

Quadtree-adaptive tsunami modelling

BY STÉPHANE POPINET

National institute of water and atmospheric research,
P.O. Box 14-901, Kilbirnie, Wellington, New Zealand

May 10, 2011

Abstract

The well-balanced, positivity-preserving scheme of Audusse et al, 2004, for the solution of the Saint-Venant equations with wetting and drying, is generalised to an adaptive quadtree spatial discretisation. The scheme is validated using an analytical solution for the oscillation of a fluid in a parabolic container, as well as the classic Monai tsunami laboratory benchmark. An efficient database system able to dynamically reconstruct a multiscale bathymetry based on extremely large datasets is also described. This combination of methods is successfully applied to the adaptive modelling of the 2004 Indian ocean tsunami. Adaptivity is shown to significantly decrease the exponent of the power law describing computational cost as a function of spatial resolution. The new exponent is directly related to the fractal dimension of the geometrical structures characterising tsunami propagation. The implementation of the method as well as the data and scripts necessary to reproduce the results presented are freely available as part of the open-source Geris Flow Solver framework.

1 Introduction

Numerical modelling of the propagation and run-up of tsunamis is an important tool for risk assessment, evacuation planning and more recently, real-time forecasting for tsunami warning systems. A common modelling strategy relies on numerical methods designed for the solution of the Saint-Venant (or non-linear shallow-water) equations [de Saint-Venant (1871)]. These two-dimensional equations are valid for a wide range of environmental problems besides tsunamis such as river flows [Liang et al. (2004), Liang and Borthwick (2009)], tides and storm surges. As a consequence a vast literature exists on numerical schemes for their solution [Audusse et al. (2004), Liang et al. (2004), LeVeque and George (2006), Walters (2006), George and LeVeque (2008), Liang and Borthwick (2009)]. In the context of tsunamis (and coastal flows more generally), two important issues related to numerical methods arise: 1) accurate handling of equilibrium solutions, and 2) robust description of wetting and drying near the coastline. These issues are particularly relevant for tsunamis (e.g. compared to river flows) as the solutions sought are typically small perturbations to the lake-at-rest equilibrium solution, and the handling of wetting and drying will determine the description of tsunami run-up. While many authors have used somewhat ad hoc schemes to deal with these issues, it is only recently that numerical methods have been derived which formally preserve the lake-at-rest condition and the positivity of the water depth in the case of wetting and drying [Audusse et al. (2004), LeVeque and George (2006)].

In addition to these requirements on numerical schemes, tsunamis are also difficult to model because they occur over a wide range of spatial scales. Propagation can take place over entire ocean basins while the destructive impact at the coast can be influenced by details as small as a few meters (e.g. sea defences). While high-resolution regular Cartesian (longitude/latitude) grids are still widely used for the solution of the Saint-Venant or Boussinesq equations [Wang and Liu (2006), Ioualalen et al. (2007), Grilli et al. (2007)], their computational cost (e.g. the wall-clock computation time) typically scales like

$$C \Delta_{\star}^{-3} \quad (1)$$

with C a constant and Δ_{\star} the ratio between the smallest and the largest length scales resolved by the Cartesian grid (i.e. a dimensionless spatial resolution). The exponent is the number of dimensions of the problem (two spatial dimensions plus time). With such a scaling practical high-resolution models requires large computational resources combined with parallel programming. It also limits the operational forecasting capability of tsunami models.

An alternative to Cartesian grids is to use fully-unstructured grids, often composed of triangular elements. This type of discretisation offers much greater flexibility and can be used to create variable-resolution meshes where, for example, static coastal details are represented at high resolution while the deep ocean uses a much coarser mesh [Walters (2006), Harig et al. (2008)]. An important issue with such a static resolution however, is that accurate propagation of tsunami waves in the deep ocean can also require high resolution, so that predicting beforehand which areas can be resolved with a coarser mesh is difficult.

One solution is to use *adaptive* methods where the resolution continuously adjusts to follow features in the flow. Combining adaptivity with unstructured meshes is possible but relatively complex and the results of such methods for tsunami modelling are only preliminary [Pranowo et al. (2008)]. An alternative is to use dynamic embedding of Cartesian meshes of different resolutions, also called *block adaptivity* [Berger and Oliger (1984)]. This has been applied successfully to tsunami modelling by George and LeVeque [George and LeVeque (2008)]. Block adaptivity works well when relatively large blocks of uniform resolution can be combined to track medium- to fine-scale features in the flow. A different approach is to keep Cartesian cells as basic discretisation volumes, while allowing neighbouring cells with a different spatial resolution. Through a generalisation of the basic discrete operators (gradient, divergence etc...), solvers can be constructed which work across resolution boundaries. This *fine-scale* adaptivity allows a more flexible tracking of flow features than block adaptivity. The price to pay is a higher overhead (compared to pure Cartesian methods) in areas of uniform spatial resolution. An elegant way to organise the resulting collection of cells is to use a *quadtree* hierarchical structure. Methods based on quadtrees (and octrees in 3D) have been successfully used to compute adaptive solutions of many problems, such as the 2D and 3D Navier-Stokes equations [Yiu et al. (1996), Popinet (2003b), Popinet (2009)], or shallow-water equations [Popinet and Rickard (2007)]. In the context of the Saint-Venant equations, the work of Liang and Borthwick using quadtrees is noteworthy [Liang et al. (2004), Liang and Borthwick (2009)] but seems to have been applied mainly to river flows.

In this article, I will describe how a quadtree-adaptive Saint-Venant solver was implemented within the Gerris Flow Solver [Popinet (2003a)] using a generalisation of the well-balanced, positivity-preserving scheme of Audusse et al. (2004). Results for several classical validation test cases will be presented as well as a realistic application to the 2004 Indian ocean tsunami. This practical case required an efficient multiscale framework for the storage and retrieval of high-resolution bathymetry data at runtime. The algorithms required for this framework as well as a performance analysis will also be presented. Finally, I will discuss how fine-scale adaptivity leads to a scaling of computational efficiency much more favourable than (1) and how this scaling is related to the *fractal dimension* of tsunami propagation.

2 Governing equations

Tsunamis are most commonly modelled using a *long-wave approximation* of the mass and momentum conservation equations for a fluid with a free-surface. In this approximation, the slopes of both the free-surface and the bathymetry are assumed to be vanishingly small. Note that these assumptions are often violated in the case of tsunamis, particularly close to shore, however and somewhat suprisingly, the long-wave approximation still gives reasonable results in practice. Simple scaling arguments show that the vertical velocity of the fluid is also vanishingly small in this context. Fluid motion can thus be described only using the vertically-averaged components of the horizontal momentum and the free-surface elevation. The resulting set of equations was first derived by Saint-Venant [de Saint-Venant (1871)] and is best seen as a system of conservation laws, which can be expressed in integral form as

$$\partial_t \int_{\Omega} \mathbf{q} d\Omega = \int_{\partial\Omega} \mathbf{f}(\mathbf{q}) \cdot \mathbf{n} d\partial\Omega, \quad (2)$$

where Ω is a given subset of space, $\partial\Omega$ its boundary and \mathbf{n} the unit normal vector on this boundary. For conservation of mass and momentum in the shallow-water context, Ω is a subset of bidimensional space and \mathbf{q} and \mathbf{f} can be written

$$\mathbf{q} = \begin{pmatrix} h \\ h u \\ h v \end{pmatrix}, \quad \mathbf{f}(\mathbf{q}) = \begin{pmatrix} h u & h v \\ h u^2 + \frac{1}{2} g h^2 & h u v \\ h u v & h v^2 + \frac{1}{2} g h^2 \end{pmatrix}, \quad (3)$$

with h the water depth, u and v the components of the horizontal velocity and g the acceleration of gravity and where we have assumed a flat bottom (bathymetry). If the bathymetry is not flat a forcing term needs to be added to equation (2) which becomes

$$\partial_t \int_{\Omega} \mathbf{q} d\Omega = \int_{\partial\Omega} \mathbf{f}(\mathbf{q}) \cdot \mathbf{n} d\partial\Omega - \int_{\Omega} h g \nabla z,$$

where z is the depth of the bathymetry. Dissipative terms (viscosity and/or friction) are also often added but they do not change the fundamental structure of the equations. This system of conservation laws is hyperbolic and admits an entropy inequality, which guarantees that the total energy $h \mathbf{u}^2/2 + g h^2/2$ behaves physically (with $\mathbf{u} = (u, v)$ the velocity vector). Another important property is that the steady state

$$\mathbf{u} = 0, \quad h + z = \text{constant}, \quad (4)$$

is a solution of this system. This simply reflects the expected property that a *lake at rest* remains so.

3 Numerical scheme

Hyperbolic systems of conservation laws often occur in physics and as such, numerical methods for their solution have been a topic of extensive research. In fact, the Saint-Venant equations (3) are closely related to the two-dimensional Euler equations for compressible gases (with the fluid depth h playing the role of the gas density). As such they are amenable to be solved using the large range of numerical methods developed to deal with flow of compressible gases; such as Godunov-type finite-volume discretisations where the fluxes on the boundary of the domain ($\partial\Omega$ in the notation above) are computed by reconstructing “upwind” solutions on the left- and right-sides of the boundary. A local Riemann problem is then solved on the boundary to select only the characteristic waves travelling in the correct direction, which gives a stable approximation of the physical flux.

In the context of the Saint-Venant equations, two important issues arise which are not usually present when dealing with compressible gases. The depth h can tend toward zero, which reflects the wetting and drying of the bathymetry. This is analogous to a vanishing gas density. Aside from questions of numerical stability, several important fundamental issues arise when this happens. Following the analogy with compressible gases, a vanishing gas density will create a *rarefaction wave* which is accompanied with a *contact discontinuity*. These waves are similar to the shock waves (i.e. hydraulic jumps for Saint-Venant) which occur when the flow becomes supercritical (i.e. the Mach or Froude number is larger than one). In both cases, the *long-wave approximation* is invalid: i.e. quantities vary abruptly across the shock wave or across the contact discontinuity. Furthermore, in the case of a vanishing density, the assumption of thermodynamic equilibrium becomes invalid (because averaging over a large-enough sample of thermodynamic states is not possible anymore). For compressible gases, an accurate solution of the rarefied-gas dynamics thus requires taking into account the non-equilibrium distribution of thermodynamic states close to the wave (for example using Boltzman’s kinetic theory of gases).

The second issue regards the numerical preservation of the lake-at-rest equilibrium solution (4). While it is “only” a numerical issue and does not involve subtle and complex considerations about the validity of the underlying equations, it has been the topic of many publications [LeVeque (1998), Zhou et al. (2001), Rogers et al. (2003)]. The issue is that the lake-at-rest equilibrium results from the balance of two terms: essentially the gradient of bathymetry and the gradient of fluid depth, which are computed numerically using independent methods. If this discretisation is done naively there is no guarantee that numerical discretisation errors will compensate one another. If the numerical schemes are consistent, these errors will vanish as spatial resolution is increased and the exact equilibrium solution will be recovered asymptotically, however it is desirable to guarantee exact equilibrium irrespective of resolution. Even more so given that for practical problems the solutions are often small perturbations around the equilibrium state (which could be swamped by noise due to numerical imbalance).

