



HAL
open science

A tree-based solver for adaptive ocean modelling

S. Popinet, G. J. Rickard

► **To cite this version:**

S. Popinet, G. J. Rickard. A tree-based solver for adaptive ocean modelling. *Ocean Modelling*, 2007, 16 (3-4), pp.224-249. 10.1016/j.ocemod.2006.10.002 . hal-01445434

HAL Id: hal-01445434

<https://hal.science/hal-01445434>

Submitted on 7 Feb 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A tree-based solver for adaptive ocean modelling

Stéphane Popinet and Graham Rickard*
National Institute of Water and Atmospheric research
PO Box 14-901 Kilbirnie, Wellington, New Zealand

October 13, 2006

Abstract

The development of an adaptive (in space *and* time) ocean model from an existing adaptive finite-volume Navier-Stokes model is described. A flexible and efficient quadtree spatial discretisation is used which requires collocation of all variables (i.e. an A-grid discretisation). We demonstrate that the use of an approximate projection method allows for implicit damping of instabilities generally associated with the A-grid, at the expense of a relatively small amount of numerical energy dissipation, while accurately preserving dispersive properties and geostrophic balance. The finite-volume formulation also maintains second-order spatial accuracy at all solid boundaries. Test cases demonstrate the efficacy of the adaptive ocean model, and the advantages it has in terms of efficient representation of multi-scale behaviour within a single model. The model is freely available as open-source code.

1 Introduction

This paper details the extension of an adaptive, finite-volume, three-dimensional, incompressible, Navier-Stokes fluid solver (the Gerris Flow Solver, Popinet (2006, 2003)) into a dynamical core able to model geophysical fluid flows. Here the focus will be on ocean modelling, but the core can be extended to the atmosphere as well (Marshall, Adcroft, Campin, Hill, and White, 2004). The ability of the model to adapt in space and time constrains the structure of the solver. Nevertheless, the tests described here show that many of the fundamental requirements of a geophysical solver can be more than adequately met by this novel adaptive solver.

Why is having adaptivity relevant? In nature, the global dynamics of oceanic or atmospheric flows are controlled by processes occurring on widely different spatial scales. In the ocean, narrow straits control the exchanges between large ocean basins, energetic but thin western boundary currents define latitudinal heat transport and a few geographically small areas control much of the deep ocean circulation through deep water formation. Getting to even smaller spatial scales, the dynamics of near-shore mixing, internal tides, sediments transport and freshwater input is of crucial importance to ecosystems and human activities but depends directly on larger-scale offshore processes (and conversely).

The regular Cartesian spatial discretisations traditionally used in geophysical models are easy to analyse mathematically and easy to implement but lead to a constant spatial resolution (or quasi-constant depending on the type of spherical-to-planar projection) independent from any physics-related scale dependence. While large increases in computational resources have allowed ever finer spatial resolutions to be reached using traditional models, the computational efficiency of the numerical schemes is still a very important issue for many applications. A logical way to decrease the computational cost is to discretise the solution using a more optimal spatial distribution of the discretisation elements. Simple applications of this idea to Cartesian meshes lead to grid nesting or stretched grids using curvilinear coordinates. In the wider field of computational fluid dynamics, finite-volume or more rarely finite-element discretisations on unstructured meshes provide almost full flexibility in the choice of local spatial resolution.

*(s.popinet|g.rickard)@niwa.co.nz

It is important to make the distinction between methods where the spatial resolution can vary but is fixed in time (and usually defined a priori by the user) and *adaptive* methods where the resolution is continuously adjusted according to the properties of the solution. Adaptive methods do not require a priori assumptions about the solution but instead use error estimates to ensure an optimal distribution of discretisation elements. In the context of geophysical fluid flows, adaptive methods have the potential to greatly improve the description of spatially restricted but time-varying processes such as: strong density gradients associated with oceanic or atmospheric fronts, eddying of western boundary currents, tidally-driven small-scale mixing/transport near coastlines etc...

A number of recent initiatives have sought to apply the more flexible techniques developed for computational fluid dynamics to geophysical fluid modelling, in recognition of the shortcomings of regular Cartesian discretisation (as well as ad hoc grid nesting). The focus has been mainly on using finite-volume or finite-element methods on statically-refined unstructured grids for barotropic coastal ocean modelling (Chen, Liu, and Beardsley, 2003; Walters, 1986; Walters, Goring, and Bell, 2001). Unstructured grids allow an accurate description of complex coastlines and the coastal ocean is a natural extension of the river flows for which these models were initially developed (Walters and Casulli, 1998). The extension of finite-element or finite-volume models to large-scale oceanic flows is non-trivial because of the need to guarantee specific properties of the large-scale flows. Geostrophic balance must be represented accurately at the discrete level and – more generally – energy must be better conserved because of the low energy dissipation in the open ocean. Cartesian methods have benefited from decades of development and describe the large-scale ocean dynamics accurately. Optimal schemes for finite-volume or finite-element formulations on unstructured grids remain a subject of active research (Chen, Liu, and Beardsley, 2003; Le Roux, Staniforth, and Lin, 1997; Le Roux, Lin, and Staniforth, 2000; Walters and Barragy, 1997).

Different adaptive methods have been applied to atmospheric and ocean modelling. For instance Blayo and Debreu (1999), Bacon, Ahmad, Boybeyi, Dunn, Hall, Lee, Sarma, Turner, Waight, Young, and Zack (2000) and Giraldo (2000) have studied adaptive techniques, and indeed Bacon et al. (2000) use their system OMEGA (The Operational Multiscale Environment Model with Grid Adaptivity) for real time hazard prediction. More recent adaptive developments include, e.g., Power, Piggott, Fang, Gorman, Pain, Marshall, Goddard, and Navon (2006) (and references therein to their adaptive finite element ocean model), Herrnstein, Wickett, and Rodrigue (2005), Barros and Garcia (2004), indicating the on-going level of activity and interest for the application of adaptive methods in geophysical modelling.

Recently, Popinet (Popinet, 2003, 2006) showed that a hierarchical tree-based spatial discretisation can efficiently provide the flexibility of fully unstructured grids while preserving the simplicity and conceptual framework of Cartesian grids. This work also showed that adapting the discretisation dynamically at each timestep was practical and only represented a small fraction of the total computational cost. The adaptive solver led to several order of magnitude savings in computational time for complex three-dimensional problems.

In this article, we propose to extend the work of Popinet (2003) to obtain an adaptive method for large-scale oceanic flows with good mass and energy conservation properties and accurate geostrophic balance representation. The overall numerical scheme will be described but we refer the reader interested in the details of the quadtree/octree discretisation and multilevel solver implementation to Popinet (2003). We will also try to reproduce previously published test cases designed to assess several properties of the numerical scheme: representation of inertia-gravity waves, geostrophic balance and bathymetry. An application illustrating the potential of the adaptive method for practical problems will also be presented.

In the following section the system under consideration is defined. In sections 3 and 4 the spatial and temporal discretisations, respectively, consistent with the Gerris flow solver are described, followed by section 5 detailing our application of the adaptive aspects implied by Gerris for oceanic flows. Section 6 presents a series of applications to two- and three-dimensional test problems, and finishes with an example of adaptivity in the context of multi-scale flow evolution (in this case that of tidally driven flow in Cook Strait, New Zealand). Conclusions are presented in section 7.

2 Problem definition

Our initial goal is to solve the Boussinesq, incompressible, hydrostatic equations of oceanic motion in a Cartesian reference frame. We will also make a “shallow water” approximation and neglect the coupling terms arising from the vertical motion (Marshall, Hill, Perelman, and Adcroft, 1997b). The equations can then be written as

$$\partial_t \mathbf{U} + \mathbf{U} \cdot \nabla \mathbf{U} + w \partial_z \mathbf{U} = \mathbf{B} \mathbf{U} + \mathbf{S} - \nabla p - g \nabla \eta$$

for the horizontal momentum,

$$\partial_t H + \nabla \cdot H \bar{\mathbf{U}} = 0$$

for continuity,

$$\partial_t T + \mathbf{U} \cdot \nabla T + w \partial_z T = S_T$$

for temperature,

$$\partial_t S + \mathbf{U} \cdot \nabla S + w \partial_z S = S_S$$

for salinity, with $\mathbf{U} \equiv (u, v)$ the horizontal components of the velocity,

$$\mathbf{B} \equiv \begin{pmatrix} 0 & f \\ -f & 0 \end{pmatrix}$$

where f is the Coriolis parameter, $H \equiv h + \eta$ where h is the depth of the ocean and η the surface elevation. \mathbf{S} , S_T and S_S contain the forcing and dissipation terms. $\bar{\mathbf{U}}$ is the depth-averaged fluid velocity. The Boussinesq hydrostatic pressure p is obtained as

$$p = \int_0^z -g \frac{\delta \rho}{\rho_0},$$

where $\delta \rho$ is the deviation from the reference density ρ_0 obtained using the equation of state $\rho \equiv \rho(T, S, p)$. The vertical component w of the velocity is obtained by vertical integration of the continuity equation as

$$w = - \int_0^z \nabla \cdot \mathbf{U} \quad (1)$$

If we assume that $\eta \ll h$ the continuity equation can be linearised as

$$\partial_t \eta + \nabla \cdot h \bar{\mathbf{U}} = 0$$

These approximations are representative of current ocean models designed for large-scale flows (as noted by Marshall, Hill, Perelman, and Adcroft (1997b) they are often incorrectly referred to as the “primitive equations” of oceanic motion). The hydrostatic assumption breaks down at scales corresponding to convective overturning in the ocean (≈ 1 km) (Marshall, Hill, Perelman, and Adcroft, 1997b). The linearised free-surface approximation can cause problems for global tracer conservation (Roullet and Madec, 2000; Campin, Adcroft, Hill, and Marshall, 2004) and becomes inconsistent when the amplitude of free-surface motion is comparable to the depth (e.g. tidal perturbations near coastlines). None of these approximations is a necessary prerequisite for the approach described in this article but they make the initial implementation simpler while preserving all the important physical properties of the oceanic system. In the future we expect to be able to follow the work of Marshall et al. to relax both the hydrostatic and linear free-surface assumptions (Marshall, Adcroft, Hill, Perelman, and Heisey, 1997a). Our starting point being a full 3D Navier–Stokes solver should make this easier.

3 Spatial discretisation

In this article, we would like to explore how well existing geophysical discretisation techniques fit within the tree-based approach presented in Popinet (2003). The tree-based discretisation was shown to be a good compromise in term of flexibility and computational efficiency. By tree-based we mean either a quadtree structure in two dimensions or an octree structure in three dimensions. Here, the discussion focuses on the quadtree structure, but all properties map equally well into the third dimension.

[Figure 1 about here.]

