

A quadtree-adaptive spectral wave model

BY STÉPHANE POPINET¹, RICHARD M. GORMAN²,
GRAHAM J. RICKARD¹ AND HENDRIK L. TOLMAN³

¹*National Institute of Water and Atmospheric research,
P.O. Box 14-901, Kilbirnie, Wellington, New Zealand*

²*National Institute of Water and Atmospheric research,
P.O. Box 11-115, Hamilton, New Zealand*

³*NOAA/NCEP/EMC Marine Modeling and Analysis Branch,
5200 Auth Road Room 209, Camp Springs, MD 20746, USA*

April 14, 2010

Abstract

A spectral wave model coupling a quadtree-adaptive discretisation of the two spatial dimensions with a standard discretisation of the two spectral dimensions is described. The implementation is greatly simplified by reusing components of the Gerris solver (for spatial advection on quadtrees) and WAVEWATCH III (for spectral advection and source terms). Strict equivalence between the anisotropic diffusion and spatial filtering methods for alleviation of the Garden Sprinkler Effect (GSE) is demonstrated. This equivalence facilitates the generalisation of GSE alleviation techniques to quadtree grids. For the case of a cyclone-generated wave field, the cost of the adaptive method increases linearly with spatial resolution compared to quadratically for constant-resolution methods. This leads to decreases in runtimes of one to two orders of magnitude for practical spatial resolutions. Similar efficiency gains are shown to be possible for global spectral wave forecasting.

Keywords: Gerris; WAVEWATCH III; Garden Sprinkler Effect; Cyclone-generated waves

1 Introduction

Spectral wave models are the primary means of forecasting global-scale, wind-forced, ocean wave fields (Hasselmann et al., 1988; Booij et al., 1999; Tolman, 1991). Initially designed for the prediction of deep-water waves they have been extended over the past decade to include some of the processes affecting shallow-water surface gravity waves such as refraction and shallow-water breaking (Komen et al., 1996). With these improvements these models are theoretically applicable – given suitable forecast wind fields – for spatial scales spanning at least four orders of magnitude: from ocean basin scales (thousands of kilometres) down to coastal scales (hundreds of metres). In practice however the computational efficiency of existing models severely limits the range of spatial scales accessible. In this paper we will investigate how adaptive methods can help lift some of these limitations.

Spectral wave models typically describe the evolution of the wave action density spectrum $N(k, \theta, \mathbf{x}, t)$ with k the wavenumber, θ the propagation direction, \mathbf{x} the spatial coordinate and t the time. The evolution equation is a four-dimensional (two spatial dimensions \mathbf{x} and two spectral dimensions (k, θ)) advection equation which can be written

$$\partial_t N + \nabla \cdot (\dot{\mathbf{x}} N) + \partial_k(\dot{k} N) + \partial_\theta(\dot{\theta} N) = S, \quad (1)$$

where $\dot{\mathbf{x}}$, \dot{k} and $\dot{\theta}$ are the advection velocities in spatial and spectral spaces respectively and S contains parameterised source terms such as wind forcing, dissipation due to breaking and non-linear wave interactions. Advection in spectral space is only used to represent shallow-water processes and is not necessary for deep-water waves (on a Cartesian grid: on a spherical grid, there still is a “refraction” term representing great-circle propagation (Hasselmann et al., 1988)). The high computational cost of the method is due to the high dimension of the wave action density spectrum field. The spectral dimensions θ and k are typically discretised using at least 20 increments each so that equation (1) requires the solution of over 400 independent advection problems in \mathbf{x} space.

In classical spectral wave models such as WAM (Hasselmann et al., 1988), WAVEWATCH IIITM (Tolman et al., 2002; Tolman, 2009) and SWAN (Booij et al., 1999; Ris et al., 1999), the two spatial dimensions \boldsymbol{x} can be discretised on a regular Cartesian grid and the spatial resolution is thus constant in space and time. This is simple but not optimal as wave fields often have varying spatial scales. This scale variability reflects a similar variability of the forcing fields such as wind (e.g. localised weather systems, fronts etc.) and bathymetry (coastlines, islands, shallow banks etc.) (Tolman, 2007; Tolman, 2008). The same limitation applies to regular latitude–longitude grids, widely used in large-scale applications, which have spatially-variable resolution arising solely as a function of latitude with no relationship to the length scales of relevant forcing fields. Some flexibility can be gained by using regular grids at various resolutions coupled using either one-way (Tolman, 1991) or two-way coupling (Tolman, 2008). Another approach is to discretise the spatial dimensions using unstructured meshes which allows good flexibility in the description of static features such as coastlines and bathymetry (Hsu et al., 2005; Qi et al., 2009).

It is important however to make the distinction between these variable-spatial-resolution methods (hereafter referred to as *static adaptivity*) and *adaptive* methods. In contrast to static adaptivity, adaptive methods also allow the spatial resolution to adjust as the solution evolves in time. While static adaptivity can reduce the computational cost of an accurate description of static features such as coastlines, it cannot do the same for dynamic features such as weather and associated wave systems. Adaptive methods are now a well-established tool in engineering, numerical analysis and theoretical fluid mechanics (Popinet, 2003; Popinet, 2009e), however their application to geophysical fluid dynamics problems is still a work-in-progress (Kurihara et al., 1995; Bacon et al., 2000; Pain et al., 2005; Popinet and Rickard, 2007; Rickard et al., 2009). In the context of spectral wave models, we are aware only of the “moving grid” approach developed for WAVEWATCH III (Tolman and Alves, 2005).

In the following, we will describe a method combining an adaptive quadtree discretisation of spatial advection (using Gerris (Popinet, 2009d; Popinet, 2009c)) with the standard spectral description implemented within WAVEWATCH III. Using a model test case of wave generation by a cyclone we will assess the gains in efficiency obtained with the adaptive method.

2 Numerical method

The implementation of the numerical method is based on coupling Gerris and WAVEWATCH III. Spatial advection ($\nabla \cdot (\boldsymbol{x} N)$ term in equation (1)) is solved using Gerris while spectral advection and source terms are computed using the corresponding routines in WAVEWATCH III. Neither Gerris nor WAVEWATCH III were modified to perform this coupling. Consequently we will only give a summary of the method and refer the reader to existing publications on Gerris and WAVEWATCH III for more details (Popinet, 2003; Popinet, 2009e; Tolman, 1991; Tolman and Chalikov, 1996; Tolman et al., 2002; Tolman, 2009).

2.1 Spatial discretisation

For two spatial dimensions Gerris discretises space using a quadtree data structure (Figure 1). This is a classical data structure in computational geometry and image processing which combines simplicity with a large degree of flexibility (Samet, 1989). This is also a hierarchical data structure where each “parent cell” has four “children cells” and so on recursively. Cells are easily and efficiently removed or added by moving up or down the hierarchy. Gerris uses finite-volume algorithms where quantities are interpreted as cell averages. In this context the quadtree data structure also has the advantage that the description of a field on each level of the hierarchy is a consistent, conservative embedding of the field description on the other levels.

All spectral components are interpreted as cell averages and are discretised on a single quadtree grid so that the (varying) spatial resolution is identical for all spectral components.




Figure 1. Quadtree data structure. (left) Geometrical description. (right) Logical description. The numbers indicate the “level” of the corresponding cell.

2.2 Advection scheme

The advection scheme in Gerris is an extension to quadtrees of the Bell-Colella-Glaz (BCG) Godunov scheme (Colella, 1990). This scheme is conservative and positive definite which guarantees positivity of the wave action densities. In contrast to other classical advection schemes (such as QUICKEST used in WAVEWATCH III), the BCG scheme is *unsplit* i.e. advection is performed simultaneously in both space dimensions and transverse fluxes are taken into account. The Godunov scheme is based on the reconstruction of the advected field on both sides of the boundary of a finite volume using only volume-averaged values. These two reconstructed values do not generally match and upwinding is then used to compute the final flux through the boundary.

Considering a tracer field $c(\mathbf{x}, t)$ with $\mathbf{x} \equiv x \mathbf{i} + y \mathbf{j}$, the leading terms of a Taylor series for the value on the boundary $(\mathbf{x} + \Delta x \mathbf{i})$ at time $t + \Delta t$ can be written

$$c(\mathbf{x} + \Delta x \mathbf{i}, t + \Delta t) = c(\mathbf{x}, t) + \Delta x \partial_x c + \Delta t \partial_t c + O(\Delta x^2, \Delta t^2). \quad (2)$$

Using the advection equation, the temporal derivative can be replaced by spatial derivatives yielding

$$c(\mathbf{x} + \Delta x \mathbf{i}, t + \Delta t) = c(\mathbf{x}, t) + (\Delta x - \Delta t u) \partial_x c - \Delta t v \partial_y c, \quad (3)$$

where $\mathbf{u} \equiv (u, v)$ is the advection velocity vector. The derivatives ∂_x and ∂_y are estimated using second-order centered operators generalised for a quadtree (Popinet, 2003). The formula for the boundary $(\mathbf{x} + \Delta y \mathbf{j})$ of the finite volume is derived in a similar way. It is clear that any of the four boundaries of the square finite volume are shared with each of the four neighbouring finite volumes, so that independent boundary values can be computed by applying (3) to either a cell or its neighbour. Once these two values have been computed, the final boundary value is obtained by simple upwinding (Colella, 1990). The advection flux through a given boundary is then just the product of the boundary length, the boundary value, the corresponding component of the advection velocity and the timestep.

The extension of this scheme to quadtrees is relatively straightforward. Care must be taken so that fluxes at the boundary between a “fine” and a “coarse” cell are consistent. This is easily achieved by computing the total flux at the boundary as the sum of each of the fine cell fluxes. A difficulty arises in the case of a general advection velocity field because the same consistency must be maintained for the (boundary-centered) velocity field used to compute the fluxes. In the case of the spatially-constant advection velocity fields which interest us here, consistency is trivially assured however.