3.1 Spatial and temporal discretisation

A recent scheme which combines all these requirements was analysed in detail by Audusse et al. (2004) and constitutes the basis for the scheme used in the present work. Following Audusse et al, the one-dimensional version of the spatially second-order accurate scheme can be summarised as

$$\Delta x_i \partial_t \mathbf{q}_i + \mathcal{F}_l(\mathbf{q}_{i,r}, \mathbf{q}_{i+1,l}, z_{i,r}, z_{i+1,l}) - \mathcal{F}_r(\mathbf{q}_{i-1,r}, \mathbf{q}_{i,l}, z_{i-1,r}, z_{i,l}) = \mathbf{S}_{ci},$$

where i is the index of the one-dimensional array of control volumes and the l and r subscripts indicate that the value is reconstructed on either the left- or right-side of the boundary between control volumes. The combined fluxes and source terms \mathcal{F}_l and \mathcal{F}_r are defined as

$$\begin{aligned} \mathcal{F}_l(\mathbf{q}_i, \mathbf{q}_{i+1}, z_i, z_{i+1}) &= \mathcal{F}(\mathbf{q}_{i+1/2-}, \mathbf{q}_{i+1/2+}) + \begin{pmatrix} 0 \\ \frac{1}{2} g h_i^2 - \frac{1}{2} g h_{i+1/2-}^2 \end{pmatrix}, \\ \mathcal{F}_r(\mathbf{q}_i, \mathbf{q}_{i+1}, z_i, z_{i+1}) &= \mathcal{F}(\mathbf{q}_{i+1/2-}, \mathbf{q}_{i+1/2+}) + \begin{pmatrix} 0 \\ \frac{1}{2} g h_{i+1}^2 - \frac{1}{2} g h_{i+1/2+}^2 \end{pmatrix}, \end{aligned}$$

where

$$\mathbf{q}_{i+1/2-} = \begin{pmatrix} h_{i+1/2-} \\ h_{i+1/2-} u_i \end{pmatrix}, \quad \mathbf{q}_{i+1/2+} = \begin{pmatrix} h_{i+1/2+} \\ h_{i+1/2+} u_{i+1} \end{pmatrix},$$

and

$$\begin{aligned} h_{i+1/2+} &= \max(0, h_{i+1} + z_{i+1} - z_{i+1/2}), \\ h_{i+1/2-} &= \max(0, h_i + z_i - z_{i+1/2}), \end{aligned}$$

with

$$z_{i+1/2} = \max(z_i, z_{i+1}).$$

Audusse et al show that this particular choice of reconstruction of the bottom topography leads to a scheme which preserves the lake-at-rest condition (is “well-balanced”) and ensures that the fluid depth is always positive.

The additional source term \mathbf{S}_{ci} is necessary to ensure balance for the second-order scheme and is written

$$\mathbf{S}_{ci} = \begin{pmatrix} 0 \\ \frac{1}{2} g (h_{i,l} + h_{i,r}) (z_{i,l} - z_{i,r}) \end{pmatrix}. \quad (5)$$

The numerical fluxes $\mathcal{F}(\mathbf{q}_{i+1/2-}, \mathbf{q}_{i+1/2+})$ are obtained using an approximate Riemann solver. For this study, I chose the HLLC (Harten-Lax-van Leer Contact) solver which also considers contact waves and is thus able to deal with wet-dry transitions [Harten et al. (1983), Liang and Borthwick (2009)].

The spatially second-order scheme requires the computation of second-order accurate left and right states (l and r subscripts) for \mathbf{q} and z . To ensure stable solutions, this is done using slope-limited extrapolations such that

$$\mathbf{q}_{i,r} = \mathbf{q} + \frac{1}{2} \psi(r) \nabla \mathbf{q}_i \Delta x_i, \quad (6)$$

where $\nabla \mathbf{q}_i$ is the cell-centred estimate of the gradient and $\psi(r)$ is a slope limiter, function of the ratio

$$r = \frac{\mathbf{q}_{i+1} - \mathbf{q}_i}{\mathbf{q}_i - \mathbf{q}_{i-1}}.$$

Several choices are possible for ψ . Within the Total Variation Diminishing (TVD) family of limiters, the most-dissipative is the “minmod” limiter and the least-dissipative the “superbee” limiter. As discussed in more detail in what follows, I found that while the minmod limiter is indeed very stable and gives reasonable results for dissipative systems (i.e. problems typically dominated by explicit bottom friction), it tends to be over-dissipative for realistic tsunami modelling. Solutions using the superbee limiter on the other hand tend to suffer from excessive sharpening and oscillations of wave fronts. The “Sweby” limiter with $\beta = 1.5$ seems to be a good compromise between dissipation and stability [Sweby (1984)].

The extension of this scheme to two spatial dimensions is obtained by using a MUSCL-type unsplit discretisation combined with a predictor-corrector time-stepping scheme. This can be summarised as:

1. Predictor step

$$\mathbf{q}_{ij}^{n+1/2} = \mathbf{q}_{ij}^n - \frac{\Delta t}{2} \mathbf{F}_{ij}^n + \frac{\Delta t}{2} \mathbf{S}_{cij}$$

2. Corrector step

$$\mathbf{q}_{ij}^{n+1} = \mathbf{q}_{ij}^n - \Delta t \mathbf{F}_{ij}^{n+1/2} + \Delta t \mathbf{S}_{cij}$$

with

$$\begin{aligned} \mathbf{F}_{ij}^n = & \Delta x_{ij} \mathcal{F}_l(\mathbf{q}_{ij,r}^n, \mathbf{q}_{i+1j,l}^n, z_{ij,r}, z_{i+1j,l}) - \Delta x_{ij} \mathcal{F}_r(\mathbf{q}_{i-1j,r}^n, \mathbf{q}_{ij,l}^n, z_{i-1j,r}, z_{ij,l}) + \\ & \Delta y_{ij} \mathcal{F}_b(\mathbf{q}_{ij,t}^n, \mathbf{q}_{ij+1,b}^n, z_{ij,t}, z_{ij+1,b}) - \Delta y_{ij} \mathcal{F}_t(\mathbf{q}_{ij-1,t}^n, \mathbf{q}_{ij,b}^n, z_{ij-1,t}, z_{ij,b}), \end{aligned}$$

where the t and b subscripts denote the “top” and “bottom” reconstructed values. This scheme is explicit in time and thus needs to verify the Courant-Friedrich-Levy (CFL) condition for stability

$$\Delta t < C \min \left[\min_{ij} \left(\frac{\Delta x_{ij}}{|u_{ij}| + \sqrt{g h_{ij}}} \right), \min_{ij} \left(\frac{\Delta y_{ij}}{|v_{ij}| + \sqrt{g h_{ij}}} \right) \right],$$

with C a constant equal to 0.5 for the MUSCL-type unsplit discretisation. This constraint is enforced in practice by choosing Δt at each timestep so that the inequality is verified. For tsunamis, the Froude number

$$F_r = \frac{|\mathbf{u}|}{\sqrt{g h}},$$

is small in most of the domain (i.e. away from the coastline), so that the time-step is usually limited by the speed of gravity waves in deep water. The current algorithm uses a global time-step for all cells. While the varying cell size of adaptive methods could be used to ensure a roughly constant value of the stability threshold (i.e. so that $\Delta x \propto \sqrt{g h}$), we will show below that an accurate description of the propagating wave requires high resolution also in deep water. Another way of improving the time-stepping efficiency would be to use a local time-step. This is relatively simple to do for hyperbolic systems [Berger and Oliger (1984)] but has not been implemented in our method yet.

To summarise, the overall scheme is second-order accurate in space and time, preserves the positivity of the water depth and the lake-at-rest condition and is volume- and momentum-conserving.

3.2 General orthogonal curvilinear coordinates

When considering tsunami propagation over large distances ($\gtrsim 1000$ km), the curvature of the Earth's surface needs to be taken into account. This is often done by rewriting the differential form of the Saint-Venant equations in spherical coordinates [Williamson et al. (1992)]. One of the drawbacks of this approach is that doing so somewhat obscures the conservative nature of the system of equations, so that a discretely-conservative formulation of the resulting system is not straightforward. Another drawback is that longitude/latitude spherical coordinates lead to a very inhomogeneous distribution of spatial resolution when considering the entire sphere (i.e. scales of order 10,000 km).

One of the goal of the framework developed within Gerris is to combine quad/octree discretisations with general curvilinear orthogonal coordinates. Within this framework, one can easily define arbitrary curvilinear orthogonal mappings applied to the different solvers available in Gerris. One could for example use longitude/latitude spherical coordinates (a particular case of curvilinear orthogonal coordinates) to solve the Saint-Venant equations on part of the sphere, while using the same solver combined with the “cubed sphere” orthogonal coordinates [Rančić et al. (1996)] to solve the Saint-Venant equations on the entire sphere.

General orthogonal coordinates are a complex topic which cannot be covered extensively within the scope of this article. In what follows I will only attempt to summarise how the discretisation of the Saint-Venant equations is generalised to orthogonal curvilinear coordinates within the Gerris framework.

Let us consider two orthogonal directions on a curved two-dimensional manifold, denoted by indices λ and θ . The orthogonal curvilinear metric can be defined using a scaling factor along each dimension, noted $m_\lambda(\lambda, \theta)$ and $m_\theta(\lambda, \theta)$ respectively. Using these notations, the system of conservation laws (2) can be rewritten in discrete form

$$m_\lambda m_\theta \partial_t \mathbf{q} + \frac{1}{\Delta_\lambda} \sum_{\lambda} m_\theta \mathbf{f}_\lambda(\mathbf{q}) + \frac{1}{\Delta_\theta} \sum_{\theta} m_\lambda \mathbf{f}_\theta(\mathbf{q}) = \mathbf{S}_g, \quad (7)$$

where the sums are computed over the faces of the cells in each respective direction, and \mathbf{S}_g is a system-specific *geometric source term* resulting from integration of the Christoffel symbols applied to vector quantities. Note that this formulation is independent from the details of both the orthogonal curvilinear metric and the system of conservation laws.

For the Saint-Venant equations one can show that \mathbf{S}_g takes the form

$$\mathbf{S}_g = \begin{pmatrix} 0 \\ -\frac{g}{2} h^2 \partial_\lambda m_\theta + h f_G v \\ -\frac{g}{2} h^2 \partial_\theta m_\lambda - h f_G u \end{pmatrix}, \quad (8)$$

with

$$f_G = v \partial_\lambda m_\theta - u \partial_\theta m_\lambda$$

One way to check that this formulation is indeed correct is to apply it to a specific orthogonal curvilinear metric. For example, the spherical coordinates metric is defined by

$$m_\lambda = R \cos \theta, \quad m_\theta = R, \quad \partial_\lambda m_\theta = 0, \quad \partial_\theta m_\lambda = -R \sin \theta,$$

with R the radius of the sphere. We can recover the differential form for (7) by making Δ_λ and Δ_θ tend to zero. Together with the definition of Saint-Venant fluxes in (3), we then get

$$\begin{aligned} d_t h + \frac{h}{m_\lambda m_\theta} [\partial_\lambda (m_\theta u) + \partial_\theta (m_\lambda v)] &= 0, \\ d_t u - \frac{f_G}{m_\lambda m_\theta} v + \frac{g}{m_\lambda} \partial_\lambda h + \frac{g h}{2 m_\lambda m_\theta} \partial_\lambda m_\theta &= 0, \\ d_t v + \frac{f_G}{m_\lambda m_\theta} u + \frac{g}{m_\theta} \partial_\theta h + \frac{g h}{2 m_\lambda m_\theta} \partial_\theta m_\lambda &= 0, \end{aligned}$$

where we have introduced the notation

$$d_t = \partial_t + \frac{u}{m_\lambda} \partial_\lambda + \frac{v}{m_\theta} \partial_\theta$$

Using the definition of the spherical-coordinates metric then gives

$$\begin{aligned} d_t h + \frac{h}{R \cos \theta} [\partial_\lambda u + \partial_\theta (v \cos \theta)] &= 0, \\ d_t u - \frac{u v}{R} \tan \theta + \frac{g}{R \cos \theta} \partial_\lambda h &= 0, \\ d_t v + \frac{u^2}{R} \tan \theta + \frac{g}{R} \partial_\theta h - \frac{g h}{2 R} \tan \theta &= 0, \end{aligned}$$

which is indeed the standard form for the Saint-Venant equations in spherical coordinates [Williamson et al. (1992)].