In the quadtree discretisation, individual cells are classical Cartesian square finite volumes but they are organised hierarchically rather than as an array. Each cell can be the parent of up to four children cells; this is the process at the heart of adaptivity. The resulting recursive structure is a tree (Figure 1) and is easily represented using memory pointers. While this discretisation is not as flexible as a fully-unstructured mesh, it does not require additional geometric information and as importantly preserves the simplicity of the discretisation given by orthogonal Cartesian discretisation volumes.

[Figure 2 about here.]

A finite-volume approach (where the solution is averaged on the control volume defined by each cell) imposes an important restriction on the type of discretisation permitted on a quadtree. Let us consider a C-grid discretisation where the quadtree hierarchy is centred on the pressure (Figure 2). When refining the grid, the new control volumes for each velocity component do not necessarily delimit a full subset of the coarser control volumes (the single-hashed area in Figure 2.b is covered by the coarser control volume but not by any of the finer control volumes). It is thus not clear how to consistently and conservatively define the refined values of the momentum components as well as fluxes at coarse–fine boundaries. If one relaxes the finite-volume constraint and uses finite differences instead this problem does not occur but the important conservation properties of finite-volume schemes are lost. A solution would be to maintain independent staggered quadtree hierarchies but this greatly complicates the computation of terms coupling variables defined on different hierarchies. A collocated discretisation (Arakawa A-grid) eliminates this specific problem and is used in the present work in order to allow the use of adaptivity.

Collocated A-grid discretisations are usually avoided in ocean models because of the strong computational modes induced by pressure–velocity decoupling as well as the associated degradation of the dispersive properties of gravity waves. Ocean models traditionally use a staggered Arakawa C-grid or B-grid. C-grids are preferred for finer-scale ocean modelling due to their superior discretisation of the dispersion properties of gravity waves. B-grids show superior discretisation of large-scale inertial waves because they use collocated definitions of the velocity components which allows the Coriolis terms to be computed without any need for interpolation. It is important to note that both B-grids and C-grids suffer from spatial computational modes. In the case of C-grids the computational mode arise from the need to interpolate the values of each component of the velocity to the location of the other component in order to compute the Coriolis term. The resulting computational mode can be problematic for relatively coarse-resolution global ocean modelling (Adcroft, Hill, and Marshall, 1999). B-grids suffer from a pressure computational mode caused by the collocation of the velocity components. Neither discretisation is thus entirely mode-free and filters are often used to keep the computational modes in check. More complex discretisations have been proposed which attempt to solve this problem (C-D grid of Adcroft, Hill, and Marshall (1999)).

Several notable exceptions that use an A-grid have relied on higher-order discretisations (Dietrich, 1997) or other forms of built-in filtering (Kar, 1999) in order to control computational modes. More generally, the pressure–velocity decoupling problem of collocated schemes had also been recognised early in the history of incompressible flow solvers (Harlow and Welch, 1965) and more recently solutions have been developed (Lai, Bell, and Colella, 1993; Rider, 1995; Almgren, Bell, and Crutchfield, 2001). Building on this previous work, we propose in section 5 a semi-implicit barotropic solver on a collocated A-grid.

Vertical discretisation has been the topic of a large number of recent studies (Bleck, 2002; Adcroft and Campin, 2004; Song and Hou, 2006). While our work relies on horizontal collocation, it does not impose restrictions on the vertical discretisation. In section 5.4 we describe a z -level vertical discretisation with an accurate finite-volume description of solid boundaries.

4 Temporal discretisation

The time discretisation of the equations of motion can theoretically be chosen independently from the spatial discretisation although in practice we will see that some combinations of time-discretisations are much easier to implement using specific spatial schemes.

Advancing the equations of motion in time requires the solution of three independent sub-problems:

1. Pressure update through the barotropic equation.
2. Computation of advection terms in the momentum equation.
3. Computation of Coriolis terms in the momentum equation.

Time-explicit discretisation of the barotropic equation is simple but stability requires the resolution of the fast external gravity waves which imposes a strong restriction on the maximum timestep. This problem has been addressed in two ways. The barotropic equation can be advanced independently from the momentum equation using a stable, shorter timestep (split-explicit schemes (Shchepetkin and McWilliams, 2005)). Changing the discretisation to a fully- or semi-implicit scheme removes the constraint on the timestep altogether (Dukowicz and Smith, 1994) and allows the use of a single timestep for the whole scheme.

Both schemes have advantages and drawbacks. An implicit discretisation requires the resolution of a spatially-coupled linear system at every timestep. This can be expensive if naive algorithms are used, however modern techniques such as multigrid or pre-conditioned conjugate gradients can be very efficient for this type of problem (arguably more efficient than the multiple timesteps of a split-explicit solution method). Implicit schemes are stable but do not guarantee a physically-meaningful aliasing of the unresolved fast gravity waves. While a split-explicit technique should describe correctly all the external gravity wavelengths, the synchronisation step necessary to reconcile the fast pressure solution and the slow momentum solution adds complexity (Shchepetkin and McWilliams, 2005). As a result, for both schemes the impact of the coupling of the fast gravity waves with the slow momentum equations is problem-dependent and is unclear in most cases. In this work we use a semi-implicit discretisation with a multigrid solver for the linear system.

A common way to compute the advection term of the momentum equation is to use a leap-frog scheme. This scheme is simple to implement and has the theoretical advantage of providing a purely non-diffusive solution. This solution is not guaranteed to be oscillations-free (monotonic) however and the leap-frog discretisation introduces a temporal computational mode. Some form of explicit filtering (added diffusion) or implicit filtering (e.g. damping provided by the implicit discretisation of the barotropic equation (Dukowicz and Smith, 1994)) is used in practice, thus cancelling the non-diffusive property of the initial scheme. More recently, ocean modellers have started to use more sophisticated advection schemes initially developed for the solution of compressible flows. These schemes are stable and non-oscillatory even in the presence of strong discontinuities. These properties are obtained through a judicious construction of built-in diffusivity. They are also more computationally expensive than simpler schemes. In our previous work (Popinet, 2003) we have successfully used the incompressible version of a two-time-levels, predictor-corrector, unsplit scheme developed by Bell, Collella and Glaz (BCG) (Bell, Colella, and Glaz, 1989) and this is what we make use of in this work.

The Coriolis term can also be discretised using a leap-frog scheme. An alternative is to use an Adams-Bashforth (AB) discretisation which can be adapted to any underlying time-discretisation (leap-frog, predictor-corrector (Marshall, Adcroft, Hill, Perelman, and Heisey, 1997a) or semi-implicit). The AB scheme uses the previous two time-levels to extrapolate the value of the Coriolis term at the time-level required to ensure second-order accuracy in time. In the case of an AB discretisation of the Coriolis terms coupled with the barotropic equation, the stability criterion is non-trivial. We present in Appendix A a detailed analysis of the stability properties of an AB discretisation of the Coriolis terms coupled with a semi-implicit discretisation of the barotropic equation. The analysis is difficult due to the third time-level introduced by the AB scheme. We show that the complex behaviour of inertia-gravity waves gives rise to a stability criterion which can be quite restrictive for this combination.

A simpler and more easily analysed alternative is to use a semi-implicit discretisation of the Coriolis terms, which fits well with a semi-implicit treatment of the barotropic equation. The resulting two-time-levels scheme is easy to analyse and is unconditionally stable. If the velocity components are collocated (e.g. A- or B-grid) the semi-implicit discretisation leads to a compact linear system which can be inverted analytically (see section 5). In the case of a staggered discretisation of the velocity components (C-grid), the semi-implicit discretisation of the Coriolis terms leads to a globally-coupled

sparse linear system which needs to be solved numerically at each timestep. This greatly complicates the solution procedure and is the main reason why semi-implicit discretisations of Coriolis terms are avoided on C-grids.

5 A quadtree, adaptive, semi-implicit scheme

Taking into account the previous discussion and keeping in mind that our starting point is the adaptive incompressible solver presented in Popinet (2003) we chose to use a two-time-levels, semi-implicit discretisation of both the barotropic and Coriolis terms. If the components of the velocity are collocated, the discrete time evolution of the barotropic momentum can be written

$$\frac{\mathbf{U}^{n+1} - \mathbf{U}^n}{\Delta t} = \mathbf{B}[(1 - \alpha)\mathbf{U}^n + \alpha\mathbf{U}^{n+1}] - g\nabla[(1 - \alpha)\eta^n + \alpha\eta^{n+1}] + \mathbf{S}^{n+\frac{1}{2}}, \quad (2)$$

where Δt is the timestep, n is the time level and α controls the degree of implicitness. All the other source terms (advection terms, viscous diffusion, forcing) are contained in $\mathbf{S}^{n+1/2}$. Similarly the continuity equation can be written

$$\frac{\eta^{n+1} - \eta^n}{\Delta t} + \nabla \cdot h[(1 - \theta)\mathbf{U}^n + \theta\mathbf{U}^{n+1}] = 0, \quad (3)$$

where θ is another free parameter. If α and θ are set to $1/2$, the discretisation is a second-order accurate Crank-Nicolson scheme. It is also unconditionally stable and conserves the total energy.

When expressed in this form, the semi-implicit schemes for the momentum and the continuity equations are coupled through the Coriolis terms. A simpler uncoupled second-order accurate approximation can be obtained through a splitting technique. If we pose

$$\hat{\mathbf{U}} \equiv \mathbf{U}^{n+1} + \Delta t g \nabla [\alpha' \eta^{n+1} - (1 - \alpha') \eta^n],$$

where α' is another control parameter, we can rewrite (2) as

$$(\mathbf{I} - \alpha \Delta t \mathbf{B}) \hat{\mathbf{U}} + \alpha \Delta t^2 g \mathbf{B} \nabla [\alpha' \eta^{n+1} - (1 - \alpha') \eta^n] = [\mathbf{I} + (1 - \alpha) \Delta t \mathbf{B}] \mathbf{U}^n - 2(1 - \alpha') \Delta t g \nabla \eta^n + \Delta t \mathbf{S}^{n+\frac{1}{2}},$$

where \mathbf{I} is the identity matrix. The second term on the left-hand side is clearly second-order in time, moreover if α' is set to $1/2$ the prefactor is only the (small) relative change in surface-height gradient. This leads to the second-order accurate, explicit approximation for $\hat{\mathbf{U}}$

$$(\mathbf{I} - \alpha \Delta t \mathbf{B}) \hat{\mathbf{U}} = [\mathbf{I} + (1 - \alpha) \Delta t \mathbf{B}] \mathbf{U}^n - 2(1 - \alpha') \Delta t g \nabla \eta^n + \Delta t \mathbf{S}^{n+\frac{1}{2}} + \mathcal{O}(\Delta t^2). \quad (4)$$

