Numerical simulations of internal solitary waves interacting with uniform slopes using an adaptive model

Graham Rickard [a,*], Joanne O'Callaghan [a], and Stéphane Popinet [a]

[a] *National Institute of Water and Atmospheric Research, P.O. Box 14-901, Kilbirnie, Wellington, New Zealand*

Abstract

Two-dimensional, non-linear, Boussinesq, non-hydrostatic simulations of internal solitary waves breaking and running up uniform slopes have been performed using an adaptive, finite volume fluid code "Gerris". It is demonstrated that the Gerris dynamical core performs well in this specific but important geophysical context. The "semi-structured" nature of Gerris is exploited to enhance model resolution along the slope where wave breaking and run-up occur. Comparison with laboratory experiments reveals that the generation of single and multiple turbulent surges ("boluses") as a function of slope angle is consistently reproduced by the model, comparable with observations and previous numerical simulations, suggesting aspects of the dynamical energy transfers are being represented by the model in two dimensions. Adaptivity is used to explore model convergence of the wave breaking dynamics, and it is shown that significant cpu memory and time savings are possible with adaptivity.

Key words: Internal wave breaking; modelling.

1 Introduction

Breaking internal solitary waves (ISWs) are considered to be important contributors to the energy required for mixing in the coastal zone (see, e.g., the

* Corresponding author. Tel.: +64 4 386 0311; fax: +64 4 386 2153.
Email addresses: g.rickard@niwa.co.nz (Graham Rickard), j.ocallaghan@niwa.co.nz (Joanne O'Callaghan), s.popinet@niwa.co.nz (Stéphane Popinet).

review by Helfrich and Melville (2006) and references therein). This is supplemented by recent observations (Bourgault et al 2007) of a train of ISWs impacting a coastline; integral to the study is a model comparison since it is noted that the general nature and predictability of ISWs remain largely unquantified. Furthermore, parameterisation of such energetic overturning for large scale models is hampered without detailed process scale modelling allied to field work. To that end this paper considers how high resolution modelling of such processes might be efficiently obtained using adaptive simulation techniques.

In a previous paper (Popinet and Rickard, 2007) an extension of an existing adaptive fluid code "Gerris" (Popinet, 2003 and 2006) for use with geophysical flows was described. The main issue was to demonstrate (via a series of test cases) that an "A-grid" formulation in Gerris comprising spatially collocated variables that allows for the use of adaptivity would not seriously compromise geophysical requirements (in particular geostrophic balance); the nature of the spatial operations ("approximate projection method") was shown to introduce a small amount of damping to the inertia-gravity waves while leaving the large scale flows relatively unhindered. Popinet and Rickard (2007) showed that Gerris was more than competitive in this regard in comparison to other non-standard ocean models. The test cases also demonstrated how the finite volume grid structure provided a significant advantage when it came to resolution and representation of the interaction with boundaries.

Although Popinet and Rickard (2007) showed the potential for a three dimensional geophysical extension of Gerris, the adaptivity was limited to the horizontal plane only, with depth projection through vertical z-level stacked layers. Furthermore, most of the test cases focussed on the horizontal geophysical responses of the model. In this paper the aim is to look in more detail at the vertical properties, in particular the integration of Boussinesq accelerations into the model in the context of non-hydrostatic simulations of internal solitary wave (ISW) breaking as an overall "assessment" of Gerris in the spirit suggested by Berntsen et al (2006). In this paper adaptivity operates in both the horizontal and vertical, removing the limitation of stacked layers.

Adaptive methods have long been recognised as potentially allowing efficient access to the multi-scale nature of geophysical flows. Examples for structured models include Blayo and Debreu (1999), and Burchard and Beckers (2004) and Hanert et al (2006) for vertical processes, Behrens (2005) in the context of atmospheric flows, and for unstructured models Piggott et al (2008) and references therein. The examples and tests in Piggott et al (2008) illustrate the adaptive potential, including both two and three dimensional flows, and the issues described (particularly geophysical scale pressure interacting with buoyancy) are those considered for the "semi-structured" Gerris here and by Popinet and Rickard (2007).

The approach we will take with Gerris is that of Bourgault and Kelley (2004) (as advocated by Berntsen et al, 2006) in that a non-hydrostatic model intended for use at coastal scales first be validated against existing observational data, in this instance the laboratory experiments of Helfrich (1992) and Michallet and Ivey (1999). Confidence in a model at these scales provides for a potential benchmark when assessing model performance at geophysical scales.

The motivation for this paper is to demonstrate that the collocated Gerris A-grid can indeed function readily with vertical Boussinesq accelerations and viscosity included. For the inter-model comparisons a "standard run" (see later) is used in order to allow reasonable coverage of the parameter space. In these standard runs the semi-structured Gerris grid is used to provide for enhanced (but fixed) resolution along the sloping surface where boundary layers are anticipated to evolve. A series of simulations will also be used to explore the resolution dependence of the model solutions (both fixed and adaptive grids), and to quantify the relative efficiency of the adaptivity. The comparisons show that Gerris represents an efficient modelling tool for the investigation of the nature of bolus generation, propagation, and eventual decay as a result of the ISW breaking.

2 Problem definition and model tests

In this study the coordinates are (x, y, z, t), where (x, y) and z are the horizontal and vertical coordinates, respectively, and t is time. With the Boussinesq approximation, the incompressible, non-hydrostatic equations in a Cartesian reference frame become,

$$\partial_t \boldsymbol{U} + \boldsymbol{U} \cdot \boldsymbol{\nabla} \boldsymbol{U} = -\frac{1}{\rho_0} \boldsymbol{\nabla} p - \frac{\rho'}{\rho_0} \boldsymbol{g} + \nu \nabla^2 \boldsymbol{U} \tag{1}$$

for momentum,

$$\boldsymbol{\nabla} \cdot \boldsymbol{U} = 0 \tag{2}$$

for continuity,

$$\partial_t \rho + \boldsymbol{U} \cdot \boldsymbol{\nabla} \rho = K \nabla^2 \rho \tag{3}$$

for density $\rho(x, y, z, t) = \rho_0 + \rho'(x, y, z, t)$ for constant ρ_0, with $\boldsymbol{U} \equiv (u, v, w)$ the components of the velocity, and the vector $\boldsymbol{g} = (0, 0, -1)g$ with g the constant acceleration due to gravity.

The initial model density $\rho(x, y, z, t = 0) = \rho_0 + \rho'(x, y, z)$ is set with $\rho_0 = 1000.0\,\mathrm{kgm}^{-3}$, and

$$\rho'(x, y, z) = \frac{\Delta\rho}{2} \left(1 + \tanh \left[\frac{z - z_i - \zeta(x)}{\Delta h} \right] \right), \tag{4}$$

where,

$$\zeta(x) = 2a_i \,\mathrm{sech}^2\left(\frac{x}{2W}\right), \qquad a_i\, W^2 = \frac{4(d_- d_+)^2}{3(d_- - d_+)}, \tag{5}$$

where x is the distance from the left hand edge of the domain. With $\zeta = 0$, equation 4 establishes a pycnocline of characteristic width Δh centred about a distance z_i below the top of the domain, with $\Delta\rho$ defining the change in density across the pycnocline. For finite ζ a perturbation of amplitude $2a_i$ is added (see the sketch in Figure 1). The aim is to establish an ISW of amplitude $a_0 = a_i$ propagating toward the slope. The profile mimics the initial conditions in the Helfrich (1992) and Michallet and Ivey (1999) experiments, and uses functional forms considered by Segur and Hammack (1982), Bogucki and Garrett (1993), and Bourgault and Kelley (2003).

Insert Figure 1 here

The numerical scheme used by the adaptive model to solve these equations is described in Appendix A. Most of the details have been presented previously (e.g. Popinet (2003), and Popinet and Rickard (2007)); however, the introduction of the Boussinesq terms interacting with viscosity in this particular context represents a new extension of the adaptive model, and so it is relevant to summarise the numerics.

For the present experiments, no normal flux conditions apply to the side and bottom surfaces. No-slip boundary conditions apply on the bottom boundary (both the level run-up of length $\mathrm{L_R}$ in Fig. 1, and the slope), while free slip conditions apply to the left hand boundary. There is no bottom friction. The top boundary is a rigid lid. Unlike Bourgault and Kelley (2007) we do not consider the impact of a sidewall friction that they are able to use in their laterally-averaged model. To that end our runs are purely two dimensional (i.e., the "free-slip side walls" as described by Bourgault and Kelley (2007)).

2.1 Model "standard run"

To be able to explore the parameter space observed by Helfrich (1992) and Michallet and Ivey (1999), and to be able to provide comparison with other model solutions, a "standard run" (see below) is defined. Adaptivity will be used later to explore convergence properties of our solutions. The impact of the use of adaptivity in Gerris will be more completely explored in a companion paper examining the lock-release experiment (O'Callaghan et al, 2009, pers comm).

For reference, a "standard run" will use constant viscosity and diffusivity values of $1.0 \times 10^{-6}\,\mathrm{m^2 s^{-1}}$, and will not use adaptivity. Spatially, most of the domain will be covered with a level of resolution of 7, i.e., the grid spacing

(the same in the x and z directions) will be "H" the depth of the fluid (see Figure 1) divided by 2^7, so that in this case with H = 0.15 m, the grid spacing is 1.17×10^{-3} m. At this level of resolution, Gerris has grid spacing comparable to that used by Berntsen et al (2006) and Bourgault and Kelley (2007). In anticipation of boundary layers along the slope, the semi-structured Gerris grid is exploited to give a level of resolution of 9 there. This spatial distribution is fixed in time for the standard runs.

The timestep in Gerris is dynamically adjusted in order to verify the CFL condition to satisfy advection scheme stability. A maximum timestep of 1.5×10^{-3} s is used, and for all but the runs with the finest grids suffices. For example, for the adaptive runs described later that allow maximum levels of resolution of 10, 11, and 12, the timestep drops to around 80, 60, and 40 per cent, respectively, relative to the maximum timestep, generally over the last 30 per cent of the test run time reflecting the increase of activity during wave breaking allied to on-going grid refinement.

Gerris maintains second order accuracy for spatial operators by a consistent grid mapping from the high resolution grid along the slope to the relatively coarser grid covering most of the domain (see Popinet, 2003 for details). An example of how the grid is laid out next to the sloping boundary is illustrated in Fig. 2. Note that the lower boundary in Fig. 2 is that given by the line showing the slope. Gerris uses a "cut-cell" technique to determine the best boundary fit within a cell, and boundary fluxes are calculated normal to this boundary (see A.3).

Insert Figure 2 here

The standard model test runs are initialised using the parameters describing model experiment (d) in Table 1, in particular having a slope of 0.217, and an initial pycnocline depth of $z_i = 0.024$m below the upper rigid lid of the model.

2.2 Model test : transport coefficient variation

Confidence in our standard model set-up was obtained from a convergence test where the initial grid resolution is the same, but the viscosity and diffusivity coefficients are changed. Here, the diffusivity and viscosity coefficients are set equal to one another, giving a Prandtl number (the ratio of viscosity to diffusivity) of 1 for the present set of experiments. This will go some way to demonstrating that the implicit viscous and diffusive effects from the numerical discretisation are not dominating the overall model solution at the transport values used in the majority of the runs reported.

Figure 3 shows density contours at time 7.5 s into each run for different values

of the diffusivity and viscosity coefficients. As the coefficients are reduced from Fig. 3(a) to (d) it is apparent that the model solution at this time is different in each case. In Fig. 3(a) the relatively large coefficients overly diffuse the density profile over the time of the ISW propagation. Figure 3(b) is the reference case, showing contours comparable to Bourgault and Kelley (2004) (see their Figure 10), with lighter fluid being drawn down towards the slope as the wave breaks. Upstream, downslope flow has compressed the pycnocline onto the slope itself; this downstream flow then separates where it meets the upslope flow of denser water, resulting in the circulation driving the lighter fluid down toward the slope.

Insert Figure 3 here

A further reduction in the diffusivity and viscosity coefficients results in Fig. 3(c). The first thing to note is that the pycnocline to the left of the breaking wave is tighter than that in Fig. 3(b); even in our standard case there has been some diffusion acting to spread the initial profile. The upstream pycnocline is also tighter and further down the slope than Fig. 3(b), suggesting that the downslope flow has been relatively intensified. It is also clear that the top of the breaking wave is now able to overshoot, forming a loop of lighter material ahead of the upslope flow of dense material. As a result, the lighter fluid being drawn down behind the wave appears more restricted.

The final reduction in the diffusivity and viscosity coefficients gives the contours in Fig. 3(d); the upstream downslope flow now appears even greater, such that the upstream pycnocline extends even further down the slope. It would seem that the overshoot of lighter fluid is being drawn by this down slope flow under the wave in such a way that instead of the upslope flow of denser fluid recirculating clockwise, it is now forced up and around counter-clockwise instead. The overall change in the solution from Fig. 3(c) to (d) shows that the model is still responding to changes in the coefficients, rather than being dominated by the numerical truncation errors.

2.3 Model test : impact of adaptivity

Using adaptivity, solution convergence for a specific case of a model initialised as the standard run can be explored. Now each of the model initial conditions are the same, but the maximum allowable level of resolution is changed for each run. This strategy assesses solution dependence on the grid resolution. Adaptivity is based on resolution of selected spatial scales, e.g., vorticity norms, tracer gradients, as quantified by equations (A.13) and (A.14), respectively. As a run proceeds, these spatial scales are diagnosed, and if deemed to be under-resolved, a grid refinement then operates until either the spatial

scales are resolved or the maximum allowable resolution has been reached. The scheme also allows for coarsening so as not to waste resources as gradients relax (except along solid boundaries where resolved scales must not change). The refinement preserves the space and time second order accuracy of the model. The parameters determining this time-adaptive behaviour then become extra model convergence parameters, in this case the maximum allowable resolution and the parameter ϵ in the adaptivity criteria, where equations A.13 and A.14 define ϵ in section A.4 in the Appendix. For the majority of adaptive runs, a single value of $\epsilon = 5.0 \times 10^{-2}$ is used, and the adaptive criteria are here applied to the density and vorticity. In cases of multiple adaptative criteria, cells are refined whenever they violate either of the refinement criteria, and coarsened whenever all the coarsening criteria are met.