Some care also needs to be taken when applying the well-balanced scheme of section 3.1 to the generalised system (7). In the case of a lake-at-rest discretised using curvilinear orthogonal coordinates, the discrete equilibrium condition can be written (in one dimension)

$$\frac{g}{2} (m_l h_{i,r}^2 - m_r h_{i,l}^2) = S_{ci} + S_g,$$

where the geometric source (8) can be discretised as

$$S_g = \frac{g}{4} (h_{i,r}^2 + h_{i,l}^2) (m_l - m_r).$$

The definition of S_{ci} then need to be generalised to maintain balance. A possible choice is

$$S_{ci} = \begin{pmatrix} 0 \\ \frac{1}{4} g (m_l + m_r) (h_{i,l} + h_{i,r}) (z_{i,l} - z_{i,r}) \end{pmatrix}.$$

This formulation reduces to the Cartesian scheme (formula (5)) for $m_l = m_r = 1$ and to the first-order scheme for $z_{i,l} = z_{i,r}$.

3.3 Adaptivity

This scheme has been implemented within the Gerris Flow Solver framework [Popinet (2003b), Popinet (2003a)]. Gerris was initially designed to solve the incompressible Euler equations with complex boundaries [Popinet (2003b)] but has since been extended to the linearised shallow-water equations [Popinet and Rickard (2007)], multiphase incompressible Navier-Stokes equations [Popinet (2009)], spectral wave models [Popinet et al. (2010)] as well as coupled problems including electrohydrodynamics [Lopez-Herrera et al. (2011)]. Gerris uses a quadtree (octree in 3D) spatial discretisation which allows efficient adaptive mesh refinement. An example of the quadtree structure is represented in Figure 1 together with its logical (tree) representation. This tree can be conveniently described as a “family tree” where each *parent* cell can have zero or four *children* cells.




Figure 1. An example of quadtree discretisation (left) together with its logical representation (right). The level of the cells in the tree are also given.

An important parameter is the *level* of a given cell in the tree. The root cell has level zero by convention and the level increases by one for each successive generation in the tree. Details of the actual implementation in Gerris can be found in [Popinet (2003b)]. A further simplification is obtained by constraining the quadtree so that the level of adjacent cells cannot differ by more than one. This reduces the number of configurations for neighbouring cells to two: 1) neighbouring cells are on the same level: this is identical to a regular Cartesian grid discretisation. 2) two fine cells are neighbouring the same coarse cell (Figure 2).




Figure 2. Computation of fluxes for the only quadtree-specific cell configuration.

To implement the numerical scheme described previously, we consider quantities as averages over each cell of the quadtree, in a classical finite-volume sense. For case 1) the implementation reduces to the standard regular Cartesian grid scheme. For case 2) one has to be careful to ensure consistent discretisation of the fluxes between the fine and coarse cells (Figure 2). The flux on the left side of coarse cell C needs to be computed as the sum of the fluxes on the right sides of C_1 and C_2 . Computing the flux on the right side of C_1 (and C_2) requires reconstructed values on either side of the fine/coarse cell face (e.g. u_r and u_l). The value of u_r is obtained using the standard Cartesian scheme (i.e. equation (6)). To maintain second-order accuracy, the value of u_l needs to be computed while also taking into account the transverse value of the slope-limited gradient in cell C . This can be summarised as

$$u_l = u|_C - \frac{1}{2} \psi(r_x) \nabla_x u|_C \Delta x + \frac{1}{4} \psi(r_y) \nabla_y u|_C \Delta y. \quad (9)$$

In order to refine or coarsen the mesh dynamically, one needs to choose a refinement criterion. This choice is not trivial and depends on both the details of the numerical scheme (which control the discretisation errors) and the physics of the problem considered (which control how discretisation errors will influence the accuracy of the solution). Gerris provides a variety of refinement criteria which can be flexibly combined depending on the problem. For this study I have used only a simple criterion based on the gradient of the free surface elevation $\zeta = z + h$. With this criterion, cells are refined whenever

$$|\nabla \zeta| \Delta > \epsilon, \quad (10)$$

where $\Delta = \Delta x = \Delta y$ is the cell size and ϵ is an adjustable parameter. This parameter is best interpreted as an estimate of the maximum discretisation error on field η for a spatially first-order method. Although our method is formally second-order for smooth fields, for tsunamis we expect that the errors will be dominated by the discretisation of sharp fronts (e.g. wave fronts, shocks and contact discontinuities near coastlines), where the method is first-order.

When a cell is refined the values of fields in the newly-created control volumes need to be initialised. By default, this is done using the same slope-limited interpolations as in equations (6) and (9). These relationships are formally second-order and also guarantee conservation of quantities (i.e. the average of the four interpolated values is equal to the value of the parent coarse cell). The bathymetry field z as well as the water depth h receive special treatment as explained in the next section.

Cells also need to be coarsened (i.e their children are destroyed) when a high spatial resolution is not required anymore. This is done using the criterion matching (10)

$$|\nabla \zeta| \Delta < \frac{\epsilon}{\beta},$$

where β is a “hysteresis factor” which prevents cycles of refinement/coarsening between successive time-steps. A typical value for β is 2. When cells are coarsened the values of fields are set by default to the averages over their children, which also ensures conservation of quantities.

3.4 Adaptive terrain reconstruction

Managing potentially very large bathymetry/topography databases is an important aspect of practical tsunami modelling. These databases often combine data from very different sources both in term of accuracy and spatial resolution: for example ranging from coarse (kilometric) - scale gravimetric estimates for the deep ocean to metre-resolution lidar data for coastal areas. Datasets also overlap and use different sampling strategies.

For static-mesh methods database access is limited to the preprocessing (or mesh generation) phase of the simulation, however for adaptive-mesh methods the topography database needs to be accessed every time the mesh is refined, so that the finest discretisation elements always define the bathymetry using the finest available data. This places strict requirements on database efficiency.

Appendix Zhou et al. (2001) describes a database system which meets the following requirements:

- Fast access to spatially-averaged topographic data at multiple scales,
- Small memory footprint even for very large (terabytes) databases (i.e. most data structures are stored on-disk),
- No restriction on the sampling strategy of the initial topographic datasets.




Figure 3. Bathymetry reconstruction. The bathymetry is drawn with thick black lines and the water level with blue lines. (a) Initial discretisation. (b) The central cell has been refined.

An important issue to address when refining the mesh is how reconstructing the bathymetry affects volume conservation and/or the preservation of lake-at-rest equilibrium. A simple 1D example of the problem is illustrated in Figure 3. The initial bathymetry, discretised with three cells of equal size, is represented by the thick black line in Figure 3.a. The “lake-at-rest” constant water level is represented with the blue line. Figure 3.b illustrates what happens when the central cell is refined (i.e. split into two finer cells). To preserve the lake-at-rest equilibrium, one cannot change the water level (blue line on both figures). With this constraint, conserving the volume of fluid when refining the mesh would impose the condition: $z_c = (z_1 + z_2)/2$. This condition would of course need to hold recursively for each level of successive refinement. In practice this would mean that reconstructing the bathymetry in a given cell would need to be done in a “bottom-up” manner, starting from the finest level (e.g. the level below which no bathymetry data is available). Note that this would need to be done “in advance” not knowing whether a given coarse cell was going to ever be refined (which depends on the details of the simulation further on). Such preprocessing of the bathymetry dataset is not unlike the kd-tree technique we presented in Appendix Zhou et al. (2001). A very important difference however is that such bottom-up preprocessing would only ensure volume conservation for the quadtree mesh used to do the preprocessing. In fact this would mean that bathymetry preprocessing would need to be redone every time the coordinate system associated with the mesh changed: for example if the mesh was shifted slightly in position or when a different projection was used. This is clearly not a practical option, as creating bathymetry databases for large datasets is a time-consuming process.

The compromise we have made is to choose either to ensure volume conservation or lake-at-rest equilibrium but not both. Volume conservation is ensured simply by: 1) reconstructing the bathymetry in each cell using the standard averaging algorithm described previously, 2) reconstruct the water level so that the total volume is conserved (dashed blue lines in Figure 3.b). This is done by distributing the deficit/excess volume between the two cells in proportion to their relative volumes. For lake-at-rest equilibrium, the water level is simply set to that of the parent cell. Note that choosing whether to preserve lake-at-rest equilibrium or volume is not trivial. For tsunami modelling however, we have found that preserving the lake-at-rest equilibrium rather than volume gives good results in practise.

4 Results

In this section, I will consider a test case which admits an analytical solution, a classical benchmark case with reference experimental data and an application to a full-scale tsunami. The data and code necessary to reproduce each case (including graphs etc...) is available on the Gerris web site [Popinet (2010a), Popinet (2010b), Popinet (2010c)].

4.1 Oscillations in a parabolic container

Analytical solutions for the shallow-water equations describing the linearly-damped oscillations of a liquid in a parabolic container have been derived by Sampson et al. [Sampson et al. (2006), Liang and Borthwick (2009)]. The linear friction term can be written

$$\begin{pmatrix} 0 \\ -\tau h u \end{pmatrix}$$

with τ the linear friction coefficient, and is added to the right-hand side of (3).

Obtaining the analytical solution relies on the demonstration that in the case of a parabolic container, the water surface is an inclined plane and the horizontal component of the water velocity is time-varying but spatially constant. Assuming a parabolic topography given by

$$z(x) = h_0 \left(\frac{x}{a} \right)^2,$$

the horizontal component of the velocity is given by [Sampson et al. (2006)]

$$u_0(t) = B e^{-\tau t/2} \sin s t,$$

where τ is the linear damping coefficient, B controls the initial surface slope and

$$\begin{aligned} p &= \sqrt{8 g h_0 / a^2}, \\ s &= \frac{\sqrt{p^2 - \tau^2}}{2}, \end{aligned}$$

with g the acceleration of gravity. The position of the water surface is given by

$$\zeta(x, t) = \frac{a^2 B^2 e^{-\tau t}}{8 g^2 h_0} \left[-s \tau \sin 2 s t + \left(\frac{\tau^2}{4} - s^2 \right) \cos 2 s t \right] - \frac{B^2 e^{-\tau t}}{4 g}$$

$$-\frac{e^{-\tau t/2}}{g} \left(B s \cos s t + \frac{\tau B}{2} \sin s t \right) x.$$

Aside from the availability of an analytical reference solution, this test case is interesting from a numerical point of view because wetting and drying occurs at the points where the planar water surface intersects the parabolic topography. Numerical errors due to an inappropriate treatment of the wetting-drying singularity are likely to contaminate the entire solution over the course of several water surface oscillations. Similarly any numerical imbalance between the topographic source terms and slope gradients will affect the solution when it reaches equilibrium.

For numerical validation, we use the following parameters: $L = 10000$ m, $h_0 = 10$ m, $a = 3000$ m, $\tau = 10^{-3} \text{ s}^{-1}$, $B = 5$ m/s (see also [Liang and Borthwick (2009), Liang and Marche (2009)]). Figure 4 illustrates the result of the numerical simulation at $t = 1500$ seconds for the surface profile. Figure 5 displays the evolution of the horizontal velocity u_0 computed numerically and analytically. The numerical value of u_0 is computed as the spatial-average of the (theoretically constant) horizontal velocity.




Figure 4. Solution at $t = 1500$ seconds. Six levels of refinement.