The resulting 2×2 system is easily inverted analytically. If we now pose

$$\mathbf{U}^* \equiv \hat{\mathbf{U}} + (1 - \alpha') \Delta t g \nabla \eta^n, \quad (5)$$

equation (3) can be rewritten

$$\eta^{n+1} - \theta \alpha' \Delta t^2 g \nabla \cdot h \nabla \eta^{n+1} = \eta^n - \Delta t \nabla \cdot h [(1 - \theta) \mathbf{U}^n + \theta \mathbf{U}^*]. \quad (6)$$

Equations (4), (5) and (6) are now entirely decoupled. The full solution procedure can be summarised as follows:

1. Compute $\hat{\mathbf{U}}$ using \mathbf{U}^n , η^n , $\mathbf{S}^{n+\frac{1}{2}}$ and (4)
2. Compute \mathbf{U}^* using $\hat{\mathbf{U}}$, η^n and (5)
3. Compute η^{n+1} using \mathbf{U}^n , \mathbf{U}^* , η^n and (6)
4. Compute $\mathbf{U}^{n+1} = \mathbf{U}^* - \alpha' \Delta t g \nabla \eta^{n+1}$

The overall scheme is very similar to the projection method (also known as “pressure-correction” method) often used to solve incompressible flow problems (Chorin, 1968; Peyret and Taylor, 1985) or equivalently oceans with a rigid-lid approximation (Marshall, Adcroft, Hill, Perelman, and Heisey, 1997a). In a first phase a provisional velocity field \mathbf{U}^* is obtained by adding to \mathbf{U}^n all the (explicit) source terms of the momentum equation with the exception of the pressure-gradient term. In a second phase the pressure compensating the eventual divergence of \mathbf{U}^* is obtained through the solution of a Poisson-like equation and the corresponding pressure gradient term is added to \mathbf{U}^* in order to obtain a non-divergent \mathbf{U}^{n+1} .

The semi-implicit scheme presented here differs slightly from this outline because a linear free-surface approximation is made rather than a stricter rigid-lid approximation. The free-surface equation (6) (equivalent to the pressure equation for incompressible flows) is then of Helmholtz-type ($x + a\nabla^2 x = \dots$) rather than Poisson-type ($\nabla^2 x = \dots$). In both cases, the spatial coupling induced by the Laplacian operator results in a sparse linear system. In the case of incompressible flows, the solution of this system largely dominates the computational cost of the overall scheme. Much effort has thus gone into efficient solution procedures for this type of systems. Multigrid techniques in particular can achieve an optimal $\mathcal{O}(n)$ asymptotic computational cost (where n is the number of unknowns) (Brandt, 1982; Bramble, 1993).

As noted by ocean modellers (Dukowicz and Smith, 1994), the linear system described by equation (6) is easier to solve than the Poisson problem of a rigid-lid approximation. This is due to the η^{n+1} term on the left-hand side of (6) which increases the weight of the diagonal when Δt is small enough and thus leads to a better-conditioned linear system. If Δt is large, the added diagonal term becomes small compared to the Laplacian and the solution converges toward the incompressible Poisson solution, with a corresponding stiffness of the linear system. We expect that the multilevel quadtree Poisson solver we developed in Popinet (2003) will perform well in both limits.

5.1 Exact and approximate projections

Up to this point we have only assumed that the velocity components are collocated. The discrete gradient and divergence operators appearing in (4), (5) and (6) will depend on the details of the spatial discretisation. To simplify the argument we will consider a one-dimensional problem with $g = h = 1$. The discrete divergence operator can then be expressed as

$$\nabla_i \cdot u = \frac{u_{i+1/2} - u_{i-1/2}}{\Delta x},$$

which is also a consistent finite-volume discretisation (when $u_{i+1/2}$ and $u_{i-1/2}$ are interpreted as fluxes through the boundaries of the control volume centred on u_i). For a collocated scheme, u is defined only at integer locations. The half-integer values can be approximated to second-order using a linear interpolation from cell centres. This gives the second-order accurate discretisation

$$\nabla_i \cdot u = \frac{u_{i+1} - u_{i-1}}{2\Delta x}. \quad (7)$$

Similarly the discrete gradient operator for the collocated surface-height can be expressed as

$$\nabla_i \eta = \frac{\eta_{i+1} - \eta_{i-1}}{2\Delta x}. \quad (8)$$

The discrete Laplacian appearing in (6) results from the combination of these operators as

$$(\nabla \cdot \nabla)_i \eta \equiv \nabla_i \cdot (\nabla_i \eta) = \frac{\eta_{i+2} - 2\eta_i + \eta_{i-2}}{4\Delta x^2}. \quad (9)$$

The corresponding stencil is twice as wide as the initial discretisation and results in two discrete solutions for odd and even η_i s only coupled through the boundary conditions. This is another way to express the existence of a computational mode for a collocated discretisation.

Aside from the inherent instability of a scheme containing computational modes (Lai, Bell, and Colella, 1993; Kar, 1999), the wide stencil of the Laplacian greatly complicates the solution of (6)

when a non-uniform resolution is used (Popinet, 2003). The standard compact Laplacian

$$\nabla_i^2 \eta \equiv \frac{\eta_{i+1} - 2\eta_i + \eta_{i-1}}{2\Delta x^2}$$

is also second-order and could be used instead in (6), however this would mean that the discrete Laplacian used is not anymore the convolution of the discrete divergence operator (7) and discrete gradient operator (8). The difference is $\mathcal{O}(\Delta x^2)$ and can be written

$$\mathcal{F}_i(\eta) \equiv (\nabla \cdot \nabla)_i \eta - \nabla_i^2 \eta = \frac{\eta_{i+2} - 2\eta_{i+1} + 2\eta_i - 2\eta_{i-1} + \eta_{i-2}}{4\Delta x^2}.$$

Equation (6) then becomes

$$\eta_i^{n+1} - \theta \alpha' \Delta t^2 \nabla_i^2 \eta^{n+1} - \theta \alpha' \Delta t^2 \mathcal{F}_i(\eta^{n+1}) = \eta_i^n - \Delta t \nabla_i \cdot [(1 - \theta)u^n + \theta u^*].$$

Replacing the wide-stencil Laplacian with the compact-stencil Laplacian would thus be equivalent to adding a $\theta \alpha' \Delta t^2 \mathcal{F}_i(\eta^{n+1})$ term to the left-hand-side of equation (6). This term can be seen as a spatial filter on η .

The initial wide-stencil Laplacian is constructed so that the discretisation of the continuity equation (6) guarantees volume conservation i.e. the discrete gradient of η^{n+1} applied to \mathbf{U}^* guarantees the discrete non-divergence (in the finite volume sense described in section 5.2) of \mathbf{U}^{n+1} . This is the definition of an *exact* projection (Peyret and Taylor, 1985). Replacing the wide-stencil Laplacian with a compact stencil guarantees volume conservation only in an asymptotic manner (the error on the discrete divergence is of order Δx^2). This defines an *approximate* projection (Rider, 1995; Almgren, Bell, and Crutchfield, 2001). An important consequence is that the total energy is not exactly preserved anymore. This is clearly apparent in the form of the \mathcal{F}_i error which acts as a dissipative term on the free-surface height (and thus on the potential energy). In the case of incompressible flows, this dissipative term has been shown to have the beneficial side-effect of almost eliminating the pressure oscillations associated with a collocated discretisation (Rider, 1995). Moreover, the relaxation of strict energy conservation has only a small effect on the quality of the solutions obtained and the transport of momentum and tracers can still be formulated as to guarantee exact conservation (Almgren, Bell, Colella, Howell, and Welcome, 1998). In section 6.1 we study whether these conclusions still apply in the case of a linear free-surface ocean formulation.

5.2 Advection and conservation properties

The advection terms appearing in the momentum and tracer equations (contained in $\mathbf{S}^{n+1/2}$ in (2)) are computed using a conservative finite-volume predictor-corrector scheme. If we consider a tracer c contained in a cubic cell \mathcal{C} of boundary $\partial\mathcal{C}$ the integrated advection term can be written

$$\int_{\mathcal{C}} A^{n+\frac{1}{2}} \equiv \int_{\mathcal{C}} [\mathbf{U} \cdot \nabla c]^{n+\frac{1}{2}} + w^{n+\frac{1}{2}} \partial_z c^{n+\frac{1}{2}} = \int_{\mathcal{C}} \nabla \cdot (\mathbf{U}c)^{n+\frac{1}{2}} + \partial_z (wc)^{n+\frac{1}{2}} = \int_{\partial\mathcal{C}} c^{n+\frac{1}{2}} (\mathbf{U}^{n+\frac{1}{2}} \cdot \mathbf{n} + wn_z) \quad (10)$$

where $\mathbf{n} = (n_x, n_y, n_z)$ is the unit normal to the boundary and we have used the Stokes theorem and the non-divergence (incompressibility) condition $\nabla \cdot \mathbf{U} + \partial_z w = 0$. A discretised form can be written as

$$\int_{\mathcal{C}} A_{i,j,k}^{n+\frac{1}{2}} \equiv \Delta^2 \left[\begin{aligned} & (uc)_{i+\frac{1}{2},j,k} - (uc)_{i-\frac{1}{2},j,k} + \\ & (vc)_{i,j+\frac{1}{2},k} - (vc)_{i,j-\frac{1}{2},k} + \\ & (wc)_{i,j,k+\frac{1}{2}} - (wc)_{i,j,k-\frac{1}{2}} \end{aligned} \right]^{n+\frac{1}{2}} \quad (11)$$

where Δ is the size of the cell. Values of the tracer and normal components of the velocity are needed on cell faces at time level $n + 1/2$. We use a Godunov procedure where face values are extrapolated in space and time from time- and cell-centred values using a second-order Taylor series expansion. This

scheme as been described in detail in (Popinet, 2003) and (Bell, Colella, and Glaz, 1989) and we will just summarise its properties here. Upwinding is performed in order to reconcile values extrapolated on both sides of the faces and standard limiters can be used for fields with sharp discontinuities. The resulting scheme is the second-order, non-oscillatory, unsplit upwind scheme of Bell, Colella and Glaz (Bell, Colella, and Glaz, 1989). Stability is ensured independently of an explicit dissipation term, provided the CFL number is smaller than one.

It is clear from (11) that local conservation of c is guaranteed only if

$$u_{i+\frac{1}{2},j,k} - u_{i-\frac{1}{2},j,k} + v_{i,j+\frac{1}{2},k} - v_{i,j-\frac{1}{2},k} + w_{i,j,k+\frac{1}{2}} - w_{i,j,k-\frac{1}{2}} = 0$$

which is the discrete expression of the non-divergence condition. The face velocities constructed using the Godunov scheme do not necessarily verify this condition. In our original scheme for incompressible flows an exact projection is applied to the face-based normal velocity field at time $n + 1/2$ to restore this condition. The resulting advection scheme is thus exactly conservative.

