Ocean Modelling 73 (2014) 30–44

Contents lists available at ScienceDirect

Ocean Modelling

journal homepage: www.elsevier.com/locate/ocemod

Assessment of spurious mixing in adaptive mesh simulations of the two-dimensional lock-exchange

H.R. Hiester [a,*], M.D. Piggott [b,c], P.E. Farrell [b,d], P.A. Allison [b]

[a] Center for Ocean-Atmospheric Prediction Studies, Florida State University, 2000 Levy Avenue, Building A, Suite 292, Tallahassee, FL 32306-2741, USA
[b] Applied Modelling and Computation Group, Department of Earth Science and Engineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK
[c] Grantham Institute for Climate Change, Imperial College London, South Kensington Campus, London SW7 2AZ, UK
[d] Center for Biomedical Computing, Simula Research Laboratory, Oslo, Norway

ARTICLE INFO

Article history:
Received 5 May 2013
Received in revised form 25 September 2013
Accepted 10 October 2013
Available online 30 October 2013

Keywords:
Lock-exchange
Diapycnal mixing
Adaptive mesh
Finite-element methods

ABSTRACT

Numerical simulations are used to evaluate the impact of adaptive meshes on the two-dimensional lock-exchange flow. In particular, the diapycnal mixing is quantified through calculation of the background potential energy. The choice of metric, which guides the mesh adapt, is fundamental to the success of an adaptive mesh simulation. The performance of different Hessian-based metrics is assessed and cases that both outperform and underperform, compared to fixed mesh simulations, are evaluated. The differences in performance result from the different forms of the metric and the extent to which smaller-scale fluctuations can influence the adapted mesh. The best performing metric produces levels of diapycnal mixing that are comparable to high resolution fixed mesh simulations that use one to two orders of magnitude more mesh vertices. Comparison of the mixing with the numerical simulations of Özgökmen et al. (2007) also demonstrates the validity of the adaptive mesh simulations.

1. Introduction

The range of temporal and spatial scales of ocean flows is vast, differing from hours to centuries and metres to thousands of kilometres. The ocean is also full of transient features that can change in both size and/or location; examples include algal blooms, dense water overflows and mesoscale eddies. In an ocean model how much, when and where to place numerical resolution, both spatial and temporal, must be considered and cannot necessarily be predicted *a priori*. Adaptive meshes, which coarsen or refine depending on the evolution of the flow, support efficient use of available computational resources and, in principle, do not require an extensive *a priori* knowledge of the dynamics (e.g. Behrens, 1998; Jacobs et al., 2013; Munday et al., 2010; Popinet and Rickard, 2007).

Using an adaptive mesh adds another layer of numerical complexity to a model. The performance of such meshes and the implications for the computed flow dynamics therefore require careful consideration. Adaptive mesh techniques have been used relatively widely in computational fluid dynamics (Baker, 1997; Cao, 2005; Frey and Alauzet, 2005; Remacle et al., 2005; Speares and Berzins, 1997; Venditti and Darmofal, 2003), with the use of adaptive meshes in ocean modelling still under development (Piggott et al., 2009). For structured meshes, studies include the application and extension of a quadtree based adaptive structured mesh Navier–Stokes solver (Gerris) to ocean flows (Popinet and Rickard, 2007) and investigation of a general adaptive structured mesh tool (Blayo and Debreu, 1999). For unstructured meshes, the studies have focused predominantly on the shallow-water equations (Behrens, 1998; Bernard et al., 2007; Remacle et al., 2006) with limited applications in three dimensions (Munday et al., 2010; Piggott et al., 2008; Power et al., 2006). Modelling flows that involve diapycnal mixing and demand good representation of both horizontal and vertical dynamics presents an important extension for unstructured adaptive mesh modelling. The lock-exchange presents an excellent test case with which to assess the potential for the use of adaptive meshes in these types of system. It incorporates simple boundary and initial conditions yet produces a complex transient and turbulent flow that includes diapycnal mixing.

The lock-exchange is a classic laboratory-scale fluid dynamics problem that has been the subject of many theoretical, experimental and numerical studies (e.g. Benjamin, 1968; Cantero et al., 2007; Hallworth et al., 1996; Härtel et al., 2000; Keulegan, 1958; Özgökmen et al., 2009a; Shin et al., 2004; Simpson, 1987) and has been used previously in the assessment of non-hydrostatic ocean models (Berntsen et al., 2006; Fringer et al., 2006). A flat-bottomed tank is separated into two sections by a vertical barrier. One section, the lock, is filled with the source fluid which is of different density to the ambient fluid that fills the second section.

* Corresponding author. Tel.: +1 850 645 0666.
E-mail address: hhiester@coaps.fsu.edu (H.R. Hiester).

1463-5003
http://dx.doi.org/10.1016/j.ocemod.2013.10.003

As the barrier is removed, the denser fluid collapses under the lighter. Two gravity currents form and propagate in opposite directions, one above the other, along the tank. Shear instabilities at the interface between the source and ambient fluid can result in the formation of characteristic Kelvin–Helmholtz billows (or weaker Holmboe waves) which lead to enhanced turbulence and mixing (Holmboe, 1962; Simpson and Britter, 1979; Smyth et al., 1988; Strang and Fernando, 2001; Thomas et al., 2003). This initial stage, when the system is in the gravity current regime, is referred to here as the *propagation stage*. Once the gravity current front(s) reach the end wall, the system enters a different regime, with the fluid 'sloshing' back and forth across the tank, which is referred to here as the *oscillatory stage*. In this stage the system is initially turbulent, and shear instability, internal waves and interaction with the end walls can all enhance mixing between the fluids of different densities. Eventually the system becomes less active and the motion subsides. Mixing of the fluid continues, but at a significantly slower rate than the previous two phases.

The accurate representation of diapycnal mixing in a numerical model is a major challenge as the governing processes occur across multiple scales and the cascade of energy can terminate at scales well below those represented by the mesh resolution. In order to represent these processes, parameterisations are commonly employed (e.g. Özgökmen et al., 2009b; Xu et al., 2006). Whilst a single constant value of the viscosity or diffusivity may be specified in a numerical model (which can be considered the most basic form of parameterisation), the discretisation method can introduce additional (positive or negative) numerical viscosity and/or diffusivity which can result in too little or too much mixing (Griffies et al., 2000; Legg et al., 2008). It is important, therefore, to assess the representation of the mixing in simulations that use adaptive meshes.

The two-dimensional lock-exchange is simulated with the non-hydrostatic model, Fluidity-ICOM (Applied Modelling and Computation Group, 2011). Fluidity-ICOM is a finite-element model that can use both structured and unstructured meshes and has integrated adaptive mesh capabilities for use with unstructured meshes. Simulations are performed here on both fixed and adaptive meshes. The two-dimensional lock-exchange is considered as, by neglecting the three-dimensional dynamics, complexity is removed from the system, allowing the model effects to be studied without the distraction of three-dimensional features and with a smaller computational demand.

Previous ocean modelling studies that use adaptive meshes have, for example, adapted the mesh to the vorticity field, field-based Hessians, solution discontinuities or truncation errors (Bernard et al., 2007; Blayo and Debreu, 1999; Munday et al., 2010; O'Callaghan et al., 2010; Popinet and Rickard, 2007; Remacle et al., 2005). More complex methods exist, in particular goal-based techniques that utilise the model adjoint to form the metric (e.g. Power et al., 2006; Venditti and Darmofal, 2003). These approaches are particularly useful as they provide a robust estimate of the error in a solution diagnostic but they require an adjoint to the forward model. In Fluidity-ICOM, the meshes are adapted to selected solution fields and information about the fields is incorporated into an error metric via the Hessians of these fields. The metric also includes user-defined solution field weights. The specific form of the metrics are such that they provide a bound for the interpolation error of the solution under a selected norm (e.g. Frey and Alauzet, 2005). The mesh is, therefore, adapted in an attempt to control this error. In general, the ability of the adapted mesh to represent the flow will depend on the suitability of the error measure and, hence, the metric formed.

Here, three Hessian-based metrics are considered: the absolute metric, M_∞ (Frey and Alauzet, 2005), the relative metric, M_R (Castro-Díaz et al., 1997), and the p-metric with $p = 2$, M_2 (Chen et al., 2007), which are derived from consideration of the L_∞, relative L_∞ and L_p norms of the interpolation error, respectively. In relation to M_∞, M_R includes a scaling by the local magnitude of the field and M_2 a scaling by the determinant of the local Hessian. A background potential energy diagnostic, which gives a measure of the diapycnal mixing, is used to quantitatively assess the simulations (Winters and D'Asaro, 1996). The Froude number (non-dimensional front speed) is also discussed. This second diagnostic was used extensively in a previous assessment of adaptive mesh Fluidity-ICOM simulations with the M_∞ metric (Hiester et al., 2011). The investigation demonstrated that, in order to obtain good performance of the adaptive mesh simulations compared to the fixed mesh benchmarks, a solution field weight that varied exponentially between the boundaries and domain interior was required. This demanded additional user input, which in this context, it is preferable to minimise.

The two key issues to be addressed here are the performance of the adaptive mesh simulations relative to those on a fixed mesh and the influence, if any, of the metric on the adaptive mesh simulations. The paper is organised as follows: Sections 2 and 3 describe the physical lock-exchange set-up, Fluidity-ICOM and the adaptive mesh techniques employed. Section 4 introduces the diagnostics. Section 5 presents and discusses the results from the numerical simulations, comparing them to one another and previously published results. Finally, Section 6 closes with the key conclusions of this work.

2. Physical set-up

2.1. Governing equations and parameters

The system is governed by the Navier-Stokes equations under the Boussinesq approximation, a linear equation of state and the thermal advection-diffusion equation:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u} = -\nabla p - \frac{\rho}{\rho_0} g\mathbf{k} + \nabla\cdot(\bar{\bar{\nu}}\nabla\mathbf{u}), \tag{1}$$

$$\nabla\cdot\mathbf{u} = 0, \tag{2}$$

$$\rho = \rho_0 + \Delta\rho = \rho_0(1-\alpha(T-T_0)), \tag{3}$$

$$\frac{\partial T}{\partial t} + \mathbf{u}\cdot\nabla T = \nabla\cdot(\bar{\bar{\kappa}}_T\nabla T), \tag{4}$$

with $\mathbf{u} = (u, v, w)^T$: velocity, p: pressure, ρ: density, ρ_0: background density, g: acceleration due to gravity, $\bar{\bar{\nu}}$: kinematic viscosity, T: temperature, T_0: background temperature, $\bar{\bar{\kappa}}_T$: thermal diffusivity, α: thermal expansion coefficient and $\mathbf{k} = (0, 0, 1)^T$. The model considered here is two-dimensional and consequently variation in the cross-stream (y) direction is neglected.

The diffusion term, $\nabla\cdot(\bar{\bar{\kappa}}_T\nabla T)$ in Eq. (4), is neglected in the Fluidity-ICOM simulations. However, the discretised system can still act as if a diffusion term were present, leading to spurious diapycnal mixing. This diffusion can be attributed to the numerics and occurs because, fundamentally, the numerical solution is an approximation to the true solution. It will be referred to here as *numerical diffusion* and it is preferable to minimise its effect. By removing the diffusion term, one level of parameterisation of the system is removed. This allows the response of the fixed and adaptive meshes and a comparison of the inherent numerical diffusion to be made more readily without the need to distinguish between diapycnal mixing due to parameterised diffusion and that inherent in the system. Fixed and adaptive mesh simulations with the diffusion term included were analysed in Hiester (2011) where the best performing adaptive mesh simulations (the same as discussed here) were found to perform as well as the second highest resolution fixed mesh.

Table 1
Physical parameters for the lock-exchange set-up.