An example is shown in Fig. 4. Standard run density contours at time 7.5 s into the run are shown in Fig. 4(a). An adaptive run that is initialised exactly as for the standard run, but is allowed to adapt up to a maximum level of resolution of 12 (based on the local density gradient and/or the vorticity), is shown at the same time in Fig. 4(b). The difference between these frames is contoured in Fig. 4(c). By comparing Figs. 4(a) and (b), it is apparent that the adaptive run results in some different features, notably the tightening of the contours for the downslope flow as it is forced to rotate around the denser fluid, the extended counter-clockwise roll-up of lighter fluid, and the sharper finger to its right. Figure 4(c) contours the density differences; these differences range from -16 to $10\,\mathrm{kgm}^{-3}$ relative to the constant background density of $1000\,\mathrm{kgm}^{-3}$, so there are percentage changes of around 1% or so. Since the majority of contours between Figs. 4(a) and (b) are similar, it would seem that most of these changes represent a shift in position of features, rather than gross amplitude errors. Regardless, it is clear that the standard run is different when adaptivity in time is used.

Insert Figure 4 here

To illustrate the relationship between the adaptivity and the dynamics, Fig. 5 contours the resolution levels underlying the solution in Fig. 4(b). The range of levels chosen for this solution range from 7 (the coarsest) to 12 (the finest). The dynamically evolving resolution is determined by the refinement criteria, and the maximum allowed level of resolution. For this instance, these criteria combine to cover the wave breaking region with levels 8 to 10, with most of the domain remaining at the original value of 7. The level 10 regions are further extended to cover the boundary layers along the slope to resolve the evolving flows.

Insert Figure 5 here

To better quantify the convergence consider the above experiment run out to

7.5 s using adaptive runs that allow for maximum levels of refinement of 10, 11, and 12. In these runs the initial grid is at level 7 nearly everywhere, except for close to the solid boundary where the maximum level is set. At each timestep the adaptive criteria are checked, and if violated grid refinement takes place. The maximum differences in density and speed at time 7.5 s for the adaptive sequence for $(10 - \text{standard}, 11 - 10, 12 - 11)$ were $(16.6, 3.5, 2.5)\,\text{kgm}^{-3}$ and $(4.8, 1.06, 0.6)\,\text{cms}^{-1}$, respectively. The runs were still converging as the maximum level of resolution was increased, but there has been a relatively significant change between refinement levels 10 and 11.

Adaptivity is only useful if it proves to be more efficient than the equivalent non-adaptive run initialised everywhere with the finest resolution obtained in the adaptive run, and if the adaptive and constant resolution solutions converge. To that end constant resolution versions have been run as above to time 7.5 s using initial uniform levels of refinement of 5 to 10. For these two dimensional runs, the total cpu memory increases by 4 for each increase in level, here 1 Mb for level 5 up to 1002 Mb for level 10. For the single cpu workstations used in this study, levels 11 and 12 at regular resolution exceeded the resources, and could not be run. However, cpu timings showed that for each level change, the cpu increases by around 5, going from 0.007 days for level 5 to 10.3 days for level 10, so that levels 11 and 12 would require at least 50 and 250 days, respectively, to get to time 7.5 s.

In comparison, the adaptive run with maximum level of 10 requires 0.79 days to reach time 7.5 s, and uses a maximum of 45 Mb of memory. For refinements up to 11 and 12, the cpu times and maximum memories are 2.2 and 4.9 days, and 57 and 81 Mb, respectively. For each maximum level change, the cpu time increases by between 2 and 3, not only due to memory increases but also because the timestep starts to decline to satisfy the advective CFL condition for the finest resolution. In comparison to its constant resolution counterpart (recalling the values for 11 and 12 are conservative estimates), the adaptive run uses factors of 22, 70, 198 and 13, 23, 50 less for memory and cpu time, respectively, for levels 10, 11, and 12. In terms of total numbers of grid cells, the adaptive runs use maximums of 116974, 145203, and 197613 for levels 10, 11, and 12, respectively, factors of 23, 76, and 222 times less than the respective constant resolution runs, reflecting the memory changes.

So, if the adaptive run at a given level is accessing all the scales of its constant resolution counterpart, then it's clear that the adaptivity will provide great cpu time efficiency, here increasing up to a factor of 50. These gains are case dependent, and will increase as the volume of fine scales reduces compared to the domain volume. As Fig. 5 shows, the high resolution region is mainly confined to the wave breaking zone.

Without an analytic solution, model consistency is checked using convergence.

The maximum density and speed differences for $(6-5)$ down to $(10-9)$ for the constant resolution runs were
$(34.0, 28.0, 20.0, 16.0, 8.0)\,\mathrm{kgm}^{-3}$ and $(14.0, 11.0, 8.0, 3.7, 3.2)\,\mathrm{cms}^{-1}$
for density and speed, respectively. As with the adaptive sequence above, convergence is clear. Finally, taking the level 10 constant resolution solution as a reference, the differences between it and the standard run and the level 10, 11, and 12 adaptive runs are
$(15.0, 5.0, 4.0, 4.0)\,\mathrm{kgm}^{-3}$ and $(4.7, 1.05, 1.05, 1.5)\,\mathrm{cms}^{-1}$
for density and speed, respectively. Again, convergence is clear, but some saturation as the adaptive models are accessing some finer scales.

Could we have started with a relatively low level of resolution and then allowed adaptivity to fill in evolving scales as required? In general no, as the initial conditions and the boundaries need to be adequately resolved at the outset; adaptivity cannot in general compensate for initially unresolved scales implied by the boundaries or initial stratification. For a uniform medium at rest with regular boundaries, this strategy works well (see Popinet (2006) for examples), and maximises the adaptive gain. In these ISW simulations, under resolution of the initial stratification will result in cumulative errors later in time. The resolution of the standard run is perhaps a compromise between adequately capturing the initial scales, and spanning the parameter space in reasonable time.

In summary, these model tests show that;

(1) the implicit viscosity and diffusivity of the numerical discretisation is not dominating the overall solution for values of $1 \times 10^{-6}\,\mathrm{m}^2\mathrm{s}^{-1}$ for the viscosity and diffusivity for the standard run;
(2) differences of around 1 per cent remain in the density as a result of under resolution when not using the adaptivity in time for the standard run;
(3) adaptivity can provide factors of up to 50 times saving in cpu time relative to the equivalent constant resolution run at the maximum resolution of the adaptive run. This scaling arises here because the wave breaking volume is comparatively compact, and the efficiency of the adaptive grid scheme in Gerris.

To explore a wider parameter space, model results using the standard run configuration are compiled and discussed. To provide comparison with the Helfrich (1992) and Michallet and Ivey (1999) observations, and the numerical results of Bourgault and Kelley (2007), slopes as shallow as 0.034 relative to the standard run's of 0.217 will be simulated. The shallowest slope requires a domain 5 times larger and needs to run for at least 5 times as long as the standard run in order to capture the wave breaking. From a resource point, adaptive runs may have been feasible, but given that the percentage differences might be around 1 per cent, and that the wider context is testing the Gerris

dynamical core, the standard run configuration will be presented. As we shall see, experimentally observed turbulence at the time of wave breaking does not occur in the simulations, and so there is clear scope for future adaptive runs to capture finer scales and the third dimension.

3 Model Results

In this section, model results from the experiments listed in Tables 1 and 2 are described. The intention is to compare Gerris to the laboratory experiments of Helfrich (1992) and Michallet and Ivey (1999). Comparisons with the laterally-averaged non-hydrostatic model of Bourgault and Kelley (2007) (as described in Bourgault and Kelley (2004)) are also made. Each run in this section is initialised as per the standard run in terms of grid resolution, with no adaptivity, and with diffusivity equal to viscosity of $1.0 \times 10^{-6}\,\mathrm{m^2 s^{-1}}$.

The phenomena of interest are the wave breaking process and subsequent turbulent vortex ("bolus") generation, bolus propagation, and the energetics resulting from the wave breaking, in particular what proportion of the incoming energy flux is reflected (and hence how much is left for mixing processes up-slope).

3.1 Wave breaking turbulence

In the Michallet and Ivey (1999) experiments they identify instability associated with three-dimensional motion. This instability and subsequent turbulence is apparent in the sequence of frames in their Figure 3, from the time that the initial wave breaks until later as the bolus runs upslope. Helfrich (1992) similarly identifies an initial patch of wave instability (or "break point") from shadowgraphs, and consequently defines the position of this break point to be the "maximum offshore location of the initial patch of turbulence".

Gerris does not produce such instability or turbulence as the initial wave breaking takes place, in common with Bourgault and Kelley (2004). Even as the model resolution and viscosity and diffusivity are varied (see Figs. 3 and 4 for example), the pycnocline to the left of the breaking wave remains relatively coherent, with none of the clear filamentary structure seen in Michallet and Ivey (1999). The most likely explanation – as proffered by Bourgault and Kelley (2004) – is that the two dimensional models are missing important three-dimensional processes.

This is not to imply that the model pycnocline is not "dynamic". Indeed, as the

slope is reduced and the number of boluses increases, a complex structure of vortices and density develops from which each bolus emerges (not shown). As the overturning and wave breaking processes diminish so does the pycnocline complexity, slowly returning toward a stable, final structure.

We identify a model break point by comparison with Figures 3c and 4a in Michallet and Ivey (1999), and their observation that "the first sign of breaking seems to be a gravitational instability, following the separation of the flow in the lower layer". Diagnosis of model flow vectors and vorticity allied to density contours shows when the lower layer flow first separates from the slope, leading to the formation of a clear counter-clockwise circulation beneath the wave; this we identify as our break point. Beyond this point, the counter-clockwise vortex intensifies and enlarges, such that at later times a bolus is formed; this follows the sequence of Figures 3c to 3e and 4a to 4c in Michallet and Ivey (1999).

3.2 Bolus Identification

Using model experiment 2 in Table 1 as an example, Figure 6 shows a snapshot at time 37.5 seconds into the run. In Fig. 6(a) for density contours there are three distinct and discrete features with core densities labelled as 1016.0 and 1028.0 kgm^{-3} above that of the surrounding material. Following Helfrich (1992), each of these counts as a bolus since they exist beyond the undisturbed pycnocline-slope intersection (indicated by the horizontal line in the model domain in each frame). Note that the vertical scale is stretched four times relative to the horizontal for clarity.

Insert Figure 6 here

Figure 6(b) shows horizontal flow, with positive (up-slope flow) contours within each bolus (with peak values of 4.0 and 5.0 cms^{-1}), and negative (down-slope return flow) contours extending from just above each bolus to the model domain rigid lid at the top (indicated by the 2.0 cms^{-1} contour). This shows that these density anomalies are not stationary. Lastly, Fig. 6(c) contours vorticity, showing discrete structures for each bolus comprising a positive vorticity "cap" indicating local counter-clockwise circulation (labelled with the 1.34 s^{-1} contours) overlying negative vorticity locally clockwise circulation (the $-3.34\,\mathrm{s}^{-1}$ contours).

However, the Helfrich (1992) experiments showed that boluses are generated below the undisturbed pycnocline-slope intersection as a result of the ISW breaking process. In Fig. 6 signatures of a fourth bolus are apparent within the bound of the pycnocline-slope intersection; there is discrete density (but not as obvious as those further up the slope indicated by the 1032 kgm^{-3} contour), there is upslope flow of around 3.0 cms^{-1} with return downslope flow

above, and a consequent positive vorticity cap with less clear negative vorticity below. This is typical of the model bolus structure below the pycnocline-slope intersection – and particularly so for the emergence of subsequent boluses after the first one – in that the bolus signatures tend to be less distinct.

Nevertheless, by identifying each new bolus in terms of its density, upslope flow, and vorticity, it is possible to build up a picture of their history from close to initial wave breaking until they are well up the slope. An example of such a time history is plotted in Fig. 7 for model experiment 2 showing the evolution of four boluses moving upslope, the first identified with crosses, the second squares, the third and fourth triangles and diamonds, respectively. Each bolus trajectory is plotted in terms of its position relative to the pycnocline-slope intersection ($x/X_{BP} = 0$) as a function of time. The character of these trajectories compare well with those described and plotted by Helfrich (1992) for his Figure 8, in particular the fact that each bolus is generated for $x/X_{BP} < 0$ (the horizontal dashed line), that subsequent boluses first appear further up the slope than their predecessors, and that boluses can be overtaken and absorbed by boluses following on from behind.

Insert Figure 7 here

In summary, bolus identification in these two dimensional model runs can be done using density (relative to local background values), zonal velocity, and/or vorticity signatures. Bolus identification below the undisturbed pycnocline-slope intersection is undoubtedly messy, especially for any subsequent boluses, but by tracking backwards from where they emerge identification can be made. These techniques have been used to compile the bolus statistics for the model experiments in this paper.