In the case of our hydrostatic, linear free-surface ocean model, the vertical component of the velocity w is obtained diagnostically through the vertical continuity condition (1) which is just another way to express the non-divergence condition. Noting that w appears only in the advection terms of the equations of motion and is thus needed only at cell faces, we can write a discretised finite-volume expression of (1) as

$$w_{i,j,k+\frac{1}{2}} = w_{i,j,\frac{1}{2}} + \sum_{k'=1}^{k'=k} u_{i-\frac{1}{2},j,k'} - u_{i+\frac{1}{2},j,k'} + v_{i,j-\frac{1}{2},k'} - v_{i,j+\frac{1}{2},k'} \quad (12)$$

where $w_{i,j,1/2}$ is the boundary condition for w at the bottom of the ocean. The discrete non-divergence condition is thus trivially satisfied at all levels k and the advection scheme is locally conservative.

As noted in other studies (Roulet and Madec, 2000; Griffies, Pacanowski, Schmidt, and Balaji, 2001), this does not guarantee global conservation however. This is due to the linear free-surface approximation which allows (small) fluctuations in the volume of each discretised vertical column of fluid. As a result, the vertical velocity at the top of the column obtained using (12) is non-zero (it is equal to the speed of the vertical free-surface motion). This generates either an outgoing tracer flux at the top of the column which is lost to the system (because the free-surface layer is not represented) or an extrapolated incoming flux which may not match the required boundary conditions. While these fluxes should be small when the linear approximation is valid, nothing guarantees that they balance over the whole system and strict global conservation is not ensured.

5.3 Adaptivity

The original incompressible flow solver Gerris uses a fully three-dimensional adaptive octree discretisation. Each dimension is treated in exactly the same way and refinement or coarsening occur simultaneously on all three dimensions (i.e. adaptivity is isotropic). A direct extension of this approach to oceanic flows is possible but poses several problems. Isotropic adaptivity means that the aspect ratio of cells is independent of the resolution. Taking a typical coarse oceanic discretisation with an horizontal resolution of 9.6 km and vertical resolution of 16 m, refining six times would lead to a cell with an horizontal resolution of 150 m and a vertical resolution of 25 cm. This is probably too high and would lead to a very large number of “layers” in the vertical. Non-isotropic refinement would allow the vertical resolution to be changed independently from the horizontal resolution. Such a capability would require substantial changes in the Gerris code base, however. For this first implementation we chose not to adapt the vertical resolution.

The model can then be seen as a fixed number of quadtree-discretised layers of constant thickness stacked vertically. The solution of the barotropic equation, hydrostatic pressure and diagnostic vertical velocity require local vertical integration over the whole depth of the water column. This is very simple when the discretised control volumes are aligned vertically i.e. when the horizontal resolution is independent of depth. We also add this constraint for our initial implementation. This is probably a more important restriction than the constant-thickness-layers approach because surface and deep currents are expected to have significantly different structures and associated spatial scales. Relaxing this constraint is certainly possible however and this will be the subject of future work.

When refining or coarsening the mesh, care must be taken to conservatively interpolate values from coarse to fine or fine to coarse cells. Simple averaging guarantees a conservative fine to coarse interpolation. For coarse to fine interpolation, values are interpolated linearly using the cell-centred gradient computed on the coarse cell. This also guarantees exact mass and momentum conservation.

Gerris provides a flexible framework for specifying adaptive refinement criteria. Several criteria can be specified simultaneously (such as vorticity, tracer gradients) and the threshold for refinement/coarsening as well as minimum and maximum number of refinement levels can themselves be functions of space and time. A limit on the total size of the simulation can also be set. When reached, the code automatically redistributes the cells in order to globally minimise the refinement cost given by the user-defined criteria.

We also note that Gerris does not use time-adaptivity in the sense that a single global timestep is used irrespective of the spatial resolution. Subcycling in time depending on the level of refinement is possible in principle but leads to complex synchronisation problems to ensure the discrete non-divergence of the global velocity field (Almgren, Bell, Colella, Howell, and Welcome, 1998). In practice the finest grid cells often represent as much as half the total number of cells. Thus, an ideal subcycling implementation would only halve the computational cost compared to simple global time-stepping.

5.4 Representation of boundaries

Traditional height-coordinate ocean models on Cartesian grids usually represent the bottom shape of ocean basins using a “staircase” topography. The coastlines are also often restricted to following the edges of the square discretisation elements. This is a simple but rather crude approach which has been shown to substantially affect the accuracy of the solution (Adcroft, Hill, and Marshall, 1997; Dupont, 2001).

A terrain-following vertical coordinate can be used instead of the height-coordinate (Haidvogel, Wilkin, and Young, 1991; Song and Haidvogel, 1994). This theoretically ensures an accurate representation of the bottom topography (but does not improve the coastline representation). In practice the σ -coordinate transformation causes the surfaces of constant pressure (which are nearly horizontal) to intersect the vertical discretisation at large angles near steep topography. This causes large errors in the discretisation of the horizontal pressure gradients which lead to strong spurious currents near basin boundaries (Mellor, Ezer, and Oey, 1994; Mellor, Oey, and Ezer, 1998). Substantial research has been done to minimise this problem but it is still a fundamental limitation of terrain-following models (Shchepetkin and McWilliams, 2003; Ezer, Arango, and Shchepetkin, 2002; Griffies, Boning, Bryan, Chassignet, Gerdes, Hasumi, Hirst, Treguier, and Webb, 2000).

Fully-unstructured meshes combined with a finite-element or finite-volume approach provide full flexibility for the discretisation of complex domains. This is the main reason for their widespread use in engineering problems. This flexibility comes at the expense of complexity and increased computational cost however. The sub-problem of mesh generation is itself far from trivial, especially in three dimensions (Bern and Eppstein, 1992; Gorman, Piggott, Pain, de Oliveira, Umpleby, and Goddard, 2006). We also note that most of the unstructured finite-volume or finite-element techniques recently developed for 3D geophysical fluid modelling (Casulli and Walters, 2000; Chen, Liu, and Beardsley, 2003; Bacon, Ahmad, Boybeyi, Dunn, Hall, Lee, Sarma, Turner, Waight, Young, and Zack, 2000; Fringer, Gerritsen, and Street, 2006) can be seen as layered 2D discretisations. This solves the problem of coastline representation but does not improve the representation of the bathymetry (a standard terrain-following coordinate is used for the vertical discretisation). A notable exception is the fully three-dimensional model currently developed by Ford, Pain, Piggott, Goddard, de Oliveira, and Umpleby (2004) which uses unstructured tetrahedral meshes.

The quadtree discretisation does not solve the boundary representation problem, although it could be used to increase the resolution near a staircase representation of the topography and coastline. A staircase representation is not the only choice however. A variety of approaches have been developed which can combine the simplicity of a Cartesian discretisation with an accurate representation of complex boundaries: “embedded boundaries” (Peskin, 1972; Ye, Mittal, Udaykumar, and Shyy, 1999), “cut-cell techniques” (Quirk, 1994; Berger and LeVeque, 1991; Almgren, Bell, Colella, and Marthaler, 1997) and for ocean models the “shaved cell” and “partial cell” approach of Adcroft, Hill, and Marshall

(1999). In our previous work (Popinet, 2003) we have demonstrated how a “cut-cell technique” can be applied to an octree discretisation. The same technique is applied here and we will just give a general overview.

The technique relies on the finite-volume formulation of the system of equations. In a manner analogous to the derivation in (10) all the conservation laws can be recast as flux integrals over the boundaries of the discretisation elements. If the (cubic) cell is cut by a solid boundary the advection term in (11) can be written as the more general form

$$\int_C A_{i,j,k}^{n+\frac{1}{2}} \equiv a \Delta^3 A_{i,j,k}^{n+\frac{1}{2}} = \Delta^2 \left[\begin{aligned} & (suc)_{i+\frac{1}{2},j,k} - (suc)_{i-\frac{1}{2},j,k} + \\ & (svc)_{i,j+\frac{1}{2},k} - (svc)_{i,j-\frac{1}{2},k} + \\ & (swc)_{i,j,k+\frac{1}{2}} - (swc)_{i,j,k-\frac{1}{2}} \end{aligned} \right]^{n+\frac{1}{2}} \quad (13)$$

where $s_{i+1/2,j,k}$ and a are the surface (resp. volume) fractions of the face (resp. cell) occupied by the fluid. All the discrete gradient, divergence and Laplacian operators can be generalised in a similar manner. This formulation is also used to compute consistent second-order-accurate fluxes between cells at different levels of refinement (see Popinet (2003) for a detailed description). Second-order accuracy in space is maintained for the barotropic equation in complex domains. A “merged cell” technique is used to avoid the CFL restriction caused by very-small cut cells (Popinet, 2003; Quirk, 1994).

6 Numerical results

The semi-implicit scheme ($\theta = \alpha = \alpha' = 1/2$) has been implemented as part of the Gerris flow solver. Most of the routines are common to both the ocean solver and the incompressible flow solver. In particular, the 2D multilevel Poisson algorithm is used with minor variations to solve the Helmholtz equation for the linear free-surface. The BCG advection scheme is also unchanged. The only routines strictly specific to the ocean model are the equation of state and the vertical integrations required for the hydrostatic pressure and vertical velocity.

In the following we present both two-dimensional test cases designed to highlight important properties and/or shortcomings of the model in the context of ocean modelling, as well as three-dimensional examples of realistic applications. The 2D simulations use the 3D model with a single layer in the vertical. We also try to give quantitative estimates rather than just a qualitative description of the model results.

6.1 Geostrophic adjustment problem

The approximate projection method has been shown to give good results for incompressible flows, however it has not been tested in the context of linearised inertia-gravity waves we consider here. The two main issues regard the stability of the computational mode associated with a collocated A-grid discretisation and the strength of the energy damping induced by the approximate projection.

As a first test case we consider the geostrophic adjustment problem studied by Dupont (2001) and Le Roux, Staniforth, and Lin (1997). A Gaussian bump

$$\eta(x, y) = \eta_0 e^{-\frac{x^2+y^2}{R^2}}$$

is initialised in a 1000×1000 km, 1000 m deep square basin. A reduced gravity $g = 0.01$ m/s is used to approximate a 10 m-thick stratified surface layer. The corresponding geostrophic velocities on an f -plane are initialised as

$$\begin{aligned} u(x, y) &= \frac{2g\eta_0 y}{f_0 R^2} e^{-\frac{x^2+y^2}{R^2}}, \\ v(x, y) &= -\frac{2g\eta_0 x}{f_0 R^2} e^{-\frac{x^2+y^2}{R^2}}, \end{aligned}$$

where f_0 is the Coriolis parameter. Following Dupont we set $f_0 = 1.0285 \times 10^{-4} \text{ s}^{-1}$, $R = 100 \text{ km}$, $\eta_0 = 599.5 \text{ m}$ which gives a maximum geostrophic velocity of 0.5 m/s .

