Adaptive vertical layering in TELEMAC-3D

Chris J. CAWTHORN, Sébastien E. BOURBAN, Chris T. MEAD, Matthew J. WOOD, Michael .S. TURNBULL
Hydrodynamics and Metocean Group, HR Wallingford
Howbery Park, Wallingford
Oxfordshire, OX10 8BA, UK
Correspondence to: s.bourban@hrwallingford.com

***Abstract*— Many processes in environmental hydraulics exhibit sharp spatial gradients of some physical variable(s) in a small localised part of the overall water column. Examples of this include spreading of dense or buoyant plumes and thermal or saline stratification in reservoirs. In this paper, we demonstrate a robust adaptive mesh redistribution (AMR) method coded for TELEMAC-3D. The AMR method aims to capture these sharp gradients without requiring an excessive number of mesh layers or any prior knowledge of the flow structure.**

Rather than increasing the number of mesh planes in regions of sharp spatial gradients, the idea of mesh redistribution is to maintain a fixed number of planes that move in response to the local solution structure. The movement of the planes is governed by a diffusion equation; an approach that is discussed in Ref. [1]. This approach is similar to that used in the popular GETM software (Ref. [2]). Mesh plane elevations linked to gradients in tracer concentration only are discussed in this paper, although the extension to include velocity shear and/or bathymetry in the equations governing plane placement is expected to be straightforward.

We present preliminary results demonstrating that the AMR method can adapt to relatively thin tracer plumes without the increased mesh resolution that would be required with some form of sigma mesh. Comparisons are drawn with an alternative approach in which plane elevations are specified by the user based on some *a priori* knowledge of the flow structure. The AMR method, which requires neither prior information about the flow nor user input, can be seen to give very similar results for the spreading of dense and buoyant plumes.

I. Introduction

Whether in the form of velocity shear layers or rapid changes in saline or thermal stratifications, thin layers with large gradients form an important part of many hydrodynamic processes. Such sharp spatial gradients can pose a troublesome problem for the numerical modelling of fluid flows. In order to accurately capture such layers, it is necessary to have a mesh that is sufficiently well resolved. In the context of TELEMAC-3D, which uses a set of identical two-dimensional meshes (referred to as 'planes') stacked vertically (see Ref. [3] for details), one must aim to decrease the spacing between mesh planes in regions of sharp vertical variations.

The spacing between mesh planes in TELEMAC-3D can be reduced in two ways. The simplest method is to use many horizontal planes in the TELEMAC-3D mesh. Even with a standard sigma-mesh (for which the planes are equispaced), this can provide adequate resolution of important thin layers if a large number of planes are used. However, this dramatically increases the cost of the computation. It is also wasteful in the sense that some areas of the mesh will inevitably have fine vertical resolution where capturing the fine detail of the dynamics is unnecessary.

A slightly more sophisticated approach would be to modify the CALCOT subroutine to place the layers at specific positions in the water column, carefully chosen based on the expected solution behaviour. Although this can be an effective technique, it is limited by the fact that the user must have some prior knowledge of the flow structure before modifying the layer positions. This means that such an approach must usually be made in an iterative manner: gradually adjusting the plane positions and re-running the simulation with the goal of converging on some 'optimal' configuration.

The aim of this paper is to introduce an *automatic* mesh layering approach to TELEMAC-3D, which we refer to as adaptive mesh redistribution (AMR). The idea is to devise an algorithm that moves the plane positions based on certain aspects of the local solution structure. For example, one might wish to concentrate layers in regions of sharp velocity gradient, tracer gradient or arbitrary linear combinations of these two variables.

Adaptive mesh refinement itself is not a new idea. In fact, the general formulation presented here was introduced by Winslow in 1966 [1]. Since that time, AMR methods have been substantially refined, but the general principles remain the same. We use a version of Winslow's "variable diffusion" approach described in more detail in [4], with a few refinements of our own designed specifically to work on the type of layered mesh used by TELEMAC-3D.

An outline of the AMR scheme implemented in TELEMAC-3D, as well as the underlying mathematics, is given in section II. Some simulation results are presented in section III, where comparisons are drawn with a simple sigma mesh. Finally, a discussion of the results and some possible extensions follows in section IV.

II. Adaptive Mesh Redistribution Scheme

The type of adaptive mesh refinement that we have introduced to TELEMAC-3D is based on the variational formulation used originally by Winslow [1]. In this section, we give a brief overview of the variational approach, and describe its specific implementation in TELEMAC-3D.

