^{\eta} \mathbf{u}\,dh \tag{10}$$

The equations for $\mathbf{u}'$ have the standard from and are solved on the baroclinic timestep. Different methods may be employed to to solve for $\overline{\mathbf{u}}$.

Can we filter the external mode entirely?

Since the external mode is associated with vertical displacements of a full water column, it can be filtered entirely through an assumption of a rigid lid, $\eta \equiv 0, w(z = 0) \equiv 0$. This amounts to assigning the external gravity waves an infinite phase velocity, and thus equilibrating them instantly to the surface pressure. There are several problems with this approach.

- As with any filtering assumption, the corresponding initial value problem is replaced with a boundary value problem, in this case to recover $\overline{\mathbf{u}}$ from a surface pressure or stream function. Boundary value problems are generally messy to parallelize, especially for the ocean surface where boundaries are irregular.

- In the equatorial region, the timescales associated with gravity waves and Rossby waves are comparable, so that it may not be correct to filter gravity waves.

- The fresh-water flux through the surface is given by the upper kinematic boundary condition:

$$\frac{\partial \eta}{\partial t} = w + Q_w - \mathbf{u} \cdot \nabla \eta \tag{11}$$

  The rigid lid precludes such a flux. Fresh-water fluxes must be represented as a *virtual salt flux*, promoting a host of other problems. Relaxing the assumption $w = 0$ adds a second elliptic equation to solve.

Implicit free surface

Next consider removing the rigid lid approximation, allowing $\eta \neq 0$. Since solving the external mode on the same timestep as the internal modes would violate CFL, use implicit time-stepping, which is unconditionally stable.

Implicit schemes also involve solving an elliptic equation. In MOM the iterative *conjugate gradient* method is used. There are no irregular boundaries here, simplifying the problem.

Data dependencies in the conjugate gradient method

Simple relaxation schemes involve nearest-neighbour communication, but converge slowly. The conjugate gradient method accelerates this by computing at each step, the optimal vector in phase space along which to converge.

Unfortunately, computing the direction involves a global sum.

Deficiencies in the implicit free surface

- Implicit methods are quite diffusive. Problems requiring an accurate (rather than merely numerically stable) solution of the free surface (tides) may be poorly served.

- There is a poorly understood interaction of the implicit free surface with the polar filter.

- “Checkerboarding” of the Laplacian operator.

- Global communication at each step in the iteration means that the method scales poorly.

- Non-iterative elliptic solvers, such as multi-grid methods, are not load-balanced for parallel architectures.

Explicit free surface

The final approach considered here is to solve the surface mode explicitly, but on a vastly shorter timestep. This is very simple to formulate: all contributions from the baroclinic modes are assumed to be slowly varying, and thus kept constant over a baroclinic timestep. Turbulent momentum exchange over this period may also safely be kept constant. The problem reduces to a simple stepping forward in time of $(\overline{\mathbf{u}}, \eta)$. Communication is nearest-neighbour only.

Advantages of the explicit free surface

- It is flexible: the time-stepping scheme can be tailored to the problem. Further the phase speed of the external gravity wave can be doctored (“reduced gravity”).

- The timestep is a fixed factor smaller than the baroclinic timestep, irrespective of horizontal resolution.

- Scaling behaviour on parallel systems is better than the implicit and rigid lid schemes.

- Straightforward to implement in generalized co-ordinates, unstructured grids.

Scaling behaviour of MOM on the T3E-900

For a 1.5° model (242×130):

- Implicit: 68 Mflops/PE barotropic solver, 48 Mflops/PE baroclinic solver and tracers, 55 Mflops/PE total. Scaling 60% at 4 rows/PE.

- Explicit: 80 Mflops/PE barotropic solver, 48 Mflops/PE baroclinic solver and tracers, 63 Mflops/PE total. Scaling 80% at 4 rows/PE.

Conclusions

Low order models benefit from being formulated in terms of a series of balance assumptions that reduce the number of prognostic equations. In the limit, atmospheric and oceanic dynamics could in principle be formulated in terms of a single prognostic variable, the *potential vorticity*, and a balance model that allows us to recover the mass and momentum fields from it.

As we move to higher resolutions, it becomes less easy to justify balance models, and models tend to solve more independent prognostic equations.

Happily, these are also the algorithms that lend themselves best to parallel formulation.

In short...

Nature does not vectorize, it parallelizes!