Gravitational acceleration	g	$10\ \mathrm{m\,s^{-2}}$
Kinematic viscosity	$\bar{\nu} = \nu\bar{\bar{\mathbf{I}}}$	$10^{-6}\bar{\bar{\mathbf{I}}}\ \mathrm{m^2\,s^{-1}}$
Thermal expansion coefficient	α	$10^{-3}\ ^\circ\mathrm{C}^{-1}$
Background temperature	T_0	$0\ ^\circ\mathrm{C}$
Domain height	$H = 2h$	$10^{-1}\ \mathrm{m}$
Reduced gravity	$g' = g\frac{\Delta\rho}{\rho_0} = -g\alpha\Delta T$	$10^{-2}\ \mathrm{m\,s^{-2}}$
Buoyancy velocity	$u_b = \sqrt{g'H}$	$10^{-3/2}\ \mathrm{m\,s^{-1}}$
Grashof number	$Gr = \left(\frac{h\sqrt{g'h}}{\nu}\right)^2$	1.25×10^6
Buoyancy period	$T_b = 2\pi N_\infty^{-1} = \frac{2\pi}{\sqrt{g'/H}}$	$2\pi 10^{1/2}\ \mathrm{s}$

The values for g, $\bar{\nu}$, α and T_0 are given in Table 1, following the values of Härtel et al. (2000) and Hiester et al. (2011). Note, when (3) is substituted into (1), the buoyancy term $\rho/\rho_0 g\mathbf{k}$ becomes $(1-\alpha(T-T_0))g\mathbf{k}$ and hence buoyancy forcing due to the temperature perturbation is included but no value of ρ_0 needs to be specified.

2.2. The domain, boundary conditions and initial conditions

The domain is a two-dimensional rectangular box, $0 \leqslant x \leqslant L$, $L = 0.8$ m, $0 \leqslant z \leqslant H$, $H = 0.1$ m. Initially, dense cold water, with temperature perturbation $T - T_0 = -0.5\ ^\circ$C, fills one half of the domain, $x < L/2$, and light warm water, with temperature perturbation $T - T_0 = 0.5\ ^\circ$C, fills the other half, $x \geqslant L/2$. At $t = 0$ s, $\mathbf{u} = \mathbf{0}\ \mathrm{m\,s^{-1}}$ everywhere.

At the end walls, $x = 0, L$, a free-slip, no normal flow condition, $u = 0\ \mathrm{m\,s^{-1}}$, is applied. At the bottom boundary, $z = 0$, a no-slip condition, $\mathbf{u} = \mathbf{0}\ \mathrm{m\,s^{-1}}$, is applied. At the top boundary, $z = H$, a free-slip, no normal flow condition, $w = 0\ \mathrm{m\,s^{-1}}$, is applied. Gravity currents at both no-slip and free-slip boundaries can therefore be considered in one simulation which is particularly useful for the comparison of the Froude numbers, Section 5.3.

3. Computational approach

3.1. Discretisation

The velocity and pressure fields are discretised using a continuous Galerkin finite-element formulation (Piggott et al., 2008, 2009). Linear basis functions are used for both fields and the loss of LBB stability is overcome through the use of a pressure filter (Piggott et al., 2009). A node-centred control-volume advection scheme with a Sweby limiter is used for discretisation of the temperature field (LeVeque, 2002; Sweby, 1984; Wilson, 2009). A semi-implicit, Crank–Nicolson scheme is used to advance the equations in time, with a time step of $\Delta t = 0.025$ s and two non-linear Picard iterations. For further details of these methods see the cited references and references therein.

The simulations are run for 500 s. This allows both the propagation stage and the oscillatory stage to be simulated, Section 5.1. By the end of the time period, the system is expected to reach a less active state, with a significantly reduced or near zero mixing rate, Section 5.2 (Özgökmen et al., 2007). Time will be scaled by the buoyancy period $T_b = 2\pi N_\infty^{-1}$, where $N_\infty = \sqrt{g'/H}$ is the buoyancy frequency, Table 1 (Özgökmen et al., 2007); 500 s corresponds to a scaled time of $t/T_b = 25.2$.

3.2. Fixed meshes

The lock-exchange configuration is run using four different fixed meshes. The meshes are generated with Gmsh (Geuzaine and Remacle, 2009). The meshes produced have triangular elements and are structured in both the horizontal and vertical, Fig. 1. The fixed meshes are distinguished by the length of an element edge, $|\mathbf{v}|$, in the horizontal and vertical with $|\mathbf{v}| = 0.002$, 0.0005, 0.00025 and 0.000125 m. The simulations that use each of these meshes are labelled F-coarse, F-mid, F-high1 and F-high2, respectively. The number of vertices in each mesh is given in Table 2.

3.3. Adaptive meshes

The adaptive mesh capabilities in Fluidity-ICOM are for use with unstructured meshes, Fig. 1 (Applied Modelling and Computation Group, 2011). The process used to adapt the mesh can be divided into three main steps: metric formation, which determines how to adapt the mesh; mesh optimisation, the process of altering the mesh based upon the metric; and interpolation of the fields from the pre- to post-adapt mesh. Each step will be introduced separately, Sections 3.3.1, 3.3.2 and 3.3.3, respectively. Here, the particular focus is on metric formation. In general, more information can be found in Applied Modelling and Computation Group (2011) and the cited references. It is also noted that Fluidity-ICOM treats all input meshes in the same manner and uses an unstructured data structure to represent both structured and unstructured meshes, hence the key distinction is between fixed and adaptive meshes.

3.3.1. Metric formation and criteria for adapting the mesh

In Fluidity-ICOM, a metric, represented by a symmetric positive definite tensor, is constructed (George and Borouchaki, 1998). This metric allows information about the system state to be contained in a form that can be used to guide the mesh optimisation step, Section 3.3.2. More specifically, given a metric, M, the aim of the mesh optimisation step is to form a mesh, $\mathcal{M}$, with edges, $\mathbf{v}$, such that

$$||\mathbf{v}||_M = \sqrt{\mathbf{v}^T M \mathbf{v}} = 1, \quad \forall \mathbf{v} \in \mathcal{M}. \tag{5}$$

That is to say, all edges in the mesh have unit length when measured with respect to the metric, M. The metric can be viewed as the continuous analogue of the mesh, describing both the shape and size of the elements (Loseille and Alauzet, 2011a). The choice of metric is, therefore, fundamental to the way in which the mesh adapts and where mesh resolution will be placed.

Three metrics are considered here each of which is based on the Hessian of a solution field(s), H (matrix of second-order derivatives), and a user-defined weight, ϵ, that can vary spatially and/or temporally. The form of each metric is motivated by interpolation error theory and they are chosen such that, for the exact Hessian, the metrics provide a bound for the interpolation error of the solution field under a selected norm.

The first metric, M_∞, is given by

$$M_\infty(\mathbf{x}) = \frac{|H(\mathbf{x})|}{\epsilon(\mathbf{x})}, \tag{6}$$

(e.g. Frey and Alauzet, 2005; Pain et al., 2001), where $|H(\mathbf{x})|$ is a modified Hessian:

$$|H(\mathbf{x})| = Q(\mathbf{x})^T |\Lambda(\mathbf{x})| Q(\mathbf{x}), \quad |\Lambda(\mathbf{x})|_{ij} = \begin{cases} |\lambda_i(\mathbf{x})| & i = j \\ 0 & i \neq j \end{cases} \tag{7}$$

with $\{\lambda_i\}$ the eigenvalues of the Hessian and Q the corresponding matrix of normalised eigenvectors. Information about both the magnitude and direction of the curvature of the field is therefore included, via $|\Lambda|$ and Q, respectively, and facilitates the formation of anisotropic elements. If this metric is used and the adaptivity criteria, Eq. (5), is satisfied then, for an exact Hessian, a bound for the

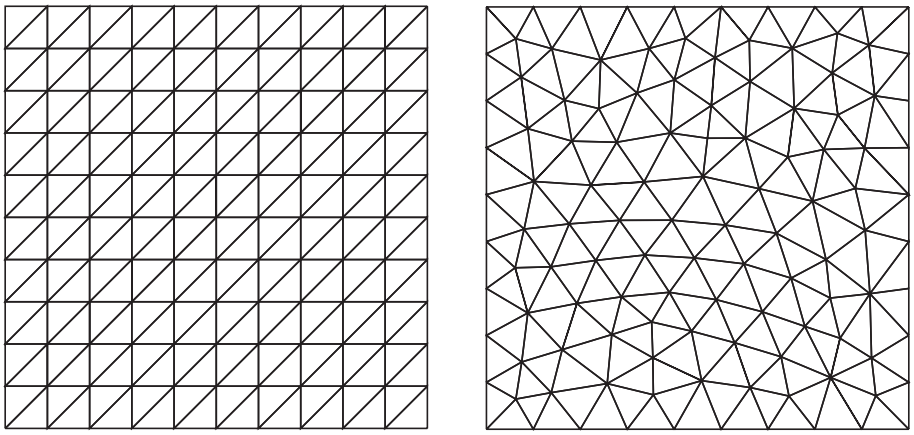

(a) Structured mesh (b) Unstructured mesh

Fig. 1. Example of a structured triangular mesh and an unstructured triangular mesh in a unit box.

Table 2
Number of vertices in the fixed meshes with the horizontal and vertical element edge length $|\mathbf{v}|$.

$\|\mathbf{v}\|$	Vertices	Reference
0.002	2.0451×10^4	F-coarse
0.0005	3.21801×10^5	F-mid
0.00025	1.283601×10^6	F-high1
0.000125	5.127201×10^6	F-high2

interpolation error is provided for a mesh element, Ω_e, under the L_∞ norm (Frey and Alauzet, 2005). In practice, areas with a high curvature of a field (large second-order derivatives) and therefore larger eigenvalues, will demand refinement of the mesh, Eqs. (5)–(7). Reducing the solution field weight will also promote more mesh refinement. Conversely, lower curvature and/or a larger solution field weight will demand coarsening of the mesh.

The second metric, M_R, has the form

$$M_R(\mathbf{x}) = \frac{1}{\epsilon(\mathbf{x})}\frac{|H(\mathbf{x})|}{\max(|f(\mathbf{x})|, f_{\min})} = \frac{M_\infty}{\max(|f(\mathbf{x})|, f_{\min})}, \tag{8}$$

(Castro-Díaz et al., 1997) is obtained through similar consideration of a relative interpolation error, with $f_{\min}$ specified by the user to avoid division by zero. The refinement or coarsening of the mesh is still guided by the curvature of the field. However, a scaling by the local magnitude of the field is now included in the metric.

The final metric is obtained by consideration of the interpolation error in the L_p norm, $p \in [1, \infty)$. The general metric, denoted M_p, has the form

$$M_p(\mathbf{x}) = \frac{1}{\epsilon(\mathbf{x})}(\det(|H(\mathbf{x})|))^{-\frac{1}{2p+n}}|H(\mathbf{x})| = (\det(|H(\mathbf{x})|))^{-\frac{1}{2p+n}}M_\infty, \tag{9}$$

(Chen et al., 2007; Loseille and Alauzet, 2011b), where n is the spatial dimension of the problem. Since $\det|H| = \prod_i |\lambda_i|$, a scaling by a measure of the magnitude of the curvature of the field is included in the metric. The extent to which $\det|H|$ influences the metric is determined by the choice of p. As p is reduced, smaller scales are given more weight in the metric and as a result are better represented (Loseille and Alauzet, 2011b). In the limit $p \to \infty$, M_∞ is recovered. The work of Loseille and Alauzet (2011b) shows that the influence of smaller scales in the metric rapidly decreases as p increases and their good results for $p = 2$ motivates the use of this value here. Hence, the third and final metric is given by

$$M_2(\mathbf{x}) = \frac{1}{\epsilon(\mathbf{x})}(\det(|H(\mathbf{x})|))^{-\frac{1}{6}}|H(\mathbf{x})| = (\det(|H(\mathbf{x})|))^{-\frac{1}{6}}M_\infty. \tag{10}$$

In Fluidity-ICOM, the user chooses which solution fields a metric will be formed for and, therefore, which fields the mesh will adapt to. If the user chooses to adapt to multiple solution fields, a metric, M_f, is formed for each chosen solution field, f. The final metric, M, is then obtained from a superposition of the metrics for individual fields $M = \bigcap_f M_f$ (Castro-Díaz et al., 1997). The user must also specify minimum and maximum edge lengths and this information is included through a restriction on the eigenvalues of $|H|$ (e.g. Pain et al., 2001). In addition, the user can provide an upper and/or lower bound on the number of mesh vertices. If the adaptive algorithm is configured appropriately, this bound should not be reached.

3.3.2. Mesh optimisation

Given a metric, the aim of the mesh optimisation step is to satisfy the criteria, Eq. (5) and thereby optimise the mesh for the current system state. The mesh is modified through a series of local topological and geometrical operations which, in two dimensions in Fluidity-ICOM, are performed using the algorithms of Vasilevskii and Lipnikov (1999). The operations include edge-collapsing, edge-splitting, edge-swapping and vertex-movement. More details and diagrams can be found in Pain et al. (2001), Piggott et al. (2009) and Vasilevskii and Lipnikov (1999).