The resulting advection scheme is monotonic, conservative, second-order in space and time and unconditionally stable for Courant-Friedrichs-Lewy (CFL) numbers smaller than one i.e.

$$\Delta t < \min(\Delta x/u, \Delta y/v) \quad (4)$$

2.3 Time stepping

For a wave field, the spatial advection velocity is given by

$$\dot{\mathbf{x}} \equiv \mathbf{c}_g + \mathbf{U}, \quad (5)$$

with $\mathbf{c}_g = c_g(k) (\cos \theta, \sin \theta)$, $c_g(k)$ the group velocity and \mathbf{U} the current velocity. For depths large compared to the wavelength, c_g is given by $\sqrt{g/k}$ with g the acceleration of gravity, so that long waves travel faster than short waves. The CFL stability condition (4) will then impose smaller timesteps for long waves than for short waves and the computational cost can be reduced by advecting each of the discretised action densities $N(k_i)$ with a different timestep, as is done in WAVEWATCH III. Of course the components will then need to be synchronised before computing the source terms (and intra-spectral advection, which is not implemented here). This leads to the following time-stepping algorithm

Algorithm 1. *Time stepping (global timestep Δt_g)*

1. **for each** k_i (advection cycle within Gerris)
 1. Compute maximum timestep Δt_a as $\min(\Delta x/u, \Delta y/v)$ where the minimum is taken over all the (variable resolution) cells of the quadtree, and $u = v = c_g(k_i)$
 2. $t' \leftarrow t$
 3. **while** $t' < t + \Delta t_g$ (advection subcycle with timestep Δt_a)
 1. **for each** θ_j
 1. Initialize $\dot{\mathbf{x}} = c_g(k_i) (\cos \theta_j, \sin \theta_j)$
 2. Solve $\partial_t N(k_i, \theta_j, \mathbf{x}, t) + \nabla \cdot \dot{\mathbf{x}} N(k_i, \theta_j, \mathbf{x}, t) = 0$ using the BCG advection scheme
 3. If GSE alleviation is required solve equation (19) with $F = N(k_i, \theta_j, \mathbf{x}, t)$
 2. Adapt the mesh
 3. $t' \leftarrow t + \Delta t_a$
2. **for each leaf cell** \mathbf{x} (source term cycle within WAVEWATCH III)
 1. Compute initial dynamic timestep Δt_n
 2. $t' \leftarrow t$
 3. **while** $t' < t + \Delta t_g$ (source term subcycle with dynamic timestep Δt_n)
 1. **for each** k_i and θ_j
 1. Compute the source term S and the derivative D of its diagonal terms
 2. Update actions

$$N(k_i, \theta_j, \mathbf{x}, t') \leftarrow N(k_i, \theta_j, \mathbf{x}, t) + \frac{S(k_i, \theta_j, \mathbf{x}, t')}{1 - D(k_i, \theta_j, \mathbf{x}, t') \Delta t_n}$$

2. Revise the dynamic timestep Δt_n

$$3. t' \leftarrow t + \Delta t_n$$

$$3. t \leftarrow t + \Delta t_g$$

The source term cycle uses a dynamic timestepping method (Tolman, 1992) implemented by the existing WAVEWATCH III source term code.

Note that mesh adaptation is performed within the advection subcycle. This is necessary to ensure that the adaptive mesh properly follows the structures of individual spectral components as they are advected. Performing adaptivity in the outer loop (advection cycle) could result in fast components (i.e. requiring several subcycles) moving from high-resolution to low-resolution areas within the advection subcycles, with the associated consequence of numerical diffusion of fine-scale structures.

2.4 Adaptivity

Several criteria can be used to adjust the local spatial resolution and Gerris provides a simple interface for user-defined criteria. A criterion is typically formulated as a cost function combined with a threshold. The mesh is adapted by “tagging” cells on the finest level of the quadtree hierarchy (the “leaf cells” of the quadtree) for which the value of the cost function is larger than the threshold value. Conversely “non-leaf” cells are tagged if the value of their cost function is smaller than the threshold multiplied by a “hysteresis factor” (a number smaller than one).

All the tagged leaf cells are then refined i.e. four new children cells are created in the hierarchy. The values of variables such as action densities need to be initialised in these new cells. To do so conservatively, the default is to compute the new values using a simple second-order-accurate interpolation of the form

$$N(\mathbf{x} + \Delta\mathbf{x}) = N(\mathbf{x}) + \nabla N(\mathbf{x}) \cdot \Delta\mathbf{x}, \quad (6)$$

where $N(\mathbf{x})$ and $\nabla N(\mathbf{x})$ are the value and gradient at the center of the cell being refined and $\Delta\mathbf{x}$ is the relative position of the center of the child cell i.e. $\Delta\mathbf{x} = (\pm \Delta x, \pm \Delta y)/4$. This guarantees that the sum of children values is equal to four times the parent value which ensures conservation.

Coarsening proceeds in a similar manner by deleting the children of all the tagged non-leaf cells. The values of the variables for parent cells are always set to the average values of their children so that conservation is maintained after deletion. The “hysteresis factor” mentioned above minimises the occurrence of “cyclic cases” when a cell having just been coarsened then violates the adaptation criterion at the next timestep and needs to be (re)refined.

For a given adaptation subcycle, the number of cells being created or deleted is typically small compared to the total number of cells and the overall computational cost of the adaptation algorithm is negligible.

In this article we have used an adaptation criterion based on an a posteriori estimate of the “truncation error” of the second-order spatial discretisation. Given a field F , to third-order accuracy the Taylor expansion of F around \mathbf{x} can be written

$$F(\mathbf{x} + \mathbf{r}) = F(\mathbf{x}) + \mathbf{r} \mathbf{F}'(\mathbf{x}) + \frac{1}{2} \mathbf{r}^T \mathbf{F}''(\mathbf{x}) \mathbf{r}, \quad (7)$$

with the gradient \mathbf{F}' and Hessian \mathbf{F}'' given by

$$\mathbf{F}' \equiv \begin{pmatrix} \frac{\partial F}{\partial x} \\ \frac{\partial F}{\partial y} \end{pmatrix}, \quad (8)$$

$$\mathbf{F}'' \equiv \begin{pmatrix} \frac{\partial^2 F}{\partial x^2} & \frac{\partial^2 F}{\partial x \partial y} \\ \frac{\partial^2 F}{\partial x \partial y} & \frac{\partial^2 F}{\partial y^2} \end{pmatrix}. \quad (9)$$

If we now assume that F is known exactly at the center of each cell and that the numerical scheme is exactly second-order accurate, to third-order accuracy the error in the numerical representation of F will be bounded by

$$\max_{\mathbf{x}} \left\| \frac{1}{2} \mathbf{r}(\mathbf{x})^T \mathbf{F}''(\mathbf{x}) \mathbf{r}(\mathbf{x}) \right\|, \quad (10)$$

with $\mathbf{r}(\mathbf{x})$ the relative position of any point contained within a cell of center \mathbf{x} . This just expresses the fact that – at leading order – the *curvature* of F controls the error. For square cells each coordinate of \mathbf{r} is within $[-\Delta x/2, \Delta x/2]$ and a simplified upper bound can be expressed as

$$\max_{\mathbf{x}} \frac{\Delta x^2}{8} \|\mathbf{F}''(\mathbf{x})\|. \quad (11)$$

Of course the exact Hessian of F is not known a priori, but we can compute an a posteriori estimate using finite differences. A difficulty arises however because the standard finite-differences operators we are using are only second-order accurate in the general case (which is what led to our error estimate in the first place). If we naively try to use these operators to estimate \mathbf{F}'' , discretisation errors will be comparable to the value sought. A solution would be to use higher-order operators to estimate \mathbf{F}'' , however constructing such operators on quadtrees is difficult (because the wide stencils required lead to the multiplication of specific cases dependent on local mesh sizes). A simpler solution is obtained by noting that the standard, centered, finite-difference, curvature operator on regular Cartesian grids

$$\frac{\partial^2 F}{\partial x^2} \approx \frac{F_{i-1} - 2F_i + F_{i+1}}{\Delta x^2} \quad (12)$$

is in fact third-order accurate (the second-order errors cancel out due to symmetries). By construction the generalised operator we use on quadtrees reduces to the regular operator for regions of the quadtree where resolution is constant (see Section 4.1 of (Popinet, 2003)) and can thus be used to compute an estimate of \mathbf{F}'' to third-order accuracy in the regions of constant resolution (see Figure 2). When there is a transition in spatial resolution, the accuracy of the generalised operator decreases to second-order and something needs to be done to obtain consistent estimates of \mathbf{F}'' . A simple solution in these regions is to approximate \mathbf{F}'' with its value computed at a coarser level of refinement. For any cell, it is always possible to find an ‘‘ancestor’’ whose neighbors have identical resolution so that \mathbf{F}'' can be estimated to third-order accuracy. In most cases it is sufficient to consider only the parent of a cell in a transition region. This means of course that in the transition regions some of the information (on the finer cells) is discarded when estimating \mathbf{F}'' , however we have found in practice that this scheme leads to robust and accurate error estimates while being simple to implement.

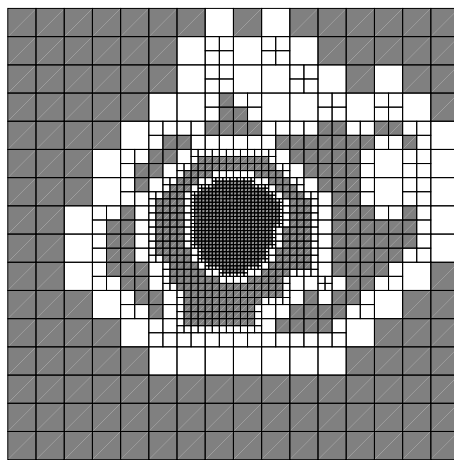


Figure 2. Example of quadtree mesh used for the cyclone problem of section 3.2. For the constant-resolution regions shown in gray the generalised finite-difference curvature operator is third-order accurate, whereas it is only second-order accurate in the transition regions shown in white.