In the context of the linearised shallow-water equations, the geostrophic balance is an exact solution which should be preserved by the numerical method. In practice, this would require an exact numerical balance between terms computed very differently: the pressure gradient and the Coriolis terms in the momentum equation. If this numerical balance is not exact, the numerical solution will adjust toward numerical equilibrium through the emission of gravity-wave noise which should not affect the stability of the solution. This problem is thus a good test of both the overall accuracy of the numerical scheme and its stability properties when dealing with inertia-gravity waves. We note in particular that a standard A-grid discretisation would develop a strong computational-mode instability in this case. Also, as studied by Le Roux, Lin, and Staniforth (2000), an inappropriate choice of finite-element basis functions will result in growing gravity-wave noise.

[Figure 3 about here.]

[Figure 4 about here.]

Figures 3 and 4 summarise the results obtained when running the geostrophic adjustment problem on a 64×64 uniform grid with a timestep $\Delta t = 1000 \text{ s}$. The maximum error on the height field (Figure 3) is small even after 18 days. After a strong initial transient corresponding to the emission of gravity waves, the error reaches a minimum at day 3 and then slowly grows with time with modulations due to the reflections of the initial gravity waves on the domain boundaries. As illustrated on figure 4, this growth is not due to any instability of the solution but to the slow decrease of the maximum amplitude of the Gaussian bump due to numerical energy dissipation.

Numerical energy dissipation is itself a problem and it is important to quantify it. As shown by Sadourny (1975) it is possible to formulate a strictly enstrophy- and energy-preserving finite-difference model on a C-grid. A finite-volume A-grid formulation using a semi-implicit time-discretisation and an exact projection would also be exactly energy-preserving (Dupont, 2001), however it would have a problematic computational mode as discussed previously. Dupont (2001) quantifies the energy conservation properties of various schemes using a non-linear variant of the geostrophic adjustment problem. A β -plane, $f = f_0 + \beta y$ is used and the advection terms are included in the momentum equation. No explicit dissipation is added. As in Dupont (2001) we chose $\beta = 1.607 \times 10^{-11} \text{ m}^{-1} \text{ s}^{-1}$. The geostrophic eddy moves slowly westward through the emission of Rossby waves and southward due to the non-linear advection terms. The resulting evolution of the total energy is shown on figure 5. For our method, the slow decrease in the total energy is due both to the dissipation of potential energy induced by the approximate projection operator and to the dissipative properties of the BCG upwind advection scheme. Another run with the advection terms switched off (figure 5, square symbols) confirms that the dissipation induced by the approximate projection operator dominates the total dissipation. The results however compare favourably with the finite-element formulations tested by Dupont which all show significantly larger energy dissipation.

[Figure 5 about here.]

6.2 Coastally-trapped wave

To test the ability of the method to represent coastlines accurately, we reproduce the *coastally-trapped wave* test case of Curchitser (1999). The simulation is initialised with the exact solution of Lamb (1993). A constant-depth ocean on an f -plane is bounded by a circular coastline. The parameters are the same as used in IMCS (2006): $h = 1500 \text{ m}$, $g = 3.92 \times 10^{-2} \text{ ms}^{-2}$, $f_0 = 8.34 \times 10^{-5} \text{ s}^{-1}$, radius of the ocean $r = 600 \text{ km}$. The simulation runs for three wave periods (5.2465 days) with a timestep of 1199 seconds for all simulations. The linearised shallow-water equations (i.e. using a linearised free-surface and neglecting advection of momentum) are solved with no added explicit dissipation. The error on the final height field is estimated using the correlation technique of Curchitser (1999). The correlation function C is defined as:

$$C(\theta_0) = \frac{\int ML(r, \theta - \theta_0) \text{d}x \text{d}y}{\int L^2 \text{d}x \text{d}y},$$

where $L(r, \theta)$ is the analytical solution of Lamb and $M(x, y)$ is the numerical solution. The phase error and the correlation error are then computed by finding the value θ_{\max} of θ_0 maximising C .

Several runs were performed for different spatial resolutions (by varying the number of quadtree levels from 5 to 8) and using a constant or adaptive refinement strategy. The final solutions obtained with 7 levels of refinement for both a constant and an adaptive resolution are illustrated on Figure 6. The vorticity adaption criterion (Popinet, 2003) was used to follow the waves as they travelled around the coastline.

Table 1 summarises the results for Gerris and reproduces the results obtained using FVCOM (Chen, Liu, and Beardsley, 2003) and ROMS (Shchepetkin and McWilliams, 2005). FVCOM is a finite-volume solver using an unstructured triangular discretisation of the circular ocean. The constant resolution version of Gerris shows a convergence and accuracy which is better than FVCOM. The adaptive version shows a similar convergence and accuracy with a slight decrease of accuracy for high-resolution due to accumulating interpolation noise when refining/coarsening the adaptive grid (this leads to slightly larger amplitudes close to the coastline which explains the correlation coefficients larger than one). ROMS is using a *staircase* representation of the coastline and consequently suffers from much larger phase errors than both Gerris (using the cut-cell technique) and FVCOM. The ROMS solution is also very noisy compared to the solutions illustrated in Figure 6. We also note that because both FVCOM and ROMS rely on an explicit barotropic solver, they used a timestep of around 250 seconds on the coarsest mesh (proportionally less on finer meshes).

An often-heard argument against the use of variable-resolution meshes for ocean modelling is the occurrence of spurious inertial and gravity waves reflections/scattering at refinement boundaries (Griffies, Boning, Bryan, Chassignet, Gerdes, Hasumi, Hirst, Treguier, and Webb, 2000). We believe one of the causes of spurious reflections is the use of discrete spatial operators which do not preserve second-order spatial accuracy when the resolution changes. The operators used in Gerris are designed to be second-order accurate everywhere (Popinet, 2003) and we do not expect spurious reflections to occur at refinement boundaries. To test whether this is indeed the case, we repeated the coastally-trapped wave test case on a mesh where half the circular domain is resolved with double the resolution (9.375 km) of the other half (18.75 km). The results are given at the end of Table 1 and are close to those obtained on a mesh with a constant resolution equal to that of the coarser half of the domain (18.75 km), which confirms that no significant spurious reflections occur at refinement boundaries.

[Figure 6 about here.]

[Table 1 about here.]

6.3 Wind-driven circular ocean

Dupont presents another interesting test case for the representation of a circular coastline (Dupont, 2001): the wind-driven circular ocean circulation. For a constant-depth ocean on an f -plane with a linear bottom friction, the linearised shallow-water equation admits the following stationary solution for the surface-height:

$$\eta(r, \theta) = \frac{Wf}{RgH\kappa} \left[\frac{R^2}{8} + \frac{r^2}{4} \left(\frac{\kappa}{f} \sin 2\theta - 1 \right) \right],$$

where R is the radius of the ocean, g the acceleration of gravity, f the Coriolis parameter, H the ocean depth, κ the linear friction coefficient and W the gradient of the wind forcing:

$$\tau_x(r, \theta) = \frac{Wr \sin \theta}{R}.$$

We follow Dupont and study the convergence with spatial resolution of the normalised error in η defined as:

$$E(\hat{\eta}) = \frac{\int |\hat{\eta} - \eta| \, dx \, dy}{\int \eta \, dx \, dy} \sqrt{\frac{\int dx \, dy}{\int \eta^2 \, dx \, dy}},$$

where $\hat{\eta}$ is the model solution. The model is run to convergence with $f = 10^{-4} \text{ s}^{-1}$, $W = 10^{-4} \text{ m}^2\text{s}^{-2}$ and $\kappa = 10^{-3} \text{ s}^{-1}$. Figure 7 illustrates the solution obtained with 5 levels of refinement.

The convergence of the normalised error with spatial resolution is plotted on Figure 8 together with the results obtained by Dupont for a streamfunction-vorticity C-grid model and the Hua-Thomasset finite-element model (HT) (Hua and Thomasset, 1984). The HT model gives the best result of all the finite-element models tested by Dupont for this particular test case. An unstructured triangular mesh is used for the finite-element model.

[Figure 7 about here.]

[Figure 8 about here.]

The C-grid model shows much larger errors and only first-order convergence due to the stair-step representation of the circular coastline. Both the HT model and Gerris show second-order convergence. The errors obtained with Gerris are also substantially smaller than that obtained with the HT model.

6.4 Flow over a Gaussian bump

The test cases presented this far are only two-dimensional. An interesting three-dimensional configuration was used by Adcroft et al. (Adcroft, Hill, and Marshall, 1999; Haidvogel, Wilkin, and Young, 1991) to validate their *shaved cells* and *partial step* representations of topography. A Gaussian bump is placed in a periodic channel bounded by solid walls with a slip boundary condition. The width of the channel is 320 km and its length 640 km. The 3D “primitive equations” of section 2 are solved with the parameters listed in Table 2.

[Table 2 about here.]

No explicit diffusion is added either on the density or momentum. The shape of the Gaussian bump is given by

$$h(x, y) = H - H_0 \exp\left(-\left[(x - x_0)^2 + (y - y_0)^2\right]/L^2\right),$$

with $x_0 = 200$ km and $y_0 = 150$ km.

The constant barotropic inflow u_0 is imposed by adding a spatially-constant body force in the x -direction. After 10 days the flow pattern illustrated on Figure 9 is obtained. The same contour intervals as in Adcroft, Hill, and Marshall (1999) are used for the representation. The result compares very well with Figure 14 and 18 of Adcroft et al. In particular the solution is as free of high-frequency noise as that obtained by Haidvogel, Wilkin, and Young (1991) with the spectral SPEM model (reproduced in Figure 14 of Adcroft et al.). The vorticity maximum in the detached vortex is well captured and does not show the excessive damping of the “full step” representation of topography (Figure 18(c) of Adcroft et al.). The Gerris solution is also cleaner than that obtained by Adcroft et al. with either the “piecewise linear” or “partial step” methods.

[Figure 9 about here.]

6.5 Tidally-driven flow in Cook Strait

We conclude with an example of adaptive barotropic flow with a complex coastline and bathymetry and a large ratio between the smallest and largest scales. Cook Strait, the stretch of water between the North and South Islands of New Zealand is a complex, highly energetic marine environment. The Marlborough Sounds – a drowned mountain river system at the northern end of the South Island – have a very convoluted coastline and provide a good test of the robustness of the treatment of embedded solid boundaries.