3.3.3. Interpolation from the pre- to post-adapt mesh

Once the mesh optimisation stage has been performed, solution fields have to be interpolated between the pre- and post-adapt meshes. The interpolation methods available in Fluidity-ICOM fall into two categories. The first is referred to as *consistent interpolation* (Applied Modelling and Computation Group, 2011). With this method, for each vertex in the post-adapt mesh, the element in the pre-adapt mesh in which it would be contained is identified. The solution field is then evaluated at the vertex in the post-adapt mesh using the finite-element basis functions of the containing element in the pre-adapt mesh. Consistent interpolation is bounded (for linear basis functions) but is non-conservative and is only well-defined for continuous function spaces. The second method uses the intersection of the pre- and post-adapt meshes to form a supermesh. The fields are then interpolated via the supermesh using Galerkin projection (Farrell et al., 2009; Farrell and Maddison, 2011). By construction, it is conservative, but is not necessarily bounded. Any overshoots or undershoots in the solution field that occur are corrected, essentially by diffusing the deviation from boundedness. The diffusion introduced in this approach is minimal when compared with consistent interpolation

(Farrell et al., 2009). This bounded, minimally diffusive, conservative method will be referred to as *bounded Galerkin projection*.

Different methods for interpolation from the pre- to post-adapt mesh have a less significant impact on the adaptive mesh simulations than that of the metric (Hiester et al., 2011; Hiester, 2011). The majority of simulations presented here use consistent interpolation for both the velocity and temperature fields as, for this numerical configuration, it provides a faster method than bounded Galerkin projection (Hiester et al., 2011). The final adaptive mesh simulations considered for the comparison with Özgökmen et al. (2007), Section 5.5, use consistent interpolation for the velocity field and bounded Galerkin projection for the temperature field as improved results for the initial set-up have been obtained with this combination (with a reduction in the mixing of approximately 7% at later times, Hiester, 2011).

3.3.4. Adaptive mesh parameters and solution field weights

The meshes are adapted every ten time steps. This choice of adapt frequency provides a balance between being sufficiently frequent so as to prevent features propagating out of the regions of higher mesh resolution and hence deteriorating the solution but not so frequent as to notably increase the computational overhead (cf. Hiester et al., 2011, and Section 3.4). The minimum and maximum edge lengths are set to 0.0001 m and 0.5 m, respectively and the maximum number of vertices is set to 2×10^5, which is comparable to the medium resolution fixed mesh, Table 2.

The meshes are adapted to the horizontal velocity field, vertical velocity field and the temperature field with solution field weights denoted ϵ_u, ϵ_v and ϵ_T, respectively. For M_∞ two sets of solution field weights are considered, Table 3, following the values of Hiester et al. (2011). The first set are spatially constant. The second set has spatially constant values of ϵ_v and ϵ_T and a value of ϵ_u that varies exponentially in the vertical such that the value at the top and bottom boundaries is two orders of magnitude smaller than that at the centre of the domain, Table 3. This spatial variation is necessary to obtain Froude numbers in adaptive mesh Fluidity-ICOM lock-exchange simulations with M_∞ that are in good agreement with high resolution fixed mesh and other published values (Hiester et al., 2011). The simulations that use these adaptive mesh configurations are denoted M_∞-const for constant solution field weights and M_∞-var for spatially varying solution field weights.

For M_R, simulations are run with weights of 0.1, 0.05 and 0.01 for temperature, horizontal velocity and vertical velocity. These correspond to a 10%, 5% and 1% bound for the relative interpolation error. In order to avoid division by zero $f_{\min} = 1 \times 10^{-5}$, Eq. (8). This value determines the minimum value of the fields that will scale the metric and is selected to allow a wide range for the velocity and temperature fields. These combinations are summarised in Table 4 and simulations that use these adaptive mesh configurations are denoted M_R-loose, M_R-mid and M_R-tight.

For M_2, three sets of solution field weights are tested. The first set, M_2-loose, reflects the values used in the simulations with M_∞, with the ratio of ϵ_u to ϵ_T remaining similar. Qualitative observation of simulation M_2-loose shows a coarse mesh and a diffusive solution. This motivates the formation of a second set of solution field weights, M_2-mid, with a reduction in size of ϵ_u, ϵ_v and ϵ_T. Finally, analysis of the background potential energy and Froude number diagnostics for the first two sets motivates the testing of a third set, M_2-tight, with further reductions in the solution field weights. In this third set, the vertical velocity field weight is reduced in order to determine if an increase in resolution can be obtained at the free-slip boundary and, hence, an improvement in the free-slip Froude number (cf. Hiester et al., 2011). The temperature weight is also halved for $t/T_b > 1.76$ to determine whether this leads to a further reduction in the diapycnal mixing at later times. These combinations are summarised in Table 5.

Table 3
Solution field weights used in the formation of M_∞, Eq. (6). For the case M_∞-var, ϵ_u has an exponential variation in the vertical with $\epsilon_u = \epsilon_0 + (\epsilon_{H/2} - \epsilon_0)(e^{100z} - 1)/(e^{100H/2} - 1)$ for $z \leqslant H/2$ and $\epsilon_0 + (\epsilon_{H/2} - \epsilon_0)(e^{100(H-z)} - 1)/(e^{100H/2} - 1)$ for $H/2 < z$ where $\epsilon_{H/2} = 0.001$ and $\epsilon_0 = 0.0001$ (Hiester et al., 2011). The mesh is adapted every ten time steps.

ϵ_u	ϵ_v	ϵ_T	Reference
0.001	0.001	0.025	M_∞-const
0.0001–0.001	0.001	0.025	M_∞-var

Table 4
Solution field weights used in the formation of M_R, Eq. (8). The mesh is adapted every ten time steps.

ϵ_u, ϵ_v and ϵ_T	Reference
0.1	M_R-loose
0.05	M_R-mid
0.01	M_R-tight

Table 5
Solution field weights used in the formation of M_2, Eq. (10). The mesh is adapted every ten time steps. The M_2-tight case has two values for ϵ_T, the former is used at earlier times and the latter for $t/T_b > 1.76$.

ϵ_u	ϵ_v	ϵ_T	Reference
0.0001	0.0001	0.002	M_2-loose
0.00005	0.00005	0.0005	M_2-mid
0.00005	0.00001	0.0005/0.000025	M_2-tight

3.4. Comparison of computational cost

In general, the number of vertices in the mesh will be taken as a gauge of the computational demand associated with a simulation. It is considered an appropriate measure when comparing the fixed and adaptive mesh Fluidity-ICOM simulations. The number of vertices is a useful measure of computational demand as it is machine independent and also gives an indication of the size of the problem. This does not account for the model scaling, either with the number of vertices in serial or the number of processors (and the number of vertices) in parallel. The run time of the simulation presents a measure of computational demand which incorporates these effects and offers a complementary measure to the number of vertices but is machine dependent and is not pursued here.[1]

The cost of the mesh adapt must also be considered. Fluidity-ICOM adaptive mesh simulations of the lock-exchange that use M_∞ with consistent interpolation require less time than one model time step for metric formation, mesh optimisation and interpolation (Hiester et al., 2011). If bounded Galerkin projection is used the time required was found to increase to approximately two time steps. Simulation M_2-mid was also profiled as a part of this investigation and the mesh adapt required a similar proportion of time to the simulations that use M_∞ (Hiester, 2011). In parallel, the overhead of adaptivity is relatively small with the overall cost of the adaptive step being dominated by the serial algorithm (Gorman et al., 2009).

4. Diagnostics

The background potential energy provides a measure of diapycnal mixing and is the main diagnostic used for analysis here,

[1] Simulations were run at both the Imperial College High Performance Computing centre and the UK National Supercomputing Service, HECToR. The number of processors used varied from 1 to 384.

Section 4.1. The Froude number is also calculated providing an additional diagnostic comparison, Section 4.2.

4.1. Background potential energy

The background potential energy is the potential energy of the minimum energy state (or reference state) that can be obtained by adiabatic redistribution of the system (Winters et al., 1995; Winters and D'Asaro, 1996). Most crucially, for a closed system, changes to the reference state caused by diapycnal mixing correspond to increases in the background potential energy (Winters et al., 1995). Denoting the vertical coordinate in the reference state z_*, the background potential energy, E_b, is given by

$$E_b = \int_\Omega \rho g z_* \, dV, \tag{11}$$

where Ω is the domain. z_* is calculated using the method of Tseng and Ferziger (2001), where a probability density function is constructed for the density (or here temperature) field and then integrated to give z_* (cf. Hiester, 2011).

The background potential energy is decomposed further to account for changes in E_b that may occur due to non-conservation of the fields through the use of a non-conservative advection scheme and consistent interpolation. Following Ilıcak et al. (2012), ρ and z_* are partitioned into a spatial mean and a perturbation: $\rho = \overline{\rho} + \rho'$ and $z_* = \overline{z_*} + z'_*$, where

$$\overline{\rho} = \frac{1}{V}\int_\Omega \rho \, dV \quad \text{and} \quad \overline{z_*} = \frac{1}{V}\int_\Omega z_* \, dV. \tag{12}$$

E_b then becomes

$$E_b = \underbrace{g\overline{\rho}\,\overline{z_*}\int_\Omega dV}_{\overline{E_b}} + \underbrace{g\int_\Omega \rho' z'_* \, dV}_{E'_b}, \tag{13}$$

where $\overline{E_b}$ changes due to changes in mass and E'_b changes due to diapycnal mixing (Ilıcak et al., 2012).

The values will be presented as a change in E'_b, normalised by the initial value of E_b:

$$\frac{\Delta E'_b(s)}{E^0_b} = \frac{E'_b(s) - E'_b(s=0)}{E_b(s=0)}, \tag{14}$$

where $s = t/T_b$ or, for a closer analysis of the propagation stages, $s = X/H$ with X the position of the no-slip front. It is noted that whilst E_b depends on density and hence ρ_0, as the values are normalised, once again no value of ρ_0 is required (cf. Section 2.1). The typical behaviour of the background potential energy is presented in Section 5.2.

4.2. Froude number

The Froude number, $Fr = U/u_b$, is the ratio of the front speed, U, to the buoyancy velocity, u_b, Table 1. After an initial acceleration, the gravity current fronts travel at a constant speed until the end walls exert an influence or viscous forces begin to dominate (Cantero et al., 2007; Härtel et al., 1999; Huppert and Simpson, 1980). As the span-averaged three-dimensional front speeds have comparable values to two-dimensional values, the Froude number may be compared between two-dimensional simulations and both three-dimensional simulations and experiments (e.g. Cantero et al., 2007; Härtel et al., 2000). High resolution fixed mesh Fluidity-ICOM simulations compare well with published values (Hiester et al., 2011) and are used here as the benchmark for comparison. The speeds with which the no-slip and free-slip fronts propagate along the domain are calculated from the model output and are used to give the corresponding no-slip and free-slip Froude numbers (Hiester et al., 2011).

5. Results

5.1. Qualitative results

The simulations exhibit dynamics that are typical of the lock-exchange, Fig. 2. Two gravity currents form and propagate in opposite directions along the tank with Kelvin–Helmholtz billows developing at the interface. Once the gravity currents hit the end walls they are reflected and the fluid begins to 'slosh' back and forth across the tank. In this second oscillatory regime, internal waves and interaction with the end walls further increase the complexity of the flow. Subsequently, the system becomes increasingly less active and the motion subsides.

The adaptive meshes coarsen or refine according to the evolution of the flow. During the propagation stages, the meshes refine along the boundaries, at the temperature interface and in and around the billows, Figs. 3–5. The meshes generated via the different metrics refine or coarsen as would be expected. Simulations that use M_∞ refine in regions with the greatest curvature and coarsen elsewhere. Simulations that use M_R also include refinement in regions where the magnitude of the fields is small. Finally, simulations that use M_2 refine in the regions with the greatest curvature, but also capture curvatures and hence features over a wider range of scales. A user can, therefore, consider *a priori* which form of the metric would be most suitable for the simulated system and the dynamics to be represented.