For the spectral wave model, this approach could be used to control the truncation error of each of the spectral components. As a single mesh is used for all spectral components, the adaptation criterion would need to aggregate all the estimated truncation errors: for example taking the maximum of the truncation error over all spectral components. In this article we have used a simpler (but potentially less accurate) approach which does not require the estimation of truncation errors for each component. We chose to control the truncation error on the significant wave height H_s , thus aggregating all the spectral components before computing the adaptation criterion (rather than aggregating all the truncation errors). Note that because mesh adaptation is performed within the advection subcycle of Algorithm 1, the individual spectral components used to compute H_s for adaptivity are not necessarily defined at the same time (some have already been advected while others have not). While the resulting value for H_s is only consistent at the end of the subcycle (after all components have been advected) we do not expect this approximation to have a significant impact on the quality of adaptation.

2.5 “Garden sprinkler effect” alleviation

The discretisation of the direction of propagation of waves (i.e. θ_j) leads to the problem known as “garden sprinkler effect” (GSE) (Booij and Holthuijsen, 1987; Tolman, 2002; Chawla and Tolman, 2008). Packets of waves propagating in discrete directions eventually separate and lead to an unphysical, fragmented wave field. This problem can be avoided by increasing the directional resolution $\Delta\theta$ sufficiently. A lower bound on the directional resolution necessary can be obtained by considering two wave packets, originating from the same source but travelling in directions separated by $\Delta\theta$ (Booij and Holthuijsen, 1987). After having travelled a distance L , the size of the simulation domain, these two packets will be separated by $L \Delta\theta$. If this separation distance is smaller than the spatial resolution Δx , no GSE will be apparent. This leads to the criterion

$$\Delta\theta < \Delta x/L. \tag{13}$$

Keeping in mind the cost of high spectral resolution mentioned earlier, this is a very restrictive criterion particularly for high spatial resolutions. This has motivated the development of “GSE alleviation” methods which seek to minimise GSE while using reasonably low directional resolution. WAVEWATCH III implements two such alleviation methods: anisotropic diffusion (Booij and Holthuijsen, 1987) and spatial filtering (Tolman, 2002). Spatial filtering was developed by Tolman as a more efficient alternative to anisotropic diffusion; the justification being that a time-explicit discretisation of the diffusion operator leads to an unacceptable constraint on the timestep (Tolman, 2002) (we will show below that this argument is in fact incorrect).

In the context of this article, we would like to extend these alleviation techniques to the quadtree spatial discretisation. The spatial filtering method of Tolman uses a stencil specific to regular Cartesian grids and as such is difficult to generalise to quadtree grids. The anisotropic diffusion method on the other hand can be discretised using standard, conservative, diffusion flux operators which we have already implemented within Gerris to solve diffusion equations (Popinet, 2009e). This consideration – together with the observation that “spatial filtering” or “spatial averaging” methods are often strictly equivalent to some form of diffusion – motivated the discussion which follows. We will demonstrate that the spatial filtering method of (Tolman, 2002) is equivalent – to third-order accuracy – to the anisotropic diffusion method of (Booij and Holthuijsen, 1987) with different diffusion coefficients.

The spatial filtering technique applied to a field F can be expressed (equation (A.5) of (Tolman, 2002))

$$F_{\text{avg}}(\mathbf{x}) = \frac{1}{3}F(\mathbf{x}) + \frac{1}{6} \sum_{n=1}^4 F(\mathbf{x} + \mathbf{r}_n), \tag{14}$$

with

$$\begin{aligned} \mathbf{r}_1 &= \mathbf{s}' + \mathbf{n}', \\ \mathbf{r}_2 &= -\mathbf{s}' + \mathbf{n}', \\ \mathbf{r}_3 &= -\mathbf{s}' - \mathbf{n}', \\ \mathbf{r}_4 &= \mathbf{s}' - \mathbf{n}' \end{aligned} \tag{15}$$

and

$$\begin{aligned} \mathbf{s}' &= \alpha \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix}, \\ \mathbf{n}' &= \beta \begin{pmatrix} -\sin \theta \\ \cos \theta \end{pmatrix}. \end{aligned} \quad (16)$$

Using the Taylor expansion (7) in (14) gives the third-order accurate expression

$$\begin{aligned} F_{\text{avg}}(\mathbf{x}) &= \frac{1}{3} F(\mathbf{x}) + \frac{1}{6} \sum_{n=1}^4 \left(F(\mathbf{x}) + \mathbf{r}_n \mathbf{F}'(\mathbf{x}) + \frac{1}{2} \mathbf{r}_n^T \mathbf{F}''(\mathbf{x}) \mathbf{r}_n \right) \\ &= F(\mathbf{x}) + \frac{1}{6} \mathbf{F}'(\mathbf{x}) \sum_{n=1}^4 \mathbf{r}_n + \frac{1}{12} \sum_{n=1}^4 \mathbf{r}_n^T \mathbf{F}''(\mathbf{x}) \mathbf{r}_n \\ &= F(\mathbf{x}) + \frac{1}{12} \sum_{n=1}^4 \mathbf{r}_n^T \mathbf{F}''(\mathbf{x}) \mathbf{r}_n \end{aligned} \quad (17)$$

Replacing the \mathbf{r}_n with their values in (15) and simplifying gives

$$\begin{aligned} F_{\text{avg}}(\mathbf{x}) &= F(\mathbf{x}) + \\ &\quad \frac{1}{3} (\alpha^2 \cos^2 \theta + \beta^2 \sin^2 \theta) F_{xx} + \\ &\quad \frac{1}{3} (\alpha^2 \sin^2 \theta + \beta^2 \cos^2 \theta) F_{yy} + \\ &\quad \frac{2}{3} (\alpha^2 - \beta^2) \cos \theta \sin \theta F_{xy}, \end{aligned} \quad (18)$$

which can be rewritten

$$\frac{F_{\text{avg}}(\mathbf{x}) - F(\mathbf{x})}{\Delta t} = D_{xx} F_{xx} + D_{yy} F_{yy} + 2 D_{xy} F_{xy}, \quad (19)$$

with

$$\begin{aligned} D_{xx} &\equiv D_{ss} \cos^2 \theta + D_{nn} \sin^2 \theta, \\ D_{yy} &\equiv D_{ss} \sin^2 \theta + D_{nn} \cos^2 \theta, \\ D_{xy} &\equiv (D_{ss} - D_{nn}) \cos \theta \sin \theta, \end{aligned} \quad (20)$$

and

$$\begin{aligned} D_{ss} &\equiv \frac{\alpha^2}{3 \Delta t}, \\ D_{nn} &\equiv \frac{\beta^2}{3 \Delta t}. \end{aligned} \quad (21)$$

This is a first-order-in-time discretisation of the anisotropic diffusion equation of Booij and Holthuijsen (but with different diffusion coefficients D_{ss} and D_{nn}). To third-order accuracy the spatial-filtering scheme is thus formally equivalent to a first-order time discretisation of the anisotropic diffusion equation. This scheme is stable provided $\Delta t < \Delta x^2 / D_{ss}$ and $\Delta t < \Delta x^2 / D_{nn}$ which can be simplified as

$$\begin{aligned} \alpha^2 &< 3 \Delta x^2, \\ \beta^2 &< 3 \Delta x^2. \end{aligned} \quad (22)$$

For the spatial filtering scheme α and β are set to (Tolman, 2002)

$$\begin{aligned} \alpha &\equiv \alpha_s \Delta c_g \Delta t, \\ \beta &\equiv \alpha_n c_g \Delta \theta \Delta t, \end{aligned} \quad (23)$$

with $\Delta c_g \equiv (\gamma - \gamma^{-1}) c_g / 2$. After simplification the stability constraints then becomes

$$\begin{aligned} \Delta t &< \frac{\sqrt{3} \Delta x}{\alpha_s \Delta c_g}, \\ \Delta t &< \frac{\sqrt{3} \Delta x}{\alpha_n c_g \Delta \theta}, \end{aligned} \tag{24}$$

which is similar to the CFL stability condition (4) but with CFL numbers of $2\sqrt{3}/(\alpha_s(\gamma - \gamma^{-1}))$ and $\sqrt{3}/(\alpha_n \Delta \theta)$ respectively. For typical values of α_s (1.5), α_n (1.5), γ (1.1) and $\Delta \theta$ ($2\pi/24$) these CFL numbers are larger than one and do not restrict the timestep compared to (4).

The anisotropic diffusion equation (19) is then easily discretised on the quadtree grid using the finite-volume, conservative, diffusion operators already implemented in Gerris (Popinet, 2009e).

2.6 Implementation

The Gerris source code is designed as a reusable, object-oriented library of functions which facilitates the implementation of new models. The implementation of the adaptive wave model only involved writing code for the *advection subcycle* of Algorithm 1 and for the time-integration of the anisotropic diffusion equation (19). Advection itself was performed using appropriate calls to the existing functions implementing the BCG advection scheme. All the other algorithms described or used in this article: mesh creation and adaptation, error estimates, visualisation etc. are not specific to the wave model and were simply reused. Gerris is also parallelised with dynamic load-balancing for adaptive simulations and these functionalities extend transparently to the adaptive spectral wave model via the object-oriented code structure.

While WAVEWATCH III is not designed as an extensible framework for developing models, it was relatively simple to isolate the function calls necessary to initialise and compute the source terms S . Although these functions are typically designed to operate on the whole spatial array of action densities, initialising a dummy WAVEWATCH III model containing only a single grid point allowed us to circumvent this particular problem. The (Fortran) WAVEWATCH III function computing source terms can then be called directly by the (C) Gerris function adding source terms, provided the Gerris data structures describing the wave spectrum at a given spatial location are first converted into their WAVEWATCH III (Fortran arrays) equivalents (and conversely after they have been updated by the WAVEWATCH III function). This “glue code” is implemented as a dynamically-loaded module which does not require any modifications or recompilation of either Gerris or WAVEWATCH III.