A. Variational principle for mesh redistribution in 1D

Suppose that the computational domain is represented in one dimension by n nodes $\{\xi_i : i = 1,n\}$, with

$$0 = \xi_1 < \xi_2 < \ldots < \xi_n = 1$$

and that the real domain (for example, the interval $a < x < b$) is then represented at the points $\{x_i : i = 1,n\}$. We then define a one-to-one mapping, $X(\xi)$, such that

$$X(\xi_i) = x_i\,, \quad \text{with} \quad X(0) = a \quad \text{and} \quad X(1) = b.$$

The variational approach is to find the mesh map $X(\xi)$ that minimises a functional of the form

$$E[X(\xi)] = \int \omega(X(\xi))|X'(\xi)|^2 d\xi \tag{1}$$

where a dash (') indicates differentiation of a function with respect to its argument. The *monitor function*, defined here by $\omega(x)$, is a positive definite function that in general depends on the structure of the solution that is to be calculated using the AMR approach. Typically, one might want to focus the mesh resolution in regions where the gradient of the function $f(x)$ is large. In such a case, an appropriate monitor function would be

$$\omega(x) = \sqrt{1 + a\left(\frac{df}{dx}\right)^2} \tag{2}$$

where a is a tuning parameter that will be discussed later.

The Euler-Lagrange equation for $X(\xi)$ associated with minimising the functional (1) is

$$\frac{d}{d\xi}\left[\omega(X(\xi))\frac{dX}{d\xi}(\xi)\right] = 0 \tag{3}$$

subject to the boundary conditions $X(0) = a$, $X(1) = b$.

The adaptive mesh refinement method therefore consists of solving a diffusion equation for the mesh node positions at each timestep. The diffusion coefficient is spatially-varying, and depends upon the current solution. This makes it possible to attract nodes to regions of interest, such as where the solution has large gradients.

B. Implementation in TELEMAC-3D

The 'variable diffusion' method described above has been implemented in TELEMAC-3D, and has been released as part of version 6.1. It is accessible by using the keyword *MESH TRANSFORMATION = 5,* and functions as an additional option in the CALCOT subroutine. It works by using an iterative (Gauss-Seidel) approach to solve the diffusion equation (3) on each vertical line of nodes in the three-dimensional mesh.

The method is currently implemented to follow only the gradient of tracer 1, using the monitor function shown in (2), with $f(x)$ representing the concentration of tracer 1. The extension to consider other physical variables and higher derivatives ought to be, in principle, a straightforward task.

In order to make the adaptive layering scale-free, the tuning parameter in (2) is chosen independently on each vertical line of nodes to be

$$a = a(x,y) = \frac{C}{\max|f'(x)|^2} \tag{4}$$

where the maximum runs over each of the nodes on the vertical line. This choice means that a very large solution gradient in one part of the mesh will not affect the mesh layering in another region with smaller (yet still significant) gradient. The constant parameter, C, can be tuned to increase or decrease the sensitivity of mesh layer positions to tracer concentration gradient. Large values of C can produce too large a deformation of the mesh planes, resulting in numerical instability. If C is too small, however, then the planes movement will not result in sufficient resolution of sharp solution gradients. In practice, we have found values of C ranging from 10 to 100 to be a good compromise between these two extremes for all of the examples studied.

Strong horizontal variation in the plane positions of a 3D mesh in TELEMAC-3D can have a destabilising effect on the simulation. In order to reduce the horizontal variation of layer positions on a local scale, the monitor function is smoothed using a simple low-pass filter in two dimensions before solving the diffusion equation (3) for layer positions.

The presence of maxima or minima in tracer concentration in the interior of the water column (as opposed to extrema at the free surface or bottom boundaries) raises a small problem for the AMR method. Because such extrema have low gradients, the AMR scheme will attempt to move solution points away from any local maxima or minima. When the solution is interpolated onto the new layer positions, this can change the position and magnitude of the extremum, resulting in a form of numerical diffusion. In the TELEMAC-3D implementation, we have attempted to eliminate this problem by first locating any local extrema in each water column, and ensuring that such points must feature in the new mesh configuration. This ensures that no interpolation takes places at extreme points, so the solution magnitude there cannot be diminished.

III. Some Examples

We now illustrate the automatic mesh redistribution method in TELEMAC-3D by considering some simple examples. In each case, the layer positions are modified according to the tracer concentration gradient alone.