The most obvious contrast between the meshes is for those produced with M_R compared to those produced with M_∞ and M_2. With M_R there are several regions where the mesh appears to be unnecessarily refined leading to an increase in the number of vertices, Fig. 6. These regions correspond to areas of the domain where the velocity fields are near zero, Fig. 4. An increase of the parameter $f_{\min}$, which determines the minimum allowed value of the field by which the metric can be scaled, Eq. (8), would lead to a reduction in resolution in the regions where the velocity field is near zero and, for this case, where the mesh was unnecessarily refined. The temperature perturbation is zero at the interface and the increase in resolution due to the smaller value of the field in this region is more desirable.

The meshes produced with M_2 place more resolution into and around the Kelvin–Helmholtz billows, than those produced with M_∞, Figs. 3 and 5. The latter also results in a coarser mesh in regions of the domain further from the interface, which is again reflected in the number of vertices in the mesh, Fig. 6. At later times, as the interface becomes more diffuse and the system less active, the simulations that use M_2 retain more structure in the mesh. There is also refinement to a mid-resolution (i.e. not very coarse or very fine) over a greater area of the domain than for the simulations with M_∞.

The adaptive meshes that use M_R here have at least three to four times more vertices in the mesh than the simulations with M_∞ and M_2 and reach the maximum number of vertices specified, Fig. 6. As a result, the simulations that use M_R were terminated early due to the increased run times. All the adaptive mesh simulations use fewer vertices than the middle resolution fixed mesh (F-mid, Table 2). Those simulations that use M_∞ and M_2 have, in general, a comparable number of vertices to the coarsest fixed mesh (F-coarse, Table 2), which is two orders of magnitude fewer vertices than the highest resolution fixed meshes considered, F-high1 and F-high2, Table 2. The relative performance of the simulations is now considered with respect to the quantitative diagnostics.

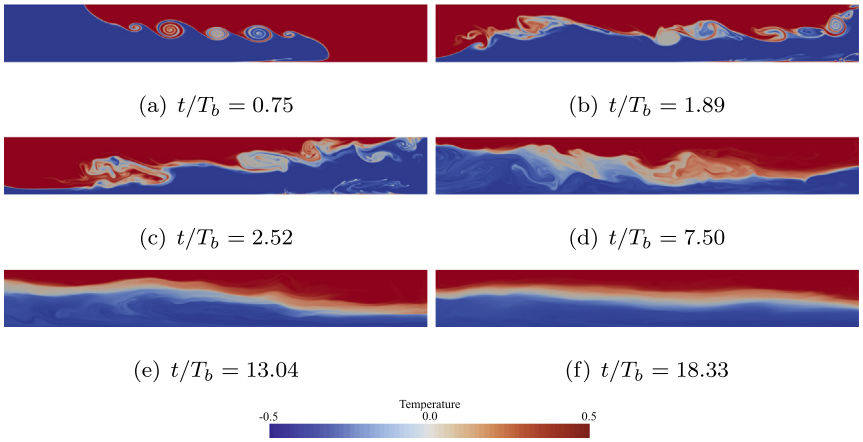

Fig. 2. Temperature distribution for the lock-exchange over scaled time, t/T_b, for the fixed mesh simulation F-high1, Table 2. Both the early propagation stages (a), with characteristic Kelvin–Helmholtz billows that enhance the diapycnal mixing, and the later oscillatory stages (b–f), during which the motion and diapycnal mixing diminish, are presented. The other simulations, with both fixed and adaptive meshes, present similar dynamics (not shown).

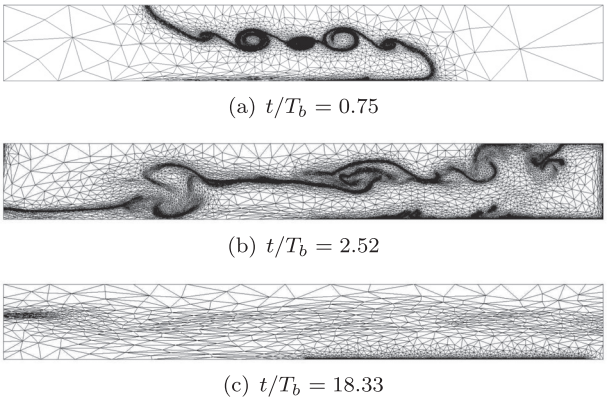

Fig. 3. Mesh at three scaled times, t/T_b, for adaptive mesh simulation M_∞-const, Table 3. Mesh resolution is increased around features such as the Kelvin–Helmholtz billows and the temperature interface. At later times, the mesh coarsens as the temperature interface becomes more diffuse and the system less active. At $t/T_b = 18.33$ there is a thin region of increased resolution along the bottom boundary where a thin layer of dense fluid over run by the no-slip head begins to rise as thin filaments into the lighter fluid above. M_∞-var, Table 3, has a similar distribution of mesh resolution in the interior of the domain and increased resolution at the boundaries (not shown).

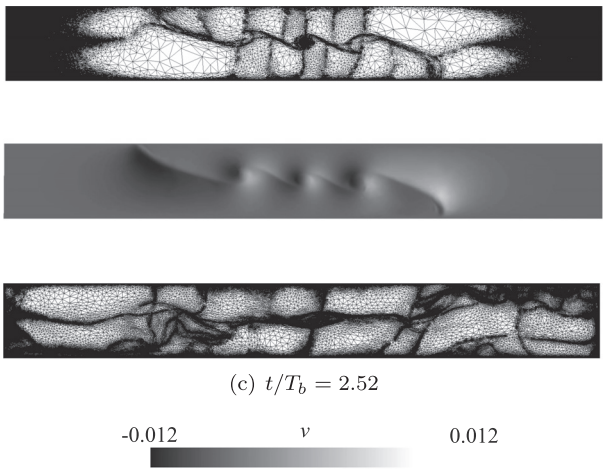

Fig. 4. Mesh (a,c) and vertical velocity, v (b), at scaled times, t/T_b, for adaptive mesh simulation M_R-mid, Table 4. Key features of the system are resolved by the mesh. There is also increased resolution in regions where the magnitude of the fields is small, even if the curvature is not large. Regions where the mesh is highly refined appear black in (a,c) due to the high density of mesh vertices. Other simulations with M_R have a similar distribution of mesh resolution (not shown).

5.2. Background potential energy

5.2.1. Fixed mesh simulations

The fixed mesh simulation F-high1 is used to demonstrate the behaviour of the potential energy, kinetic energy and background potential energy perturbation, Fig. 7. As the two gravity currents form and propagate across the domain the potential energy decreases through exchange with the kinetic energy of the system and loss to diapycnal mixing. The background potential energy perturbation, E_b', increases as diapycnal mixing takes place. As the fraction of the domain occupied by the gravity currents increases and there is more diapycnal mixing along the lengthening interface, E_b' increases more rapidly. The free-slip and no-slip fronts reach the end wall at $t/T_b \approx 1.25$ and $t/T_b \approx 1.75$, respectively. As the currents run up against the end walls, the potential energy increases, the kinetic energy decreases and the mixing rate (rate at which E_b' changes) continues to increase.

During the first oscillation, $t/T_b \approx 3$–7, the diapycnal mixing is still vigorous and is further enhanced by internal waves and interaction with the end walls. During the second oscillation,

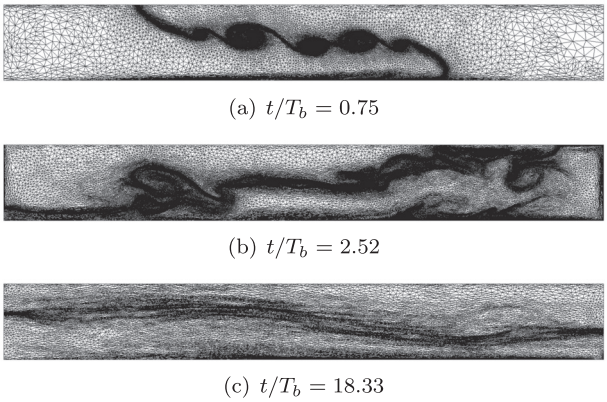

Fig. 5. Mesh at three scaled times, t/T_b, for adaptive mesh simulation M_2-mid, Table 5. Key features of the system are resolved by the mesh, with increased resolution in regions with both stronger and weaker curvature, for example ahead of the no-slip front or around the temperature interface at later times. Other simulations with M_2 have a similar distribution of mesh resolution (not shown).

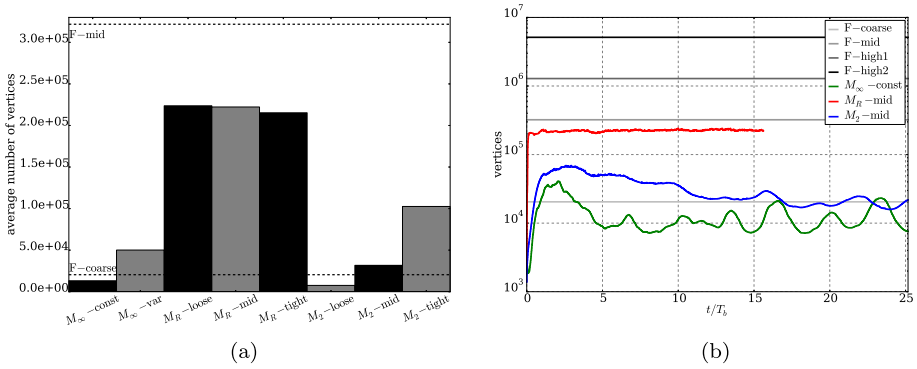

Fig. 6. Comparison of the number of vertices in the meshes over time for fixed and adaptive mesh simulations, Tables 2–5. In (a) the number of vertices are over the total time period, $0 < t/T_b < 25.2$. However, for simulations that use M_R, truncated ranges are used, with a maximum value $t/T_b = 15.1$, as the simulations are terminated before $t/T_b = 25.2$ due to increased run times. Dashed lines representing the number of mesh vertices in the two coarsest fixed meshes are also included. In (b) the variation of selected adaptive mesh simulations are presented. For a given metric, the values for the other adaptive mesh simulations follow a similar pattern but with different magnitudes. All of the adaptive mesh runs use fewer vertices than the second highest resolution fixed mesh, F-mid, and most that use M_∞ or M_2 are of a comparable magnitude to the coarsest resolution fixed mesh, F-coarse.

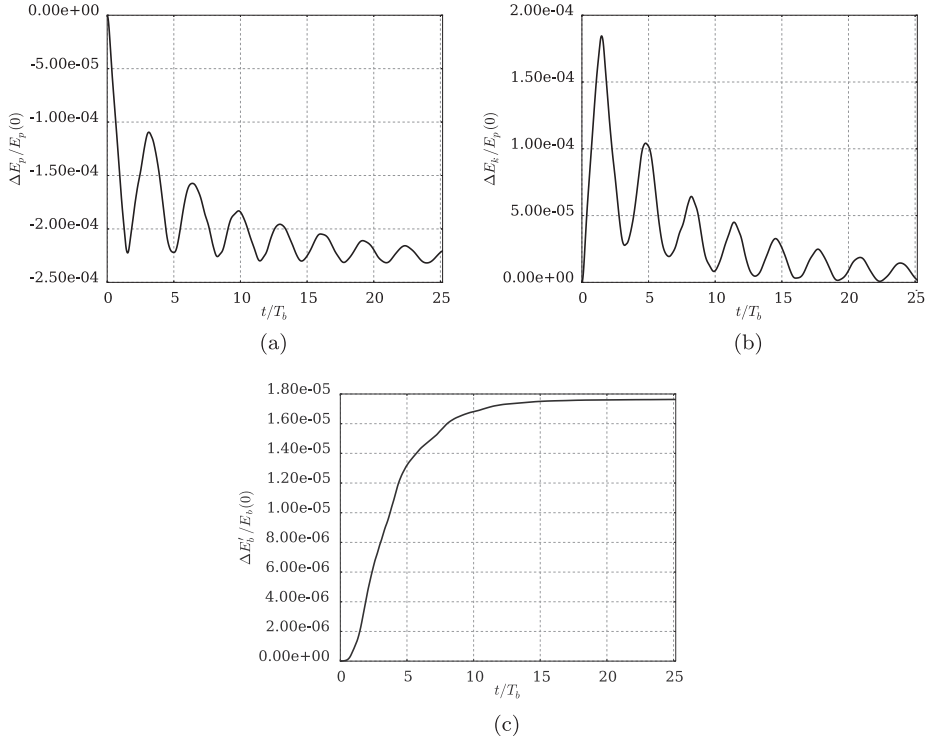

Fig. 7. Comparison of the scaled potential energy, $E_p/E_p(0)$, $E_p = \int_\Omega \rho g z dV$, the scaled kinetic energy, $E_k/E_p(0)$, and the scaled background potential energy perturbation, $E_b'/E_b(0)$, over scaled time, t/T_b, for fixed mesh simulation F-high1, Table 2. The values of E_p and E_k oscillate in accordance with the system. The value of E_b' increases corresponding to diapycnal mixing in the system. At later times, as the mixing subsides, E_b' tends to a near constant value.