3 Results

3.1 Garden sprinkler effect

We first reproduce the GSE test case of (Booij and Holthuijsen, 1987) and (Tolman, 2002) as a verification that the diffusion formulation (19) is effective at alleviating the GSE when coupled with adaptive refinement. An initial wave field with a spectral distribution given by

$$N(f, \theta) \propto \exp\left(-\frac{(f - f_\mu)^2}{2 f_\sigma^2}\right) \cos^2(\theta - \theta_\mu), \tag{25}$$

propagates in deep water during five days. The initial significant wave height is given by

$$H_s(x, y) = H_0 \exp\left(-\frac{x^2 + y^2}{2 r_\sigma^2}\right). \tag{26}$$

The parameters match those of (Tolman, 2002): $f_\mu = 0.1$ Hz, $f_\sigma = 0.02$ Hz, $\theta_\mu = 30^\circ$, $H_0 = 2.5$ m, $r_\sigma = 150$ km. The source terms are not included and the simulation reduces to spatial advection with or without GSE alleviation so that WAVEWATCH III routines are not used for this test case.

Figure 3 summarises the results obtained using the spectral wave model implemented within Gerris. This figure should be compared with the corresponding figures (1 and 5) in Tolman, 2002. Aside from minor differences close to the domain boundaries the two sets of simulations closely match. The reference solution 3.(b) is well reproduced. An almost identical GSE is obtained when GSE alleviation is turned off (Figure 3.(a)) and disappears as α_n and α_s are increased (Figures 3.(c) and (d)). Although the agreement between Figures 3.(c) and (d) and Figures 5.(a) and (b) of (Tolman, 2002) is excellent, it is only so for values of α_n and α_s which are exactly double those of Tolman (this is also the case for the cyclone-generated waves in Figure 7). This is somewhat frustrating as solving this discrepancy would provide further confirmation that the analysis of section 2.5 is indeed correct (and that the anisotropic diffusion and spatial filtering schemes are strictly equivalent in practice as well as in theory). We carefully checked the derivation in section 2.5, its implementation within Gerris as well as the implementation of spatial filtering within WAVEWATCH III and ruled out any trivial implementation error. A possible explanation could be the increased numerical diffusion induced by the lower-order discretisation of the spatial interpolation operators used in the spatial filtering implementation.




Figure 3. Significant wave height for the GSE test case of Tolman, 2002. Results for the adaptive model at $t = 5$ days, contours at 0.1 metre intervals. Axis in kilometres. Unless otherwise specified the directional resolution is $\Delta\theta = 15^\circ$ and the maximum spatial resolution is $\Delta x_{\max} = 100$ km. (a) No GSE alleviation. (b) No GSE alleviation, $\Delta\theta = 3^\circ$ and $\Delta x_{\max} = 25$ km. (c) GSE alleviation, $\alpha_n = \alpha_s = 1.5$. (d) GSE alleviation, $\alpha_n = \alpha_s = 3$.

Figures (a), (b), (c) and (d) should be compared with figures 1.(a), 1.(b), 5.(a) and 5.(b) of (Tolman, 2002) respectively.

The main difference between this set of simulations and those in (Tolman, 2002) is that rather than using a constant spatial resolution, the resolution is adapted according to the a-posteriori-error adaptivity criterion of section 2.4, applied to H_s . To maintain a sufficient resolution of the weak final wave field the threshold for refinement (i.e. the maximum truncation error on H_s) was set to 2.5 mm. A typical evolution of the wave field and corresponding mesh is given in Figure 4.

The parameters and scripts necessary to reproduce this test case are available on the Gerris web site (Popinet, 2009b).

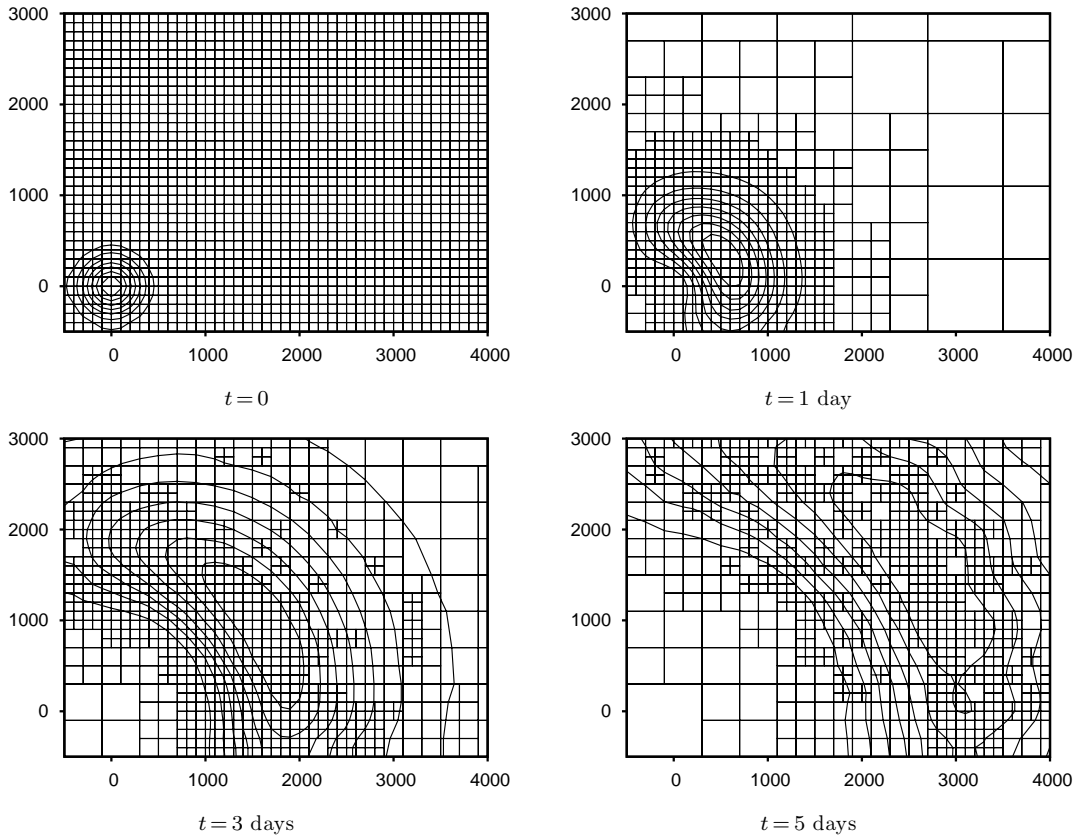


Figure 4. Evolution with time of the significant wave height H_s and corresponding adaptive mesh. Seven contours equally distributed between the minimum and maximum values of H_s at each time. The parameters are those of Figure 3.(d). The a posteriori error adaptivity criterion of section 2.4 is used for H_s with a threshold of 2.5 mm.

3.2 Cyclone-generated wave field

We now consider the evolution of a wave field generated by a localised weather pattern: a model cyclone. The domain is a square centered on the origin, of size 3328 km and described using Cartesian coordinates (curvature is not taken into account). A simplified radially-symmetric southern-hemisphere ‘‘Holland’’ cyclone model (Holland, 1983) is used for the wind field, consisting of clockwise rotating winds (with no radial component), with wind speed of the form

$$v(r, t) = v_{\max}(t) \left(\frac{r_{\max}}{r} \right)^2 \exp\left(2 \left(1 - \frac{r_{\max}}{r} \right) \right), \quad (27)$$

with r the distance to the center of the cyclone, t the time and $r_{\max} = 100$ km. The maximum velocity was set to a constant value after an initial linear spinup, i.e.

$$v_{\max}(t) = v_0 \min(1, t/t_0), \quad (28)$$

with $v_0 = 50$ m/s and $t_0 = 25$ hours.

The cyclone moves ‘‘south’’ (down) at a constant velocity of 555 km/day starting from coordinates (0, 1110) km. Water is uniformly deep so that the only source terms (computed using WAVEWATCH III) are: wind-wave interaction, nonlinear wave-wave interactions and dissipation through whitecapping. We note that WAVEWATCH III also includes an optional wind correction for atmospheric stability but this was not used in this study. The initial wave growth (starting from zero) was described using the ‘‘linear input’’ model of (Cavaleri and Malanotte-Rizzoli, 1981) as implemented in WAVEWATCH III.

This scenario was run using both Gerris/WAVEWATCH III and WAVEWATCH III only. Both adaptive and constant resolution simulations were performed with Gerris. For adaptive simulations the a-posteriori-error criterion of section 2.4 was used on both the wind field (with a threshold of 0.2 m/s) and the significant wave height (with a threshold of 0.1 metre). The global timestep Δt_g was set to 15 minutes. WAVEWATCH III used a constant spatial resolution of $\Delta x = 13$ km, which was also applied (uniformly) in constant-resolution Gerris simulations, and as the finest resolution limit (Δx_{\max}) for adaptive Gerris runs. The spectral grid consisted of 25 logarithmically-spaced frequencies $f_n = 1.1 f_{n-1}$, with $f_1 = 0.04$ Hz. In most cases 24 directional bins were used (i.e. $\Delta\theta = 15^\circ$).

Figure 5 illustrates the typical evolution of the significant wave height with time obtained using both the constant- and adaptive-resolution Gerris runs (left and centre column respectively). The simulation reproduces the shift between the extrema in significant wave height and in maximum wind speed (Figure 6) previously observed in radar measurements and simulations (Young and Burchell, 1996; Young, 2006). There is a local minimum of significant wave height located to the north of the eye, while the maximum in significant wave height is located to the south-east of the cyclone’s eye (for this clockwise Southern Hemisphere cyclone). This is a result

of the “extended fetch” mechanism proposed by King and Shemdin (King and Shemdin, 1978) which preferentially enhances waves propagating with velocities close to the translation velocity of the storm. Aside from the stronger GSE for the constant resolution at $t = 48$ hours (bottom-left corner of Figure 5), the agreement between the constant- and adaptive-resolution results is excellent.