3.3 First Bolus Evolution

Typical evolution of the first bolus from close to its emergence until propagation well upslope is illustrated in Figs. 8 and 9, again for model experiment 2 in Table 1. Figure 8 shows contours of density, while Figure 9 combines contour lines of vorticity with velocity vectors (the latter at a relatively coarser level of resolution of 6 for clarity) to show the circulation. In each figure the frames (a), (b), (c), and (d) are at times 22.5, 27.0, 30.0, and 37.5 seconds into the run, respectively. The size of each frame is proportional to the height of the bolus at that time (with the bolus heights detailed in the Fig. 8 caption for reference). Frames (a) and (b) are for bolus positions downslope of the undisturbed pycnocline-slope intersection, while frames (c) and (d) are upslope.

Insert Figure 8 here Insert Figure 9 here

The initial model bolus tends to have three elements clearly delineated in the density and flow fields. This overall structure is contained by the strong downslope flow ahead of the bolus that separates at its front, then passes over the top of the bolus, with re-circulation at the back of the bolus. The three bolus elements comprise two counter-clockwise rotating cells (at the back and front), with a clockwise rotating cell between them. Note that the velocity vectors have been thinned for clarity. However, the vorticity contours are based on the full resolution, and the tightening of contours close to the slope show the presence of boundary layers.

Dense upslope flowing fluid is drawn into the back of the bolus, and entrains with lighter fluid being circulated around the left-most vortex. This fluid then circulates and mixes within the core of the bolus. Note that this mixing process continues to dilute the density within the bolus as it propagates upslope, such that the maximum density in frame (a) of $1047.0\,\mathrm{kgm}^{-3}$ steadily reduces to $1017.0\,\mathrm{kgm}^{-3}$ by frame (d) in Fig. 8.

A characteristic bolus height and base length can be obtained from the evolving bolus structure. The height is the maximum distance from the slope up to where the down slope flow vectors intersect the density contours at the bolus top, locating the transition between the core bolus and background densities. As this density difference reduces as the bolus propagates upslope, the transition sharpness declines, but the bolus top is still indentifiable via flow vectors. Similarly, bolus base length can be defined in terms of the distance between the upward flow at the bolus leading nose to downward flow at the bolus rear edge. Again, these locations can be correlated with density gradients to identify the bounds of a single bolus. The solid lines on frames (b) in Figs. 8 and 9 illustrate these definitions. It is clear that the overall dimensions of the bolus change, with general shrinkage in both height (going from 0.03 m in frame (a) to 0.0144 m by frame (d)) and base length (from 0.05 m in (a) to 0.03 m in (d)) as the bolus progresses upslope.

It is also clear that the bolus three cell structure changes as it moves upslope. It would seem that the cell at the front of the bolus shrinks (most evidently in density), and that the left-most vortex seems to extend over the middle vortex such that at the later times the typical cap structure in vorticity is obtained. This is perhaps in response to the change in flow field. Note that in frames (a) and (b) in Fig. 9 the velocity vectors show that there is downslope flow into the nose of the bolus. However, in frames (c) and (d), where the bolus is now moving beyond the undisturbed pycnocline-slope intersection, this downslope flow disappears to be replaced by upslope flow. The bolus structure at these later times is consistent with that seen in other models (e.g. Härtel et al (2000) and Venayagamoorthy and Fringer (2007)), and similar to gravity current head structures described by Simpson (1972, 1997).

3.4 Model bolus validation

The bolus properties examined by Helfrich (1992) in his experiments for a parameter space dependent on the initial densities, pycnocline depths, and slope angle, can be diagnosed from Gerris to provide for a more quantitative evaluation of the model. While particular experiments of Helfrich (1992) are not modelled directly, the observational space in terms of the parameter λ will be explored, where

$$\lambda = (kL)^{-1}, \qquad k^{-1} = \sqrt{\frac{4(d_+ d_-)^2}{3(d_+ - d_-)a_0}}, \tag{6}$$

where L, d_+, and d_- are defined in Fig. 1, and a_0 is the maximum displacement of the incoming wave. All of the model experiments in Tables 1 and 2 will be used and plotted against their reference label, i.e., 1-4 and a-i. Note that experiments 1-4 and a-d are parallel in every regard except the initial depth of the pycnocline; this then changes the respective value of λ.

Figure 10 plots breaking location in terms of a_0/d_{BP} as a function of λ for comparison with Helfrich's (1992) Figure 6. Both show a tendency for an increase of a_0/d_{BP} as λ decreases. For the larger range of λ used here, the model suggests that the breaking location actually saturates as λ increases. The model values for a_0/d_{BP} are typically 0.2 or so higher than those found in the experiments. As described earlier, identifying the break point in the model is difficult due to the lack of an initial patch of instability seen in the experiments.

Insert Figure 10 here

The number of boluses in each experiment as a function of λ is shown in Fig. 11 for comparison with Helfrich's (1992) Figure 9. The model captures the dependence with λ well, showing a typical increase in the number of boluses as λ decreases. Furthermore, the multiple-bolus characteristics shown in Fig. 7 fit that seen and discussed by Helfrich (1992) with respect to his Figure 8. Note that the model results do not suggest an absolute scaling with respect to λ. Rather, for each sequence of "similar" experiments, i.e., 1-4, a-d, e-i, the scaling with λ indeed holds, but the number of boluses for a given λ can vary.

Insert Figure 11 here

The initial height of the first bolus in Fig. 12 compares favourably with Helfrich's (1992) result in his Figure 10, both covering the range 1 to 2 over the values of λ sampled. The agreement is relatively good given the uncertainties in estimating the time and then the actual first bolus height (see the definitions and error estimates quoted by Helfrich).

Insert Figure 12 here

Frames (a)-(c) in Fig. 13 show the first bolus aspect ratios (H/l_b) as a function of distance for the model experiments for comparison with Figure 11 in Helfrich (1992). The major consistent difference across the models is their tendency to have the aspect ratio somewhat larger than observed; Helfrich shows values in the range 0.2 to 0.4 (given his estimated uncertainties), whereas the model aspect ratios are typically in the range 0.4 to 0.8. Helfrich also notes a tendency for the aspect ratio to increase with distance; this general tendency is not so apparent in the model, but there are certainly individual instances where this does occur.

Insert Figure 13 here

The decay of the first bolus height with distance is shown in Fig. 13(d)-(f) for comparison with Helfrich's (1992) Figure 12. The near linear decay observed by Helfrich is reproduced by nearly all of the model experiments. In some of the model experiments (particularly in Fig. 13(f)) the bolus height is seen to increase at early times in the bolus evolution; again uncertainties in timing of measurements between the observations and models may contribute to these differences.

3.5 Energy Reflectance

As a final consideration, comparisons with energy reflectance results from Bourgault and Kelley (2007) are made, specifically their free slip sidewall results shown in their Figure 6. Bourgault and Kelley (2007) covered all of the Michallet and Ivey (1999) experiments in compiling their results. In this study we have focussed on specific experiments labelled e-i in Table 2 corresponding to experiments 4, 8, 12, 7, and 13, respectively, from Table 1 in Michallet and Ivey (1999); results from the latter are reproduced in Table 3. The Gerris wave amplitudes (a_0) in Table 2 match fairly well, although they tend to typically be a few percent less than measured. The comparison for the characteristic length L_w shows that the model responded a little more erratically, with slight overestimates for shorter L_w but with relatively larger underestimates for longer L_w (experiments h and i).

For direct comparison with the model results of Bourgault and Kelley (2007), the energy reflectance "R" was diagnosed, where

$$R = E_R/E_0, \tag{7}$$

where E_0 and E_R (as defined by equations 7 and 8 in Bourgault and Kelley (2007)) are the energy fluxes measured at the base of the slope for the initial

incoming and first reflected waves, respectively. The reflectance provides an important measure of the likely amount of mixing induced by the breaking of these waves, and scaling their behaviour for parameterisation in larger and climate scale models remains relevant. To validate their model, Bourgault and Kelley (2007) compared the dependence of their model R in terms of the Iribarren number (Iribarren and Nogales (1949) ξ where,

$$\xi = s/\sqrt{(a_0/L_w)}, \tag{8}$$

where s is the slope, for the Michallet and Ivey (1999) experiments. Bourgault and Kelley (2007) found their free slip sidewall model R values are typically 0.1 larger than the equivalent experimental values. They showed that by including no slip sidewalls close agreement with the laboratory experiments could be obtained. In terms of the free slip sidewall runs, Bourgault and Kelley (2007) performed a curve fit to their results and obtained a parameterisation of the form,

$$R = 1 - e^{-\xi/\xi_0}, \tag{9}$$

where a least squares fit returns $\xi_0 = 0.78 \pm 0.02$.

The results from the Gerris experiments are summarised in Figure 14 showing R as a function of ξ.

Insert Figure 14 here

The Gerris results for R show that experiments e, f, g, and h have reflectance estimates close to those of Bourgault and Kelley (2007) (the bold curve), and with a similar dependence on ξ. The only anomaly is experiment i, which falls closer to the Michallet and Ivey (1999) R value.

For comparison, Fig. 15 plots the individual energy fluxes as a function of time used to calculate the reflectances for model experiments e to i. The bold curve for experiment f in Fig. 15(b) compares favourably with that from Figure 3 in Bourgault and Kelley (2007), and indeed the value obtained for R is close to that found in their model. Note that the character of each energy flux curve changes with slope, the density jump $\Delta\rho$, and the depth of the pycnocline z_i, so that the incoming energy flux and first return pulse can be relatively close or widely separated in time from each other.

Insert Figure 15 here

It is encouraging that the tendency for increasing R with ξ is reproduced by Gerris and generally in close agreement with Bourgault and Kelley (2007) model results. The only outlier is experiment i. In terms of differences, it is clear that the time history of energy flux for experiment i has quite a narrow time separation between the incoming and first return pulse relative to the other experiments. Experiment i also has a significantly larger value of

λ (equation 6) than all the other experiments (see Tables 1 and 2). It is clear, however, that further runs will be needed to cover the full range of Michallet and Ivey (1999) experiments in order to assess and understand how anomalous the experiment i reflectance result is.

4 Conclusion

Multi-scale, non-hydrostatic simulations of geophysical flows represent a significant computational challenge. In this paper the potential of an adaptive (in space *and* time) model, Gerris, is explored, in terms of validation against laboratory scale experiments, and against other models benchmarked in a similar fashion. Gerris combines a tree-based structure with a finite-volume discretisation for efficient adaptivity, thereby allowing for local refinement over fine scale features without the need for excessive refinement over less dynamic parts of the domain. The ability of Gerris to handle largely horizontal geophysical processes was demonstrated in Popinet and Rickard (2007). This paper extends that work by looking in more detail at vertical processes.

The flexible grid structure of Gerris allows for selective refinement in the model domain. Here this property is exploited to allow for increased refinement along the sloping boundary to better resolve evolving boundary layers. For most of the runs in this paper this spatially non-uniform grid is fixed in time (for our so-called "standard runs") in order to be able to span the parameter space of interest, and to allow comparison with previously published model output. Similarly, the diffusivity and viscosity coefficients have been restricted for comparison purposes. All the runs in this paper were done on single CPU workstations, but note that Gerris is fully parallelised.

For the adaptive runs, efficiency can be measured by comparison with a version comprising constant resolution equal to the finest resolution of the adaptive run. By this measure the adaptive run requires some 10 to 50 times less cpu time, and up to 22 to 200 times less memory (dependent on maximum allowed resolution) than the constant resolution run for the breaking waves. This gain is clearly problem dependent, as it will be proportional to the volume fraction in which the finest scales arise.

The general dynamical features of internal solitary waves breaking on uniform slopes reported by Helfrich (1992) and Michallet and Ivey (1999) are reproduced by Gerris, especially given the experimental uncertainties, and the fact that the model simulations are two-dimensional and use a rigid lid rather than a free surface as a top boundary condition. The absence of turbulent generation at the time of wave breaking in the models has been attributed to the lack of a third dimension (and so will require future simulations). Nevertheless, the

overall correspondence between the model and observations suggests that the major energy conversion processes involved in wave breaking leading to bolus generation can be largely captured in two dimensions. Some of the detailed diagnostics show that the model gets the general time evolution correct, but that amplitudes can be wrong (for instance, in terms of the bolus aspect ratio).

For the tests in this paper we have limited ourselves to two dimensions, and so have not further explored the issues of side wall interactions considered by Bourgault and Kelley (2007). The latter find that they need to introduce side wall friction effects into their model in order to reproduce the experimentally obtained energy reflectance values of Michallet and Ivey (1999). Furthermore, Bourgault and Kelley (2007) find that the Helfrich (1992) reflectance values are 0.1 to 0.4 smaller again than Michallet and Ivey (1999); the reasons for this are unclear, but Bourgault and Kelley (2007) suggest that again side-wall effects have an influence, as well as the different Helfrich (1992) tank geometry. Given these uncertainties should we expect to be able to reproduce the Helfrich (1992) bolus evolution results with our present numerical experiments? Perhaps not, but the similarities between the model and experiment suggest that perhaps the lateral effects are not dominating the wave breaking process; nevertheless, more complete answers will have to wait for fully three dimensional simulations.

Aspects of bolus evolution are revealed by the model. The boluses are generated below the undisturbed pycnocline-slope intersection, and while the bolus is on this part of the slope, the model suggests a three cell structure to the bolus in which denser up-slope flowing material is mixed into the bolus interior with the lighter down-slope flowing fluid. As the bolus moves upslope toward the undisturbed pycnocline-slope intersection and beyond it, the bolus changes from three to two cells, eventually forming a structure similar to typical gravity current heads. This bolus evolution is consistent across the range of model parameters explored.