$t/T_b \approx 7$—10, diapycnal mixing still occurs but at a reduced rate. Subsequently, the system becomes increasingly less active and the diapycnal mixing subsides. While the potential energy and kinetic energy oscillate in accordance with the system, the background potential energy perturbation constantly increases (or tends to a near constant value) as diapycnal mixing continually occurs within the system (or tends to zero), demonstrating the diagnostic utility of this quantity. Due to the reduction of the mixing rate to zero (or near zero) the simulated time period (up to $t/T_b = 25.2$) is sufficient to obtain a fair representation of the enhanced mixing behaviour of the system. Diapycnal mixing will continue beyond this time at a significantly reduced rate.

As the diffusion term is neglected here, the diapycnal mixing is attributable to numerical diffusion. As the fixed mesh resolution increases, the amount of diapycnal mixing decreases indicating that the higher resolution meshes have a lower numerical diffusion, Fig. 8. The fixed mesh simulations provide a useful set of benchmarks for comparison of the adaptive mesh simulations. As all other numerical components of the model remain the same for the fixed and adaptive mesh simulations, the impact of the adaptive mesh can also be focused on more readily.

5.2.2. Adaptive mesh simulations

During the propagation stages, the adaptive mesh simulations reproduce the general mixing trends of the fixed meshes, with an increasing mixing rate as the gravity currents propagate further across the domain, Fig. 8. With the exception of those that use M_R, the adaptive mesh simulations can present comparable mixing

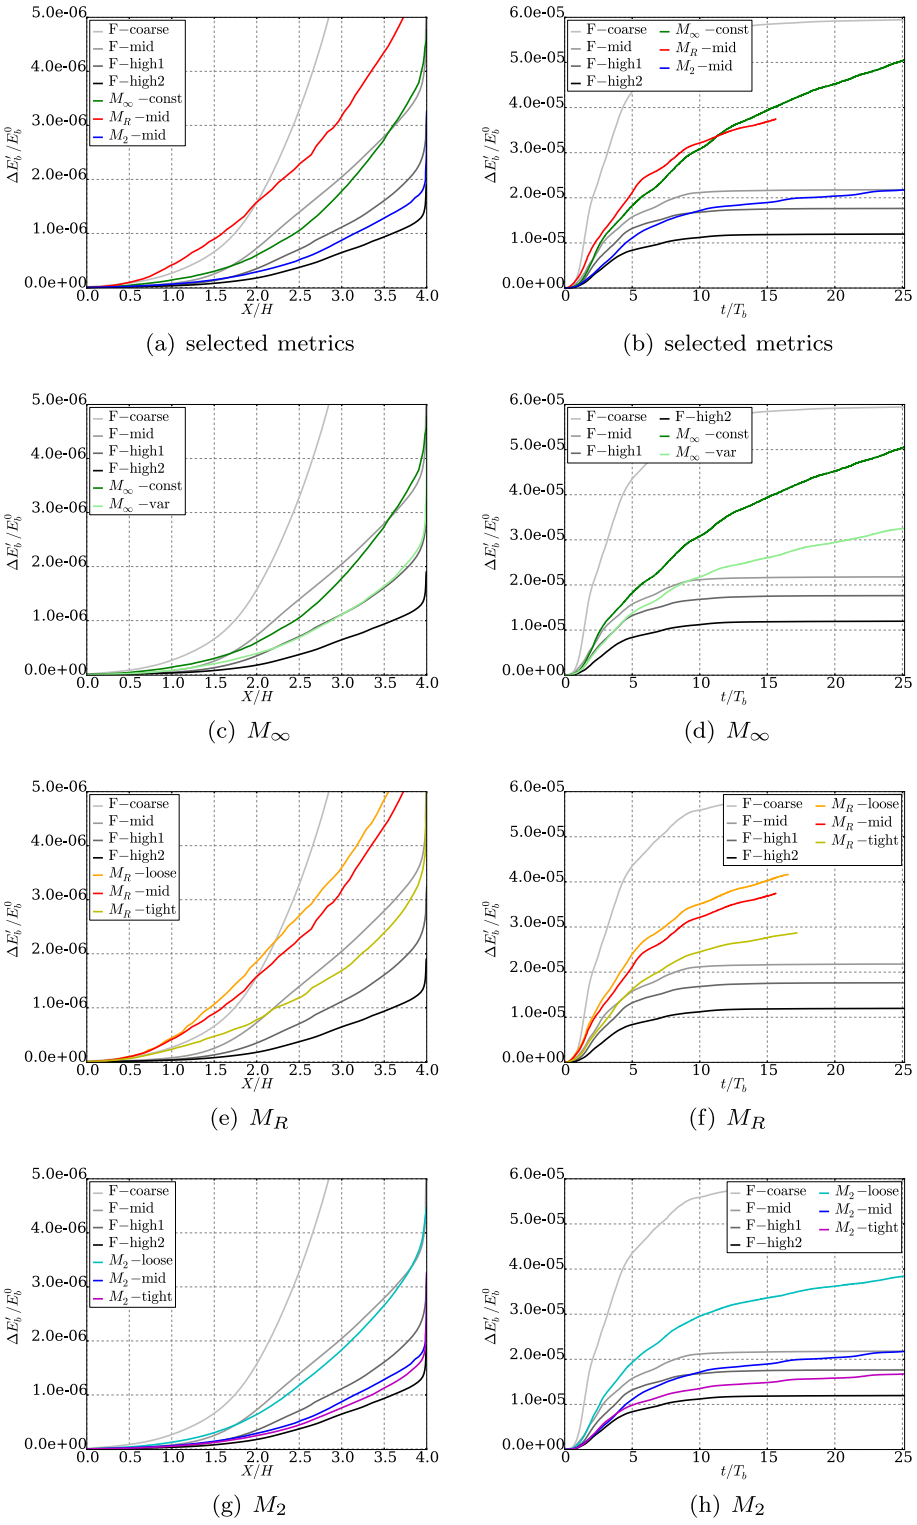

Fig. 8. Comparison of the normalised change in the background potential energy perturbation, E_b', Eq. (13), for the fixed and adaptive mesh simulations, Tables 2–5. On the left the values are plotted with the scaled position of the no-slip front, X/H, for the propagation stages. On the right the values are plotted with scaled time, t/T_b, over all time. The simulations that use M_R are terminated early due to increased run times. The simulations that use M_2 have the best performance, with trends most similar to the higher resolution fixed meshes and the least mixing compared to other adaptive meshes with comparable or larger numbers of mesh vertices, Fig. 6.

to the fixed mesh simulations that have at least one order of magnitude more vertices in the mesh. During the oscillatory stages, diapycnal mixing occurs in the simulations that use M_∞ and M_R over all time resulting in a constantly increasing value of E_b', whereas, for all but the coarsest fixed mesh simulations, this quantity tended to a near constant value.

In general, the adaptive mesh simulations that use M_2 perform the best, Fig. 8. These simulations can produce trends that are the most similar to that of the fixed meshes, with a decrease in the mixing rate at later times, and a comparable magnitude of E_b' to the fixed meshes that have at least one order of magnitude more vertices. The improved performance of simulations that use M_2

can be attributed to better representation of a range of scales than that obtained with M_∞ and M_R. This is particularly evident at later times, when the system is less active and the interface more diffuse, leading to fields with weaker curvatures, Figs. 3 and 5. These points are now considered in more detail, beginning with discussion of the simulations that use M_∞, followed by those that use M_R and finally those that use M_2.

During the propagation stages, the simulations that use M_∞, M_∞-const and M_∞-var, have comparable levels of diapycnal mixing to fixed mesh simulations F-mid and F-high1, respectively, Fig. 8. During the early oscillatory stages ($2.5 < t/T_b < 5$) the values of E_b' begin to increase above the fixed mesh values and the number of vertices begins to decrease, Fig. 6. Diapycnal mixing then continues at all times, unlike the higher resolution fixed mesh cases. The system dynamics and mesh behaviour are strongly coupled. During the oscillatory stages, the system becomes less active, with the kinetic energy decreasing overall (not shown but similar to the trend for F-high1 presented in Fig. 7) and the temperature field becomes more diffuse. This results in weaker curvature of the fields. As the metrics returned with M_∞ will aim to focus resolution into regions with high curvature, a coarser mesh is produced during these stages, resulting in a higher numerical diffusion. This, in turn, further increases the diapycnal mixing and damping of the dynamics, again resulting in weaker curvatures, a coarser mesh and increased numerical diffusion.

For the simulations that use M_R, as the solution field weights decrease, the mixing decreases. The values of E_b' obtained for M_R-tight are of the same magnitude as simulation M_∞-const for the propagation stages and approximately 10–20% larger than M_∞-var in the oscillatory stages. The number of vertices used in these simulations increases significantly (over 400% on average) compared to the simulations that use M_∞ and reaches the maximum number of mesh vertices specified for the adaptive mesh (2×10^5), Fig. 6. Snapshots of the mesh suggest that the resolution is not necessarily used effectively in the simulations with M_R, leading to worse performance than the simulations that use M_∞. The additional parameter $f_{\min}$ will influence the extent of the mesh refinement and, if increased, may be expected to result in meshes with fewer vertices, potentially more appropriately placed. Furthermore, increasing the maximum number of mesh vertices may lead to more refinement in critical regions and reduce the mixing. However, the increased diapycnal mixing with increased mesh resolution, when compared to simulations with M_∞, indicates that this metric does not perform well for the lock-exchange and further investigation of M_R is not pursued here.

The simulations that use M_2 perform the best of the adaptive mesh simulations and, as for those that use M_R and M_∞, have a decrease in diapycnal mixing as the solution field weights decrease. Simulation M_2-loose uses a comparable number of vertices to simulation M_∞-const and produces comparable or smaller values of $\Delta E_b'$ than simulation M_∞-const, Fig. 6.

During the propagation stages and the earlier oscillatory stages, $t/T_b < 10$, the values of $\Delta E_b'$ for simulation M_2-mid fall between those of the two highest resolution fixed mesh simulations, F-high1 and F-high2, Fig. 8. Subsequently, in simulation M_2-mid, the diapycnal mixing continues at a reduced rate with a trend that is more similar to the fixed mesh runs than the adaptive mesh simulations with M_∞ or M_R, whilst using just over half the number of vertices used in simulation M_∞-var and twice that of simulation M_∞-const. The final value of $\Delta E_b'$ for M_2-mid is the same as simulation F-mid and approximately two-thirds the value for M_∞-var, overall presenting a comparable level of diapycnal mixing to a fixed mesh with at least one order of magnitude more vertices and a fixed mesh with almost two orders of magnitude more vertices at early times $t/T_b < 10$, when the system is more active and the dynamics more complex.

For M_2-tight, the diapycnal mixing is reduced by approximately 10% compared to simulation M_2-mid during the propagation stages and uses approximately 250% the number of vertices used by simulation M_2-mid. By the end of the time period, simulation M_2-tight has values of $\Delta E_b'$ within 5% of the high resolution fixed mesh simulation F-high1 whilst using one order of magnitude fewer vertices. Simulation M_2-tight therefore offers an improvement in E_b' over M_2-mid but has an increased computational cost.

The diapycnal mixing behaviour and distribution of vertices indicate that it is the ability of M_2 to increase resolution even when the curvature is weaker that allows the improved representation of the field and the reduction in the diapycnal mixing. Snapshots of the mesh for simulation M_∞-var and M_2-mid show higher resolution of the billows, particularly at their centre, and also extending away from the billow edges, Figs. 3 and 5. As the fluid in the billow begins to mix and the fields homogenise, the curvature of the fields is reduced (particularly in the temperature field). The smaller-scale variations in the fields are not captured adequately in the simulation with M_∞ but are given more weight in M_2 and, hence, are better represented. Furthermore, during the oscillatory stages, the simulations that use M_∞ have much coarser resolution in the majority of the domain than simulation M_2-mid, Figs. 3 and 5. It is not surprising, therefore, that simulation M_2-mid behaves more like the higher resolution fixed meshes and demonstrates that M_2 provides a better guide of where the mesh resolution is needed.