Figure 5. Time evolution of the (spatially constant) horizontal velocity. Seven levels of refinement.

The results are visually very close (similarly to those presented in Liang and Borthwick (2009)) and it is interesting to further quantify the errors. We define the following relative error norms

$$\begin{aligned}\|h\|_1 &= \frac{\int_0^T \int_{-L/2}^{L/2} |h(x,t) - z(x) - \zeta(x,t)| dx dt}{h_0 \int_0^T \int_{-L/2}^{L/2} dx dt}, \\ \|h\|_2 &= \frac{1}{h_0} \sqrt{\frac{\int_0^T \int_{-L/2}^{L/2} [h(x,t) - z(x) - \zeta(x,t)]^2 dx dt}{\int_0^T \int_{-L/2}^{L/2} dx dt}}, \\ \|h\|_{\max} &= \frac{\max_{x,t} |h(x,t) - z(x) - \zeta(x,t)|}{h_0}, \\ \|u_0\|_2 &= \frac{1}{B \int_0^T dt} \sqrt{\int_0^T \left[u_0(t) - \frac{\int_{-L/2}^{L/2} u(x,t) h(x,t) dx}{\int_{-L/2}^{L/2} h(x,t) dx} \right]^2 dt}.\end{aligned}$$

with $T = 6000$ seconds the total duration of oscillations. Figures 6 and 7 illustrate the convergence with spatial resolution of the error norms on the free surface position and velocity respectively. The estimation of convergence orders are obtained through a linear regression of the data points (using a logarithmic scale for the ordinate). Consistent convergence is obtained for both quantities. The convergence is close to second-order for the velocity and better than first order for the position. Further examination reveals that the largest errors on position are localised close to the wetting-drying contact points. This leads to the first-order convergence observed for the $\|h\|_{\max}$ norm. This is expected since the second-order Godunov scheme degrades to first-order in the vicinity of discontinuities. These results are comparable to the results presented in Liang & Marche, 2009 (note however that our norms take the time history into account). This test case is part of the Gerris test suite [Popinet (2010a)].




Figure 6. Convergence with spatial resolution of the relative error norms on the free surface position.




Figure 7. Convergence with spatial resolution of the relative error norm on the horizontal velocity.

4.2 Monai tsunami benchmark

A classical tsunami benchmark problem is the “Monai tsunami benchmark”, the second benchmark problem proposed at the “third international workshop on long-wave run-up models” [IWLWRM (2004)]. This benchmark aims to reproduce the results of a wave tank experiment designed to study the 1993 Okushiri tsunami in Japan. Available experimental data include time-series of water level at selected locations as well as videos of the water surface. The topography is relatively complex and focuses on a narrow valley (near the village of Monai) where a run-up of 32 metres was observed. A small island also sits in front of the valley. A digital terrain model, matching the physical model used in the wave tank experiment is available as input for the benchmark, as well as the input wave profile.

Following the benchmark guidelines, we consider a numerical simulation area 5.448 m long and 3.402 m wide. Symmetry boundary conditions are used on the side walls (i.e. purely reflective boundaries), the topography is reconstructed dynamically using the digital terrain model data provided. The vertical wall at the right-hand-side of the domain is approximated as a very steep slope (i.e. a wall with a width of a single grid cell). The time-series of water level elevation is imposed at the left-hand-side boundary. Adaptivity is used with the gradient criterion applied to the water surface elevation H with a threshold value of 10^{-3} . The maximum resolution is set to ~ 7 mm and the minimum resolution to ~ 5 cm. A quadratic bottom friction term

$$-C_f |\mathbf{u}| \begin{pmatrix} 0 \\ u \\ v \end{pmatrix}$$

is added to the right-hand-side of (3), with a friction coefficient C_f set to 10^{-3} .

Figures 8 and 9 illustrate the evolution of the free surface elevation as well as the corresponding adaptive mesh. The initial wave amplitude is negative which causes an initial drying out at the coast as well as drying out of the isthmus connecting the coast with the island ($t = 10, 12$ and 14 s). At $t = 14$ s, the positive wave front has already travelled past the island. The wave front is also clearly visible in the adaptive mesh as it is associated with large free-surface gradients which are tracked by the adaptivity criterion. At $t = 16$ s, the wave has flooded both the coastline and the island and is being reflected away from the shore. Subsequent frames show multiple wave reflections being tracked by the adaptive mesh. In all frames, the moving wet-dry “contact line” is also tracked at high-resolution by the adaptive mesh as the corresponding contact discontinuity also induces large free-surface elevation gradients. These results are very

similar to the results of LeVeque & George (Figure 4 of LeVeque and George (2006)) obtained on regular Cartesian meshes. The results are also practically identical to those of a control simulation run on a constant resolution mesh.




Figure 8. Evolution of the free surface elevation (left column) and corresponding adaptive mesh (right column). Contour lines of the topography are represented. The areas in white in the left column are dry. The wet areas are coloured according to free surface elevation relative to the unperturbed water surface. A "jet" colour scale is used with a maximum value (dark red) of $+5$ cm and a minimum value (dark blue) of -5 cm.

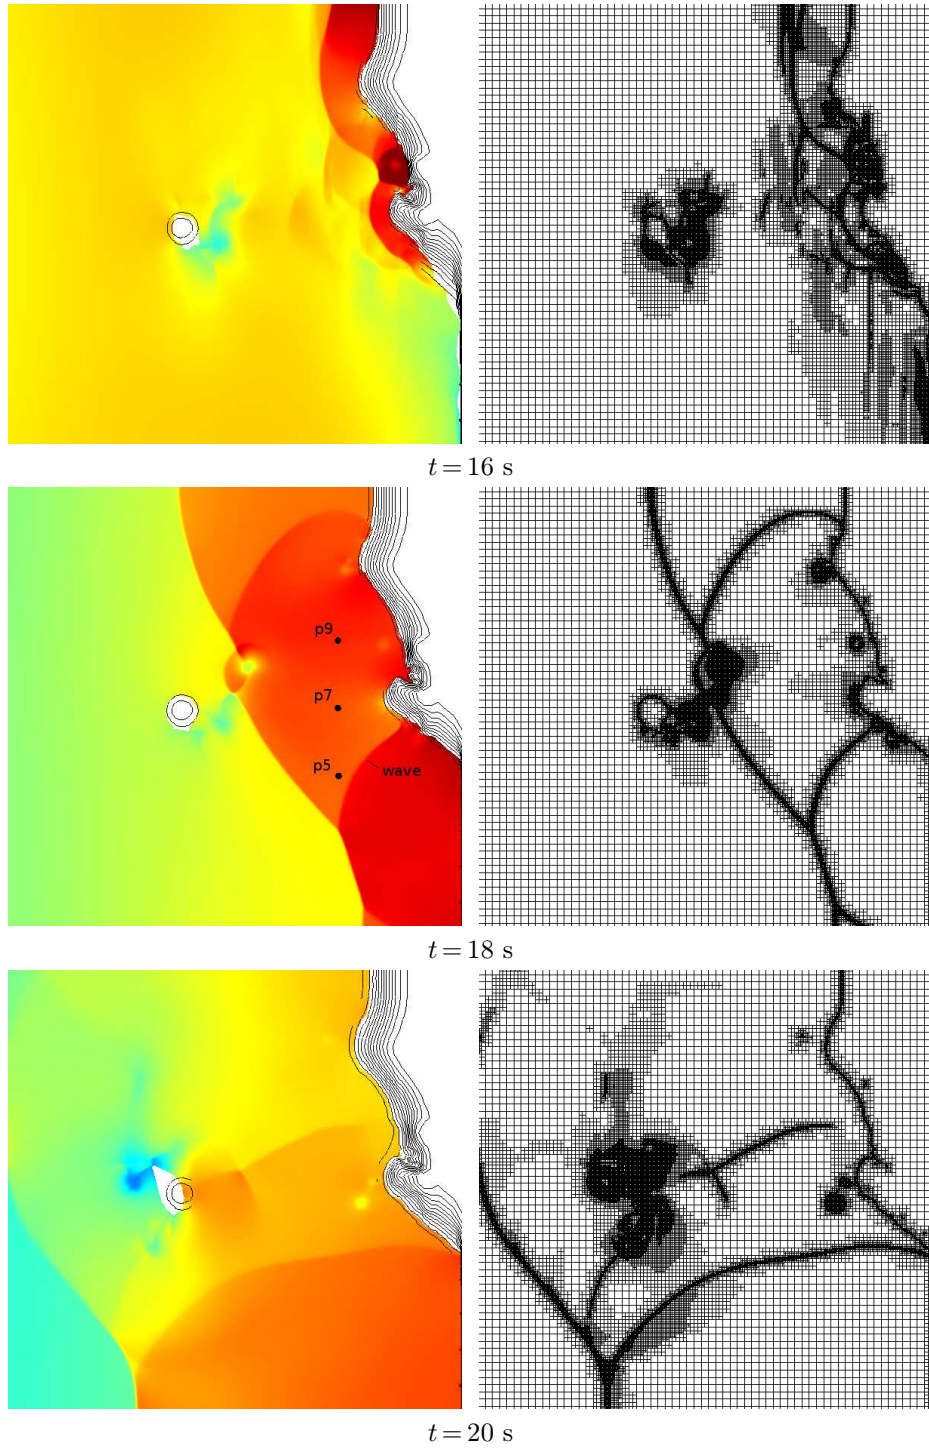


Figure 9. Evolution of the free surface elevation (continued from Figure 8).

Comparison of the time-series of water surface elevation at experimental measurement locations give a more quantitative assessment of the results. Figures 10, 11 and 12 give the measured and calculated time-series at experimental probe positions 5, 7 and 9 (on a transect parallel

to the coast and approximately mid-way from the coast and the island, see Figure 9, $t = 18$ s). We also reproduce the results of LeVeque and George (2006) and Anuga [Nielsen et al. (2005)] for comparison. The numerical results are all comparable and agree both for the amplitude and phase of the different waves. The arrival time of the initial wave as well as its amplitude are well predicted by the models. The (weaker) wave reflected off the right-hand-side wall (visible in the lower-right quadrant at $t = 18$ s in Figure 9) is measured by probe 7 at $t \sim 19$ s and is also visible in the models (which tend to under predict its amplitude). The same wave is also measured by probe 9 at $t \sim 20$ s and is visible in the Gerris simulation (again with a lower amplitude relative to the background surface elevation). Note also the initial noise in the experiment which explains the non-zero elevation in the experimental data for $t < 10$ s. The agreement is satisfactory overall and confirms that for similar spatial resolutions and solution algorithms (Godunov shock-capturing schemes) but different spatial discretisations, all methods give comparable results.




Figure 10. Time-series of free-surface elevation measured and calculated at the location of probe 5.




Figure 11. Time-series of free-surface elevation measured and calculated at the location of probe 7.




Figure 12. Time-series of free-surface elevation measured and calculated at the location of probe 9.

Finally Figure 13 illustrates the evolution of the ratio between the number of discretisation elements of the constant resolution (~ 7 mm) control simulation and the (time-dependent) number of elements used by the adaptive simulation. In principle, this ratio represents the efficiency gain (in terms of number of unknowns) of an adaptively refined mesh compared to a uniform mesh of the highest resolution. Up to $t = 10$ s a coarse mesh is sufficient to resolve the free surface which leads to savings in total number of discretisation elements of a factor larger than 50. As the wave field becomes more complex the ratio decreases and stabilises around 10. Note that these factors are resolution-dependent. We look in more detail at this issue in the next section.




Figure 13. Evolution of the ratio between the number of elements of a constant-resolution control simulation and the number of elements of the adaptive simulation.