In terms of Gerris as a geophysical tool, it is encouraging that Gerris performs well in comparison to other non-hydrostatic geophysical models such as that of Bourgault and Kelley (2004, 2007), and the Massachusetts Institute of Technology general circulation model (Marshall et al., 1997) and the Bergen Ocean Model (Berntsen (2000)) which are both discussed in Berntsen et al (2006) in terms of wave breaking tests. Adaptivity is presently enabled by collocation of prognostic variables, making the Gerris dynamical core unusual with regard to other geophysical models. The tests in this paper and previously (Popinet and Rickard (2007)) continue to suggest that Gerris and its adaptivity represents an efficient alternative to simulate multi-scale, non-hydrostatic flows.

With isotropic adaptivity, there will be no vertical stacked layers in general. Unresolved processes such as vertical mixing, convection etc are parameterised

locally down each vertical column in many global scale models. The loss of this vertical structure means parameterisation amendment will be required in order to convert purely vertical schemes to more general forms. Further, selection rules will be needed to separate when and where the adaptive model is actually resolving or not these small scale processes. The challenge will be to efficiently incorporate such parameterisations without endangering other adaptivity gains (as discussed by Pain et al (2005) for unstructured models).

Whether adaptivity coupled with a pressure solver for non-hydrostatic flows can maximise the available computational resources needs further study. A stumbling block is likely to be the convergence properties of the pressure solver in the presence of excessive grid aspect ratios, since typical geophysical systems have horizontal scales much larger than in the vertical. This issue has recently been discussed by Matsumura and Hasumi (2008) who suggest that there are potential acceleration techniques that can make non-hydrostatic multi-grid Poisson inversions competitive with their hydrostatic equivalents; this provides further encouragement as a multi-grid pressure solver is integral to Gerris.

In this paper, density is treated as the only independent tracer. Future tests with Gerris will need to examine more closely an oceanic equation of state linking potential temperature and salinity tracers to the in-situ density. Furthermore, as the model is applied to larger geophysical systems, the potential for inaccuracies in the discretisation of the horizontal hydrostatic pressure gradient (especially across a grid that is adapting in the vertical) may exist and will need to be quantified. In the present laboratory scale experiments the influence of such errors is small; maintaining accuracy will become more critical for larger systems.

Acknowledgements

This work was funded by the Foundation for Research, Science, and Technology, New Zealand, contract C01X0701. We are also grateful to two anonymous referees for their constructive critique leading to a greatly improved manuscript.

Appendix

A Numerical discretisation of adaptive model

The details of the numerical scheme are described in Popinet (2003) and we will only give a summary of the overall procedure.

A.1 Temporal discretisation

A staggered in time discretisation of the density and pressure leads to the following formally second-order accurate time discretisation of the momentum, continuity, and density equations (1), (2) and (3), respectively,

$$\frac{\boldsymbol{U}_{n+1}-\boldsymbol{U}_n}{\Delta t}+\boldsymbol{U}_{n+\frac{1}{2}}\cdot\boldsymbol{\nabla}\boldsymbol{U}_{n+\frac{1}{2}}=-\frac{1}{\rho_0}\boldsymbol{\nabla}p_{n+\frac{1}{2}}-\frac{\rho'_{n+\frac{1}{2}}}{\rho_0}\boldsymbol{g}+\frac{\nu}{2}\nabla^2\left(\boldsymbol{U}_n+\boldsymbol{U}_{n+1}\right), \tag{A.1}$$

$$\frac{\rho'_{n+\frac{1}{2}}-\rho'_{n-\frac{1}{2}}}{\Delta t}+\boldsymbol{\nabla}\cdot\left(\rho'_n\boldsymbol{U}_n\right)=0, \tag{A.2}$$

$$\boldsymbol{\nabla}\cdot\boldsymbol{U}_n=0. \tag{A.3}$$

This system is further simplified using a classical time-splitting projection method (Chorin, 1968) to give,

$$\frac{\boldsymbol{U}_\star-\boldsymbol{U}_n}{\Delta t}+\boldsymbol{U}_{n+\frac{1}{2}}\cdot\boldsymbol{\nabla}\boldsymbol{U}_{n+\frac{1}{2}}=-\frac{\rho'_{n+\frac{1}{2}}}{\rho_0}\boldsymbol{g}+\frac{\nu}{2}\nabla^2\left(\boldsymbol{U}_n+\boldsymbol{U}_\star\right), \tag{A.4}$$

$$\frac{\rho'_{n+\frac{1}{2}}-\rho'_{n-\frac{1}{2}}}{\Delta t}+\boldsymbol{\nabla}\cdot\left(\rho'_n\boldsymbol{U}_n\right)=0, \tag{A.5}$$

$$\boldsymbol{U}_{n+1}=\boldsymbol{U}_\star-\frac{\Delta t}{\rho_0}\boldsymbol{\nabla}p_{n+\frac{1}{2}}, \tag{A.6}$$

$$\boldsymbol{\nabla}\cdot\boldsymbol{U}_{n+1}=0, \tag{A.7}$$

which requires the solution of the Poisson equation

$$\boldsymbol{\nabla}\cdot\left[\frac{\Delta t}{\rho_0}\boldsymbol{\nabla}p_{n+\frac{1}{2}}\right]=\boldsymbol{\nabla}\cdot\boldsymbol{U}_\star \tag{A.8}$$

This equation is solved efficiently using the quad/octree-based multilevel solver described in detail in Popinet (2003).

The discretised momentum equation (A.4) can be re-organised as

$$\frac{1}{\Delta t}\boldsymbol{U}_\star-\frac{\nu}{2}\nabla^2\boldsymbol{U}_\star=\frac{\nu}{2}\nabla^2\boldsymbol{U}_n-\frac{\rho'_{n+\frac{1}{2}}}{\rho_0}\boldsymbol{g}+\frac{1}{\Delta t}\boldsymbol{U}_n-\boldsymbol{U}_{n+\frac{1}{2}}\cdot\boldsymbol{\nabla}\boldsymbol{U}_{n+\frac{1}{2}}, \tag{A.9}$$

where the right-hand-side depends only on values at time n and $n+1/2$. This Helmholtz-type equation for $\boldsymbol{U}_\star$ can be solved efficiently using a variant of the multilevel Poisson solver. The resulting Crank–Nicolson discretisation of the viscous terms is formally second-order accurate and unconditionally stable.

The velocity advection term $\boldsymbol{U}_{n+\frac{1}{2}} \cdot \boldsymbol{\nabla}\boldsymbol{U}_{n+\frac{1}{2}}$ and the density advection term $\boldsymbol{\nabla} \cdot (\rho'_n \boldsymbol{U}_n)$ are estimated using the Bell–Colella–Glaz second-order unsplit upwind scheme (Bell et al 1989, Popinet 2003). This scheme is stable for CFL numbers smaller than one.

A.2 Spatial discretisation

Space is discretised using a graded quadtree partitioning (octree in three dimensions). Second order accuracy is maintained in spatial operators, including all boundaries. We refer the reader to Popinet (2003) and references therein for a more detailed presentation of this data structure.

All the variables (components of the momentum, pressure and passive tracers) are collocated at the centre of each square in 2D (resp. cubic in 3D) discretisation volume. Consistent with a finite-volume formulation, the variables are interpreted as the volume-averaged values for the corresponding discretisation volume. The choice of a collocated definition of all variables makes momentum conservation simpler when dealing with mesh adaptation (Popinet, 2003). It is also a necessary choice in order to use the Godunov momentum advection scheme of Bell, Colella and Glaz (Bell et al 1989), and it simplifies the implementation of the Crank–Nicolson discretisation of the viscous terms; however one has to be careful to avoid the classic problem of decoupling of the pressure and velocity field.

To do so, an approximate projection method (Almgren et al 2000, Popinet 2003) is used for the spatial discretisation of the pressure correction equation (A.6) and the associated divergence in the Poisson equation (A.8). In a first step the auxiliary cell-centred velocity field $\boldsymbol{U}^c_\star$ is computed using equation (A.9). An auxiliary face-centred velocity field $\boldsymbol{U}^f_\star$ is then computed using averaging of the cell-centred values on all the faces of the Cartesian discretisation volumes. When faces are at the boundary between different levels of refinement of the quad/octree mesh, averaging is performed so as to guarantee consistency of the corresponding volume fluxes (see Popinet (2003) for details).

The divergence of the auxiliary velocity field appearing on the right-hand-side of equation (A.8) is then computed for each control volume as the finite-volume

approximation

$$\boldsymbol{\nabla} \cdot \boldsymbol{U}_\star = \frac{1}{\Delta} \sum_f \boldsymbol{U}_\star^f \cdot \boldsymbol{n}^f, \tag{A.10}$$

with $\boldsymbol{n}^f$ the unit normal vector to the face and Δ the length scale of the control volume.

After solving equation (A.8), the pressure correction is applied to the face-centred auxiliary field

$$\boldsymbol{U}_{n+1}^f = \boldsymbol{U}_\star^f - \frac{\Delta t}{\rho_0} \boldsymbol{\nabla}^f p_{n+\frac{1}{2}}, \tag{A.11}$$

where $\boldsymbol{\nabla}^f$ is a simple face-centred gradient operator (consistent at coarse/fine volume boundaries, see section 4.1 of Popinet 2003). The resulting face-centred velocity field $\boldsymbol{U}_{n+1}^f$ is exactly non-divergent by construction. This velocity field is used to compute the momentum and density advection terms $\boldsymbol{U}_{n+\frac{1}{2}} \cdot \boldsymbol{\nabla} \boldsymbol{U}_{n+\frac{1}{2}}$ and $\boldsymbol{\nabla} \cdot (\rho_n' \boldsymbol{U}_n)$ so that the resulting scheme conserves mass and momentum exactly.

The cell-centred velocity field at time $n+1$ is obtained by applying a cell-centred pressure correction

$$\boldsymbol{U}_{n+1}^c = \boldsymbol{U}_\star^c - \left| \frac{\Delta t}{\rho_0} \boldsymbol{\nabla}^f p_{n+\frac{1}{2}} \right|^c, \tag{A.12}$$

where the $||^c$ operator denotes averaging over all the faces delimiting the control volume. The resulting cell-centred velocity field $\boldsymbol{U}_{n+1}^c$ is approximately divergence-free.

A.3 Representation of solid boundaries

Solid boundaries are represented using the *Cartesian cut-cell* scheme described in Popinet (2003). This scheme uses surface- and volume- fluid fractions to obtain a spatially second-order accurate representation of the solid boundaries irrespective of their orientation relative to the underlying Cartesian discretisation. This scheme has been shown in Popinet (2003), and Popinet and Rickard (2007) to compare very favourably with the simpler first-order accurate *masking* commonly used in Cartesian ocean models.

A.4 Adaptivity

The overall scheme allows for space and time-varying spatial resolution. To simplify the implementation the sizes of neighbouring cells can not vary by

more than a factor of two (this is sometimes referred to as *restricted or graded quad/octree*).

One of the advantages of the quad/octree disctretisation is that mesh refinement or coarsening are cheap and can be performed at every timestep if necessary with a minimal impact on overall performance (usually less than 1% of computation time). Interpolation of quantities on newly refined or coarsened cells is also relatively simple and is done conservatively both for momentum and density.

Several refinement criteria can be used simultaneously depending on the problem. Combinations of the following criteria have been used in this article:

Vorticity

$$\frac{|\boldsymbol{\nabla} \times \boldsymbol{U}|\Delta}{\max(|\boldsymbol{U}|)} < \epsilon, \tag{A.13}$$

with Δ the mesh size and ϵ a user-defined threshold. This criterion will ensure that a finer resolution is used in areas of high vorticity. The threshold parameter ϵ can be interpreted as the maximum angular deviation from a straight path – due to the local vorticity $\boldsymbol{\nabla} \times \boldsymbol{U}$ – of a massless particle travelling at speed $\max(|\boldsymbol{U}|)$ across a cell of size Δ. The threshold ϵ is usually set to a small value, typically 10^{-2}.

Gradient

$$|\boldsymbol{\nabla} c|\,\Delta < \epsilon, \tag{A.14}$$

with c a field variable. Note that the value of ϵ can be chosen differently if required for each adaptative criteria.

REFERENCES

Almgren, A. S., J. B. Bell, and W. Y. Crutchfield, 2000. Approximate projection methods: Part I. Inviscid analysis, *SIAM J. Sci. Comput.*, **22**, pp1139-1159.

Behrens, J., 2005. Adaptive atmospheric modeling: scientific computing at its best, *Comput. Sci. Eng.*, **7**, pp 7683.

Bell, J. B., P. Colella, and H. M. Glaz, 1989. A second-order projection method for the incompressible Navier-Stokes equations, *J. Comput. Phys.*, **85**, pp 257-283.

Berntsen, J., 2000. USERS GUIDE for a modesplit σ-coordinate numerical ocean model. Technical Report 135, Department of Applied Mathematics, University of Bergen, Johs. Bruns gt. 12, N-5008, Bergen, Norway, 48pp.

Berntsen, J., J. Xing, and, G. Alendal, 2006. Assessment of non-hydrostatic ocean models using laboratory scale problems. *Continental Shelf Research*, **26**, pp 1433-1447.

Blayo, E., Debreu, L., 1999. Adaptive mesh refinement for finite difference ocean models: first experiments. *Journal of Physical Oceanography*, **29**, pp 12391250.

Bogucki, D., Garrett, C., 1993. A simple model for the shear-induced decay of an internal solitary wave, *Journal of Physical Oceanography*, **23**, pp 17671776.

Bourgault, D., Kelley, D.E., 2003. Wave-induced boundary mixing in a partially mixed estuary. *Journal of Marine Research*, **61**, pp 553576.

Bourgault, D., and, D. E. Kelley, 2004. A laterally averaged nonhydrostatic ocean model, *Journal of Atmospheric and Oceanic Technology*, **21**, pp 1910-1924.