5.3. Froude number

The values of the no-slip and free-slip Froude number tend to near constant values as the fixed mesh resolution increases, Fig. 9. The no-slip values for the two higher resolution simulations, F-high1 and F-high2, are affected by the shedding of a billow at the

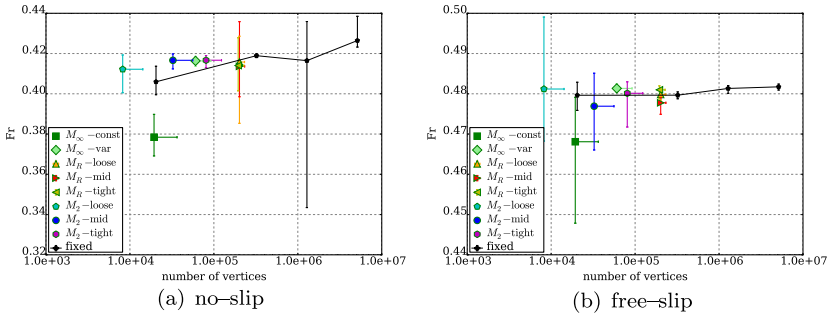

(a) no–slip (b) free–slip

Fig. 9. Comparison of the Froude number, $Fr = U/u_b$, Table 1, at the no-slip and free-slip boundaries for the fixed and adaptive mesh simulations, Tables 2–5. The markers represent the average Froude number and the vertical error bars represent the minimum and maximum values. For the adaptive mesh simulations, the number of vertices in the mesh is an average over the period between the initial time and the time when the no-slip front reaches the end wall, with the horizontal error bar extending to the maximum number of vertices in the time period. With the exception of M_∞-const, the adaptive mesh simulations are able to produce Froude numbers that are comparable to the higher resolution fixed mesh simulations.

nose of the gravity current; this results in an acceleration and deceleration of the front and is captured by the large error bars for these values (cf. Hiester et al., 2011). For F-high2 only part of the acceleration/deceleration occurs within the window over which the values of Froude number are calculated and, therefore, the average value is slightly over-estimated (Hiester, 2011).

The values of the Froude number for simulations with M_R and M_2 show good agreement with the fixed meshes and M_∞-var, Fig. 9. Simulation M_2-loose presents the best performance for the Froude number diagnostic, compared to both the fixed meshes and other adaptive meshes. The next best performing adaptive mesh simulation is M_2-mid, followed by M_∞-var. Only simulation M_∞-const significantly under performs.

An increase in boundary resolution ahead of the gravity current fronts can be seen in the mesh for simulations with M_2 and M_R, Fig. 5. It is precisely this increase in resolution, in regions where the curvature of the velocity fields is weaker, that is responsible for the increase in Froude numbers in simulation M_∞-var compared to simulation M_∞-const (Hiester et al., 2011). The region of increased resolution in simulation M_2-mid, for example, does not extend as far and does not demand as much refinement as in simulation M_∞-var, Figs. 3 and 5, but is sufficient to obtain comparable Froude numbers. The reduction in the number of vertices used in simulation M_2-mid compared to simulation M_∞-var suggests that in the latter case more refinement has occurred than was necessary. Furthermore, with M_2, the increase in resolution along the boundary is achieved without the need for spatial variation of the horizontal velocity weight, which, from the perspective of a model user, is clearly desirable. Again it is the ability of simulations with M_2 to capture variations at a range of scales that facilitates the improved performance.

5.4. Adaptive meshes with a constrained total number of mesh vertices

The adaptive mesh simulations discussed above are guided by the metric, and the number of vertices in the mesh is essentially unconstrained (in practice a maximum number of vertices is set by the user, Section 3.3.4, and, here, the meshes produced with M_∞ and M_2 do not reach this maximum, Fig. 6). Simulations that use different metrics (or even the same metric with different solution field weights) can have both a different average mesh resolution and a different distribution of mesh resolution. In order to separate the effects of these two factors, adaptive mesh simulations with a constrained number of mesh vertices are investigated. In these simulations, the number of mesh vertices is constrained by setting an upper and lower bound for the number of vertices to 2.0451×10^4, the same as the number of vertices in the coarsest fixed mesh, Table 2. The previously shown best performing M_2 metric and, for comparison, the M_∞ metric are used with the solution field weights as in simulations M_∞-const, M_2-coarse and M_2-mid. The constrained simulations are denoted by an asterisk, M_∞-const*, M_2-coarse* and M_2-mid*, respectively. This set allows comparison between both different metrics and different solution field weights. Note, the constraint on the number of mesh vertices leads to a reduction in the number of vertices for M_∞-const* and M_2-mid* compared to M_∞-const and M_2-mid and an increase for M_2-coarse* compared to M_2-coarse, Fig. 6.

The adapted mesh is subject to two constraints: the solution field weights and the bounds on the number of vertices. The adaptive mesh procedure adopted first computes the metric according to the solution field weights, as for the case with the unconstrained number of vertices. The metric is then scaled, if necessary, to coarsen or refine so that the number of vertices lies above or below the supplied lower or upper bound. This produces a mesh that attempts to meet the solution field weight criteria whilst satisfying the vertex constraint.

The meshes produced look qualitatively similar to the cases with an unconstrained number of vertices and the meshes for M_2-coarse* and M_2-mid* have a very similar distribution (not shown, cf. Figs. 3 and 5). The values of the background potential energy perturbation are also comparable between the simulations with M_2, with a difference at the end of the simulated time period of only 10% in the constrained case compared to approximately 50% in the unconstrained case, Figs. 8 and 10. Most crucially, once again, both the background potential energy and Froude number show improved performance with simulations that use M_2 over those that use M_∞, Figs. 10 and 11.

5.5. Comparison to published values

In Özgökmen et al. (2007), the two-dimensional lock-exchange is used to investigate the performance of different sub-grid-scale (SGS) models in large eddy simulations (LES) using a non-hydrostatic formulation. With this approach, the larger-scale eddies in the flow are computed and the SGS model represents the effect of smaller-scale eddies. The SGS models are found to improve the results for a given mesh resolution. As a part of the study, simulations without the SGS models are performed at a range of resolutions and the highest resolution values are taken as the benchmark solution.

Following Özgökmen et al. (2007), two Reynolds numbers $Re = 2800$ and 4300 are considered, where $Re = u_b h/\nu$, and u_b is the buoyancy velocity, h the domain half height and $\bar{\bar{\boldsymbol{\nu}}} = \nu\bar{\bar{\mathbf{I}}}$ is the kinematic viscosity, cf. Table 1. A Prandtl number $Pr = 7$ is used, where $Pr = \nu/\kappa$, where $\bar{\bar{\boldsymbol{\kappa}}} = \kappa\bar{\bar{\mathbf{I}}}$ is the thermal diffusivity which is

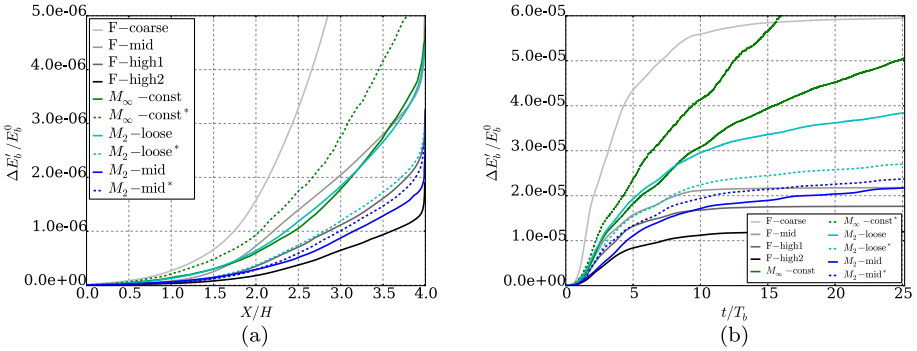

Fig. 10. Comparison of the normalised change in the background potential energy perturbation, E_b', Eq. (13), for fixed and adaptive mesh simulations, the latter both with and without a constrained number of mesh vertices (solid and dashed lines, respectively), Tables 2, 3, 5, Section 5.4. For the constrained case, denoted with an asterisk, the solution field weights are the same as the corresponding unconstrained case. In (a) the values are plotted with the scaled position of the no-slip front, X/H, for the propagation stages. In (b) the values are plotted with scaled time, t/T_b, over all time. The simulations that use M_2 perform better than those that use M_∞ in both the constrained and unconstrained cases.

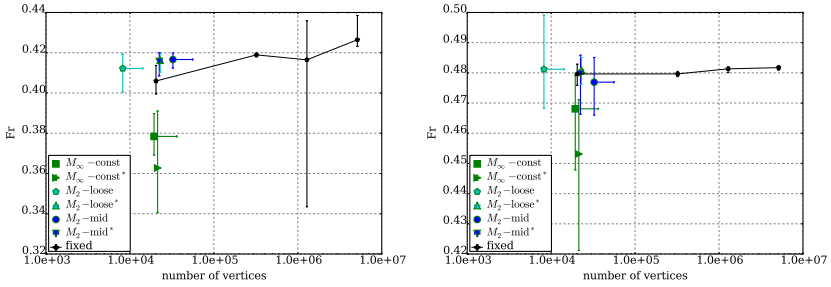

(a) no–slip (b) free–slip

Fig. 11. Comparison of the Froude number, $Fr = U/u_b$, Table 1, at the no-slip and free-slip boundaries for the fixed and adaptive mesh simulations, the latter both with and without a constrained number of mesh vertices, Tables 2, 3, 5, Section 5.4. For the constrained case, denoted with an asterisk, the solution field weights are the same as the corresponding unconstrained case. For the adaptive mesh simulations, the number of vertices in the mesh, given on the x-axis, is an average over the period between the initial time and the time when the no-slip front reaches the end wall, with the horizontal error bar extending to the maximum number of vertices in the time period. The adaptive mesh simulations that use M_2 are able to produce Froude numbers that are comparable to the higher resolution fixed mesh simulations, whereas the simulations that use M_∞ under perform.

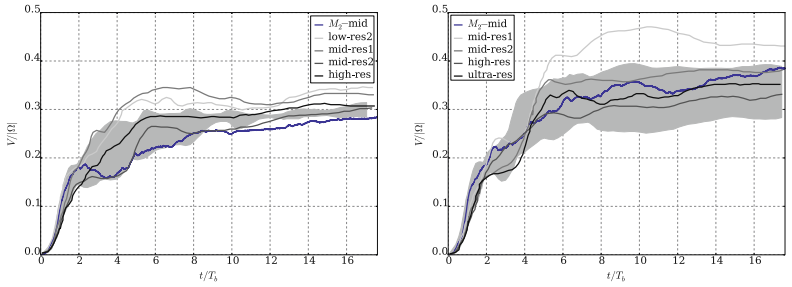

(a) $Re = 2800$ (b) $Re = 4300$

Fig. 12. Comparison of the volume fraction of the domain that contains mixed fluid over scaled time between Fluidity-ICOM adaptive mesh simulation M_2-mid and simulations of Özgökmen et al. (2007). V: volume of mixed fluid, $|\Omega|$: volume of the domain and t/T_b: scaled time. For the Fluidity-ICOM simulations, mixed fluid is classified as the fluid with temperature $-1/6 \leqslant T < 1/6$. The benchmark values (with no SGS models) from Özgökmen et al. (2007) are presented as lines and have the following degrees of freedom: low-res2 -1.08×10^4; mid-res1 -7.68×10^4; mid-res2 -1.728×10^5; high-res -2.7×10^5; and ultra-res 1.0404×10^6. The shaded area marks the range of values for the LES simulations with the SGS closure that was found to perform the best across the different mesh resolutions. The Fluidity-ICOM adaptive mesh simulations show good agreement with the Özgökmen et al. (2007) simulations and in particular the mid- to high-resolution benchmark values.

reinstated for the comparison. The values of u_b and h are as in Table 1 and the values of ν and κ are then determined from the values of Re and Pr. The domain used is shortened to be 0.5 m long to match the aspect ratio of 5 used in Özgökmen et al. (2007) and the bottom boundary condition is also changed from a no-slip to a free-slip, no normal flow condition. The adaptive mesh solution field weights are as in simulation M_2-mid, Table 5.