4.3 2004 Indian ocean tsunami

As the most destructive tsunami in recorded history, the 2004 Indian ocean tsunami has been studied extensively both using field data and numerical models [Wang and Liu (2006), Grilli et al. (2007), Ioualalen et al. (2007), Pranowo et al. (2008), Harig et al. (2008), George and LeVeque (2008)]. It is also an interesting event in the context of the present article since it combines strong near-field effects (shoaling, inundation) with far-field wave travel. The spatial scales involved thus range from tens of metres to thousands of kilometres (i.e. five orders of magnitude).

Several authors have analysed the 2004 tsunami using numerical models. Aside from differences in the numerical schemes – which can affect the solution significantly as we will see – one of the main source of uncertainty in modelling real tsunamis is the estimation of the initial surface displacement. Various schemes have been proposed to reconstruct this initial condition [Wang and Liu (2006), Grilli et al. (2007), Pietrzak et al. (2007)]. I have based the following validation on the work of Grilli et al. (2007) who used an initial tsunami source based on the staggered rupture at the Indian–Australian and Eurasian–Andaman plate boundaries. The rupture is modelled as five discrete vertical displacement fields with rupture times staggered over about 20 minutes. The vertical displacement field for each segment is modelled using Okada’s dislocation solution [Okada (1985)] (see Table 1 of Grilli et al. (2007) for details). Grilli et al optimised the distribution of vertical displacements in order to obtain the best possible fit when comparing model results with both tide gauge records and satellite altimetry.

The available bathymetry data is another important source of uncertainty when comparing model results and point measurements such as tide gauge data. Tide gauges are typically located within harbours and other complicated, semi-enclosed, fine-scale coastal features. These local bathymetry details can have an important effect on incoming tsunami waves: focusing, excitation of local resonances etc... The results presented in this section were obtained using the ETOPO1 topography data [Amante and Eakins (2008)] which has a nominal maximum spatial resolution of one nautical mile. Note that Grilli et al used a higher resolution dataset for the coast of Thailand since they were specifically interested in validating the model for this area [Ioualalen et al. (2007)].

The general orthogonal coordinates formulation is used with a longitude–latitude mapping on a domain centred on 94°E , 8°N and extending over 54° of both longitude and latitude. This is enough to cover the Bay of Bengal and part of the Indian Ocean (including the Maldives and Diego Garcia archipelagos) as well as most of Indonesia. The mesh is adapted at every time-step using the gradient criterion (10) with $\epsilon = 5$ cm. This criterion is also applied during initialisation so that the initial surface displacement is properly resolved. This criterion is only applied to “wet cells”. The maximum spatial resolution is limited to twelve levels of refinement or $54/$

$2^{12} \approx 0.8$ nautical mile. The minimum spatial resolution is set to 5 levels of refinement or $54/2^5 \approx 101$ nautical miles.

Figure 14 illustrates the wave elevation field and the corresponding adaptive mesh two hours (top row) and three and a half hours (bottom row) after the initial fault rupture. After two hours, the wave front has just reached Sri Lanka, so that the eastern side of the island is resolved at the maximum resolution (0.8 nautical mile) while the western coastline still uses a much coarser resolution (50 nautical miles). After three and a half hours, the wave front reaches the Maldives and wave reflections off the eastern coastline of Sri Lanka are also evident. These images also illustrate the dynamic terrain reconstruction which guarantees that the bathymetry (and consequently coastline) is always resolved using the maximum resolution data available (one nautical mile ETOPO1 data for this example): the Maldives, for example, are invisible on the coarse resolution mesh at two hours but are resolved at high resolution at three and a half hours.

The complex individual wave fronts due to the staggered fault rupture and initial dispersion are also tracked individually at high resolution. Some areas (e.g. the dark blue patch south of the Aceh peninsula in Figure 14 top right) have already been traversed by the wave fronts and have returned to coarser resolutions, appropriate to resolve the smoother elevation field there.




Figure 14. 2004 Indian Ocean tsunami. (left column) Wave elevation, dark red > 2 metres, dark blue < -2 metres. (right column) Spatial resolution of adaptive mesh, dark blue 101 nautical miles, dark red 0.8 nautical mile.

Figure 15 illustrates the maximum wave elevation reached over 10 hours after the initial fault rupture. Several classical features are readily apparent such as: the “waveguide” propagation along submarine ridges toward Sri Lanka and the Maldives; large wave elevations along the Aceh peninsula, Thailand, eastern Sri Lanka, India and Maldives. The results in figures 15.a and 15.b are also in qualitative and quantitative agreement with the results of Grilli et al (Figures 4 and 5 of Grilli et al. (2007)).




Figure 15. Maximum wave elevation reached over 10 hours, 1 metre interval contours. (a) Bay of Bengal, dark blue zero, dark red > 5 metres. (b) Detail near Northern Sumatra and Thailand, dark blue zero, dark red > 8 metres.

Generating a “time-averaged” map such as that of Figure 15 raises a small issue when using an adaptive discretisation. If using only the slope criterion (10) to control mesh adaptation, nothing guarantees that the fine-scale spatial variations of time-averaged quantities such as the maximum wave elevation will be properly discretised. To keep an adequate resolution of the maximum wave elevation field the “discretisation error criterion” of Popinet et al. (2010) (section 2.4) is added to (10) with a threshold (maximum discretisation error) of 5 cm. Adding this requirement will increase the number of cells compared to tracking the wave only, but will not improve the overall solution. This problem could be circumvented by using a separate mesh to store the maximum wave elevation. The spatial resolution requirements for the dynamic wave elevation and the essentially-static maximum wave elevation could then be decoupled. For the results presented here however, the overhead is around 20% which does not warrant the complication of such a decoupling (it is also clear in Figure 14 that the high-resolution areas are closely correlated with the dynamic wave elevation rather than with features in the maximum wave elevation field).

To obtain a more quantitative comparison with previous model results and field data, I have reproduced in Table 1 the observed [Kawata (2005), Yalciner et al. (2005), Yamada et al. (2005)] and modelled [Grilli et al. (2007)] run-ups at a number of locations. The agreement is satisfying overall with the exception of Kho Phi Phi and Khao Lak where the present study predicts run-ups significantly lower than both observations and model results from Grilli et al. This discrepancy could be explained by an inaccurate bathymetry in these areas (where Grilli et al used additional data from the Thai navy). It is clear in particular that Kho Phi Phi island is only crudely represented in the ETOPO1 dataset.

Locations	(Long. E, Lat. N)	Grilli et al	Present study	Field surveys
Aceh (N coast), Indonesia	(95.323,5.570)	9.33	6.16	10–11
Aceh (N coast), Indonesia	(95.284,5.556)	14.40	8.25	10–16
Aceh (W coast), Indonesia	(95.247,5.458)	16.94	17.60	24–35
Galle, Sri Lanka	(80.475,5.974)	3.23	3.16	2–3
SE coast, Sri Lanka	(81.816,7.427)	8.13	5.60	5–10
Chennai, India	(80.279,13.021)	2.43	3.01	2–3
Nagappaattinam, India	(79.740,10.865)	4.67	3.20	2–3.5
Pulikat, India	(80.333,13.383)	2.62	2.35	3.45
Kamala Bch., Phuket, Thailand	(98.275,7.973)	3.47	5.95	4.5–5.3
Patong Bch., Phuket, Thailand	(98.276,7.900)	2.48	6.20	4.8–5.5
Kho Phi Phi, Thailand	(98.777,7.739)	3.68	0.60	4.6–5.5
Khao Lak, Thailand	(98.268,8.857)	13.88	4.35	15.77

Table 1. Observed and modelled runups (metres) at several locations. The Grilli et al model results are reproduced from [Grilli et al. (2007)] (table 3, NSW model).

Comparison of model time-series and tide gauge/satellite altimetry records provide another direct validation of the model results. The 2004 tsunami was recorded by more than twenty tide gauges within the Indian ocean. Figure 16 compares model results with observed time-series for the tide gauges located within our simulation domain (and for which records were readily available in digital form [Sea Level Center (2006)]). The arrival times and wave amplitudes are predicted satisfactorily overall. The level of agreement is comparable to that obtained in previous modelling studies (Figure 7 of Grilli et al. (2007) in particular). Aside from a relatively small error in arrival time, the agreement with the Columbo tide gauge is satisfying (Figure 16.e): Columbo is located on the western coastline of Sri Lanka and predicting the tsunami correctly at this location involves accurately following the wave as it wraps around the island in a narrow band close to shore.

Figure 17 gives a comparison of the sea level anomaly observed by the JASON-1 satellite altimeter with the corresponding modelled transect. Note that the source model I have used was tuned by [Grilli et al. (2007)] to obtain a good agreement with the JASON transect. Figure 17 can be compared with Figure 6 of [Grilli et al. (2007)] which shows very similar agreement between modelled and measured transects.

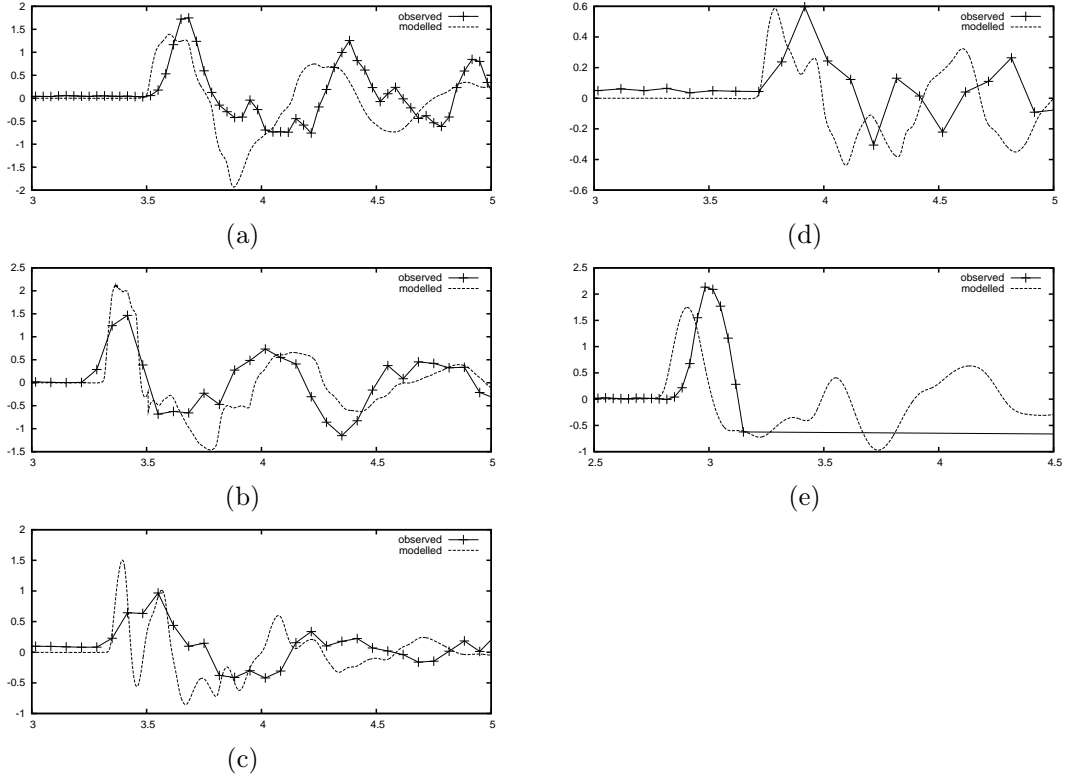


Figure 16. Observed and modelled time-series of wave elevation (metres) at different tide gauge locations. Horizontal axis is time in hours after the initial fault rupture. (a) Hannimaadhoo, Maldives. (b) Male, Maldives. (c) Gan, Maldives. (d) Diego Garcia. (e) Columbo, Sri Lanka.