A. Tracer advection over a compound slope

The first example of adaptive layering concerns a buoyant tracer released at the upstream boundary of a straight channel. The channel bed is formed from two planar slopes with different inclinations and a constant flow rate from left to right is imposed at the upstream boundary.

Fig. 1 shows the resulting plume using a sigma mesh, whilst Fig. 2 shows the same plume simulated using the adaptive layering approach described in Section II. The AMR algorithm has moved three of the four internal planes into the spreading front of the plume, whereas the plume front only occupies a single mesh layer in the sigma-mesh case. The plume is also more concentrated towards the free surface when using the AMR approach, indicating a reduction in numerical diffusion caused by the divergence of the planes in the sigma mesh.

Finally, note that the adaptive layering reduces to equispaced sigma-layering ahead of the spreading plume, where there is no tracer concentration gradient. This demonstrates that the adaptive layering will only take effect in regions of sharp tracer gradient, leaving the mesh unchanged elsewhere.

B. Dense tracer source in a straight horizontal channel

For our second example we consider a point source of dense tracer located at the bottom of a straight channel with rectangular cross-section. The channel is 1km long, 100m wide and 10m deep. A depth-averaged velocity of 1m/s is applied along the channel. At the source, dense fluid with an excess salinity of 215 parts-per-thousand is added at a rate of 0.5m^3/s. These source and flow parameters are typical of those found in studies of hypersaline discharge dispersion conducted by HR Wallingford.

TELEMAC-3D was first used to establish a steady velocity profile without the tracer source, and then this steady profile was used as an initial condition for a second computation including the dense tracer source. Fig. 3 shows a comparison of the results of simulations of this dense plume taken 240 minutes after release began.

By comparing panels (a) and (b) of Fig. 3, it is clear that the plume modelled using the adaptive mesh contains more mesh planes than that modelled with the sigma mesh, despite the sigma mesh having almost twice as many planes overall. The focussed resolution gives the resulting plume a more realistic shape. This is particularly apparent at the upstream end of the plume, where the AMR result shows a more rounded front than that predicted using a sigma mesh.

It is interesting to note that additional simulations carried out using a sigma mesh with more layers (not shown) seem to suggest a convergence towards a tracer concentration distribution very similar to that obtained using the AMR approach.

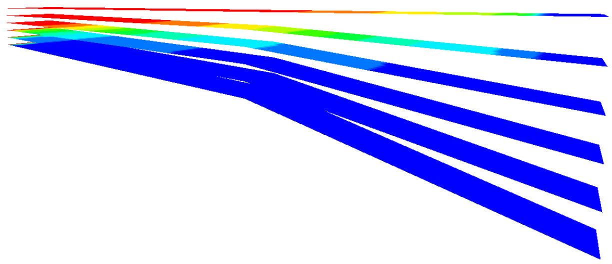

Figure 1. Buoyant tracer release over a compound slope without adaptive mesh redistribution.

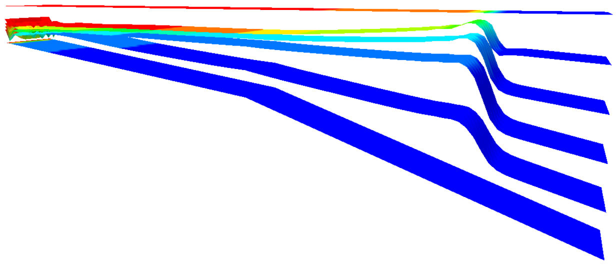

Figure 1. Buoyant tracer release over a compound slope with adaptive mesh redistribution.

Panels (c) and (d) of Fig. 3 clearly show that the horizontal spreading of tracer is strongly affected by the choice of layering strategy. The dense plume spreads further across the flow using the sigma mesh when compared to the AMR approach. We believe that this cross-flow spreading follows from a ‘blocking’ phenomenon, caused by the spurious sharp front upstream of the source. This obstructs the ambient flow, forcing fresh water around the dense plume. For the adaptive mesh, the plume occupies a smaller vertical extent, thus it has a reduced effect on the ambient flow. Dense fluid is therefore mostly swept directly downstream, with minimal cross-stream spreading. This hypothesis is supported by investigation of the velocity fields in both cases (not shown).