Bourgault, D., M. D. Blokhina, R. Mirshak, and, D. E. Kelley, 2007. Evolution of a shoaling internal solitary wavetrain, *Geophysical Research Letters*, **34**, L03601, doi:10.1029/2006GL028462.

Bourgault, D., and, D. E. Kelley, 2007. On the reflectance of uniform slopes for normally incident interfacial solitary waves, *Journal of Physical Oceanography*,

37, pp 1156-1162.

Burchard, H., and, J-M. Beckers, 2004. Non-uniform adaptive vertical grids in one-dimensional numerical ocean models, *Ocean Modelling*, **6**, pp 51-81.

Chorin, A., 1968. Numerical solution of the Navier-Stokes equations. *Mathematics of Computation*, **22** (104), pp 745-762.

Hanert, E., E. Deleersnijder, and, V. Legat, 2006. An adaptive finite element water column model using the MellorYamada level 2.5 turbulence closure scheme, *Ocean Modelling*, **12**, pp 205-223.

Härtel, C., E. Meiburg and F. Necker, 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**, pp 189-212.

Helfrich, K.R., 1992. Internal solitary wave breaking and run-up on a uniform slope, *Journal of Fluid Mechanics*, **243**, pp 133-154.

Helfrich, K.R., and W.K. Melville, 2006. Long nonlinear internal waves, *Annu. Rev. Fluid Mech.*, **38**, pp 395-425.

Iribarren, C.R., Nogales, C., 1949. Protection des ports. Section II, Comm. 4, XVIIth Int. Nav. Congress, Lisbon.

Marshall, J., A. Adcroft, C. Hill, L. Perelman, C. Heisey, 1997. A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers, *Journal of Geophysical Research*, **102(C3))**, pp 5753-5766.

Matsumura, Y., and H. Hasumi, 2008. A non-hydrostatic ocean model with a scalable multigrid Poisson solver, *Ocean Modelling*, **24**, pp 15-28.

Michallet, H., and, G. N. Ivey, 1999. Experiments on mixing due to internal solitary waves breaking on uniform slopes, *Journal of Geophysical Research*, **104**(C6), pp13467-13477.

Pain, C.C., M.D. Piggott, A.J.H. Goddard, F. Fang, G.J. Gorman, D.P. Marshall, M.D. Eaton, P.W. Power, C.R.E. de Oliveira, 2005. Three-dimensional unstructured mesh ocean modelling, *Ocean Modelling*, **10**, pp 5-33.

784 Piggott, M.D., G.J. Gorman, C.C. Pain, P.A. Allison, A.S. Candy, B.T. Mar-
785 tin, and M.R. Wells, 2008. A new computational framework for multi-scale
786 ocean modelling based on adapting unstructured meshes, *Int. J. Numer. Meth.*
787 *Fluids*, **56**, pp 10031015.

788 Popinet, S., 2003. A tree-based adaptive solver for the incompressible Euler
789 equations in complex geometries, *Journal of Computational Physics*, **190**(2),
790 pp572-600.

791 Popinet, S., 2006. The Gerris Flow Solver, http://gfs.sourceforge.net.

792 Popinet, S., and, G.J. Rickard, 2007. A tree-based solver for adaptive ocean
793 modelling, *Ocean Modelling*, **16**, p224-249.

794 Segur, H., Hammack, J.L., 1982. Soliton models of long internal waves, *Journal*
795 *of Fluid Mechanics*, **118**, pp 285304.

796 Simpson, A.E., 1972. Effects of the lower boundary on the head of a gravity
797 current, *J. Fluid Mech.*, **53**, pp 759-768.

798 Simpson, A.E., 1997. Gravity Current, Cambridge University Press.

799 Venayagamoorthy, S. K., and, O. B. Fringer, 2007. On the formation and
800 propagation of nonlinear internal boluses across a shelf break, *Journal of Fluid*
801 *Mechanics*, **577**, pp 137-159.

LIST OF FIGURES

(1) Experimental and model set-up. Horizontal dashed line is initial undisturbed centre of the pycnocline, with solid line illustrating initial conditions for model perturbed pycnocline. Slope gradient is "s", such that $s = H/L_S = d_-/L$. For reference, horizontal distance from slope bottom to left hand domain edge is L_R.

(2) Grid layout on slope for standard run (see text). Largest squares show volumes associated with coarsest grid at level of resolution of 7. As the boundary set here by the slope (the solid line) is approached, the resolution is increased by halving the grid size until the maximum resolution (here level 9) is reached. Gerris is finite volume, and accounts for cell changes in volume where cut by the boundary in order to return the correct boundary fluxes.

(3) Density contours at time 7.5 s for viscosity and diffusivity values of (a) 6.75×10^{-5}, (b) 1.0×10^{-6}, (c) 1.5×10^{-7}, and (d) $1.5 \times 10^{-8}\,\mathrm{m^2 s^{-1}}$ for model experiment "d" in Table 1 in the (x, z)-plane. The height of each frame is 0.075 m.

(4) Density contours at time 7.5 s for (a) standard run (see text), (b) adapting to a maximum level of resolution of 12, and (c) frame(b) - frame(a) in the (x, z)-plane. The height of each frame is 0.075 m.

(5) Contours of resolution levels for frame (b) in Fig. 4 from lowest resolution at level 7 to finest resolution at level 12.

(6) Contours at time 37.5 s for model experiment 2 in Table 1 of (a) density, (b) x component of velocity, and (c) vorticity in the (x, z)-plane, with contour intervals of (a) $4.0\mathrm{kgm^{-3}}$, (b) $0.01\mathrm{ms^{-1}}$, and (c) $4.67\mathrm{s^{-1}}$. The height of each frame is 0.072 m. Vertical scale is four times the horizontal scale for clarity. Vertical lines locate root cells at horizontal spacing of 0.15 m, and horizontal line inside model domain locates initial undisturbed pycnocline.

(7) Multiple bolus trajectories plotted in terms of its position relative to the pycnocline-slope intersection ($x/X_{BP} = 0$) as a function of time for model experiment 2 (see Table 1). Four boluses are shown moving upslope, the first identified with crosses, the second squares, the third and fourth triangles and diamonds, respectively. Horizontal dashed line locates horizontal position of undisturbed pycnocline-slope intersection.

(8) Density contours for first bolus for model experiment 2 in Table 1 at (times (seconds), normalised positions x/X_{BP}) (a) (22.5, -0.28), (b) (27.0, -0.07), (c) (30.0, 0.06), and (d) (37.5, 0.31), respectively. Bolus trajectory shown as crosses in Figure 7. Contour interval for density in $\mathrm{kgm^{-3}}$ for frames (a,b) is 4.0, for frame (c) is 3.0, and for (d) is 2.0. Maximum

(density, bolus height, length, speed) is (a) (1047, 0.03, 0.05, 0.11), (b) (1035, 0.022, 0.033, 0.084), (c) (1032, 0.017, 0.03, 0.07), and (d) (1017, 0.0144, 0.024, 0.047) (kgm^{-3}, m, m, ms^{-1}), respectively. Bold lines on frame (b) indicate bolus height and length.

(9) Vorticity contours and velocity vectors for first bolus for model experiment 2 in Table 1. Velocity vectors scaled proportional to the maximum speed in each frame, and the contour interval for vorticity is $5.0s^{-1}$. All other details for frames (a)-(d) are the same as in the caption for Fig. 8. Bold lines on frame (b) indicate bolus height and length.

(10) Breaking location (criterion) as a function of λ. Characters identify individual experiments 1 to 4 and a to i in Tables 1 and 2.

(11) Number of boluses in a breaking event as a function of λ. Characters identify individual experiments 1 to 4 and a to i in Tables 1 and 2.

(12) Initial height of first bolus scaled by a_0 (maximum displacement of the interface) as a function of λ. Characters identify individual experiments 1 to 4 and a to i in Tables 1 and 2.

(13) First bolus aspect ratio H/l_B (frames (a)-(c)) and first bolus height (frames (d)-(f)) as a function of upslope distance for model experiments as labelled in Tables 1 and 2. Dashed line in frames (d)-(f) indicates linear decay of bolus height with distance.

(14) Reflection coefficient R against Iribarren number ξ for the dots (model) and Michallet and Ivey (1999) experiments (crosses). Size of boxes and crosses at each point indicate estimated error bounds. The bold curve is the best fit to model by Bourgault and Kelley (2007) for the free slip sidewalls.

(15) Energy flux (W/m) against time (seconds) for model experiments (e) to (i) in Table 2.

Table 1
Initial conditions for eight test model experiments. For each of these model experiments $H = 0.15$ m, $L_R = 0.05$ m (see Figure 1), $\rho_0 = 1000\, kgm^{-3}$, and $a_i = 0.027$ m, $\Delta\rho = 47.0\, kgm^{-3}$, $\Delta h = 0.0035m$, $d_+ = z_i$, and $d_- = 0.15 - z_i$ (see equations 4 and 5). Pairs of experiments differ only in their initial z_i value, which changes subsequent a_0 and λ model values.

Model Label	slope	z_i (m)	a_0 model (cm)	λ model
1,a	0.034	0.05, 0.024	2.7, 2.6	0.053, 0.018
2,b	0.05	0.05, 0.024	2.8, 2.7	0.077, 0.026
3,c	0.128	0.05, 0.024	2.8, 2.6	0.198, 0.083
4,d	0.217	0.05, 0.024	3.0, 2.9	0.32, 0.12

Table 2
Model conditions for five selected experiment numbers (in brackets) from Michallet and Ivey (1999). For each of these model experiments H = 0.15 m, L_R = 1.0 m, (see Figure 1), $\rho_0 = 1000\, kgm^{-3}$, and $\Delta\rho$ is used in equation 4.

Model Label (Expt)	slope	z_i (m)	$\Delta\rho$ (kgm^{-3})	a_0 (cm)	L_w (cm)	ξ	λ
e (4)	0.069	0.033	20.0	2.8	31.0	0.23	0.054
f (8)	0.169	0.0255	40.0	2.6	28.0	0.55	0.098
g (12)	0.214	0.024	12.0	2.7	24.0	0.64	0.113
h (7)	0.169	0.0495	40.0	2.0	41.0	0.77	0.3
i (13)	0.214	0.06	44.0	1.8	46.0	1.08	0.64

Table 3
Experimental data for five selected experiment numbers from Michallet and Ivey (1999). Variation in experimental Iribaren number ξ is based on range of internal wave amplitudes a_0, and a 5% error estimate in L_w by Michallet and Ivey (1999).

Expt	a_0 (cm)	L_w (cm)	ξ
4	3.0	30.0	0.21 - 0.22
8	2.8	24.0	0.48 - 0.51
12	2.5 - 3.1	24.0	0.58 - 0.67
7	2.0	46.0	0.79 - 0.83
13	1.6 - 2.0	56.0	1.1 - 1.3

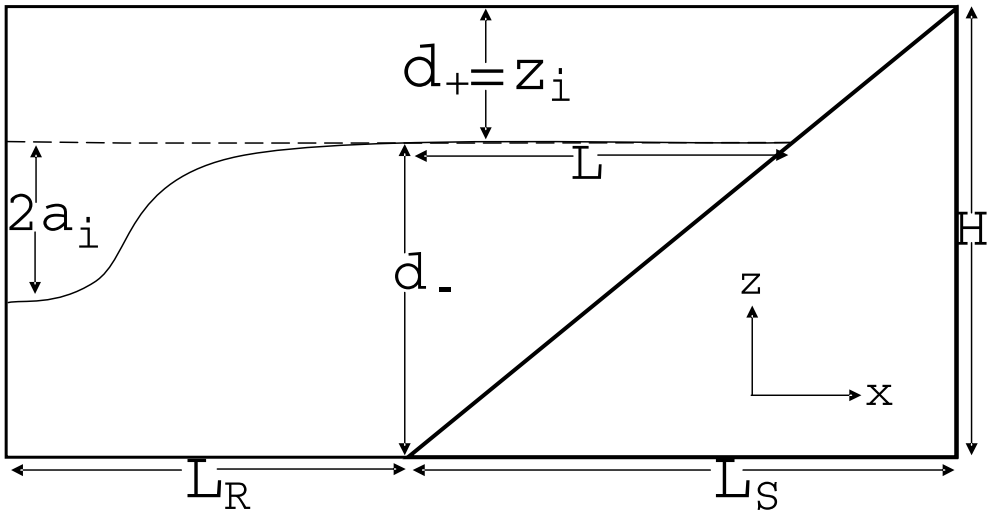

Fig. 1. Experimental and model set-up. Horizontal dashed line is initial undisturbed centre of the pycnocline, with solid line illustrating initial conditions for model perturbed pycnocline. Slope gradient is "s", such that $s = H/L_S = d_-/L$. For reference, horizontal distance from slope bottom to left hand domain edge is L_R.