The simulation domain is illustrated on Figure 10. The eastern part of the strait is relatively shallow while the western part drops off steeply. Because we only solve for the barotropic flow in this example, we limit the maximum depth to 400 meters as an approximation to a fixed-depth thermocline. A “Flather” boundary condition is used on all four boundaries (Flather and Heaps, 1975) where only the M_2 component of the surface elevation is imposed, using results from the Walters and Goring tidal model (Walters, Goring, and Bell, 2001). The main feature of the M_2 tides around New

Zealand is a single Kelvin-like wave propagating anti-clockwise around both islands. As a result, the tides on the eastern and western sides of Cook Strait are in opposition of phase with an amphidrome close to the narrowest point of the passage. The resulting tidal pressure gradient then forces strong currents both in the main passage and in some branches of the Marlborough Sounds.

[Figure 10 about here.]

The full 3D model is used to solve for the barotropic part of the solution by using only a single vertical layer. The bathymetry is described using the full 3D volume-of-fluid technique described in section 5.4. A small amount ($4 \times 10^{-6} \text{ s}^{-1}$) of linear bottom friction is added to dampen large velocities in very shallow areas. The resolution of the cells defining the coastline is fixed at 12 quadtree-levels i.e. $500 \text{ km} / 2^{12} \approx 122$ meters. The minimum resolution is set to 6 levels i.e. $\approx 7.8 \text{ km}$. Away from the coastline the resolution is adapted dynamically at every timestep according to the local vorticity. The timestep is also adapted dynamically in order to verify the CFL condition required for the stability of the advection scheme. The average timestep is approximately 40 seconds. In established regime the total number of cells is approximately constant at $\approx 140,000$ (compared to $(2^{12})^2 \approx 17,000,000$ for a regular Cartesian grid of equivalent resolution). Each timestep takes 10 seconds of CPU time on a Pentium 4 2.6 GHz processor, of which less than 5% is used to adapt the grid. The simulation required approximately 50 MBytes of memory.

Figure 11 and 12 shows part of the depth-averaged vorticity field for $t = 6$ days. The combination of strong currents and sharp headlands creates vorticity “dipoles” close to the coastline which are then advected away from the coast by the mean tidal circulation. The weak dissipation due to bottom friction in deeper water as well as the good energy conservation properties of the numerical scheme lead to long-lived eddies which interact with newer eddies produced at each tidal cycle. This leads to the complex vorticity field illustrated in the figures. This process is more clearly illustrated by the accompanying animation (Popinet, 2005).

[Figure 11 about here.]

Figure 12 displays a detail of the quadtree mesh for $t = 6$ days. High-vorticity areas are tracked with a maximum resolution of 12 levels. Using a statically refined mesh with a high resolution near the coastline decreasing offshore leads to a very different picture. Eddies are strongly damped by numerical diffusion as soon as they leave the coastline and the process of inter-tidal interactions between eddies does not occur.

While the code has been validated in the Cook Strait area for the prediction of tidal amplitudes and currents (Msadek, 2005), the high-resolution results presented here still need to be compared with available observational measurements. However, satellite imagery provide convincing first evidence that the processes of eddy generation off headlands, subsequent vorticity transport offshore and inter-tidal eddy interactions illustrated here are important in reality. The International Space Station pictures in the accompanying material (NASA, 2002a,b) show structures very similar to the small offshore eddies and vorticity layers off headlands illustrated in Figure 11.

One may argue that trying to resolve individual eddies at the scales illustrated here is not necessary as these scales are meant to be adequately represented by sub-grid turbulence models, however most turbulence models rely on assumptions (e.g. spatial isotropy) which are not verified in strongly sheared flows such as boundary layers. We believe that the high resolution made possible by adaptive techniques will be a key requirement in accurately modelling the on-shore/off-shore transport processes controlling sediment dynamics, nutrient fluxes in the marine ecosystem and bio-geochemical cycles generally.

[Figure 12 about here.]

[Figure 13 about here.]

7 Conclusion

The development of an adaptive ocean model has been described. The dynamic adaptivity – allowing for multi-scale representation within a single model that can evolve in time – is achieved through

a conservative finite-volume discretisation which requires collocation of all the geophysical variables (A-grid discretisation). It is demonstrated, however, that the extension to a free-surface ocean model of the approximate projection method overcomes the A-grid limitations at the expense of a small amount of numerical energy dissipation. Test cases clearly underline the accuracy and efficiency of this adaptive ocean model; accurate geostrophic balance, robust and accurate representation of bathymetry and coastlines are demonstrated.

As a first demonstration of the functional adaptivity, the present three-dimensional ocean model has a linear free surface, and is incompressible, hydrostatic, and Boussinesq. Here, the adaptivity is also limited to the horizontal plane only, and is projected down through a series of z -level, stacked layers in the vertical. Nevertheless, the boundaries are represented with a “cut-cell” technique, so the familiar “staircase” pitfalls of typical finite difference, z -level models are avoided (Adcroft, Hill, and Marshall, 1997). Furthermore, the range of relevant geophysical solutions for such an ocean model remains large, and therefore permit the exploration of a number of testing problems.

For the case study of coupling between a complex coastal and offshore barotropic flow, we demonstrate how the adaptive mesh results in a reduction of two orders of magnitude in the number of grid points. Moreover, the dynamic adaptivity gives the model an advantage over other models that have an unstructured spatial grid that is fixed in time, as it has the ability to track evolving fine-scale features (e.g. eddies shed from headlands, meandering frontal boundaries etc...). This is particularly important for the accurate simulation of processes which are not adequately described by current subgrid-scale models (e.g. turbulence near complex coastlines). Compared to more classical unstructured grid approaches, the “semi-structured” quadtree grid we employ is very easy to generate and – as a result – the cost of adapting the grid at every timestep is a small fraction of the total computational cost. While time-adaptivity is certainly possible on unstructured grids, it is expensive and cannot usually be applied at every timestep (Bacon, Ahmad, Boybeyi, Dunn, Hall, Lee, Sarma, Turner, Waight, Young, and Zack, 2000). In the context of geophysical fluid flows, we are only aware of two studies demonstrating time-adaptive unstructured grids (Bacon, Ahmad, Boybeyi, Dunn, Hall, Lee, Sarma, Turner, Waight, Young, and Zack, 2000; Ford, Pain, Piggott, Goddard, de Oliveira, and Umpleby, 2004).

From a user perspective, the fact that adaptive mesh generation is part of the solution procedure allows flexible control of the level of refinement, so that issues of convergence – for example – can be addressed by simply changing an input file, thereby avoiding having to do a model re-build each time the resolution is altered. The hierarchical quadtree structure is also naturally suited to multigrid algorithms which leads to robust and efficient solvers for the implicit free-surface discretisation even with arbitrarily complex coastlines.

We are currently working on generalising this approach to a non-hydrostatic, nonlinear free-surface formulation with adaptivity in all dimensions. The underlying adaptive Navier-Stokes fluid solver Gerris naturally deals with non-hydrostatic, nonlinear free surface problems, in which adaptivity in all three spatial dimensions is routinely performed. In the oceanic case, decoupling of the hydrostatic and non-hydrostatic pressures can be employed to overcome the stiffness of the pressure equation due to the large aspect ratio between horizontal and vertical scales (Marshall, Adcroft, Hill, Perelman, and Heisey, 1997a). The extension of adaptivity to the vertical coordinate could be more problematic and has not yet been studied in any significant detail by us or others. The main issue will be ensuring the accurate discretisation of the horizontal hydrostatic pressure gradients across levels of refinement. Another important general issue for adaptive solvers which has received little attention is the development of criteria for refinement to control errors in subgrid-scale parametrisation.

We are confident that these issues can be overcome and that the resulting dynamic adaptivity will be an essential component of the next generation of geophysical fluid flows models.

A Stability of Adams-Bashforth discretisation

Model simulations of the geostrophic adjustment problem tested an Adams-Bashforth scheme to discretise the Coriolis terms in time. The simulations seemed to require a timestep much shorter than expected as the simulation proceeded. This appendix confirms the need for such a short timestep, and motivates the use of the semi-implicit scheme for the Coriolis terms which is possible within the

adaptive model because of the collocated velocity components.

A simple, one-dimensional shallow water system is used to explore the stability properties of an inertia-gravity wave system using Adams-Bashforth discretisation for the Coriolis terms, viz,

$$\frac{\partial u}{\partial t} = fv - g \frac{\partial \eta}{\partial x}, \quad (14)$$

$$\frac{\partial v}{\partial t} = -fu, \quad (15)$$

$$\frac{\partial \eta}{\partial t} = -H \frac{\partial u}{\partial x}. \quad (16)$$

Taking $\mathbf{u} = \hat{\mathbf{u}} \cos kx$, $v = \hat{v} \cos kx$, $\eta = \hat{\eta} \sin kx$, analytic solutions are,

$$u = u_o \cos \omega t \cos kx, \quad v = -u_o \frac{f}{\omega} \sin \omega t \cos kx, \quad \eta = H k u_o \frac{f}{\omega} \sin \omega t \sin kx, \quad (17)$$

for each wavenumber k with frequency ω , where,

$$\omega^2 = f^2(1 + k^2 R^2), \quad R^2 = \frac{gH}{f^2}, \quad (18)$$

and u_o is a constant.

Following Marshall et al. (1997), an Adams-Bashforth method (AB2) is used to discretise the Coriolis terms, while the remaining terms are allowed to range from fully explicit to fully implicit. The discrete (in time) system can then be written,

$$u^{n+1} - u^n = a(v^n + \alpha(v^n - v^{n-1})) - b(d\eta^{n+1} + e\eta^n), \quad (19)$$

$$v^{n+1} - v^n = -a(u^n + \alpha(u^n - u^{n-1})), \quad (20)$$

$$\eta^{n+1} - \eta^n = c(du^{n+1} + eu^n), \quad (21)$$

where $a = f\Delta t$, $b = gk\Delta t$, $c = Hk\Delta t$, $\alpha = 1/2 + \chi$, $d = \theta$, and $e = 1 - \theta$, where θ is the implicitness parameter such that $0 \leq \theta \leq 1$. As in Marshall et al. (1997), the offset in the AB2 scheme χ is set to be 0.1, and is consistent with the stability of this AB2 scheme *alone*, i.e., $a \leq 1/2$ when $\chi > a^2/4 + a^4/2$. In terms of an amplification factor λ such that $\eta^{n+1} = \lambda \eta^n$, the discrete system can be reduced to the following fourth order polynomial;

$$\lambda^4(1 + cbd^2) + \lambda^3(2bcde - 2) + \lambda^2(1 + bce^2 + a^2(1 + \alpha)^2) + \lambda(-2a^2\alpha(1 + \alpha)) + a^2\alpha^2 = 0. \quad (22)$$

Stable numerical solutions arise for $|\lambda| < 1$. The fourth order system reduces to two complex conjugate pairs, one representing the computational mode which is found to be always stable for the range of parameters explored here, the other the physical mode.