This example highlights a very important fact about vertical mesh spacing in TELEMAC-3D. It shows that a crude sigma mesh may predict spurious flow patterns if insufficient mesh planes are used. The fact that the spreading of the dense plume tends to resemble that obtained using AMR as the number of sigma mesh planes are increased suggests that the predictions of the AMR method are more accurate than those obtained using a sigma mesh, even when using far fewer planes.

The implications of these observations for studies of dense discharges, though significant, are not discussed here. Interested readers are directed to Ref. [5] for a discussion of the importance of mesh plane positions in discharge studies.

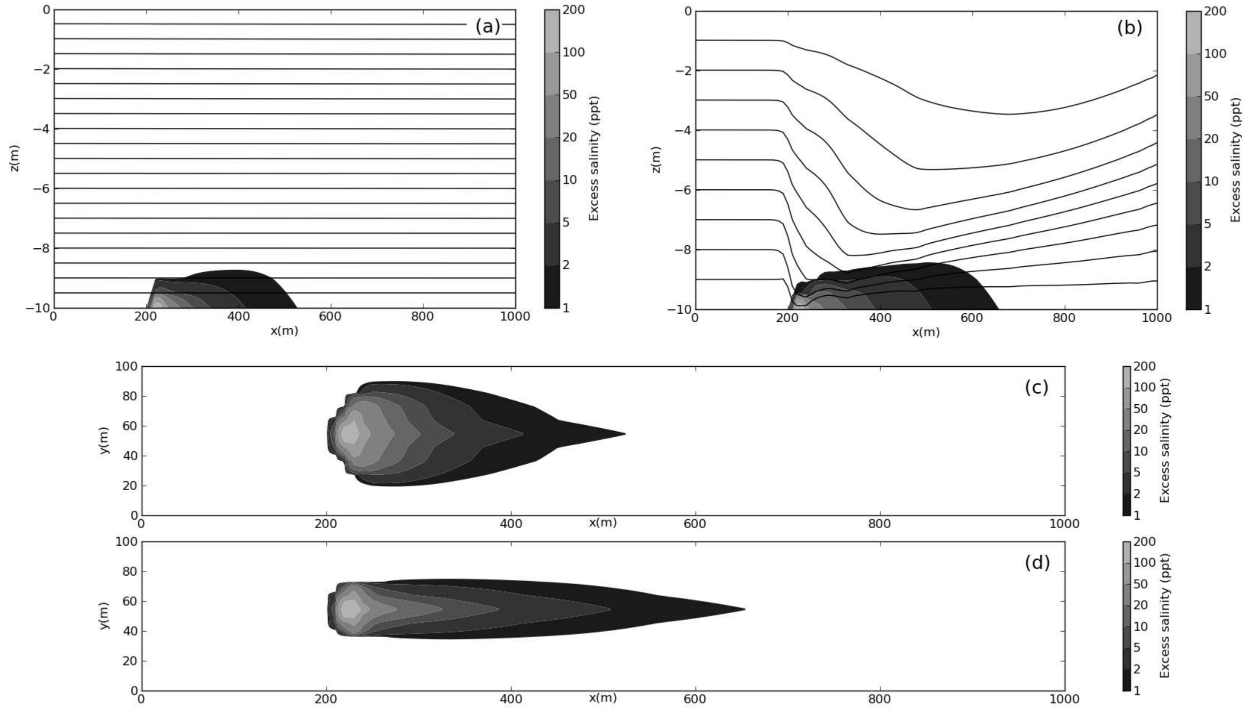

Figure 3. Dense saline source in channel flow. Shading shows the excess salinity in parts-per-thousand (ppt) 240 minutes after release began. Panels (a) and (b) show a vertical cross-section through the source point, combining the layer positions and excess salinity in the same plot. Panels (c) and (d) show the excess salinity at the bottom of the channel (i.e. on the lowest plane) for each simulation.

C. Saline lock exchange

For our final example, we turn to a simple lock-exchange problem, in which relatively dense and relatively buoyant dense fluids are initially contained side-by-side in a cuboidal container. As the denser fluid sinks, it spreads beneath the less dense fluid, driven by gravity. This example is distinct from the previous examples in that the tracer concentration (which is directly linked to the fluid density) is the only factor driving the fluid flow. In the earlier examples, the applied ambient flow was large enough to dominate the dynamics.

Fig. 4 and Fig. 5 show the solution and mesh layers for the lock exchange flow using sigma and adaptively-layered meshes, respectively. The most important feature to note is that the shape of the lock exchange current (as visualised by the colours in Figs. 4 and 5) is essentially the same in each case. This demonstrates that the AMR algorithm is not degrading the quality of the solution. In fact, the clustering of mesh planes near the spreading and receding fronts of the current ought to make the solution more accurate at these key locations.