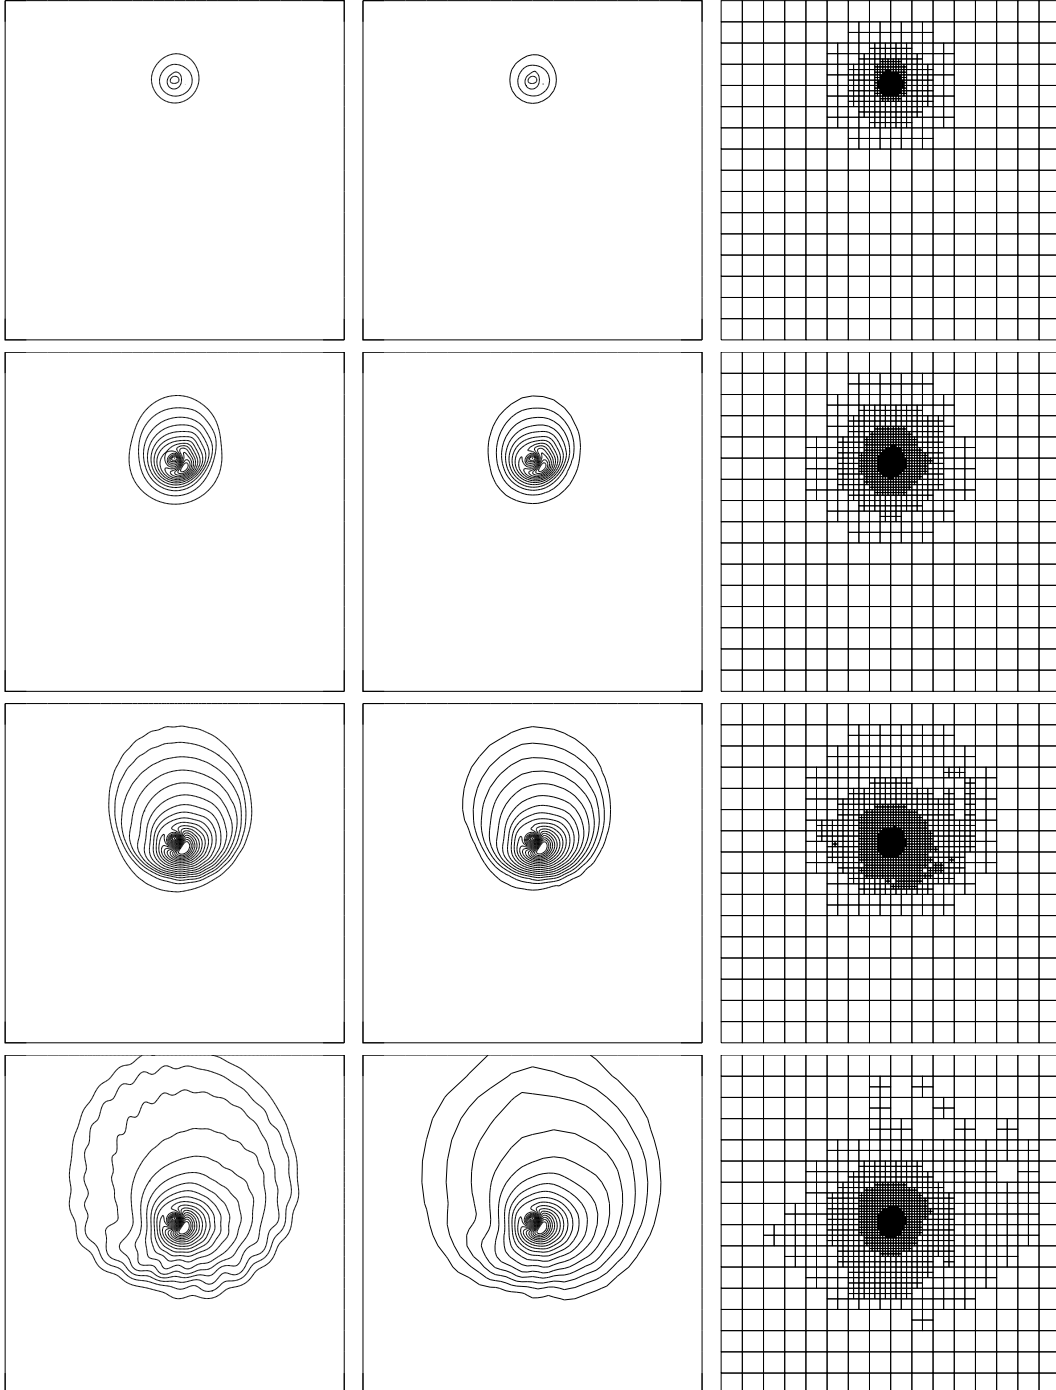


Figure 5. Evolution of a cyclone-generated wave field. Significant wave height contours at one metre interval. Rows from top to bottom correspond to $t = 12, 24, 36$ and 48 hours respectively. Left column: constant resolution $\Delta x = 13$ km. Central column: adaptive resolution, maximum spatial resolution $\Delta x_{\max} = 13$ km. Right column: adaptive mesh. GSE alleviation with $\alpha_s = \alpha_n = 4$.




Figure 6. Detail of the significant wave height (black) and wind field (red). Adaptive Gerris result. Maximum spatial resolution $\Delta x_{\max} = 13$ km. $t = 48$ hours.

To further validate these results Figure 7 gives a comparison of the results obtained using Gerris/WAVEWATCH III and WAVEWATCH III only. Agreement is excellent but the same discrepancy in the values of α_s and α_n as in section 3.1 is necessary.

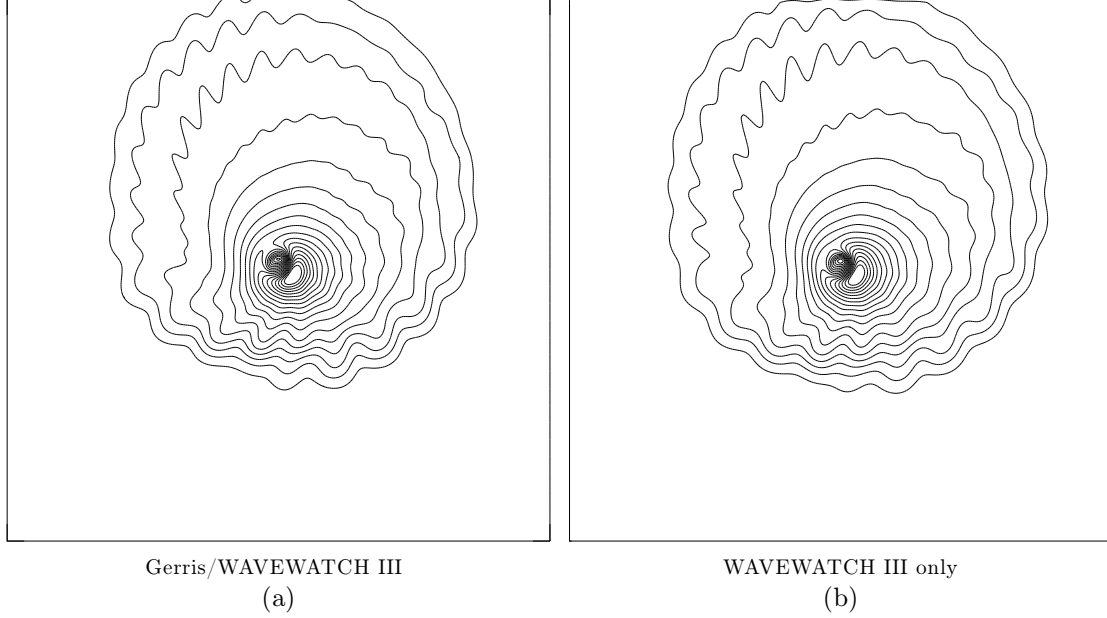


Figure 7. Comparison of solutions obtained with Gerris/WAVEWATCH III (a) and WAVEWATCH III only (b). $t = 36$ hours. Constant spatial resolution $\Delta x = 13$ km. GSE alleviation is used in both models with $\alpha_s = \alpha_n = 3$ for Gerris/WAVEWATCH III and $\alpha_s = \alpha_n = 1.5$ for WAVEWATCH III only.

Running Gerris with 15° directional resolution and without GSE alleviation, a strong Garden Sprinkler Effect develops after $t = 36$ hours (Figure 8.(a)) and makes either GSE alleviation (Figure 8.(b)) or higher directional resolution (Figure 8.(c)) mandatory. The results of Figure 8 confirm that adaptive refinement combined with the anisotropic diffusion scheme is effective at alleviating the GSE, while comparing well with the high-directional-resolution reference result.




Figure 8. Comparison of simulations with and without GSE alleviation. $t = 48$ hours. Maximum spatial resolution $\Delta x_{\max} = 26$ km. (a) Gerris, constant resolution, no GSE alleviation, $\Delta\theta = 15^\circ$. (b) Gerris, adaptive resolution, $\alpha_s = \alpha_n = 4$, $\Delta\theta = 15^\circ$. (c) Gerris, constant resolution, no GSE alleviation, $\Delta\theta = 6^\circ$.

Diffusion (or spatial filtering) is synonymous with higher numerical dissipation which could lead to the prediction of smaller maximum wave heights. The evolution of the maximum significant wave height illustrated in Figure 9 shows that GSE alleviation does not have a significant impact (somewhat counter-intuitively the predicted maximum significant wave height is actually *higher* with GSE alleviation than without). The good agreement between Gerris/WAVEWATCH III and WAVEWATCH III-only is also confirmed.