Figure 17. Comparison between JASON-1 elevation anomaly and modelled wave elevation along the satellite transect.

These comparisons demonstrate that the solver presented in this article gives results which are close to those obtained using similar methods (compressible schemes for the Saint-Venant equations) using regular longitude–latitude grids at comparable – but constant – spatial resolutions (\approx one nautical mile). Figure 18 illustrates the ratio of the number of discretisation elements required when using a constant spatial resolution versus the number of elements used by the adaptive method. The average reduction in the number of elements over the 10 hours of simulation is approximately 50. The total runtime for this simulation on a single Intel Core 2, 2.66 GHz processor was 13 hours.




Figure 18. Evolution of the ratio between the number of elements of a constant-resolution simulation and the number of elements of the adaptive simulation.

Is such a high resolution really necessary when considering processes at ocean scales? Figure 19 (left column) illustrates the effect of varying the maximum spatial resolution on time-series of wave elevation at two tide gauge locations (Male: top row and Columbo: bottom row). It is clear that the coarser-resolution simulations lead to significant numerical damping of the tsunami wave (while having little influence on the phase). This is also confirmed when examining maps of maximum wave elevation as a function of resolution (not shown). This damping is not due to local under-resolution (e.g. of bathymetry) but to numerical dissipation of the wave as it travels long distances across deep water. Other models show similar trends: for example the regular grid, linearised shallow-water model of Wang and Liu (2006) shows a significant under-estimation of wave elevation at the Male and Gan tide gauge locations when run at 2 nautical miles resolution. Note also that, while it looks like the time-series for the Male tide gauge (Figure 19 top left) is close to convergence (when comparing the 0.8 and 1.6 NM results), this does not seem to be the case for the Columbo tide gauge (Figure 19 bottom left). This suggests that still higher resolution could improve the propagation of the coastal “edge wave” around Sri Lanka.

Another source of numerical damping which seems to be rarely discussed (in the context of tsunami modelling) is the choice of limiter used in the second-order Godunov scheme. Figure 19 (right column) illustrates the increased wave dissipation induced by choosing a minmod limiter instead of a Sweby limiter. The minmod limiter is the most dissipative of the Total Variation Diminishing (TVD) limiters and is often employed in Godunov schemes for the Saint-Venant equations (particularly for river flows) due to the increased robustness brought by the increased dissipation. I have also tried using the superbee limiter, which is the least dissipative of the TVD limiters, but this lead to spurious oscillations of the wave front. The Sweby limiter was found to be the best compromise between dissipation and stability. Note how the increased dissipation induced by the minmod limiter is similar to that induced by roughly halving the spatial resolution.

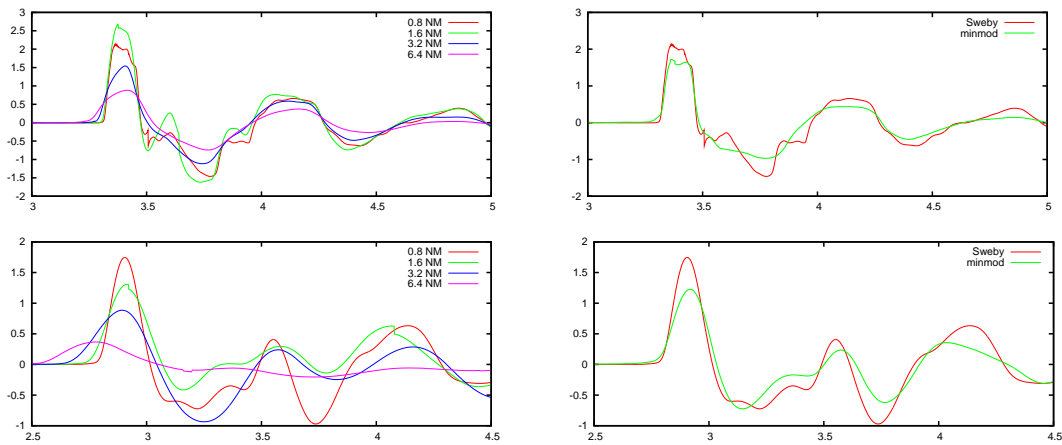


Figure 19. Modelled time-series of wave elevation (metres) at the Male (top row) and Columbo (bottom row) tide gauge locations. Left column: Influence of the maximum spatial resolution (in Nautical Miles). Right column: Influence of the choice of limiter for a maximum resolution of 0.8 NM.

Figure 18 shows that a factor of ≈ 50 in mesh size can be gained for this particular example when using the adaptive method at 0.8 NM maximum resolution. This factor is problem- and resolution-dependent. Figure 20.a illustrates how the (time-averaged) mesh size scales with maximum spatial resolution. While the constant-resolution mesh size increases quadratically with increasing resolution, the exponent in the adaptive case is only 1.4. This exponent is a measure of the “information content” of the wave elevation field and is closely related to the *fractal dimension* of this field. Indeed one of the definition of the fractal dimension is the Minkowski-Bouligand or box-counting dimension [Tricot (1995)] which involves covering a pattern with boxes and looking at how the number of boxes scales as their size decreases, which is almost what we do here. The 1.4 exponent indicates that the geometrical complexity of the wave pattern is intermediate between that of a curve (exponent 1) and a surface (exponent 2). This is consistent with the observation that the wave pattern can be seen as an ensemble of interacting “wave fronts” (i.e. a collection of curves). The shape of these curves is directly influenced by the bathymetry (particularly as they start wrapping around the coastline) and the 1.4 exponent also reflects the fractal dimension of the bathymetry/coastline (the archetypal fractal object [Mandelbrot (1967)]).

This interpretation can be further refined by looking at how the scaling exponent evolves in time (Figure 21). A first observation is that the scaling exponent generally decreases with increasing resolution. This could indicate that the asymptotic regime which defines the “true” Minkowski dimension of the field has not been reached yet, or that the field is not truly self-similar (i.e. does not strictly have a Minkowski dimension). In practice however this means that the 1.4 exponent is likely to be an upper bound on the scaling (a trend which is also apparent in Figure 20.a). The evolution in time of the scaling exponent is similar at all resolutions: we will focus on the 1.6/0.8 NM curve in the following description. In a first phase ($t < 30$ min) the scaling exponent reaches a maximum of ≈ 1.7 which reflects the geometrical complexity of the vertical displacement field created by the staggered fault ruptures (using the Okada fault model, see Figure 21.b). In a second phase ($1 \text{ hour} < t < 5 \text{ hours}$), the solution is dominated by individual wave fronts propagating mostly in deep water (Figures 14 and 21.c). This causes the exponent to decrease to less than 1.2, reflecting the quasi-unidimensional nature of the field. From 5 to 8 hours the solution contains both deep-water wave fronts and multiple wave reflections from strong interactions with coastlines (this interpretation could also explain the “bump” at 3.5 hours which corresponds to the wave front hitting the Maldives, Figure 14 bottom row). Finally at 10 hours most of the deep-water waves have left the domain or reached the coastline and the scaling exponent essentially reflects the geometrical complexity of the coastline (Figure 21.d). Coincidentally, the asymptotic exponent of 1.25 is identical to that computed by Mandelbrot for the west coast of Britain in his classic paper [Mandelbrot (1967)]. Exponents of 1.2 – 1.3 were later shown to be typical of convoluted coastlines [Carr and Benzer (1991)].

The reduction in “dimensional complexity” leads to the scale-dependent reduction in mesh size illustrated in Figure 20.b. One must be careful when trying to extrapolate these gains in mesh size into gains in running time. The quadtree mesh discretisation as well as the dynamic terrain reconstruction incur additional costs compared to a static discretisation on a regular longitude-latitude grid. In practice, we have found this overhead to be a factor 5 to 10 compared to the same numerical scheme implemented on regular grids. Taking this factor into account, it is clear from Figure 20.b that substantial gains in runtime can still be expected, particularly at high resolutions. Furthermore, the gain in *memory* requirements scales roughly as Figure 20.b (with minimal overhead from the quadtree discretisation) which can make simulations possible

at resolutions which would be impractical using regular grids (due to memory limitations).




Figure 20. (a) Evolution of the time-averaged (over 10 hours) mesh size as a function of maximum spatial resolution for constant and adaptive discretisations. (b) Ratio of constant and adaptive mesh sizes as a function of maximum spatial resolution.




Figure 21. (a) Evolution of the scaling exponent as a function of time and spatial resolutions. The scaling exponent is computed as: $\log(s_f/s_c)/\log(\Delta_c/\Delta_f)$ where s is the (instantaneous) number of cells and Δ is the maximum spatial resolution, for pairs of fine (subscript f) and coarse (subscript c) meshes. The pairs of spatial resolutions considered for each curve are given in the legend. The inset figures (b), (c) and (d) represent the spatial resolution at the different times indicated in Figure (a).

5 Conclusion

The well-balanced, positivity-preserving and conservative scheme of Audusse et al. (2004) for the solution of the Saint-Venant system of conservation laws has been successfully generalised to an adaptive quadtree discretisation and general orthogonal coordinates. Second-order global convergence, and first-order local convergence has been demonstrated for oscillations of a fluid in a parabolic container. The convergence is first-order close to the wetting-and-drying contact points, as expected for a solution with contact discontinuities.

Adaptive resolution of the classical Monai tsunami benchmark problem has also been studied. The results are comparable to those obtained using similar numerical methods on static grids but with a gain in mesh size of order 10 for the spatial resolutions considered.

A hierarchical database system designed for the retrieval and reconstruction of adaptive bathymetry has been described and analysed. This system can efficiently reconstruct multiscale representations of bathymetry at runtime, even for extremely large bathymetry datasets. The system is independent from the details of the numerical method and could also be useful for other solvers, for example using adaptive unstructured meshes.

The adaptive solver coupled with the bathymetry database system was used to simulate the 2004 Indian ocean tsunami with a ratio between the coarsest and finest resolved scales of order 100. The results were close to those previously published using similar solvers and the same (finest) spatial scales. The computation time for a given spatial resolution was shown to scale like

$$C_1 \Delta_\star^{-(\alpha+1)} \quad (11)$$

with C_1 a constant, Δ_\star the ratio between the finest and coarsest resolved scales, and $\alpha \approx 1.4$. The exponent α is closely related to the average fractal dimension of the geometrical structures characterising tsunami propagation. Although the constant C_1 will typically be larger than the corresponding constant for a regular Cartesian grid (due to the overheads of the quadtree data structure), it is clear that at sufficiently high resolutions the adaptive method will lead to increasingly large gains in computational efficiency compared to a regular Cartesian grid (for which $\alpha = 2$). Such favourable scalings were also demonstrated in a previous study on spectral wave modelling [Popinet et al. (2010)]. The efficiency of the adaptive method could be further improved by using a variable timestep depending on the local resolution. This would further decrease the overall exponent (which also includes the temporal dimension i.e. the $+1$ term in the exponent of (11)).

The method was implemented within the open-source Gerris Flow Solver framework [Popinet (2003a)] so that parallel simulations with dynamic load-balancing are also readily possible [Agbaglah et al. (2011)]. This combination of methods should provide a powerful tool for tsunami modelling which is also freely available to the community. In the same spirit, all the data and scripts necessary to reproduce the results presented in this article are available on the Gerris web site [Popinet (2010a), Popinet (2010b), Popinet (2010c)].

Acknowledgements

I would like to thank Mansour Ioualalen for sharing his analysis of the JASON-1 data.