To quantitatively assess the diapycnal mixing in the flow, Özgökmen et al. (2007) divide the temperature field into three classes, light, mixed and heavy, and compare the volume fraction of fluid in each class. Here, the mixed class is compared between the different simulations and, in the Fluidity-ICOM simulations, corresponds to fluid with temperature perturbation $-1/6 \leqslant T - T_0 < 1/6$, Fig. 12.

In general, the spread of values across resolutions and SGS methods reported by Özgökmen et al. (2007) is larger for $Re = 4300$ than $Re = 2800$. At $Re = 2800$, the M_2-mid mixed water mass volume fractions behaves most like the (second) mid-resolution (1.728×10^5 degrees of freedom) benchmark case from Özgökmen et al. (2007) with generally comparable or smaller values than this case. At $Re = 4300$, the values for M_2-mid are more similar to the Özgökmen et al. (2007) high-resolution (2.7×10^5 degrees of freedom) benchmark case at early times and the Özgökmen et al. (2007) ultra-resolution benchmark case (1.0404×10^6 degrees of freedom) at late times.

At $Re = 2800$, M_2-mid uses an average of 3.2×10^4 vertices which increases to 4.3×10^4 vertices at $Re = 4300$. In terms of degrees of freedom (which given the control volume discretisation for temperature and P1 basis functions for pressure and velocity is the equivalent to the number of vertices for the Fluidity-ICOM simulations), this places M_2-mid between the Özgökmen et al. (2007) (second) low-resolution and (first) mid-resolution benchmark simulations (1.08×10^4 and 7.68×10^4 degrees of freedom, respectively). However, the M_2-mid mixed water mass volume fractions agree well with the higher resolution Özgökmen et al. (2007) simulations which have one to two orders of magnitude more degrees of freedom. This again highlights the good performance of the adaptive mesh simulations that use the metric M_2.

6. Conclusions

Simulations of the two-dimensional lock-exchange performed with Fluidity-ICOM on fixed and adaptive meshes have been evaluated primarily by comparison of the diapycnal mixing quantified through the background potential energy perturbation, Section 4.1. The diffusion term is neglected and, therefore, any diffusion is considered numerical. Values from simulations on the fixed meshes are taken as the benchmark for comparison, with the diapycnal mixing decreasing as the mesh resolution increases. The progress of the system is categorised into two main stages: the propagation stage, when the gravity currents travel across the domain, and the subsequent oscillatory stage, where the fluid 'sloshes' back and forth across the domain, Fig. 2.

Four different resolution fixed meshes are considered with horizontal and vertical element edge lengths $|\mathbf{v}| = 0.002$, 0.0005,

0.00025 and 0.000125 and are labelled F-coarse, F-mid, F-high1 and F-high2, respectively, Table 2. Three different forms of the metric, which guides the mesh adapt, are investigated: the absolute metric, M_∞, Eq. (6), the relative metric, M_R, Eq. (8), and the p-metric (with $p = 2$), M_2, Eq. (10) (Chen et al., 2007; Castro-Díaz et al., 1997; Frey and Alauzet, 2005). All meshes adapt to the temperature, horizontal velocity and vertical velocity, Tables 3–5.

The simulations capture the key dynamics of the lock-exchange, including propagation of the fronts, Kelvin–Helmholtz billows and turbulent mixing. The adaptive mesh simulations with M_∞ and M_2 use, in general, a comparable number of vertices to the coarsest resolution fixed mesh, F-coarse, and one to two orders of magnitude fewer vertices than F-high1 and F-high2, Fig. 8. The number of vertices for simulations that use M_R is more comparable to fixed mesh simulation F-mid.

The simulations that use M_2 produce the best performance, Fig. 8. The diapycnal mixing is comparable to the highest resolution fixed meshes (F-high1 and F-high2) during the propagation stage and more active part of the oscillatory stage ($t/T_b < 10$). During the later stages, the values of the background potential energy perturbation tend towards those of the middle resolution fixed mesh, F-mid. The simulations that use M_∞ produce variable performance with respect to the mixing diagnostics. The simulation that uses M_∞ with a spatially varying solution field weight has comparable levels of diapycnal mixing to the fixed mesh simulation F-high1 during the propagation stage. During the oscillatory stage the simulations with M_∞ exhibit more diapycnal mixing than the higher resolution fixed meshes and continue to mix at all times. The simulations with M_R do not offer an improvement over the simulations with M_∞ or M_2 and use at least 1.5–2 times as many vertices, Fig. 6. Comparison of adaptive mesh simulations with a constrained number of mesh vertices further demonstrate the improved performance with M_2, Figs. 10 and 11.

The weighting given to the smaller-scale fluctuations with M_2 facilitates the formation of a more appropriate mesh, Fig. 5. This leads to improved representation of the Kelvin–Helmholtz billows during the propagation stage and of the interface during the oscillatory stage and hence better representation of the diapycnal mixing. During the oscillatory stages, due to the diapycnal mixing, the curvature in the temperature field is not as large and the system also becomes less active. This leads to a coarsening of the mesh with M_∞, which tends to favour the strongest variations, and an increase in numerical diffusion, Figs. 3 and 8. A reduction in the solution field weights at later times would require additional user intervention but has the potential to improve performance of the simulations with M_∞ as the system evolves.

With M_R, the mesh is found to refine unnecessarily in regions of the domain where the velocity fields are near zero, Fig. 4. The temperature field, however, has near zero values at or near the interface, where resolution is required. The successful use of scaling by the local field value is, therefore, highly problem and field dependent. Using the global maximum or average of the magnitude of the field to scale the Hessian offers an alternative form of M_R that has the potential to be utilised effectively in scenarios where an initially active flow diminishes over time. However, in the current form, the use of M_R is not appropriate for the lock-exchange.

The Froude numbers for the adaptive mesh simulations are also calculated. With the exception of simulation M_∞-const which uses M_∞ with spatially constant solution field weights, the values are found to be in good agreement with the higher resolution fixed meshes and hence published values Fig. 9 (Hiester et al., 2011). With simulations that use M_2 and M_R this is achieved with no need for user-defined spatial variation of the solution field weights. Furthermore, simulations with M_2 in the lock-exchange configuration of Özgökmen et al. (2007), show good agreement between the Fluidity-ICOM mixing bin values and those from Özgökmen et al. (2007), Fig. 12.

The value $p = 2$ for the M_p metric is found to work well. The successful use of M_2 demonstrated here builds on the good results obtained with M_2 in Loseille and Alauzet (2011b) by extension to a turbulent and time-varying flow. A smaller value of p would lead to a more equal weighting between the smaller- and larger-scale fluctuations and a more uniform mesh would be expected (Loseille and Alauzet, 2011b). Conversely, as p increases, the larger-scale fluctuations will become increasingly dominant and the meshes produced will become more like those for the M_∞ case (Loseille and Alauzet, 2011b). The ability to capture a range of scales will also be useful for modelling of the lock-exchange in three dimensions, where the lobe and cleft instability adds to the complexity of the flow. The extension to three dimensions offers an important and tractable avenue for future investigation which also presents the opportunity for more extensive comparison to published results from other types of model e.g. Özgökmen et al. (2009a,b).

Whilst there are many other factors which will affect the efficiency of the individual models, such as the discretisation method, the adaptive meshes are able to produce flow characteristics that are equivalent to fixed meshes with one to two orders of magnitude more vertices (or degrees of freedom). This reduction in the number of vertices presents a significant improvement in the efficiency of the simulation for the finite-element discretisation method and numerical configuration used here. Such decreases in computational demand are not limited to the discretisation method and mesh structure considered here with, for example, 80% efficiency gains for the lock-exchange problem using a quadtree finite-volume discretisation reported in O'Callaghan et al. (2010). In addition, the reduction in computational demand with the use of adaptive meshes can provide an offset against the inherent increased cost of, for example, a finite-element discretisation method on an unstructured mesh compared to a finite-difference model on a structured mesh.

The performance of the adaptive mesh is highly dependent on the choice of metric. Changing the adaptive mesh settings can and will change the solution, particularly for a turbulent system such as the lock-exchange. However, the impact is not necessarily any greater than changing the discretisation method or the resolution of a fixed mesh. The effective use of an adaptive mesh with the simple metrics used here demands consideration of the problem to which it is applied and preliminary test simulations to obtain an appropriate set of solution field weights. The ability of the simulations that use M_2 to capture the dynamics of the lock-exchange in both the propagation stages and the oscillatory stages offers a promising outlook for the application of this metric to other flow regimes. There is clear potential for the utilisation of adaptive meshes in ocean modelling and this work provides further progress towards facilitating the wider use of adaptive meshes in this field.

Acknowledgement

The authors would like to acknowledge the generous funding of Imperial College London through the Janet Watson scholarships, the Grantham Institute for Climate Change and the UK Natural Environment Research Council (Project NE/F012594/1). This research is also funded by a Center of Excellence grant from the Research Council of Norway to the Center for Biomedical Computing at Simula Research Laboratory. The support of the High Performance Computing centre at Imperial College London, www.imperial.ac.uk/ict/services/teachingandresearchservices/highperformancecomputing, and access to the UK National Supercomputing Service HECToR Cray XT4 system, www.hector.ac.uk, under the NERC Shelf Seas Consortium are greatly appreciated. Thanks must be made to the authors' colleagues in the Applied Modelling and Computational Group at

Imperial College London, in particular, Stephan Kramer and Cian Wilson, for their continued advice and to the three anonymous reviewers for their comments. H.R. Hiester would also like to thank Paul Holland and Gareth Collins for their critique of this work.

References

Applied Modelling and Computation Group, 2011. Fluidity Manual. Department of Earth Science and Engineering, South Kensington Campus, Imperial College London, London SW7 2AZ, UK. <http://amcg.ese.ic.ac.uk>, version 4.1-release Edition.

Baker, T.J., 1997. Mesh adaptation strategies for problems in fluid dynamics. Finite Elements in Analysis and Design 25 (3–4), 243–273. http://dx.doi.org/10.1016/S0168-874X(96)00032-7, ISSN 0168-874X.

Behrens, J., 1998. Atmospheric and ocean modeling with an adaptive finite element solver for the shallow-water equations. Applied Numerical Mathematics 26 (1–2), 217–226. http://dx.doi.org/10.1016/S0168-9274(97)00090-1, ISSN 0168-9274.

Benjamin, T.B., 1968. Gravity currents and related phenomena. Journal of Fluid Mechanics 31 (2), 209–248.

Bernard, P.-E., Chevaugeon, N., Legat, V., Deleersnijder, E., Remacle, J.-F., 2007. High-order *h*-adaptive discontinuous Galerkin methods for ocean modelling. Ocean Dynamics 57 (2), 109–121. http://dx.doi.org/10.1007/s10236-006-0093-y.

Berntsen, J., Xing, J., Alendal, G., 2006. Assessment of non-hydrostatic ocean models using laboratory scale problems. Continental Shelf Research 26 (12–13), 1433–1447. http://dx.doi.org/10.1016/j.csr.2006.02.014, ISSN 0278-4343.

Blayo, E., Debreu, L., 1999. Adaptive mesh refinement for finite-difference ocean models: first experiments. Journal of Physical Oceanography 29 (6), 1239–1250. http://dx.doi.org/10.1175/1520-0485(1999)029<1239:AMRFFD>2.0.CO;2.

Cantero, M.I., Lee, J.R., Balachandar, S., García, M.H., 2007. On the front velocity of gravity currents. Journal of Fluid Mechanics 586, 1–39. http://dx.doi.org/10.1017/S0022112007005769.

Cao, J., 2005. Application of a posteriori error estimation to finite element simulation of incompressible Navier-Stokes flow. Computers and Fluids 34 (8), 972–990. http://dx.doi.org/10.1016/j.compfluid.2004.09.001, ISSN 0045-7930.

Castro-Díaz, M.J., Hecht, F., Mohammadi, B., Pironneau, O., 1997. Anisotropic unstructured mesh adaption for flow simulations. International Journal for Numerical Methods in Fluids 25 (4), 475–491. http://dx.doi.org/10.1002/(SICI)1097-0363(19970830)25:4<475::AID-FLD575>3.0.CO;2-6.

Chen, L., Sun, P., Xu, J., 2007. Optimal anisotropic meshes for minimizing interpolation errors in L^p-norm. Mathematics of Computation 76, 179–204. http://dx.doi.org/10.1090/S0025-5718-06-01896-5.