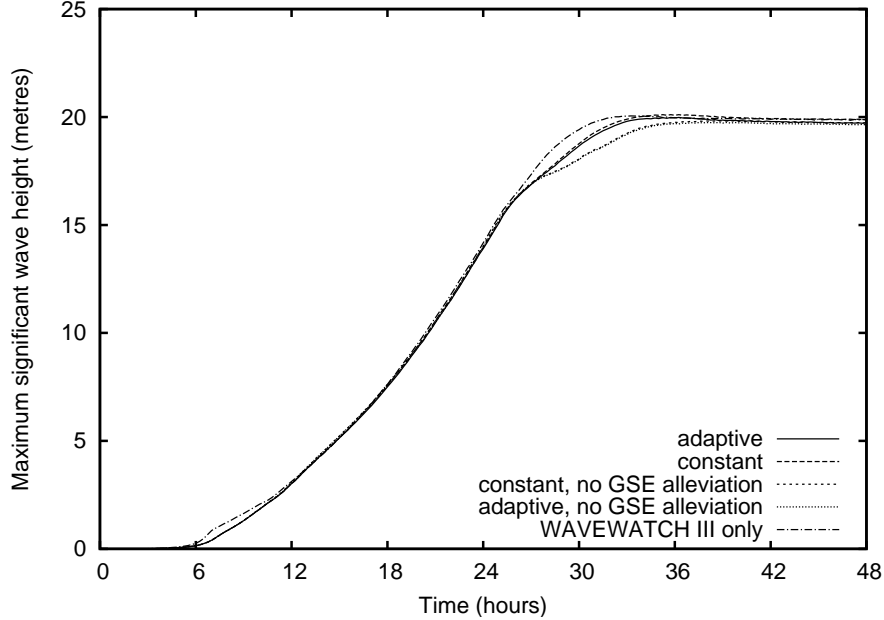


Figure 9. Evolution of the maximum significant wave height. $\Delta x_{\max} = 13$ km.

To study the influence of spatial resolution, we repeated the simulations for maximum resolutions of 26 and 6.5 km. The results are summarised in Figure 10. Finer details are indeed captured at higher resolutions close to the center of the cyclone (bottom row). As expected the GSE is stronger at higher spatial resolutions (left column, due to violation of criterion (13)). To estimate quantitatively the agreement between the constant- and adaptive-resolution results we computed the difference in significant wave height between each set of runs. Figure 11 gives the evolution of the corresponding maximum and Root-Mean-Square (RMS) difference norms as a function of time. While the maximum difference is somewhat higher than the 0.1 metre threshold set for adaptivity this nonetheless confirms that the a-posteriori-error estimate is appropriate for controlling the error (closer examination also reveals that the largest differences are associated with areas where the GSE is stronger in the constant-resolution simulation than in the adaptive-resolution simulation, as is visible in Figure 10, middle row).

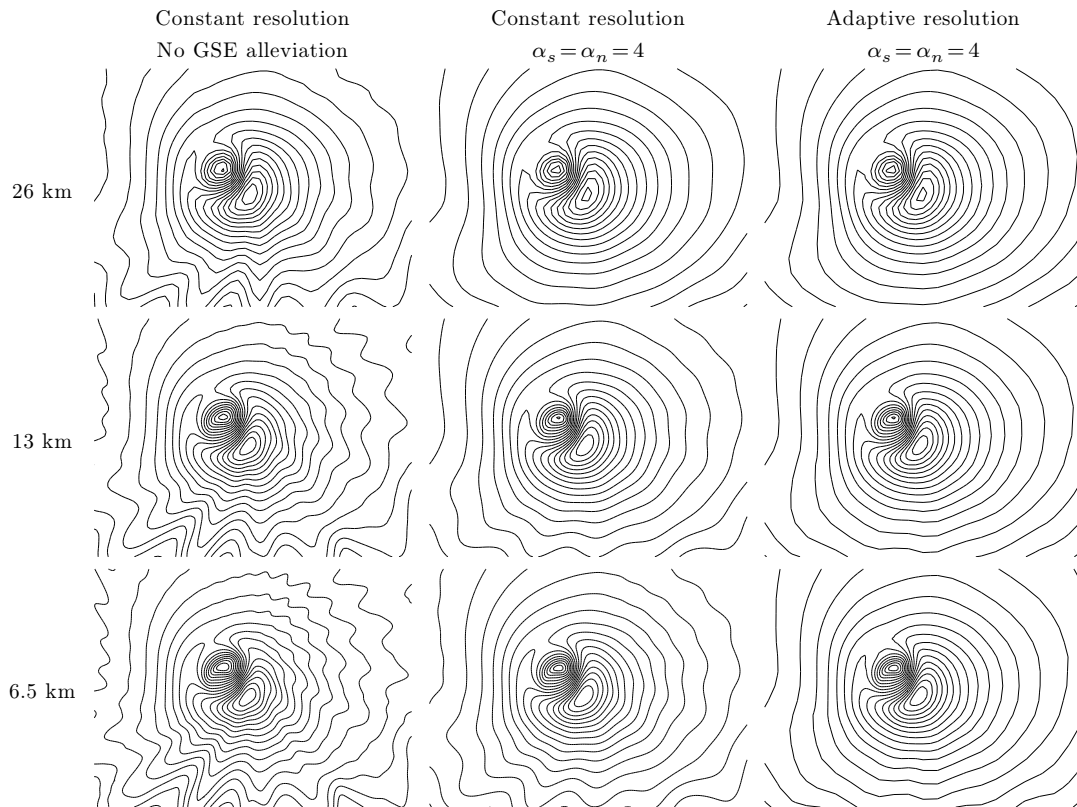


Figure 10. Detail of the significant wave height close to the center of the cyclone for varying resolutions (vertically) and discretisation schemes (horizontally). Significant wave height contours at one metre interval. $t = 48$ hours.




Figure 11. Evolution of the Root-Mean-Square (RMS) and maximum differences in significant wave height between the adaptive resolution and constant resolution simulations. $\Delta x_{\max} = 13$ km. $\alpha_s = \alpha_n = 4$. The adaptive error criterion was set to 0.1 metre.

Finally Table 1 and figure 12 give the total computing times for each simulation, run on an Intel PC 32-bits 2.6 GHz. The constant-resolution Gerris/WAVEWATCH III runs are significantly slower than the equivalent runs with WAVEWATCH III-only due to the overhead of the quadtree grid structure compared to regular arrays (within the advection and GSE-alleviation schemes).

For this particular example using adaptivity leads to gains in efficiency of factors comprised between 15 and 78 depending on resolution. The cost of constant-resolution runs increases approximately quadratically with spatial resolution; which is expected if the overall cost of the algorithm is dominated by the cost of computing source terms. In contrast the cost of the adaptive runs only increases linearly with (maximum) spatial resolution, so that the gain in efficiency also increases linearly with spatial resolution.

This change in scaling mirrors the particular structure of the solution and will be different for other problems. More specifically, this scaling can be directly related to the fractal (or information) dimension of the solution: the information content of the cyclone-generated wave field scales like the information content of a curve (rather than a surface) which then leads to a correspondingly more efficient encoding using the quadtree structure compared to the regular grid. In the next section we extend this scaling analysis to a more “typical” operational wave forecasting scenario.

Note that the scripts necessary to reproduce the results presented in this section are available on the Gerris web site (Popinet, 2009a).

Maximum spatial resolution	26 km	13 km	6.5 km
WAVEWATCH III only	180	702	3260
Gerris/WAVEWATCH III constant	315	1350	5400*
Gerris/WAVEWATCH III adaptive	12	20	42
Ratio WAVEWATCH III only/Gerris adaptive	15	35	78

Table 1. Computing times in minutes for different models and varying maximum spatial resolutions. GSE alleviation is applied with $\alpha_s = \alpha_n = 3$ (corresponding to $\alpha_s = \alpha_n = 1.5$ for WAVEWATCH III only). (*) Estimated.

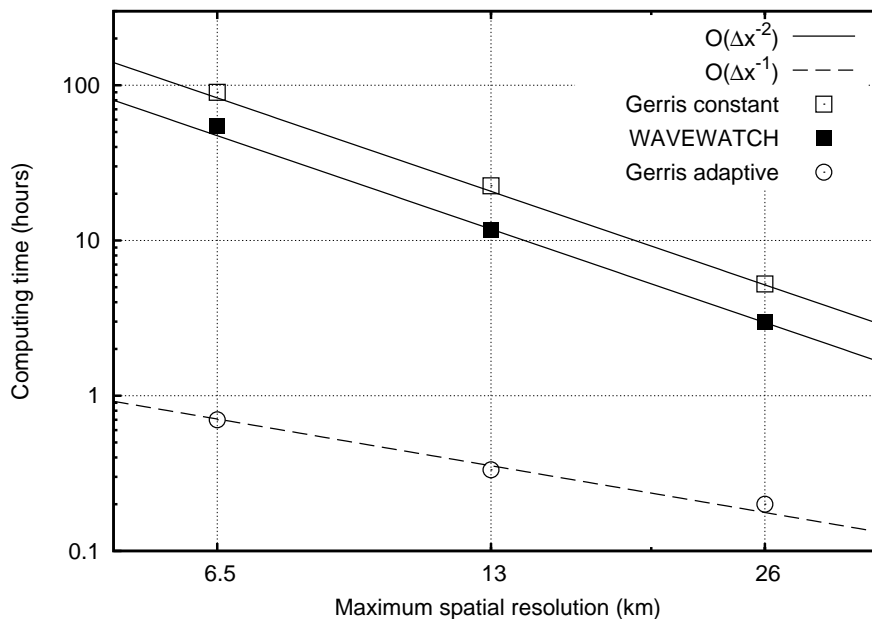


Figure 12. Computing times as functions of spatial resolution for different models.

3.3 Scaling analysis for a global forecast snapshot

Figure 13 illustrates an example of a significant-wave-height field derived from the global output at half-degree resolution from the NOAA/NCEP multigrid wave model (Chawla et al., 2007). The analysis (0-hour forecast) field for September 1st 2009, 00:00 UTC was selected arbitrarily. It is fairly typical of global wave distributions in August/September with strong weather systems and high wave amplitudes at mid-latitudes in the southern hemisphere, as well as similar (but generally weaker) systems at mid-latitudes in the northern hemisphere.