Appendix A Database system for efficient terrain reconstruction

In the general context of finite-volume methods, quantities associated with discretisation elements are local spatial averages (over the volume of the element). Thus in our case topography will also be defined as spatial averages of characteristic quantities such as altitude and/or slope. If we assume that topography is defined as (x_i, y_i, z_i) triplets corresponding to discrete point samples of the real topography, the approximate average altitude for a given area \mathcal{A} will be

$$\bar{z} \equiv \frac{1}{m} \sum_{(x_i, y_i) \in \mathcal{A}} z_i,$$

where m is the number of samples contained in \mathcal{A} . The first problem to solve is thus a *point location* problem: given a set of N points, construct the subset of m points belonging to \mathcal{A} . An optimal algorithm for point location will have a minimal cost of $O(m)$ i.e. the minimum number of operations necessary just to report the result of the search. A naive algorithm would test each point of the set and would have a cost of $O(N) + O(m)$ where N is typically much larger than m . One way to speed up the search is to first construct a *spatial index* of the set. For example, storing points only at the vertices of a regular Cartesian grid directly provides a spatial index as the logical coordinates (i.e. memory/disk location) of a sample are simply related to their spatial coordinates. Such a data structure allows optimal point location, however it severely restricts the sampling strategy of the underlying data. While raw data could be resampled on a regular grid, this would only push the problem upstream (how to perform this resampling without an efficient point location algorithm?) while introducing other issues such as aliasing and accuracy of interpolations etc...

Fortunately efficient point location algorithms are not limited to Cartesian-grid data and there is a vast literature on the topic [Preparata and Shamos (1985)]. A simple and efficient data structure for point location in k dimensions (when k is small) is the kd -tree [Bentley and Friedman (1979)]. This data structure does not place any restriction on the coordinates of the samples and allows point location queries with a number of operations of order $O(N^{1-1/k}) + O(m)$. As its name indicate, the 2d-tree ($k = 2$ for our application) is a hierarchical data structure (a binary tree) which can be constructed recursively as follows

1. alternatively set c to either the x - or y -axis,
2. find the median coordinate \bar{c} of the point set along axis c ,
3. split the point set into two subsets according to whether their c -coordinate is smaller or larger than \bar{c} ,
4. repeat for each of the two subsets.

This procedure automatically constructs a hierarchy of axis-aligned bounding boxes which contains the hierarchy of sub-trees. This allows fast traversal of the tree for point location. The 2d-tree data structure itself is then just an ordered list of the (alternating axis) \bar{c} values for each of the sub-trees.

Coming back to the initial problem of estimation of spatial averages on a given area \mathcal{A} , we see that using a 2d-tree results in a cost of order $O(\sqrt{N}) + O(m)$. While N is usually much larger than m , \sqrt{N} can often be comparable to m so that the $O(m)$ contribution is not negligible in the total cost. For example, the 2004 tsunami simulation of section 4.3 uses a topography database with $N \approx 233$ million ($\sqrt{N} \approx 15,000$) while m can vary between one (for a 0.8 nautical mile resolution) and 10,000 (for a 101 nautical miles resolution). Note also that the $O(\sqrt{N})$ and $O(m)$ costs are only order estimates and that the real cost will also include pre-factors (multiplicative constants) which will depend on the relative cost of each set of operations. In practice, computing averages (the $O(m)$ cost) involves arithmetic operations which tend to be more expensive than the memory operations required for tree traversal (the $O(\sqrt{N})$ cost).

To reduce this cost, the 2d-tree data structure is augmented with spatial averages for each of the sub-trees. To do this the following step is added to the construction algorithm above

- 2.a compute $\Sigma_z \equiv \sum z_i$ for all the points in the point set.

The augmented 2d-tree data structure is then an ordered list of (\bar{c}, Σ_z) doublets for each of the sub-trees. In practice other moments are also stored for each sub-tree (e.g. $\Sigma_x, \Sigma_y, \Sigma_{xy}, \Sigma_{xz}$ etc...) so that higher-order average moments (slopes, curvature etc...) can also be computed efficiently. Given this data structure the following recursive algorithm can be used to retrieve the average altitude \bar{z} of a given area \mathcal{A} (\bar{z} is initially set to zero)

1. if the bounding box of the sub-tree is entirely outside \mathcal{A} do nothing,

2. if the bounding box of the sub-tree is entirely contained within \mathcal{A} : add Σ_z of the sub-tree to \bar{z} , add the number of points in the sub-tree to m ,
3. otherwise repeat for each of the two sub-trees.

The average altitude is then given by \bar{z}/m . Using this algorithm, for large m , the bulk of the points used for computing the average altitude will be contained in the bounding boxes of a few (\sqrt{m}) sub-trees; only the sub-trees which intersect the boundary of \mathcal{A} will need to be traversed. The number of operations will thus scale like the number of points lying on (or close to) the boundary of \mathcal{A} i.e. $O(\sqrt{m})$. The overall algorithm will thus scale like $O(\sqrt{N}) + O(\sqrt{m})$.

For the moment we have only described a data structure suitable for storage in main memory. This would obviously limit the capacity of the database to the size of available main memory rather than available disk storage. Aside from access speed, the main constraint of disk storage compared to main memory is that data access is performed in contiguous blocks of finite size (typically 4096 bytes). To obtain good performance it is then important to try to access data in contiguous chunks with a size which is a multiple of the disk block size. To follow this rule, the traversal of the 2d-tree is stopped when the amount of memory necessary to describe all the points in a sub-tree is smaller than the disk block size. Each point is encoded using a double-precision (x, y, z) triplets i.e. 24 bytes, so that the leaves of the 2d-tree will typically contain $4096/24 = 170$ points. An on-disk 2d-tree database with N points is thus made of three files:

1. the 2d-tree index: a linear array of \bar{c} values for each “node” of the tree. This will require $2N/170$ double-precision numbers i.e. $16N/170$ bytes.
2. average values: a linear array of averages (Σ_z etc...) for each node of the tree. This will require $2Np/170$ double-precision numbers where p is the number of averages stored.
3. point coordinates: a linear array of groups of 170 (x, y, z) triplets for each leaf of the tree. This will require $3N$ double-precision numbers i.e. $24N$ bytes.

Note that the 2d-tree index (the array of \bar{c} values) is optimal in term of storage requirements because we are using a *linear binary tree* which does not require extra storage (e.g. pointers) to encode the logical structure of the tree (the index of a node in the tree is given directly by arithmetic relations, see [Gargantini (1982)]). The same is true for the on-disk array of point coordinates and average values: the logical structure of the tree is encoded implicitly in the ordering of the values. Note also that the index itself is much smaller (by a factor of 255) than the space required to store the point coordinates. The 2d-tree index for terabytes-size databases could thus easily fit in main memory. This is not strictly necessary with the current layout but will improve database performance.

Aside from its good scaling and storage properties, the database system we have described has several desirable features:

- data from any sources can be easily combined even if they overlap, have variable sampling density etc...
- there is no constraints on the initial sampling of the raw data. The only requirement is that data need to be converted to a list of (x, y, z) triplets.

A further refinement which we have not implemented yet would be to add information about the quality of samples i.e. represent samples with a (x, y, z, q) quadruplet. Spatial averages would then be defined as *weighted* averages, which would give more weight to better-quality data in areas of overlap. Note that, even without quality information, the current system will by construction give more weight to high-density samples in areas of overlap. Implementing such a quality-weighted scheme would be trivial. The main obstacle is the (non)availability of quality information for many topographic datasets.

To assess the practical performance of the system and to verify the theoretical scalings, a simple experiment was performed. A realistic terrain database was created (the ETOPO1 database used in section 4.3) which covers the surface of the globe using longitude and latitude coordinates (i.e. $-180 < x \leq 180$ and $-90 < y \leq 90$) at a resolution of one arc-minute, giving a total of approximately 233 million samples. A random set of 1000 axis-aligned query squares were then generated with random centres (within $[-180: 180] \times [-90: 90]$) and random sizes within $[0: 10]$. Figure 22 summarises the query times for each query square. Two runs were performed: one with an empty disk cache (red circles) where all the data needed to be accessed on disk and a second one for which all data were already present in the main memory disk cache (i.e. no disk access was necessary). As expected disk access is much slower (by a factor of ~ 100) than main memory access which explains the shift between the two sets of curves. In both cases the expected \sqrt{m} dependence is obtained when m is large. When m is smaller than the number of samples in the leaves of the tree (~ 170), the query time becomes constant, also as expected. In all cases averages over a million samples are retrieved in less than about one second. The same test was also performed on databases containing 23, 2.3 and 0.23 million samples to evaluate the dependency in database size (which is expected to be $O(\sqrt{N})$). This showed only a very weak dependency in N which suggests that the practical prefactor of averaging (with $O(\sqrt{m})$ scaling) is much larger than that of tree traversal.

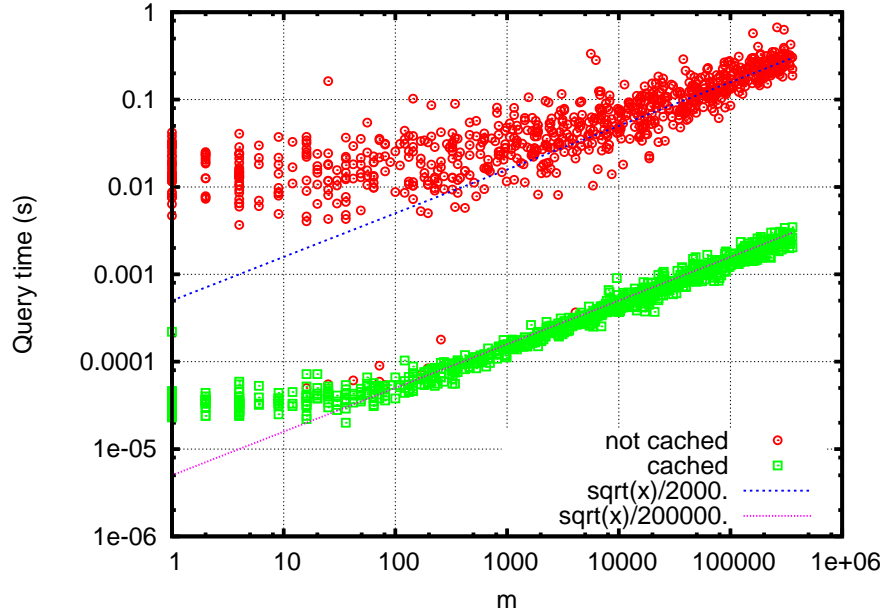


Figure 22. Query time as a function of the number of samples m used to compute local averages over a set of 1000 squares of random positions and sizes. The ETOPO1 database with 233 million samples is used. The red circles correspond to a “cold start” where no data is present in the main memory disk cache, the green squares are for a “hot start” where all data are already cached in main memory.

A special case which needs to be addressed is initialising the average depth of a cell which does not contain any depth sample. This can happen either when data sampling is very inhomogeneous (typical examples being sounding transects from survey ships or digitised contour lines) and/or when features in the flow field (e.g. a wave front) need to be resolved at a resolution higher than that of the bathymetry data. For this particular case, depth information needs to be interpolated from surrounding cells where depth is defined. To do this in a robust manner, we make use of the multiscale representation of the data encoded in the quadtree. The average depth is defined starting from the root of the quadtree (i.e. a single square which contains all the bathymetry data). This is repeated for each successively finer level of the quadtree. If the cells on each level contain depth samples, these are used to compute the average depth. Otherwise, a depth value is constructed using the average depth and slope values defined for the (coarser) parent cell. Any “data gap” (irrespective of its size) will thus be filled with a linearly-interpolated value using as much information as possible from parent cells.