In the centre of the domain, the transition between dense and light fluid (visually, the transition from red to blue) is sharper in Fig. 5 than in Fig. 4, indicating once more that the AMR method reduces vertical numerical diffusion by increasing the mesh resolution in the transition region.

At the ends of the domain, the tracer concentration is essentially constant within each water column, so the equispaced sigma-mesh is retained.

IV. Discussion

We have introduced to TELEMAC-3D a powerful method of automatically increasing mesh resolution in key portions of the water column. It is important to note that this adaptive mesh redistribution does not change the number of nodes or elements in the mesh. Furthermore, the mesh redistribution algorithm adds only a modest cost in terms of computation time.

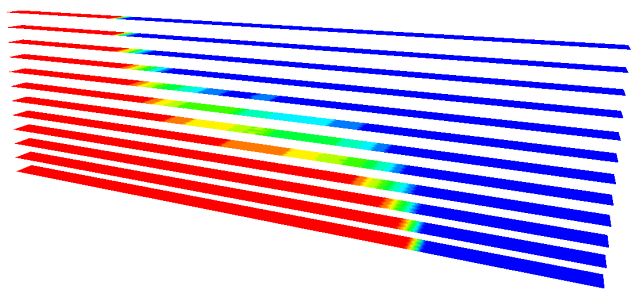

Figure 4. Lock exchange flow using a standard sigma mesh. Colours represent the salinity distribution.

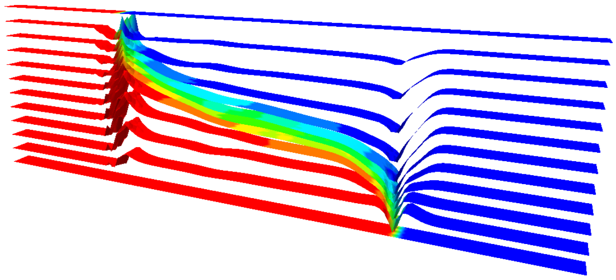

Figure 5. Lock exchange flow using the AMR algorithm. Colours represent the salinity distribution.

The AMR algorithm that is currently implemented in TELEMAC-3D release 6.1 is a first example of what will eventually become a general adaptive layering toolbox for use in TELEMAC-3D. At present, the AMR implementation allows layers to cluster where the tracer concentration gradient is largest, and can be easily generalised to focus on velocity gradient or linear combinations of velocity and tracer gradients in much the same way as other established software, notably the GETM package [2].

One can imagine cases for which it is more appropriate to consider higher spatial derivatives of physical variables. An example of this might feature relatively small deviations from a strong saline background gradient. Interesting behaviour may occur where the salinity *gradient* changes sharply, but an AMR approach based on salinity gradient alone will not cluster layers at such locations. Instead, basing the mesh redistribution on the *curvature* of the salinity function (for example) ought to provide increased resolution in regions of interest. This can be achieved by considering a different monitor function in place of (2), with the curvature operator replacing gradient. Higher derivatives still can be tracked with equally straightforward modifications.

We look forward to further development of the AMR options in TELEMAC-3D, which we hope will be carried out both at HR Wallingford and in cooperation with other open-source developers. As part of the development process, we would greatly welcome any feedback from users regarding generalisations or suggestions for improvement.

REFERENCES

[1] A. M. Winslow, "Numerical solution of the quasilinear Poisson equation in a nonuniform triangle mesh", J. Comp. Phys., vol. 1, pp. 149-172), 1966.

[2] R. Hofmeister, H. Burchard, and J.-M. Beckers, "Non-uniform adaptive vertical grids for 3D numerical ocean models", Ocean Modelling, vol. 33, pp. 70-86, 2010.

[3] J.-M. Hervouet, Hydrodynamics of free surface flows: Modelling with the finite element method, Wiley, 2007.

[4] H. Tang and T. Tang, "Adaptive mesh methods for one- and two-dimensional hyperbolic conservation laws", SIAM J. Numer. Anal, vol. 41, pp. 487-515, 2003.

[5] C. T. Mead, S. E. Bourban, C. J. Cawthorn and M. S. Turnbull, "The importance of model plane location and movement in dense discharge assessment", Proceedings of the International Symposium on Outfall Systems, Mar del Plata, Argentina, May 2011.