While some further development (mostly technical) is required before the adaptive model described here can be used to reproduce such forecasts, we can use the scaling analysis methodology described in the previous section to estimate the gains in efficiency expected from the adaptive model. To do so, the results from the NOAA forecast have been interpolated on a constant-resolution Gerris quadtree longitude/latitude grid with 1024×512 grid points corresponding to a spatial resolution of 0.35 degrees, slightly higher than the original data (Figure 13). The total number of sea-covered grid points is 313,410.




Figure 13. Significant wave height, constant resolution, 1024×512 (0.35 degrees). Contours at 1 m intervals. 313,410 grid points.

Figure 14 illustrates the truncation error estimated using equations (11) and (12). Note that this truncation error estimate is not specific to the adaptive model. It is valid for any model using a second-order-accurate representation of the wave field (including the original NOAA global wave model). This first analysis reveals a number of interesting features. First of all the truncation error is highly variable: large areas have negligible truncation errors while a significant fraction of the domain has truncation errors larger than 5 cm. High-truncation-error areas

are not necessarily correlated with high-wave-amplitude areas but rather reveal the variability in the underlying forcings. For example small features which are not obvious in the wave-height field (Figure 13) become so in the truncation-error field: the influence of islands (e.g. Maldives, Polynesian, Hawaii, Aleutian, Galapagos, Caribbean, Azores, Falkland...) is particularly clear. These arise either through grid-scale variation in model bathymetry, or through WAVE-WATCH's representation of subgrid-scale obstruction by island chains (Tolman, 2003). Further to these *static* geographic features, the signatures of *dynamic* weather systems also appear more clearly such as the weather fronts in depressions, orographic wind jets etc.




Figure 14. Truncation error, constant resolution, 1024×512 (0.35 degrees). Logarithmic color scale. Dark blue is < 1 mm, dark red is larger than 5 cm.

The initially-constant discretisation of the significant wave height field can then be coarsened or refined locally according to this error estimate. This is equivalent to applying the adaptation algorithm of section 2.4 to the significant wave height field only, to find the mesh that an adaptive model would use to continue the integration from this point. Choosing an adaptation threshold of 5 cm and setting the maximum equivalent resolution to 1024×512 (i.e. coarsening but never refining the initial field) leads to the wave field and adaptive mesh illustrated in Figures 15 and 16. Aside from the obvious coarsening of parts of the land mask, the field in Figure 15 closely matches the original field of Figure 13 while being discretised with 10 times fewer grid points.




Figure 15. Significant wave height, adaptive resolution, maximum equivalent resolution 1024×512 (0.35 degrees). Adaptation threshold for discretisation error: 5 cm. 31,207 grid points.




Figure 16. Level of adaptive refinement for Figure 15. Dark red corresponds to an equivalent resolution of 1024×512 (0.35 degrees), dark blue 64×32 (5.6 degrees).

Changing the distribution of spatial resolution obviously leads to a change in the distribution of truncation errors. This is summarised by the “cumulative area-fraction distributions” of Figure 17 which represents the fraction of the total area described with a truncation error larger than a given value. The blue curves are the distributions for a constant discretisation with three different uniform spatial resolutions. The red curves are the corresponding adaptive distributions. The curve for a 1024×512 constant resolution shows that $\sim 20\%$ of the domain is described with a truncation error smaller than 1 mm, $\sim 0.5\%$ larger than 5 cm and $\sim 0.1\%$ larger than 10 cm. This is consistent with our earlier description of Figure 14.

Using an adaptive discretisation limited to a 1024×512 maximum equivalent resolution leads to the distribution drawn using a solid red line. As cells with low truncation errors (i.e. “over-resolved cells”) are coarsened, the distribution converges toward the “ideal” adaptive distribution: all cells having the same truncation error i.e. a step function centered on the adaptation threshold. In practice the step is rounded rather than sharp due to the added constraint on the quadtree that the resolutions of neighbouring cells are either equal or differ by a factor of two (Popinet, 2003). Also, as refinement of the initial constant resolution field is not allowed, adaptivity cannot improve the truncation error of cells which are initially under-resolved and the adaptive distribution thus matches the constant distribution when the error is above the 5-cm threshold.

Note that for the cyclone test case we found that setting a maximum truncation error smaller than 10 cm was necessary to avoid excessive numerical diffusion of the solution. For the weaker wave and wind fields of the global forecast, it is likely that even a 5-cm truncation error will lead to significant numerical diffusion of the associated features. This means that a significant fraction ($\sim 0.5\%$, dark red areas in Figure 14) of the initial forecast solution is under-resolved.

Doubling the resolution of the constant-discretisation representation uniformly shifts the distributions towards lower errors by a factor of four (blue curves). A resolution of 4096×2048 (~ 0.1 degrees) is necessary to describe the wave field with a truncation error smaller than 5 cm (almost) everywhere. For this constant resolution $\sim 95\%$ of the total area is described with a discretisation error smaller than 1 mm. Increasing the maximum resolution of the adaptive discretisation does not change the distribution below the threshold significantly (aside from a slightly more rounded step as explained above) but allows refinement of under-resolved areas and thus convergence toward a truncation error of 5 cm everywhere.




Figure 17. Cumulative area-fraction distribution of truncation errors. The value of the function $f(x)$ is the fraction of the total area described with a truncation error larger than x . Adaptation threshold 0.05 m.

Finally Figure 18 gives the number of grid points for the constant and adaptive discretisations as functions of spatial resolution (a coarse-grid $512 \times 256/0.7$ degrees data point has also been added). The curves are comparable to those in Figure 12 but with a smaller exponent (0.45 instead of 1) for the adaptive discretisation. Gains in number of grid points range from factors of between 5 and 100.




Figure 18. Number of grid points as functions of spatial resolution for constant-resolution and adaptive grids. Adaptation threshold 0.05 m.

While this analysis relies on a static snapshot of a global forecast, it gives confidence that the gains in efficiency demonstrated for the time-dependent cyclone can realistically be expected also for operational wave forecasting. Also note that, while a half-degree resolution is very high for a global operational model, it is unlikely to be sufficient to accurately represent localised weather patterns such as the cyclone wind/wave field of the previous section. The analysis presented in the current section thus relies on an under-resolved model wave field which is likely to underestimate the true natural spatial variability. As such we expect a gain of a factor of 10 in number of grid points (Figure 18) to be a conservative estimate of potential gains for operational scenarios.

4 Conclusion

We have described a method coupling a quadtree-adaptive discretisation of the spatial dimensions (within Gerris) with a regular discretisation of the spectral dimensions (within WAVEWATCH III). We also demonstrated the formal and practical equivalence of GSE-alleviation methods using either anisotropic diffusion or spatial filtering. The anisotropic diffusion formulation leads to a simple implementation of the GSE-alleviation method on quadtree grids. For the case of a cyclone-generated wave field, the cost of the adaptive method increases linearly with spatial resolution compared to quadratically for constant-resolution methods. This leads to decreases in runtimes of one to two orders of magnitude for practical spatial resolutions. An a posteriori scaling analysis of a global spectral wave forecast shows that efficiency gains of this order can reasonably be expected also in an operational context.

Some further development is required to make the present Gerris/WAVEWATCH III model suitable for “real world” applications. For the most part, these are relatively straightforward implementations of model “infrastructure”, such as the handling of environmental forcing fields, model outputs and the specification of incident wave fields at the boundaries of limited-area models. Extension of the advection algorithms to handle shallow water, currents and subgrid physics are also desirable, but not expected to involve new scientific challenges. These developments will largely reuse existing components in both Gerris and WAVEWATCH III in the same spirit as in this article. Gerris also implements distributed-memory parallelism with load-balancing and initial tests suggest that this approach should extend to Gerris/WAVEWATCH III in a straightforward manner. The accurate representation of complex boundaries (Popinet, 2003; Popinet and Rickard, 2007) will also extend naturally to the discretisation of spectral wave fields near coastlines within the quadtree-adaptive framework.

In operational forecasting applications, spectral wave models are always used as part of a forecasting system, coupled at least to a 3-dimensional weather model and in some cases to other models as well. Because of the need to solve for a full directional spectrum in each spatial cell, the wave model contributes a relatively high fraction of the computational load of the system. This computational load is dominated by evaluation of source terms, and in particular the nonlinear interaction term. This is the case even when a heavily-simplified algorithm such as the Discrete Interaction Approximation is used (as was used in all simulations in the present work, as well as in most operational forecasting applications), let alone using any of the more exact techniques presently available (such as the Webb–Resio–Tracy method) (van Vledder, 2006; Webb, 1978; Resio and Perrie, 1991; Tracy and Resio, 1982). As a result, computation times scale approximately in proportion to the number of spatial grid cells used, so large gains in efficiency can be made by the use of adaptive techniques to limit the number of spatial cells required.

We have shown in this article that efficiency gains greater than an order of magnitude can be expected for realistic wave fields. These gains reflect the natural scale-separation of the processes being simulated which – in the case of wave fields – is essentially induced by the spatial scale distribution of forcing processes: for example the scale distribution (or fractal dimension) of bathymetry/coastlines, or of weather patterns (cyclones and fronts).

In an operational context, these gains could be realised by setting adaptivity criteria so that the wave model can operate within its allocated computing resources and schedule, by applying a range of spatial resolutions up to the full horizontal resolution of the forcing atmospheric model. An adaptive wave model would also be well suited to applications specifically aimed at cyclone forecasting, in which blended atmospheric products are often developed combining operational weather model outputs with additional cyclone-specific data (e.g. from satellites and aircraft) and models (Cardone and Cox, 2009). In this situation a generic adaptive model could be immediately applied to accurately simulate the associated wave conditions, with less “cyclone-specific” setup required than in existing approaches, such as the use of a nested “moving grid” (Tolman and Alves, 2005).