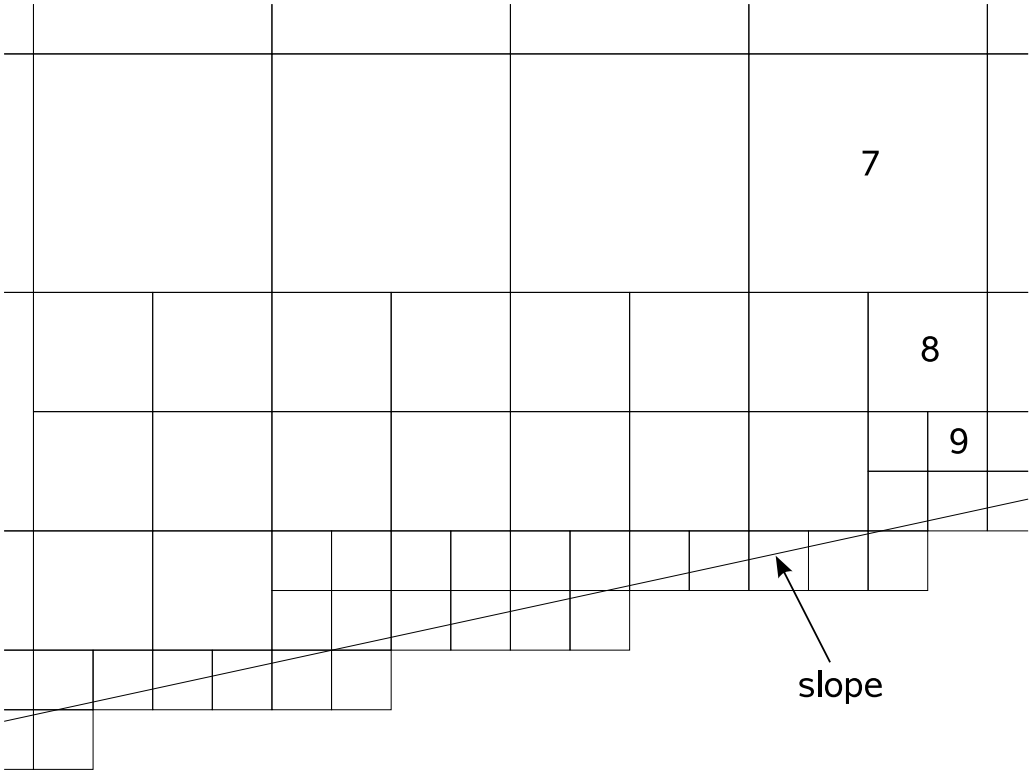

Fig. 2. Grid layout on slope for standard run (see text). Largest squares show volumes associated with coarsest grid at level of resolution of 7. As the boundary set here by the slope (the solid line) is approached, the resolution is increased by halving the grid size until the maximum resolution (here level 9) is reached. Gerris is finite volume, and accounts for cell changes in volume where cut by the boundary in order to return the correct boundary fluxes.

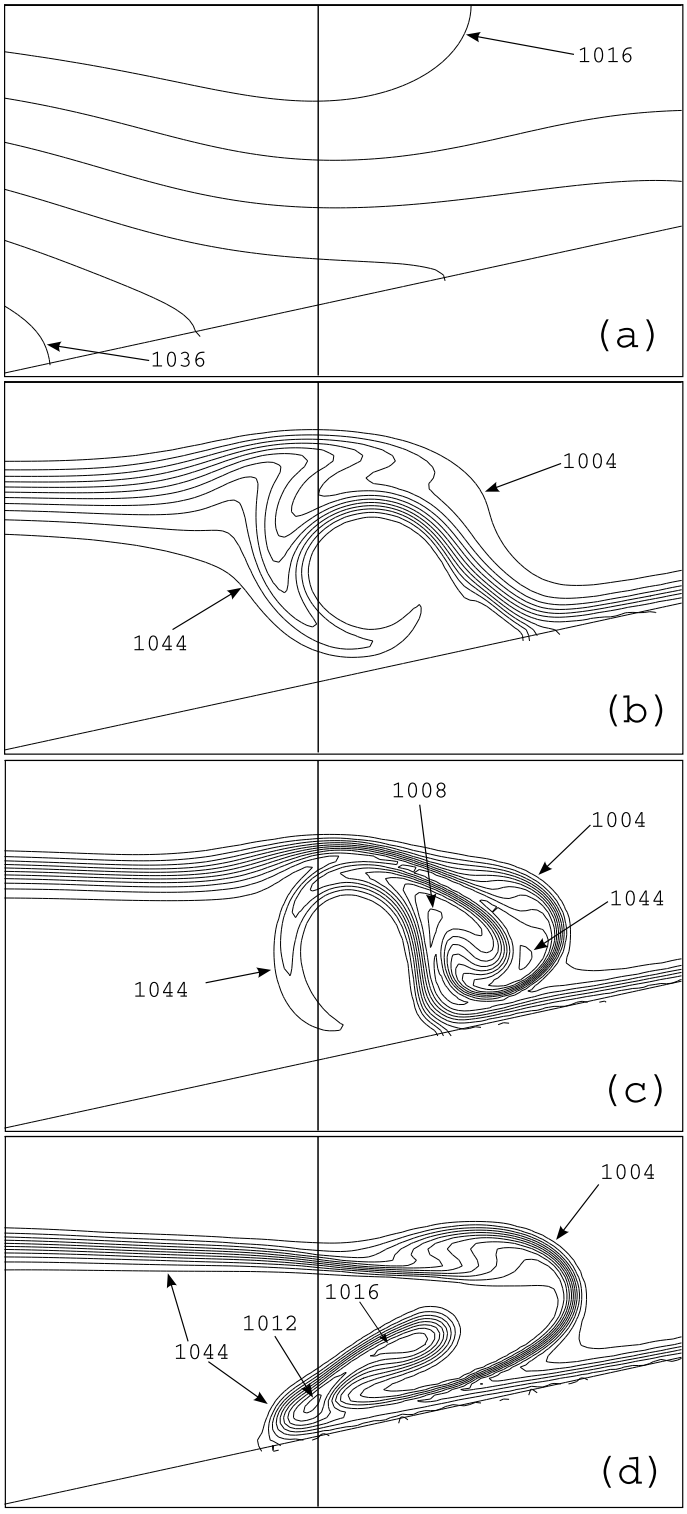

Fig. 3. Density contours at time 7.5 s for viscosity and diffusivity values of (a) 6.75×10^{-5}, (b) 1.0×10^{-6}, (c) 1.5×10^{-7}, and (d) $1.5 \times 10^{-8}\,\mathrm{m^2 s^{-1}}$ for model experiment "d" in Table 1 in the (x, z)-plane. The height of each frame is 0.075 m.

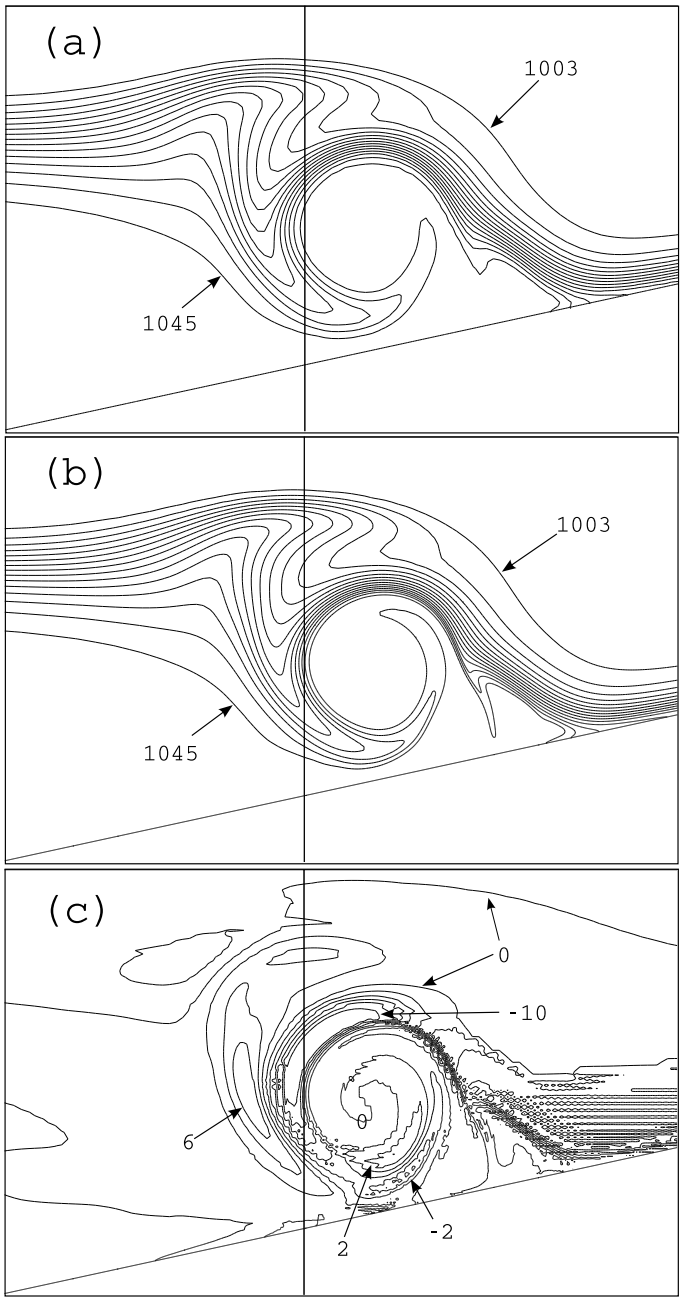

Fig. 4. Density contours at time 7.5 s for (a) standard run (see text), (b) adapting to a maximum level of resolution of 12, and (c) frame(b) - frame(a) in the (x, z)-plane. The height of each frame is 0.075 m.

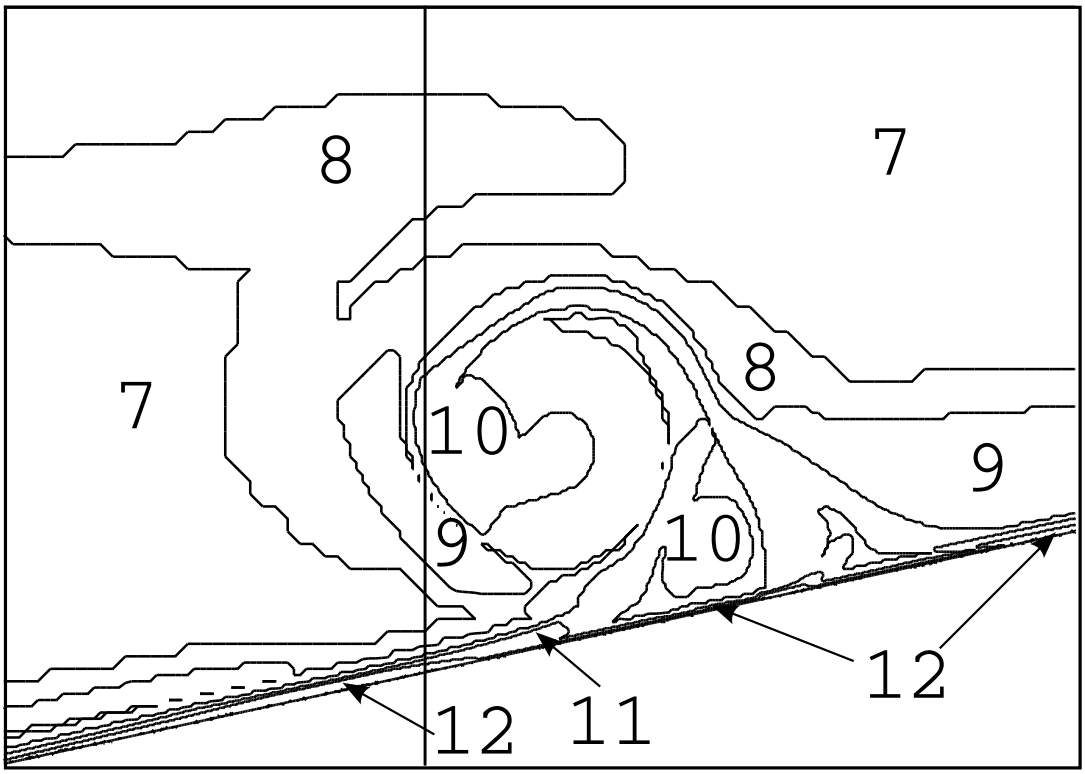

Fig. 5. Contours of resolution levels for frame (b) in Fig. 4 from lowest resolution at level 7 to finest resolution at level 12.