Farrell, P.E., Maddison, J.R., 2011. Conservative interpolation between volume meshes by local Galerkin projection. Computer Methods in Applied Mechanics and Engineering 200 (1–4), 89–100. http://dx.doi.org/10.1016/j.cma.2010.07.015, ISSN 0045-7825.

Farrell, P.E., Piggott, M.D., Pain, C.C., Gorman, G.J., Wilson, C.R., 2009. Conservative interpolation between unstructured meshes via supermesh construction. Computer Methods in Applied Mechanics and Engineering 198 (33–36), 2632–2642. http://dx.doi.org/10.1016/j.cma.2009.03.004, ISSN 0045-7825.

Frey, P.J., Alauzet, F., 2005. Anisotropic mesh adaptation for CFD computations. Computer Methods in Applied Mechanics and Engineering 194 (48–49), 5068–5082. http://dx.doi.org/10.1016/j.cma.2004.11.025, ISSN 0045-7825.

Fringer, O.B., Gerritsen, M., Street, R., 2006. An unstructured-grid, finite-volume, nonhydrostatic, parallel coastal ocean simulator. Ocean Modelling 14 (3–4), 139–173. http://dx.doi.org/10.1016/j.ocemod.2006.03.006.

George, P.L., Borouchaki, H., 1998. Delaunay Triangulation and Meshing: Application to Finite Elements. Hermes.

Geuzaine, C., Remacle, J.-F., 2009. Gmsh: a 3-D finite element mesh generator with built in pre- and post-processing facilities. International Journal for Numerical Methods in Engineering 79 (11), 1309–1331. http://dx.doi.org/10.1002/nme.2579.

Gorman, G.J., Pain, C.C., Piggott, M.D., Umpleby, A.P., Farrell, P.E., Maddison, J.R., 2009. Interleaved parallel tetrahedral mesh optimisation and dynamic load-balancing. In: Bouillard, P., Díez, P. (Eds.), Adaptive Modelling and Simulation. pp. 101–104.

Griffies, S.M., Pacanowski, R.C., Hallberg, R.W., 2000. Spurious diapycnal mixing associated with advection in a z-coordinate ocean model. Monthly Weather Review 128 (3), 538–564. http://dx.doi.org/10.1175/1520-0493(2000)128<0538:SDMAWA>2.0.CO;2.

Hallworth, M.A., Huppert, H.E., Phillips, J.C., Sparks, R.S.J., 1996. Entrainment into two-dimensional and axisymmetric turbulent gravity currents. Journal of Fluid Mechanics 308, 289–311. http://dx.doi.org/10.1017/S0022112096001486.

Härtel, C., Meiburg, E., Necker, F., 1999. Vorticity dynamics during the start-up phase of gravity currents. Societá Italiana di Fisica 22 (6), 823–833.

Härtel, C., Meiburg, E., Necker, F., 2000. Analysis and direct numerical simulation of the flow at a gravity-current head. Part 1. Flow topology and front speed for slip and no-slip boundaries. Journal of Fluid Mechanics 418, 189–212. http://dx.doi.org/10.1017/S0022112000001221.

Hiester, H.R., 2011. The application of adaptive mesh techniques to numerical simulations of gravity current flows. Ph.D. thesis, Imperial College London.

Hiester, H.R., Piggott, M.D., Allison, P.A., 2011. The impact of mesh adaptivity on the gravity current front speed in a two-dimensional lock-exchange. Ocean Modelling 38 (1–2), 1–21. http://dx.doi.org/10.1016/j.ocemod.2011.01.003.

Holmboe, J., 1962. On the Behaviour of Symmetric Waves in Stratified Shear Flows. Geophysical Publications 24, 67–113.

Huppert, H.E., Simpson, J.E., 1980. The slumping of gravity currents. Journal of Fluid Mechanics 99 (4), 785–799. http://dx.doi.org/10.1017/S0022112080000894.

Ilcak, M., Adcroft, A.J., Griffies, S.M., Hallberg, R.W., 2012. Spurious dianeutral mixing and the role of momentum closure. Ocean Modelling 45–46, 37–58. http://dx.doi.org/10.1016/j.ocemod.2011.10.003, ISSN 1463-5003.

Jacobs, C.T., Collins, G.S., Piggott, M.D., Kramer, S.C., Wilson, C.R.G., 2013. Multiphase flow modelling of volcanic ash settling in water using adaptive unstructured meshes. Geophysical Journal International 192 (2), 647–655.

Keulegan, G.H., 1958. Twelfth progress report on model laws for density currents. The motion of saline fronts in still water. National Bureau of Standards, Report (5831).

Legg, S., Jackson, L., Hallberg, R.W., 2008. Eddy-resolving modeling of overflows in Ocean Modeling in an Eddying Regime. Geophysical Monograph, vol. 177. American Geophysical Union, pp. 63–82.

LeVeque, R.J., 2002. Finite-Volume Methods for Hyperbolic Problems. Cambridge University Press.

Loseille, A., Alauzet, F., 2011a. Continuous mesh framework part I: well-posed continuous interpolation error. SIAM Journal on Numerical Analysis 49 (1), 38–60. http://dx.doi.org/10.1137/090754078.

Loseille, A., Alauzet, F., 2011b. Continuous mesh framework part II: validations and applications. SIAM Journal on Numerical Analysis 49 (1), 61–86. http://dx.doi.org/10.1137/10078654X.

Munday, D., Marshall, D., Piggott, M., 2010. Idealised flow past an island in a dynamically adaptive finite element model. Ocean Dynamics 60 (4), 835–850. http://dx.doi.org/10.1007/s10236-010-0291-5.

O'Callaghan, J., Rickard, G., Popinet, S., Stevens, C., 2010. Response of buoyant plumes to transient discharges investigated using an adaptive solver. Journal of Geophysical Research 115 (C11025). http://dx.doi.org/10.1029/2009JC005645.

Özgökmen, T.M., Iliescu, T., Fischer, P.F., Srinivasan, A., Duan, J., 2007. Large eddy simulation of stratified mixing in two-dimensional dam-break problem in a rectangular enclosed domain. Ocean Modelling 16 (1–2), 106–140. http://dx.doi.org/10.1016/j.ocemod.2006.08.006.

Özgökmen, T.M., Iliescu, T., Fischer, P.F., 2009a. Reynolds number dependence of mixing in a lock-exchange system from direct numerical and large eddy simulations. Ocean Modelling 30 (2–3), 190–206. http://dx.doi.org/10.1016/j.ocemod.2009.06.013.

Özgökmen, T.M., Iliescu, T., Fischer, P.F., 2009b. Large eddy simulation of stratified mixing in a three-dimensional lock-exchange system. Ocean Modelling 26 (3–4), 134–155. http://dx.doi.org/10.1016/j.ocemod.2008.09.006, ISSN 1463-5003.

Pain, C.C., Umpleby, A.P., de Oliveira, C.R.E., Goddard, A.J.H., 2001. Tetrahedral mesh optimisation and adaptivity for steady-state and transient finite element calculations. Computer Methods in Applied Mechanics and Engineering 190 (29–30), 3771–3796. http://dx.doi.org/10.1016/S0045-7825(00)00294-2, ISSN 0045-7825.

Piggott, M.D., Gorman, G.J., Pain, C.C., Allison, P.A., Candy, A.S., Martin, B.T., Wells, M., 2008. A new computational framework for multi-scale ocean modelling based on adapting unstructured meshes. International Journal for Numerical Methods in Fluids 56 (8), 1003–1015. http://dx.doi.org/10.1002/fld.1663.

Piggott, M.D., Farrell, P.E., Wilson, C.R., Gorman, G.J., Pain, C.C., 2009. Anisotropic mesh adaptivity for multi-scale ocean modelling. Philosophical Transactions of The Royal Society A: Mathematical, Physical and Engineering Sciences 367 (1907), 4591–4611. http://dx.doi.org/10.1098/rsta.2009.0155.

Popinet, S., Rickard, G., 2007. A tree-based solver for adaptive ocean modelling. Ocean Modelling 16 (3–4), 224–249. http://dx.doi.org/10.1016/j.ocemod.2006.10.002, ISSN 1463-5003.

Power, P.W., Piggott, M.D., Fang, F., Gorman, G.J., Pain, C.C., Marshall, D.P., Goddard, A.J.H., Navon, I.M., 2006. Adjoint goal-based error norms for adaptive mesh ocean modelling. Ocean Modelling 15 (1–2), 3–38. http://dx.doi.org/10.1016/j.ocemod.2006.05.001, ISSN 1463-5003.

Remacle, J.-F., Li, X., Shephard, M.S., Flaherty, J.E., 2005. Anisotropic adaptive simulation of transient flows using discontinuous Galerkin methods. International Journal for Numerical Methods in Engineering 62 (7), 899–923. http://dx.doi.org/10.1002/nme.1196.

Remacle, J.-F., Frazão, S.S., Li, X., Shephard, M.S., 2006. An adaptive discretization of shallow-water equations based on discontinuous Galerkin methods. International Journal for Numerical Methods in Fluids 52 (8), 903–923. http://dx.doi.org/10.1002/fld.1204.

Shin, J.O., Dalziel, S.B., Linden, P.F., 2004. Gravity currents produced by lock exchange. Journal of Fluid Mechanics 521, 1–34. http://dx.doi.org/10.1017/S002211200400165X.

Simpson, J.E., 1987. Gravity Currents in the Environment and the Laboratory. Ellis Horwood Limited.

Simpson, J.E., Britter, R.E., 1979. The dynamics of the head of a gravity current advancing over a horizontal surface. Journal of Fluid Mechanics 94 (3), 477–495. http://dx.doi.org/10.1017/S0022112079001142.

Smyth, W.D., Klaassen, G.P., Peltier, W.R., 1988. Finite amplitude Holmboe waves. Geophysical & Astrophysical Fluid Dynamics 43 (2), 181–222. http://dx.doi.org/10.1080/03091928808213625.

Speares, W., Berzins, M., 1997. A 3D unstructured mesh adaptation algorithm for time-dependent shock-dominated problems. International Journal for

Numerical Methods in Fluids 25 (1), 81–104. http://dx.doi.org/10.1002/(SICI)1097-0363(19970715)25:1<81::AID-FLD541>3.0.CO;2-0.
Strang, E.J., Fernando, H.J.S., 2001. Entrainment and mixing in stratified shear flows. Journal of Fluid Mechanics 428, 349–386.
Sweby, P.K., 1984. High resolution schemes using flux limiters for hyperbolic conservation laws. SIAM Journal on Numerical Analysis 21 (5), 995–1011. http://dx.doi.org/10.1137/0721062.
Thomas, L.P., Dalziel, S.B., Marino, B.M., 2003. The structure of the head of an inertial gravity current determined by particle-tracking velocimetry. Experiments in Fluids 34 (6), 708–716. http://dx.doi.org/10.1007/s00348-003-0611-3.
Tseng, Y.-H., Ferziger, J.H., 2001. Mixing and available potential energy in stratified flows. Physics of Fluids 13 (5), 1281–1293. http://dx.doi.org/10.1063/1.1358307.
Vasilevskii, Y., Lipnikov, K., 1999. An adaptive algorithm for quasioptimal mesh generation. Computational Mathematics and Mathematical Physics 39 (9), 1468–1486.
Venditti, D.A., Darmofal, D.L., 2003. Anisotropic grid adaptation for functional outputs: application to two-dimensional viscous flows. Journal of Computational Physics 187 (1), 22–46. http://dx.doi.org/10.1016/S0021-9991(03)00074-3, ISSN 0021-9991.
Wilson, C., 2009. Modelling multiple-material flows on adaptive unstructured meshes. Ph.D. thesis, Imperial College London.
Winters, K.B., D'Asaro, E.A., 1996. Diascalar flux and the rate of fluid mixing. Journal of Fluid Mechanics 317 (1), 179–193. http://dx.doi.org/10.1017/S0022112096000717.
Winters, K.B., Lombard, P.N., Riley, J.J., D'Asaro, E.A., 1995. Available potential energy and mixing in density-stratified fluids. Journal of Fluid Mechanics 289, 115–128. http://dx.doi.org/10.1017/S002211209500125X.
Xu, X., Chang, Y.S., Peters, H., Özgökmen, T.M., Chassignet, E.P., 2006. Parameterization of gravity current entrainment for ocean circulation models using a high-order 3d nonhydrostatic spectral element model. Ocean Modelling 14 (1–2), 19–44. http://dx.doi.org/10.1016/j.ocemod.2006.02.006.