A further advance might be to apply adaptive methods to other components of the forecast system (Popinet and Rickard, 2007), so that model resolution is matched to the scales of relevant physical processes throughout the coupled atmosphere-ocean system.

Acknowledgements

This work was partly funded by a NIWA Innovation Seed Fund project, and by the New Zealand Foundation for Research, Science and Technology under contract CO1X0812. This development was made possible by the availability of source codes for both Gerris (Popinet, 2009c; Popinet, 2009d) and WAVEWATCH III.

Bibliography

- [Bacon et al., 2000] D.P. Bacon, N.N. Ahmad, Z. Boybeyi, T.J. Dunn, M.S. Hall, P.C.S. Lee, R.A. Sarma, M.D. Turner, K.T. Waight III, S.H. Young, et al. A dynamically adapting weather and dispersion model: the operational multiscale environment model with grid adaptivity (OMEGA). *Monthly Weather Review*, 128 (7): 2044–2076, 2000.
- [Booij and Holthuijsen, 1987] N. Booij and L.H. Holthuijsen. Propagation of ocean waves in discrete spectral wave models. *Journal of Computational Physics*, 68: 307–326, 1987.
- [Booij et al., 1999] N. Booij, R.C. Ris, and L.H. Holthuijsen. A third-generation wave model for coastal regions. I- Model description and validation. *Journal of Geophysical Research*, 104 (C4): 7649–7666, 1999.
- [Cardone and Cox, 2009] V. Cardone and A. Cox. Tropical cyclone wind field forcing for surge models: critical issues and sensitivities. *Natural Hazards*, 51 (1): 29–47, 2009.
- [Cavaleri and Malanotte-Rizzoli, 1981] L. Cavaleri and P. Malanotte-Rizzoli. Wind wave prediction in shallow water: Theory and applications. *Journal of Geophysical Research*, 86 (10): 961–10, 1981.
- [Chawla and Tolman, 2008] A. Chawla and H.L. Tolman. Obstruction grids for spectral wave models. *Ocean Modelling*, 22 (1-2): 12–25, 2008.
- [Chawla et al., 2007] A. Chawla, D. Cao, V. M. Gerald, T. Spindler, and H. L. Tolman. Operational implementation of a multi-grid wave forecast system. In *10th international workshop on wave hindcasting and forecasting & coastal hazard symposium*, 2007. Paper B3.
- [Colella, 1990] P. Colella. A multidimensional second order Godunov scheme for conservation laws. *Journal of Computational Physics*, 87: 171–200, 1990.
- [Hasselmann et al., 1988] S. Hasselmann, K. Hasselmann, E. Bauer, L. Bertotti, C. V. Cardone, J. A. Ewing, J. A. Greenwood, A. Guillaume, P. A. E. M. Janssen, G. J. Komen, P. Lionello, M. Reistad, and L. Zambresky. The WAM Model - a third generation ocean wave prediction model. *Journal of Physical Oceanography*, 18 (12): 1775–1810, 1988.
- [Holland, 1983] G.J. Holland. Tropical cyclone motion: Environmental interaction plus a beta effect. *J. Atmos. Sci.*, 40: 328–342, 1983.
- [Hsu et al., 2005] T.W. Hsu, S.H. Ou, and J.M. Liau. Hindcasting nearshore wind waves using a FEM code for SWAN. *Coastal Engineering*, 52 (2): 177–195, 2005.
- [King and Shemdin, 1978] D.B. King and O.H. Shemdin. Radar observations of hurricane wave directions. In *16th International Conference on Coastal Engineering*, Hamburg, Germany, 1978. ICCE.
- [Komen et al., 1996] G. J. Komen, L. Cavaleri, M. Donelan, K. Hasselmann, S. Hasselmann, and P. A. E. M. Janssen. *Dynamics and Modelling of Ocean Waves*. Cambridge University Press, 1996.

- [Kurihara et al., 1995] Y. Kurihara, M.A. Bender, R.E. Tuleya, and R.J. Ross. Improvements in the GFDL hurricane prediction system. *Monthly Weather Review*, 123 (9): 2791–2801, 1995.
- [Pain et al., 2005] C.C. Pain, M.D. Piggott, A.J.H. Goddard, F. Fang, G.J. Gorman, D.P. Marshall, M.D. Eaton, P.W. Power, and C.R.E. De Oliveira. Three-dimensional unstructured mesh ocean modelling. *Ocean Modelling*, 10 (1-2): 5–33, 2005.
- [Popinet, 2009a] S. Popinet. Gerris – cyclone-generated wave field example, 2009a. URL <http://gfs.sourceforge.net/examples/examples/cyclone.html>.
- [Popinet, 2009b] S. Popinet. Gerris – "garden sprinkler effect" in wave model example, 2009b. URL <http://gfs.sourceforge.net/examples/examples/garden.html>.
- [Popinet, 2009c] S. Popinet. Gerris wavewatch module, 2009c. URL http://gfs.sourceforge.net/wiki/index.php/Object_hierarchy#Wavewatch.
- [Popinet, 2009d] S. Popinet. The Gerris Flow Solver, 2009d. URL <http://gfs.sf.net>.
- [Popinet, 2003] S. Popinet. Gerris: a tree-based adaptive solver for the incompressible Euler equations in complex geometries. *Journal of Computational Physics*, 190 (2): 572–600, 2003. URL <http://gfs.sf.net/gerris.pdf>.
- [Popinet, 2009e] S. Popinet. An accurate adaptive solver for surface-tension-driven interfacial flows. *Journal of Computational Physics*, 228: 5838–5866, 2009e. URL <http://gfs.sf.net/papers/tension.pdf>.
- [Popinet and Rickard, 2007] S. Popinet and G. Rickard. A tree-based solver for adaptive ocean modelling. *Ocean Modelling*, 16: 224–249, 2007. URL <http://gfs.sf.net/ocean.pdf>.
- [Qi et al., 2009] J. Qi, C. Chen, R.C. Beardsley, W. Perrie, G.W. Cowles, and Z. Lai. An unstructured-grid finite-volume surface wave model (FVCOM-SWAVE): Implementation, validations and applications. *Ocean Modelling*, 2009.
- [Resio and Perrie, 1991] D.T. Resio and W. Perrie. A numerical study of nonlinear energy fluxes due to wave-wave interactions. Part 1. Methodology and basic results. *Journal of Fluid Mechanics*, 223: 603–629, 1991.
- [Rickard et al., 2009] G. Rickard, J. O’Callaghan, and S. Popinet. Numerical simulations of internal solitary waves interacting with uniform slopes using an adaptive model. *Ocean Modelling*, 30: 16–28, 2009. URL <http://gfs.sf.net/papers/isw.pdf>.
- [Ris et al., 1999] R.C. Ris, L.H. Holthuijsen, and N. Booij. A third-generation wave model for coastal regions 2. Verification. *Journal of Geophysical Research*, 104 (C4): 7667–7681, 1999.
- [Samet, 1989] H. Samet. *Applications of Spatial Data Structures*. Addison-Wesley Publishing Company, 1989.
- [Tolman, 1991] H. L. Tolman. A third-generation model for wind waves on slowly varying, unsteady, and inhomogeneous depths and currents. *Journal of Physical Oceanography*, 21: 782–797, 1991.
- [Tolman, 1992] H. L. Tolman. Effects of numerics on the physics in a third-generation wind-wave model. *Journal of Physical Oceanography*, 22: 1095–1111, 1992.
- [Tolman, 2003] H. L. Tolman. Treatment of unresolved islands and ice in wind wave models. *Ocean Modelling*, 5: 219–231, 2003.
- [Tolman, 2009] Hendrik L. Tolman. User manual and system documentation of WAVEWATCH IIITM version 3.14. Technical Note 276, NOAA/NWS/NCEP/MMAB, 2009.
- [Tolman, 2002] H.L. Tolman. Alleviating the garden sprinkler effect in wind wave models. *Ocean Modelling*, 4 (3-4): 269–289, 2002.
- [Tolman, 2007] H.L. Tolman. Development of a multi-grid version of WAVEWATCH III. Technical Note 256, NOAA/NWS/NCEP/MMAB, 2007.
- [Tolman, 2008] H.L. Tolman. A mosaic approach to wind wave modeling. *Ocean Modelling*, 25 (1-2): 35–47, 2008.
- [Tolman and Alves, 2005] H.L. Tolman and J.H.G.M. Alves. Numerical modeling of wind waves generated by tropical cyclones using moving grids. *Ocean Modelling*, 9 (4): 305–323, 2005.
- [Tolman and Chalikov, 1996] H.L. Tolman and D. Chalikov. Source terms in a third-generation wind wave model. *Journal of Physical Oceanography*, 26 (11): 2497–2518, 1996.
- [Tolman et al., 2002] H.L. Tolman, B. Balasubramanian, L.D. Burroughs, D.V. Chalikov, Y.Y. Chao, H.S. Chen, and V.M. Gerald. Development and Implementation of Wind-Generated Ocean Surface Wave Modelsat NCEP. *Weather and Forecasting*, 17 (2): 311–333, 2002.
- [Tracy and Resio, 1982] Barbara A. Tracy and Donald T. Resio. Theory and calculations of the nonlinear energy transfer between sea waves in deep water. WES Report 11, US Army Engineer Waterways Experiment Station, 1982.

- [van Vledder, 2006] Gerbrant Ph. van Vledder. The WRT method for the computation of non-linear four-wave interactions in discrete spectral wave models. *Coastal Engineering*, 53 (2-3): 223–242, 2006.
- [Webb, 1978] D.J. Webb. Nonlinear transfer between sea waves. *Deep Sea Research*, 25: 279–298, 1978.
- [Young and Burchell, 1996] I. R. Young and G. P. Burchell. Hurricane generated waves as observed by satellite. *Ocean Engineering*, 23 (8): 761–776, 1996.
- [Young, 2006] Ian R. Young. Directional spectra of hurricane wind waves. *Journal of Geophysical Research*, 111 (C8): C08020, 2006.