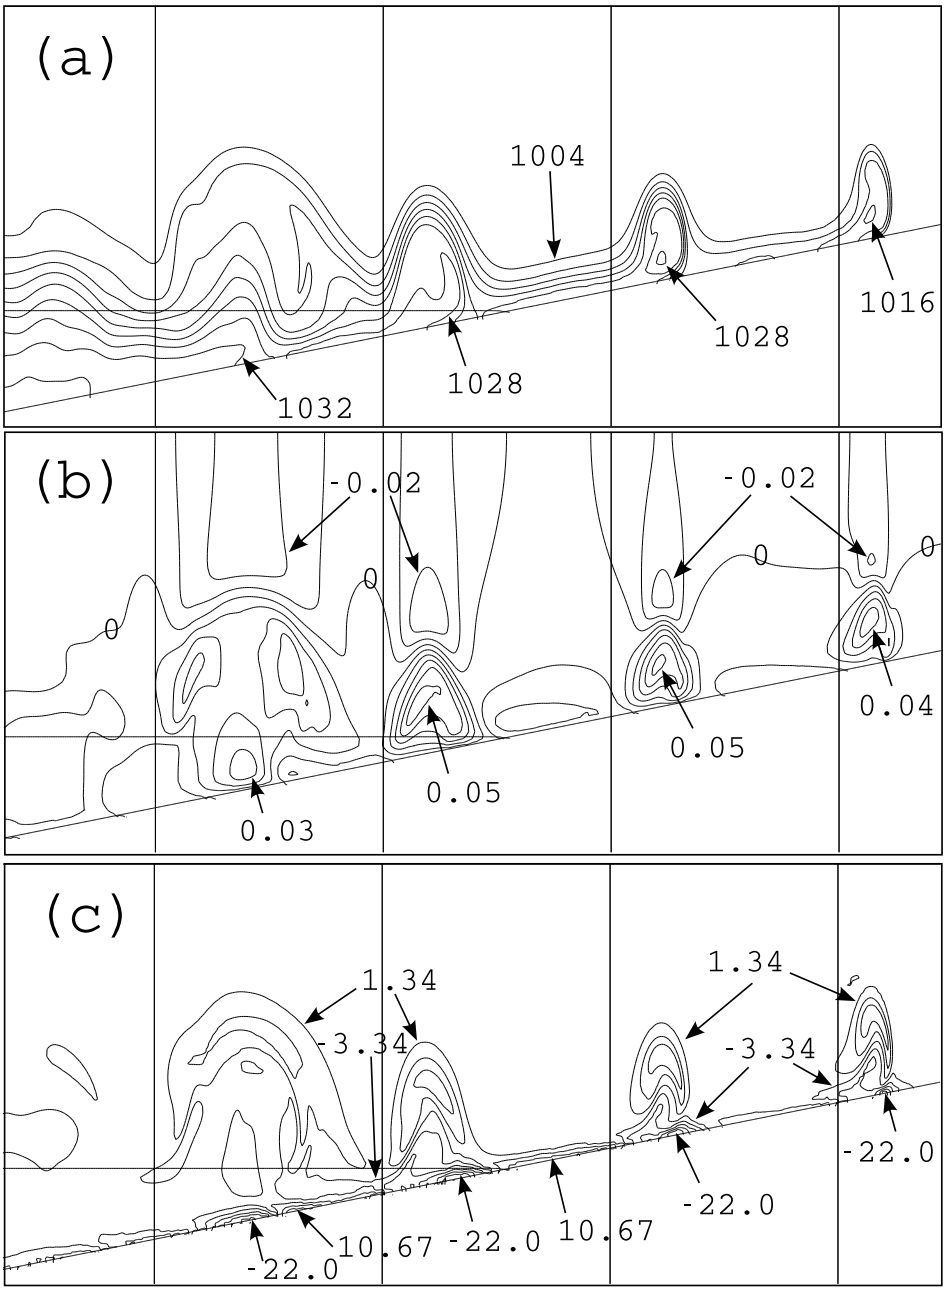

Fig. 6. Contours at time 37.5 s for model experiment 2 in Table 1 of (a) density, (b) x component of velocity, and (c) vorticity in the (x, z)-plane. The height of each frame is 0.072 m. Vertical scale is four times the horizontal scale for clarity. Vertical lines locate root cells at horizontal spacing of 0.15 m, and horizontal line inside model domain locates initial undisturbed pycnocline.

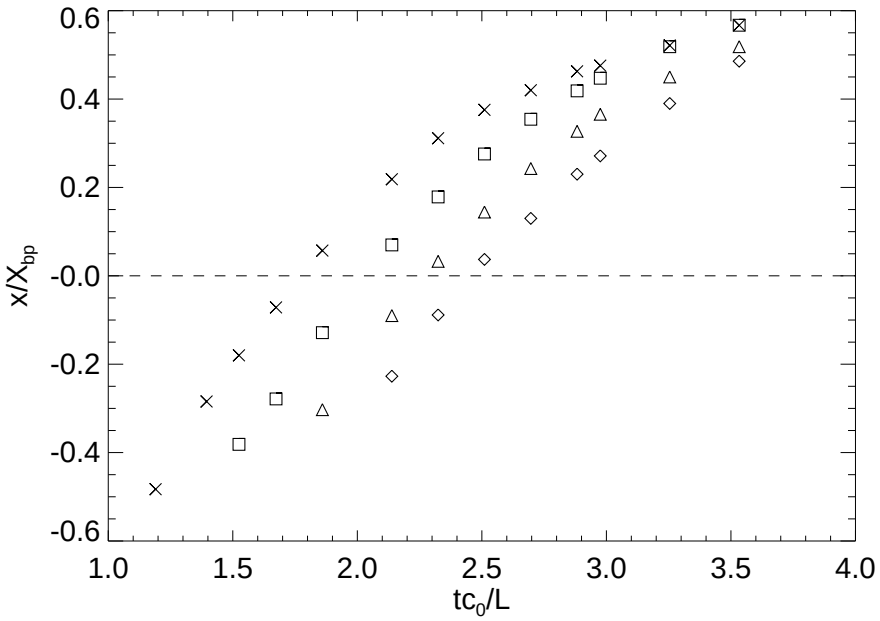

Fig. 7. Multiple bolus trajectories in time for model experiment 2 (see Table 1). Four boluses are shown moving upslope, the first identified with crosses, the second squares, the third and fourth triangles and diamonds, respectively. Horizontal dashed line locates horizontal position of undisturbed pycnocline-slope intersection.

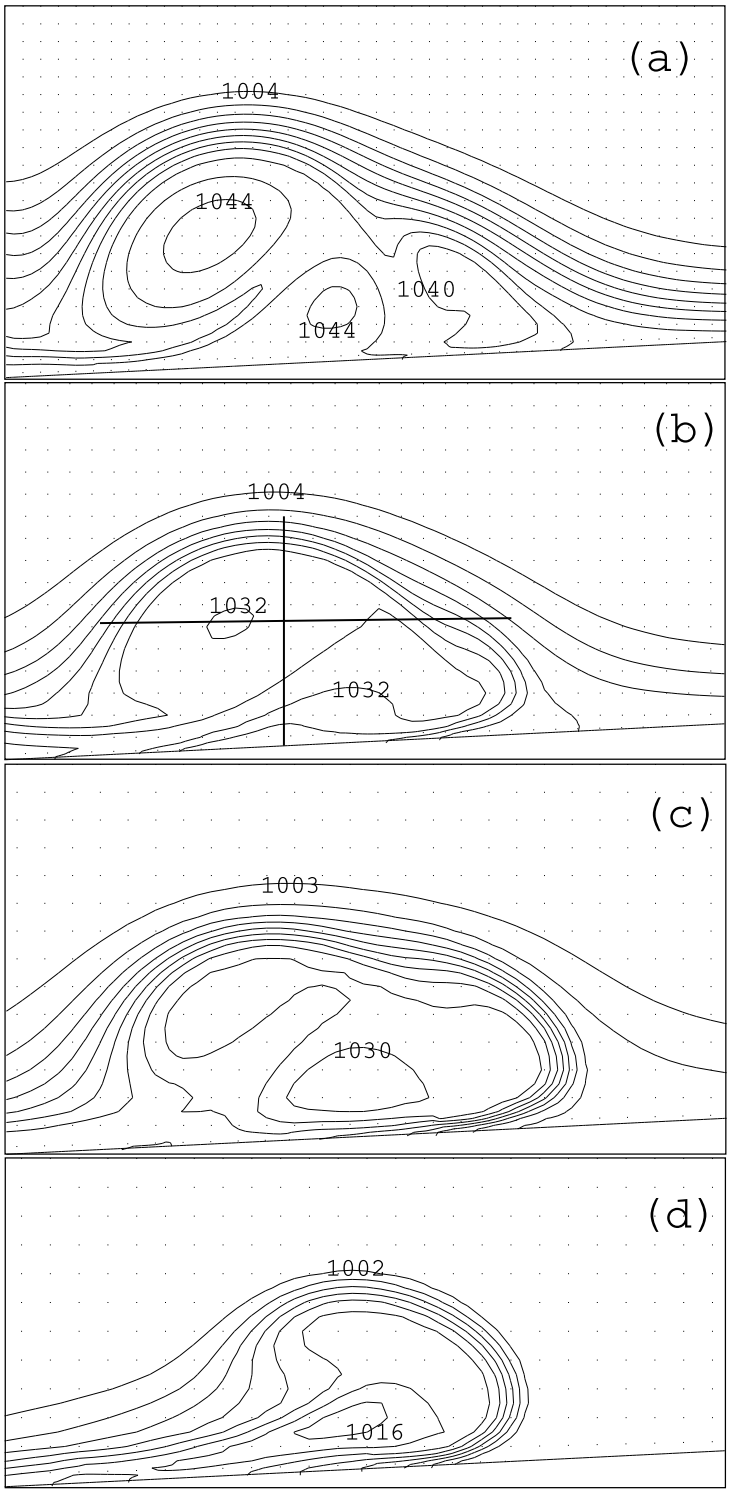

Fig. 8. Density contours for first bolus for model experiment 2 in Table 1 at (times (seconds), normalised positions x/X_{BP}) (a) (22.5, -0.28), (b) (27.0, -0.07), (c) (30.0, 0.06), and (d) (37.5, 0.31), respectively. Bolus trajectory shown as crosses in Figure 7. Contour interval for density in kgm^{-3} for frames (a,b) is 4.0, for frame (c) is 3.0, and for (d) is 2.0. Maximum (density, bolus height, length, speed) is (a) (1047, 0.03, 0.05, 0.11), (b) (1035, 0.022, 0.033, 0.084), (c) (1032, 0.017, 0.03, 0.07), and (d) (1017, 0.0144, 0.024, 0.047) (kgm^{-3}, m, m, ms^{-1}), respectively. Bold lines on frame (b) indicate bolus height and length.

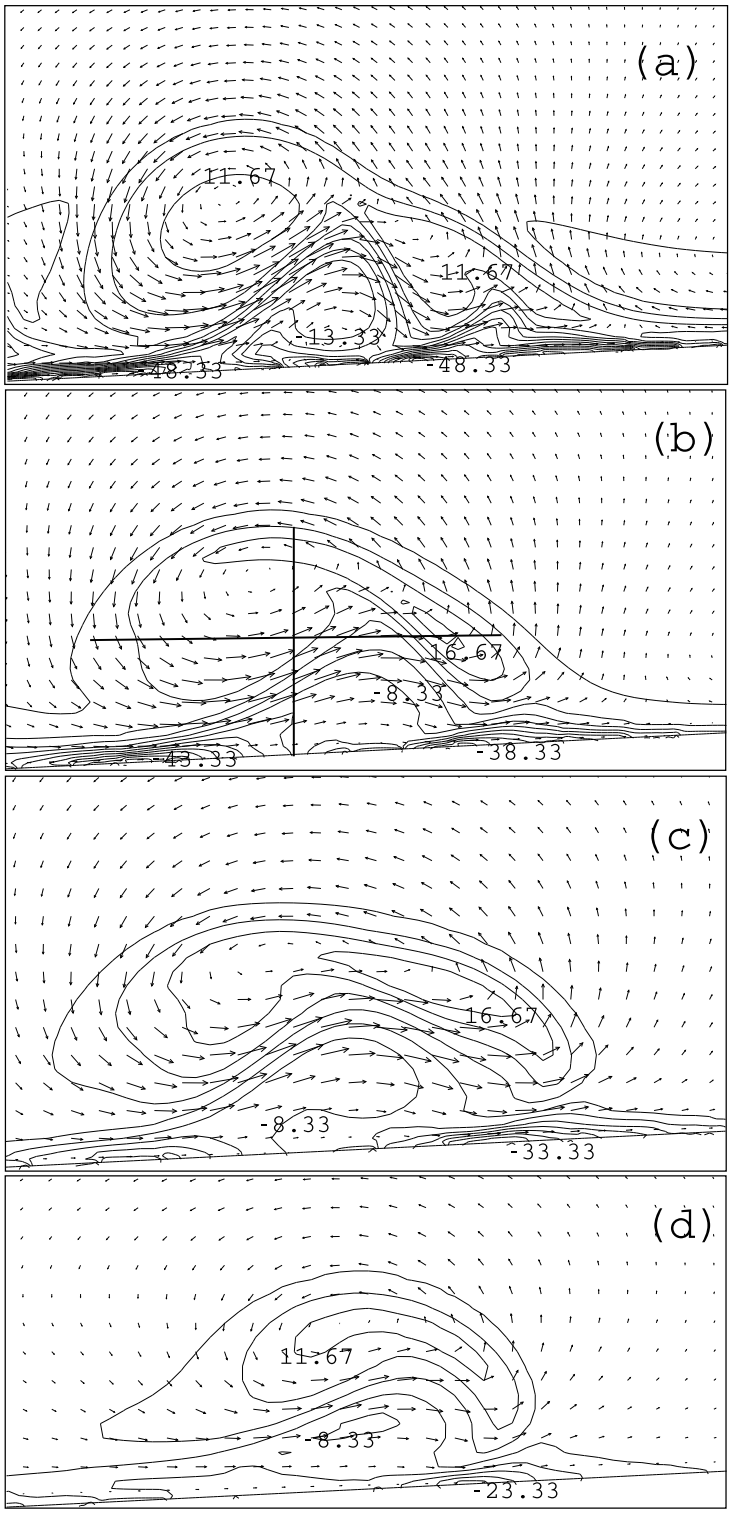

Fig. 9. Vorticity contours and velocity vectors for first bolus for model experiment 2 in Table 1. Velocity vectors scaled proportional to the maximum speed in each frame, and the contour interval for vorticity is $5.0s^{-1}$. All other details for frames (a)-(d) are the same as in the caption for Fig. 8. Bold lines on frame (b) indicate bolus height and length.