Solutions for the physical mode are plotted in Figure 14 as a function of the dimensionless numbers $f\Delta t$ and Rk . The region of stability is that below (or to the right of) each curve, and the different curves are labelled with the value of the implicitness parameter θ . The horizontal dashed line is the AB2 stability boundary in the absence of gravity waves. For pure gravity waves with $\theta \geq 0.5$ the system is always stable.

The important result is that the stability region for the full inertia-gravity wave system falls well below the AB2 line at $f\Delta t = 1/2$ for the centred-in-time $\theta = 1/2$ discrete system, particularly as Rk increases. Such Rk changes represent either an increase in the depth H of the domain, or an increase in the wavenumber k (or equivalently a reduction in the resolved wavelengths). This theoretical solution explains the need for significant reductions in the timestep as the generation of resolved, fine-scale structure in the numerical model solution occurs. Even though the separate gravity and inertia waves are theoretically stably resolved by $f\Delta t = 1/2$, their interaction as an inertia-gravity wave system significantly brings down the available region of stability.

Allowing for more implicitness in the gravity wave terms via increasing θ expands the stability range, and with $\theta = 0.6$ a permissible timestep around $f\Delta t = 1/2$ is recovered for all values of Rk . Further increases in θ allow for significant increases in $f\Delta t$ at larger values of Rk , with smaller gains for smaller Rk . However this comes at the cost of increased dissipation of gravity waves.

[Figure 14 about here.]

References

- Adcroft, A., Campin, J.-M., 2004. Rescaled height coordinates for accurate representation of free-surface flows in ocean circulation models. *Ocean Modelling* 7 (3), 269–284.
- Adcroft, A., Hill, C., Marshall, J., 1997. Representation of topography by shaved cells in a height coordinate ocean model. *Monthly Weather Review* 125 (9), 2293–2315.
- Adcroft, A., Hill, C., Marshall, J., 1999. A new treatment of the Coriolis terms in C-grid models at both high and low resolutions. *Monthly Weather Review* 127 (8), 1928–1936.
- Almgren, A., Bell, J., Colella, P., Howell, L., Welcome, M., 1998. A conservative adaptive projection method for the variable density incompressible Navier-Stokes equations. *Journal of Computational Physics* 142 (1), 1–46.
- Almgren, A., Bell, J., Colella, P., Marthaler, T., 1997. A Cartesian grid projection method for the incompressible Euler equations in complex geometries. *SIAM J. Sci. Comp.* 18 (5).
- Almgren, A., Bell, J., Crutchfield, W., 2001. Approximate projection methods: Part I. Inviscid analysis. *SIAM J. Sci. Comput.* 22 (4), 1139–1159.
- Bacon, D., Ahmad, N., Boybeyi, Z., Dunn, T., Hall, M., Lee, P., Sarma, R., Turner, M., Waight, K., Young, S., Zack, J., 2000. A dynamically adapting weather and dispersion model: The operational multiscale environment model with grid adaptivity (OMEGA). *Mon. Weather Rev.* 128, 2044–2076.
- Barros, S., Garcia, C., 2004. A Global Semi-Implicit Semi-Lagrangian Shallow-Water Model on Locally Refined Grids. *Monthly Weather Review* 132 (1), 53–65.
- Bell, J., Colella, P., Glaz, H., 1989. A second-order projection method for the incompressible Navier-Stokes equations. *J. Comput. Phys* 85 (2), 257–283.
- Berger, M., LeVeque, R., June 1991. A rotated difference scheme for Cartesian grids in complex geometries. In: *AIAA 10th Computational Fluid Dynamics Conference*. Honolulu, Hawaii, pp. 1–7.
- Bern, M., Eppstein, D., 1992. *Mesh Generation and Optimal Triangulation*. Xerox Corp., Palo Alto Research Center.
- Blayo, E., Debreu, L., 1999. Adaptive mesh refinement for finite difference ocean models: first experiments. *J. Phys. Oceanogr.* 29, 1239–1250.
- Bleck, R., 2002. An oceanic general circulation model framed in hybrid isopycnic-Cartesian coordinates. *Ocean Modelling* 4 (1), 55–88.
- Bramble, J., 1993. *Multigrid Methods*. Chapman & Hall/CRC.
- Brandt, A., 1982. Guide to multigrid development. *Multigrid Methods* 960, 220–312.
- Campin, J.-M., Adcroft, A., Hill, C., Marshall, J., 2004. Conservation of properties in a free-surface model. *Ocean Modelling* 6 (3), 221–244.
- Casulli, V., Walters, R., 2000. An unstructured grid, three-dimensional model based on the shallow water equations. *International Journal for Numerical Methods in Fluids* 32 (3), 331–348.
- Chen, C., Liu, H., Beardsley, R., 2003. An unstructured grid, finite-volume, three-dimensional, primitive equations ocean model: Application to coastal ocean and estuaries. *Journal of Atmospheric and Oceanic Technology* 20 (1), 159–186.
- Chorin, A., 1968. Numerical solution of the Navier-Stokes equations. *Mathematics of Computation* 22 (104), 745–762.
- Curchitser, E., 1999. On the transient adjustment of a mid-latitude abyssal ocean basin with realistic geometry and bathymetry. Ph.D. thesis, Rutgers University, New Brunswick, NJ.

- Dietrich, D., 1997. Application of a modified Arakawa ‘A’ grid ocean model having reduced numerical dispersion to the Gulf of Mexico circulation. *Dynamics of Atmospheres and Oceans* 27, 201–217.
- Dukowicz, J., Smith, R., 1994. Implicit free-surface method for the Bryan-Cox-Semtner ocean model. *J. Geophys. Res* 99, 7991–8014.
- Dupont, F., 2001. Comparison of numerical methods for modelling ocean circulation in basins with irregular coasts. Ph.D. thesis, McGill University, Montreal.
- Ezer, T., Arango, H., Shchepetkin, A., 2002. Developments in terrain-following ocean models: inter-comparisons of numerical aspects. *Ocean Modelling* 4 (3), 249–267.
- Flather, R., Heaps, N., 1975. Tidal computations for Morecambe bay. *Geophys. J.R. astr. Soc.* 42, 489–517.
- Ford, R., Pain, C., Piggott, M., Goddard, A., de Oliveira, C., Umpleby, A., 2004. A nonhydrostatic finite-element model for three-dimensional stratified oceanic flows. part I: Model formulation. *Monthly Weather Review* 132 (12), 2816–2831.
- Fringer, O., Gerritsen, M., Street, R., 2006. An unstructured-grid, finite-volume, nonhydrostatic, parallel coastal ocean simulator. *Ocean Modelling* 14, 139–173.
- Giraldo, F., 2000. The Lagrange-Galerkin method for the two-dimensional shallow water equations on adaptive grids. *International Journal for Numerical Methods in Fluids* 33 (6), 789–832.
- Gorman, G., Piggott, M., Pain, C., de Oliveira, C., Umpleby, A., Goddard, A., 2006. Optimisation based bathymetry approximation through constrained unstructured mesh adaptivity. *Ocean Modelling* 12, 436–452.
- Griffies, S., Boning, C., Bryan, F., Chassignet, E., Gerdes, R., Hasumi, H., Hirst, A., Treguier, A.-M., Webb, D., 2000. Developments in ocean climate modelling. *Ocean Modelling* 2 (3/4), 123–192.
- Griffies, S., Pacanowski, R., Schmidt, M., Balaji, V., 2001. Tracer conservation with an explicit free surface method for z-coordinate ocean models. *Monthly Weather Review* 129 (5), 1081–1098.
- Haidvogel, D., Wilkin, J., Young, R., 1991. A semi-spectral primitive equation ocean circulation model using vertical sigma and orthogonal curvilinear horizontal coordinates. *Journal of Computational Physics* 94 (1), 151–185.
- Harlow, F., Welch, J., 1965. Numerical calculation of time-dependent incompressible flow of fluid with free surface. *Phys. Fluids* 8, 2182–2189.
- Herrnstein, A., Wickett, M., Rodrigue, G., 2005. Structured adaptive mesh refinement using leapfrog time integration on a staggered grid for ocean models. *Ocean Modelling* 9 (3), 283–304.
- Hua, B., Thomasset, F., 1984. A noise-free finite element scheme for the two-layer shallow water equations. *Tellus A* 36, 157.
- IMCS, 2006. Coastally trapped waves in a circular basin,
<http://marine.rutgers.edu/po/index.php?model=test-problems&title=circle>.
- Kar, S., 1999. Stable centered-difference schemes, based on an unstaggered a grid, that eliminate two-grid interval noise. *Monthly Weather Review* 128 (10), 3643–3653.
- Lai, M., Bell, J., Colella, P., 1993. A projection method for combustion in the zero Mach number limit. *Proceedings of the Eleventh AIAA Computational Fluid Dynamics Conference*, 776–783.
- Lamb, H., 1993. *Hydrodynamics*. Cambridge University Press.
- Le Roux, D., Lin, C., Staniforth, A., 2000. A semi-implicit semi-lagrangian finite-element shallow-water ocean model. *Monthly Weather Review* 128 (5), 1384–1401.