Bibliography

- [IWLWRM (2004)] The third international workshop on long-wave runup models. http://isec.nacse.org/workshop/2004_cornell/bmark2.html, June 2004.
- [Sea Level Center (2006)] University of Hawai'i Sea Level Center. <http://uhslc.soest.hawaii.edu/>, 2006.
- [Agbaglah et al. (2011)] G. Agbaglah, S. Delaux, D. Fuster, J. Hoepffner, C. Josserand, S. Popinet, P. Ray, R. Scardovelli, and S. Zaleski. Parallel simulation of multiphase flows using octree adaptivity and the volume-of-fluid method. *Compte-rendus de l'Académie des Sciences, Paris*. 339: 194–207, 2011.
- [Amante and Eakins (2008)] C. Amante and B. W. Eakins. ETOPO1 1 arc-minute global relief model: procedures, data sources and analysis. Technical report, National Geophysical Data Center, NESDIS, NOAA, US Department of Commerce, Boulder, CO, 2008.
- [Audusse et al. (2004)] E. Audusse, F. Bouchut, M.-O. Bristeau, R. Klein, and B. Perthame. A fast and stable well-balanced scheme with hydrostatic reconstruction for shallow water flows. *SIAM J. Sci. Comp.*, 25 (6): 2050–2065, 2004.
- [Bentley and Friedman (1979)] J.L. Bentley and J.H. Friedman. Data structures for range searching. *ACM Computing Surveys (CSUR)*, 11 (4): 397–409, 1979. ISSN 0360-0300.
- [Berger and Olinger (1984)] M.J. Berger and J. Olinger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of computational Physics*, 53 (3): 484–512, 1984. ISSN 0021-9991.
- [Yalciner et al. (2005)] Yalciner A. C., Ghazali N. H., and Wahab A. K. A. December 26, 2004 indian ocean tsunami field survey july 09-10, 2005 at north west peninsular malaysia coast, penang and langkawi islands. Technical report, ITST of UNESCO IOC, 2005. <http://yalciner.ce.metu.edu.tr/sumatra/survey>.
- [Carr and Benzer (1991)] J.R. Carr and W.B. Benzer. On the practice of estimating fractal dimension. *Mathematical Geology*, 23 (7): 945–958, 1991.
- [de Saint-Venant (1871)] A. Barré de Saint-Venant. Théorie du mouvement non-permanent des eaux, avec application aux crues des rivières et à l'introduction des marées dans leur lit. *CR Acad. Sci. Paris*, 73: 147–154, 1871.
- [Gargantini (1982)] I. Gargantini. An effective way to represent quadrees. *Communications of the ACM*, 25 (12): 905–910, 1982. ISSN 0001-0782.
- [George and LeVeque (2008)] D.L. George and R.J. LeVeque. High-resolution methods and adaptive refinement for tsunami propagation and inundation. *Hyperbolic Problems: Theory, Numerics, Applications*, pages 541–549, 2008.
- [Grilli et al. (2007)] S.T. Grilli, M. Ioualalen, J. Asavanant, F. Shi, J.T. Kirby, and P. Watts. Source constraints and model simulation of the December 26, 2004, Indian Ocean Tsunami. *Journal of Waterway, Port, Coastal, and Ocean Engineering*, 133: 414–428, 2007.
- [Harig et al. (2008)] S. Harig, Chaeroni, W.S. Pranowo, and J. Behrens. Tsunami simulations on several scales. *Ocean Dynamics*, 58 (5): 429–440, 2008. ISSN 1616-7341.
- [Harten et al. (1983)] A. Harten, P.D. Lax, and B. Van Leer. On upstream differencing and Godunov-type schemes for hyperbolic conservation laws. *SIAM review*, pages 35–61, 1983. ISSN 0036-1445.
- [Ioualalen et al. (2007)] M. Ioualalen, J. Asavanant, N. Kaewbanjak, S.T. Grilli, J.T. Kirby, and P. Watts. Modeling the 26 December 2004 Indian Ocean tsunami: Case study of impact in Thailand. *J. Geophys. Res.*, 112: 1–21, 2007.
- [LeVeque (1998)] R.J. LeVeque. Balancing source terms and flux gradients in high-resolution godunov methods: The quasi-steady wave-propagation algorithm. *Journal of Computational Physics*, 146 (1): 346–365, 1998. ISSN 0021-9991.
- [LeVeque and George (2006)] R.J. LeVeque and D.L. George. High-resolution finite volume methods for the shallow water equations with bathymetry and dry states. *Advanced numerical models for simulating tsunami waves and runup*, 10: 43–73, 2006.
- [Liang and Borthwick (2009)] Q. Liang and A.G.L. Borthwick. Adaptive quadtree simulation of shallow flows with wet-dry fronts over complex topography. *Computers & Fluids*, 38 (2): 221–234, 2009.
- [Liang and Marche (2009)] Q. Liang and F. Marche. Numerical resolution of well-balanced shallow water equations with complex source terms. *Advances in Water Resources*, 32 (6): 873–884, 2009. ISSN 0309-1708.
- [Liang et al. (2004)] Q. Liang, A.G.L. Borthwick, and G. Stelling. Simulation of dam-and dyke-break hydrodynamics on dynamically adaptive quadtree grids. *International Journal for Numerical Methods in Fluids*, 46 (2): 127–162, 2004. ISSN 1097-0363.
- [Lopez-Herrera et al. (2011)] J. M. Lopez-Herrera, S. Popinet, and M. A. Herrada. A charge-conservative approach for simulating electrohydrodynamic two-phase flows using volume-of-fluid. *Journal of Computational Physics*, 230: 1939–1955, 2011. URL <http://gfs.sf.net/papers/ehd.pdf>.

- [Mandelbrot (1967)] B. Mandelbrot. How long is the coast of Britain? Statistical self-similarity and fractional dimension. *Science*, 156 (3775): 636–638, 1967.
- [Nielsen et al. (2005)] O. Nielsen, J. Sexton, and J. Schneider. The ANUGA tsunami inundation model. In *Proceedings of ModSIM*, pages 518–523, 2005.
- [Okada (1985)] Y. Okada. Surface deformation due to shear and tensile faults in a half-space. *Bulletin of the Seismological Society of America*, 75 (4): 1135, 1985. ISSN 0037-1106.
- [Pietrzak et al. (2007)] J. Pietrzak, J., A. Socquet, D. Ham, W. Simons, C. Vigny, R.J. Labeur, E. Schrama, G. Stelling, and D. Vatvani. Defining the source region of the Indian Ocean Tsunami from GPS, altimeters, tide gauges and tsunami models. *Earth and Planetary Science Letters*, 261 (1-2): 49-64, 2007.
- [Popinet (2003a)] S. Popinet. The Gerris Flow Solver. <http://gfs.sf.net>, 2003a.
- [Popinet (2003b)] S. Popinet. Gerris: a tree-based adaptive solver for the incompressible euler equations in complex geometries. *J. Comput. Phys.*, 190 (2): 572–600, 2003b. URL <http://gfs.sf.net/gerris.pdf>.
- [Popinet (2009)] S. Popinet. An accurate adaptive solver for surface-tension-driven interfacial flows. *Journal of Computational Physics*, 228: 5838–5866, 2009. URL <http://gfs.sf.net/papers/tension.pdf>.
- [Popinet (2010a)] S. Popinet. Gerris – Oscillations in a parabolic container test case. <http://gfs.sourceforge.net/tests/tests/parabola.html>, 2010a.
- [Popinet (2010b)] S. Popinet. Gerris – Monai tsunami benchmark. <http://gfs.sourceforge.net/examples/examples/monai.html>, 2010b.
- [Popinet (2010c)] S. Popinet. Gerris – The 2004 Indian Ocean tsunami. <http://gfs.sourceforge.net/examples/examples/tsunami.html>, 2010c.
- [Popinet and Rickard (2007)] S. Popinet and G. Rickard. A tree-based solver for adaptive ocean modeling. *Ocean Modelling*, (16): 224–249, 2007. URL <http://gfs.sf.net/ocean.pdf>.
- [Popinet et al. (2010)] S. Popinet, R. M. Gorman, G. J. Rickard, and H. L. Tolman. A quadtree-adaptive spectral wave model. *Ocean Modelling*, 34: 36–49, 2010. URL <http://gfs.sf.net/papers/wave.pdf>.
- [Pranowo et al. (2008)] W.S. Pranowo, J. Behrens, J. Schlicht, and C. Ziemer. Adaptive mesh refinement applied to tsunami modeling: Tsunaflash. In *Proceedings of the International Conference on Tsunami Warning (ICTW), Bali, Indonesia*, volume 6, 2008. URL http://pirba.ristek.go.id/ictw/e-proceedings/PDF/3.%20DMS/DMS25DE_Pranowo_Widodo_r.pdf.
- [Preparata and Shamos (1985)] F.P. Preparata and M.I. Shamos. *Computational geometry: an introduction*. Springer, 1985. ISBN 0387961313.
- [Rančić et al. (1996)] M. Rančić, R. J. Purser, and F. Mesinger. A global shallow-water model using an expanded spherical cube: Gnomonic versus conformal coordinates. *Quarterly Journal of the Royal Meteorological Society*, 122 (532): 959–982, 1996. ISSN 1477-870X.
- [Rogers et al. (2003)] B.D. Rogers, A.G.L. Borthwick, and P.H. Taylor. Mathematical balancing of flux gradient and source terms prior to using Roe’s approximate Riemann solver. *Journal of Computational Physics*, 192 (2): 422–451, 2003. ISSN 0021-9991.
- [Sampson et al. (2006)] J. Sampson, A. Easton, and M. Singh. Moving boundary shallow water flow above parabolic bottom topography. *ANZIAM J*, 47, 2006.
- [Sweby (1984)] P.K. Sweby. High resolution schemes using flux limiters for hyperbolic conservation laws. *SIAM Journal on Numerical Analysis*, 21 (5): 995–1011, 1984. ISSN 0036-1429.
- [Kawata (2005)] T. Kawata. Comprehensive analysis of the damage and its impact on coastal zones by the 2004 indian ocean tsunami disaster. Technical report, Disaster Prevention Research Institute, 2005.
- [Tricot (1995)] C. Tricot. *Curves and fractal dimension*. Springer, 1995. ISBN 0387940952.
- [Walters (2006)] R.A. Walters. Design considerations for a finite element coastal ocean model. *Ocean Modelling*, 15 (1-2): 90–100, 2006.
- [Wang and Liu (2006)] X. Wang and P.L.F. Liu. An analysis of 2004 Sumatra earthquake fault plane mechanisms and Indian Ocean tsunami. *Journal of Hydraulic Research*, 44 (2): 147–154, 2006.
- [Williamson et al. (1992)] D.L. Williamson, J.B. Drake, J.J. Hack, R. Jakob, and P.N. Swarztrauber. A standard test set for numerical approximations to the shallow water equations in spherical geometry. *Journal of Computational Physics*, 102 (1): 211–224, 1992. ISSN 0021-9991.
- [Yamada et al. (2005)] F. Yamada, H. Fernandez, J. Goff, B. Higman, and B. Jaffe. An overview of the ITST expedition to Sri Lanka. In *Proc. 5th Int. Conf. on Ocean Wave Measurement and Analysis, WAVES 2005*, 2005.
- [Yiu et al. (1996)] K. F. C. Yiu, D. M. Greaves, S. Cruz, A. Saalehi, and A.G.L. Borthwick. Quadtree grid generation: information handling, boundary fitting and CFD applications. *Computers & Fluids*, 25 (8): 759–769, 1996. ISSN 0045-7930.
- [Zhou et al. (2001)] J.G. Zhou, D.M. Causon, C.G. Mingham, and D.M. Ingram. The surface gradient method for the treatment of source terms in the shallow-water equations. *Journal of Computational Physics*, 168 (1): 1–25, 2001. ISSN 0021-9991.