- Le Roux, D., Staniforth, A., Lin, C., 1997. Finite elements for shallow-water equation ocean models. *Monthly Weather Review* 126 (7), 1931–1951.
- Lynch, D., Werner, F., 1987. Three-dimensional hydrodynamics on finite elements. Part I: Linearized harmonic model. *International Journal for Numerical Methods in Fluids* 7, 871–909.
- Marshall, J., Adcroft, A., Campin, J., Hill, C., White, A., 2004. Atmosphere-ocean modeling exploiting fluid isomorphisms. *Monthly Weather Review* 132 (12), 2882–2894.
- Marshall, J., Adcroft, A., Hill, C., Perelman, L., Heisey, C., 1997a. A finite-volume, incompressible Navier-Stokes model for studies of the ocean on parallel computers. *J. Geophys. Res* 102 (5), 753–5,766.
- Marshall, J., Hill, C., Perelman, L., Adcroft, A., 1997b. Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling. *J. Geophys. Res* 102 (5), 733–5,752.
- Mellor, G., Ezer, T., Oey, L., 1994. The pressure gradient conundrum of sigma coordinate ocean models. *Journal of Atmospheric and Oceanic Technology* 11, 1126–1134.
- Mellor, G., Oey, L., Ezer, T., 1998. Sigma coordinate pressure gradient errors and the seamount problem. *Journal of Atmospheric and Oceanic Technology* 15 (5), 1122–1131.
- Msadek, R., June 2005. Hydrodynamic tidal model of Cook Strait. Tech. rep., National Institute of Water and Atmospheric research, available at <http://gfs.sf.net/tides.pdf>.
- NASA, 2002a. International Space Station photograph. <http://eol.jsc.nasa.gov/scripts/sseop/photo.pl?mission=ISS005&roll=E&frame=16782>.
- NASA, 2002b. International Space Station photograph. <http://eol.jsc.nasa.gov/scripts/sseop/photo.pl?mission=ISS005&roll=E&frame=16780>.
- Peraire, J., Zienkiewicz, O., Morgan, K., 1986. Shallow water problems: A general explicit formulation. *International Journal for Numerical Methods in Engineering* 22 (3), 547–574.
- Peskin, C., 1972. Flow patterns around heart valves: a numerical method. *J. Comp. Phys.*, 10–252.
- Peyret, R., Taylor, T., 1985. *Computational methods for fluid flow*. Springer-Verlag, New York.
- Popinet, S., 2003. Gerris: A tree-based adaptive solver for the incompressible Euler equations in complex geometries. *Journal of Computational Physics* 190 (2), 572–600.
- Popinet, S., 2005. Animation of the vorticity field in Cook Strait. MPEG file (9.3 MB), <http://gfs.sf.net/tmp/cook.mpg>.
- Popinet, S., 2006. The Gerris Flow Solver. <http://gfs.sourceforge.net>.
- Power, P., Piggott, M., Fang, F., Gorman, G., Pain, C., Marshall, D., Goddard, A., Navon, I., 2006. Adjoint goal-based error norms for adaptive mesh ocean modelling. *Ocean Modelling* 15 (1-2), 3–38.
- Quirk, J., 1994. An alternative to unstructured grids for computing gas dynamic flows around arbitrarily complex two-dimensional bodies. *Computers and Fluids* 23 (1), 125–142.
- Rider, W., 1995. Approximate projection methods for incompressible flow: Implementation, variants and robustness. Tech. rep., Technical Report LA-UR-2000, Los Alamos National Laboratory.
- Roulet, G., Madec, G., 2000. Salt conservation, free surface, and varying levels: A new formulation for ocean general circulation models. *Journal of Geophysical Research* 105 (C10), 23927–23942.
- Sadourny, R., 1975. The dynamics of finite-difference models of the shallow-water equations. *J. Atmos. Sci* 32 (4), 680–689.
- Shchepetkin, A., McWilliams, J., 2003. A method for computing horizontal pressure-gradient force in an oceanic model with a nonaligned vertical coordinate. *J. Geophys. Res* 108 (10.1029).

- Shchepetkin, A., McWilliams, J., 2005. The regional oceanic modeling system (ROMS): A split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Modelling* 9 (4), 347–404.
- Song, Y., Haidvogel, D., 1994. A semi-implicit ocean circulation model using a generalized topography-following coordinate system. *Journal of Computational Physics* 115 (1), 228–244.
- Song, Y., Hou, T., 2006. Parametric vertical coordinate formulation for multiscale, Boussinesq, and non-Boussinesq ocean modeling. *Ocean Modelling* 11, 298–332.
- Walters, R., 1986. A finite element model for tidal and residual circulation. *Communications in Applied Numerical Methods* 2, 393–398.
- Walters, R., Barragy, E., 1997. Comparison of H and P finite element approximations of the shallow water equations. *International Journal for Numerical Methods in Fluids* 24 (1), 61–79.
- Walters, R., Casulli, V., 1998. Robust, finite element model for hydrostatic surface water flows. *Commun. Numer. Methods Eng.* 14 (10), 931–940.
- Walters, R., Goring, D., Bell, R., 2001. Ocean tides around New Zealand. *New Zealand Journal of Marine and Freshwater Research* 35, 567–579.
- Ye, T., Mittal, R., Udaykumar, H. S., Shyy, W., 1999. An accurate Cartesian grid method for viscous incompressible flows with complex immersed boundaries. *J. Comp. Phys.* 156, 209–240.

List of Figures

1	Example of quadtree discretisation and corresponding tree representation. The dotted lines on the right-hand-side define the <i>level</i> of the cells.	23
2	Staggered C-grid discretisation on a quadtree. (a) Initial coarse grid showing the control volume for the horizontal component of the velocity. (b) The cell centred on the pressure discretisation has been refined once. The union of the control volumes of the newly defined horizontal components is given by the cross-hashed area.	24
3	Evolution of the maximum error on the surface height for the geostrophic adjustment problem.	25
4	Evolution of the surface-height error field. (a) $t = 1.157$ days, (b) $t = 2.315$ days, (c) $t = 3.472$ days, (d) $t = 4.630$ days, (e) $t = 17.361$ days.	26
5	Evolution of the total energy for the non-linear geostrophic adjustment problem. The C-grid model is based on Sadourny (1975) and implemented by Dupont (2001). The finite-element formulations are those studied by Dupont. LW: Lynch and Werner (1987), LLS: Le Roux, Lin, and Staniforth (2000), PZM: Peraire, Zienkiewicz, and Morgan (1986).	27
6	Surface-height isolines for the coastally-trapped wave solution at $t = 5.2465$ days (3 wave periods). (a) Constant resolution 9.375 km. (b) Adaptive resolution, white 9.375 km, black 150 km.	28
7	Velocity and height field for the wind-driven circular ocean test case.	29
8	Convergence of the normalised error with spatial resolution for different models. Both the HT and C-grid results are reproduced from the study of Dupont (2001).	30
9	Non-dimensional depth-integrated vorticity $\hat{\zeta}/f$ at $t = 10$ days. Contour interval = 100.	31
10	Coastline and isobaths for the Cook Strait tidal model. The domain extent is 500×500 km. Both the coastline and the isobaths are drawn using their exact representation on the quadtree mesh. The light grey area indicates depths deeper than 400 metres. The black box indicates the area zoomed in on Figure 11.	32
11	Detail of the vorticity field at $t = 6$ days. The domain extent is 83×73 km. The black boxes indicate the areas represented on Figure 12 and 13.	33
12	Detail of the adaptive quadtree mesh. The coastline is statically refined at the finest level while individual eddies are tracked dynamically. The domain extent is 36×29 km. The minimum spatial resolution is 122 metres.	34
13	Detail of the barotropic velocity field in the Tory channel area of the Marlborough Sounds.	35
14	Stability boundaries for the discretisation of an inertia-gravity wave system as a function of the dimensionless numbers Rk and $f\Delta t$. Discretisation is Adams-Bashforth for the Coriolis terms, and implicit in the gravity wave terms. Stable solutions lie below and/or to the right of the curves. Each curve is labelled with the implicitness parameter θ of the gravity wave terms. The horizontal dashed line marks the stability boundary for solutions with no gravity waves.	36




Figure 1: Example of quadtree discretisation and corresponding tree representation. The dotted lines on the right-hand-side define the *level* of the cells.




Figure 2: Staggered C-grid discretisation on a quadtree. (a) Initial coarse grid showing the control volume for the horizontal component of the velocity. (b) The cell centred on the pressure discretisation has been refined once. The union of the control volumes of the newly defined horizontal components is given by the cross-hashed area.




Figure 3: Evolution of the maximum error on the surface height for the geostrophic adjustment problem.




Figure 4: Evolution of the surface-height error field. (a) $t = 1.157$ days, (b) $t = 2.315$ days, (c) $t = 3.472$ days, (d) $t = 4.630$ days, (e) $t = 17.361$ days.




Figure 5: Evolution of the total energy for the non-linear geostrophic adjustment problem. The C-grid model is based on Sadourny (1975) and implemented by Dupont (2001). The finite-element formulations are those studied by Dupont. LW: Lynch and Werner (1987), LLS: Le Roux, Lin, and Staniforth (2000), PZM: Peraire, Zienkiewicz, and Morgan (1986).




Figure 6: Surface-height isolines for the coastally-trapped wave solution at $t = 5.2465$ days (3 wave periods). (a) Constant resolution 9.375 km. (b) Adaptive resolution, white 9.375 km, black 150 km.




Figure 7: Velocity and height field for the wind-driven circular ocean test case.




Figure 8: Convergence of the normalised error with spatial resolution for different models. Both the HT and C-grid results are reproduced from the study of Dupont (2001).




Figure 9: Non-dimensional depth-integrated vorticity $\hat{\zeta}/f$ at $t = 10$ days. Contour interval = 100.




Figure 10: Coastline and isobaths for the Cook Strait tidal model. The domain extent is 500×500 km. Both the coastline and the isobaths are drawn using their exact representation on the quadtree mesh. The light grey area indicates depths deeper than 400 metres. The black box indicates the area zoomed in on Figure 11.




Figure 11: Detail of the vorticity field at $t = 6$ days. The domain extent is 83×73 km. The black boxes indicate the areas represented on Figure 12 and 13.




Figure 12: Detail of the adaptive quadtree mesh. The coastline is statically refined at the finest level while individual eddies are tracked dynamically. The domain extent is 36×29 km. The minimum spatial resolution is 122 metres.




Figure 13: Detail of the barotropic velocity field in the Tory channel area of the Marlborough Sounds.




Figure 14: Stability boundaries for the discretisation of an inertia-gravity wave system as a function of the dimensionless numbers Rk and $f\Delta t$. Discretisation is Adams-Bashforth for the Coriolis terms, and implicit in the gravity wave terms. Stable solutions lie below and/or to the right of the curves. Each curve is labelled with the implicitness parameter θ of the gravity wave terms. The horizontal dashed line marks the stability boundary for solutions with no gravity waves.

List of Tables

1	Results for the coastally-trapped waves in a circular basin. The FVCOM and ROMS results are reproduced from IMCS (2006).	38
2	Parameters for the Gaussian bump test case.	39

Model	Resolution (km)	Maximum C	Angle of max C ($^{\circ}$)
Gerris	37.5	0.9823	3.7
	18.75	0.9954	0.3
	9.375	0.9974	0.1
	4.6875	1.0000	0.0
Gerris adaptive	37.5	0.9766	3.7
	18.75	0.9904	0.2
	9.375	1.0040	0.0
	4.6875	1.0310	0.0
FVCOM	40	0.9921	-4.6
	20	0.9934	-1.6
	10	0.9993	-0.5
	5	0.9999	-0.2
ROMS	40	0.9801	-35.5
	20	0.9909	-18.2
	10	0.9971	-9.3
	5	0.9986	-4.9
Gerris half-refined	9.375/18.75	0.9965	0.3

Table 1: Results for the coastally-trapped waves in a circular basin. The FVCOM and ROMS results are reproduced from IMCS (2006).

Grid size		$256 \times 128 \times 8$
Ocean depth H	m	4500
Height of bump H_0	m	4050
Length scale of bump L	km	25
Stratification $\frac{NH}{fL}$		1.5
Barotropic in-flow u_0	cm s ⁻¹	25
Coriolis parameter f	s ⁻¹	10^{-4}
Timestep	s	600

Table 2: Parameters for the Gaussian bump test case.