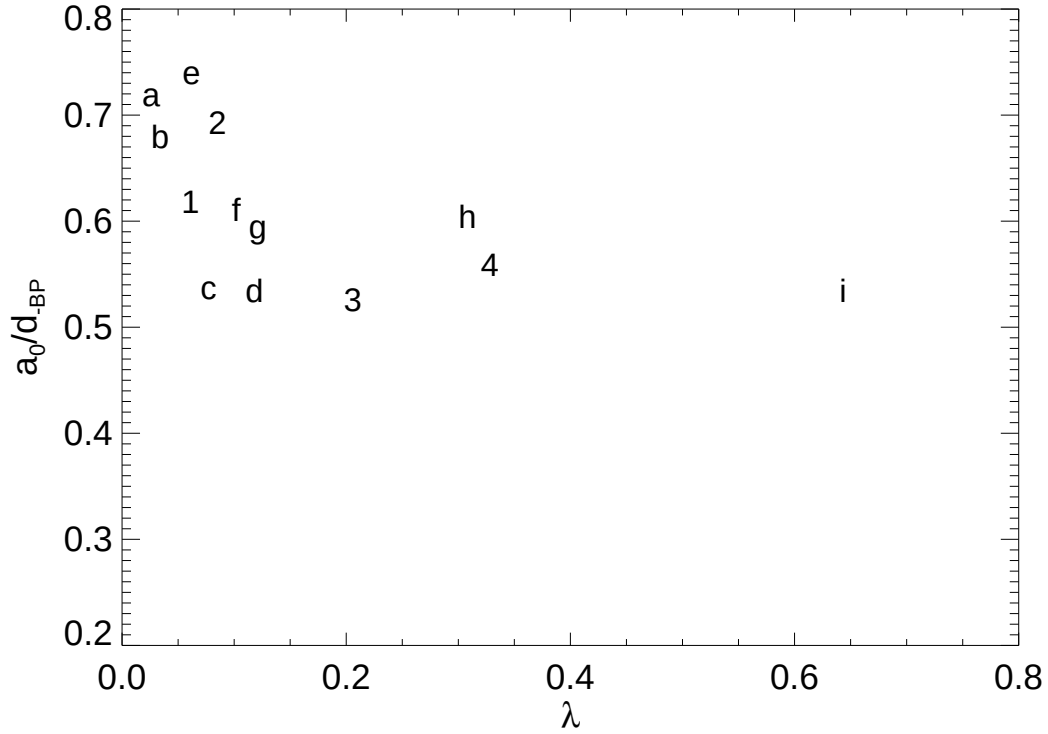

Fig. 10. Breaking location (criterion) as a function of λ. Characters identify individual experiments 1 to 4 and a to i in Tables 1 and 2.

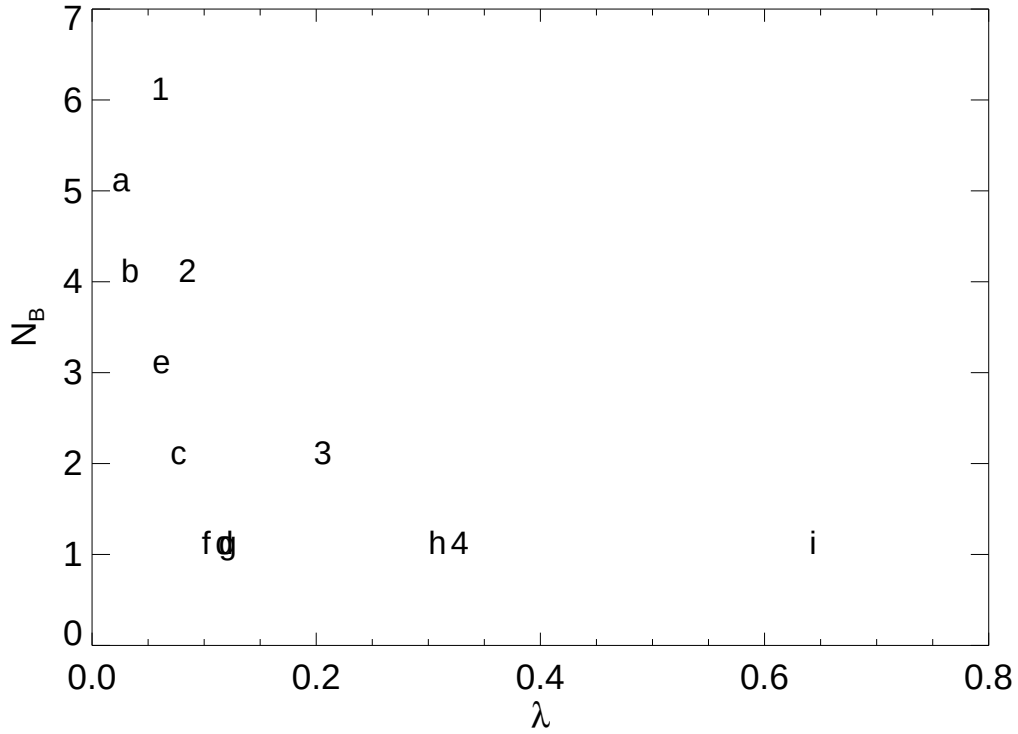

Fig. 11. Number of boluses in a breaking event as a function of λ. Characters identify individual experiments 1 to 4 and a to i in Tables 1 and 2.

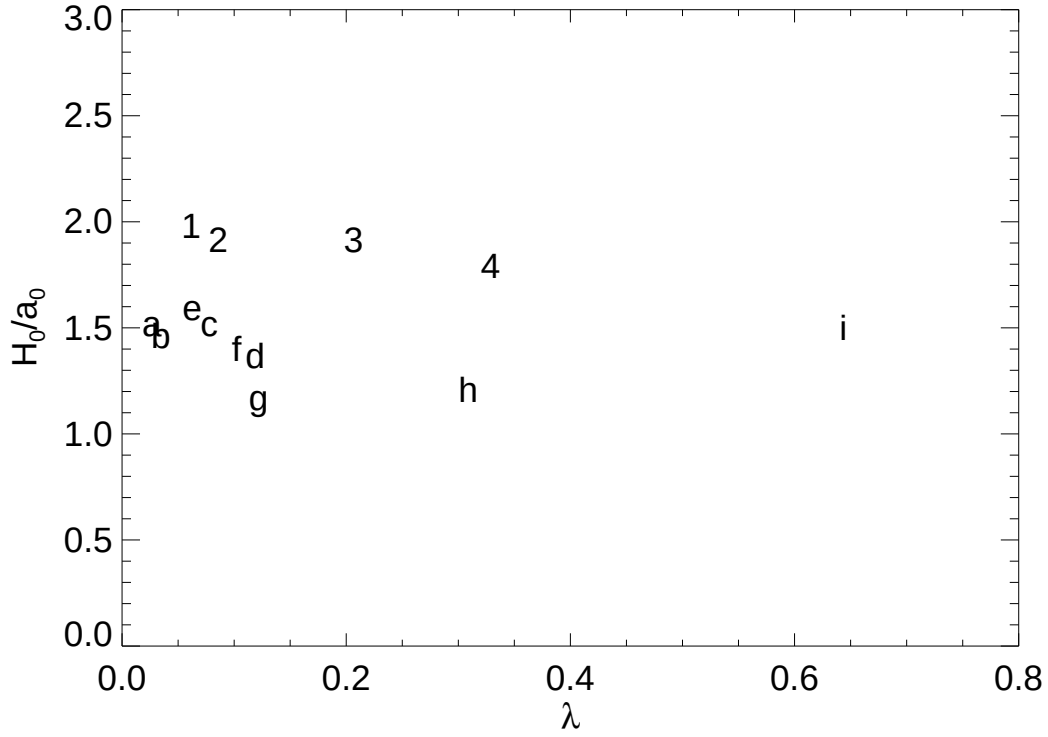

Fig. 12. Initial height of first bolus scaled by a_0 (maximum displacement of the interface) as a function of λ. Characters identify individual experiments 1 to 4 and a to i in Tables 1 and 2.

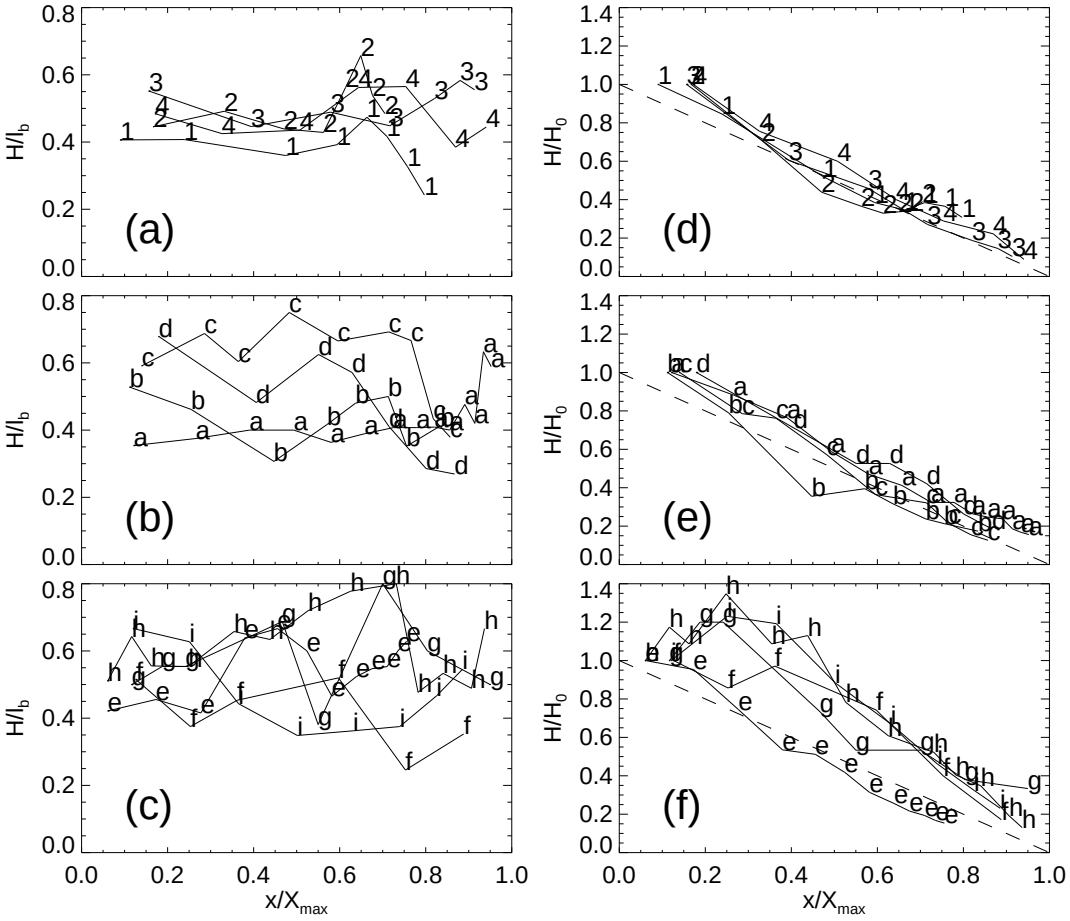

Fig. 13. First bolus aspect ratio H/l_B (frames (a)-(c)) and first bolus height (frames (d)-(f)) as a function of upslope distance for model experiments as labelled in Tables 1 and 2. Dashed line in frames (d)-(f) indicates linear decay of bolus height with distance.

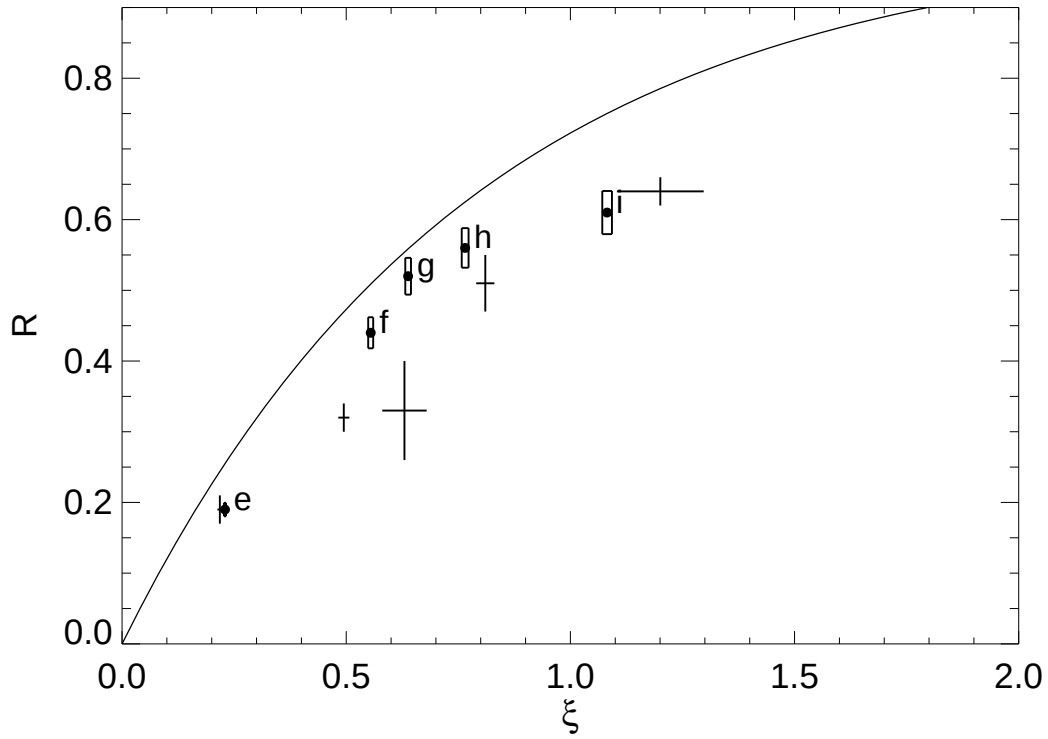

Fig. 14. Reflection coefficient R against Iribarren number ξ for the dots (model) and Michallet and Ivey (1999) experiments (crosses). Size of boxes and crosses at each point indicate estimated error bounds. The bold curve is the best fit to model by Bourgault and Kelley (2007) for the free slip sidewalls.

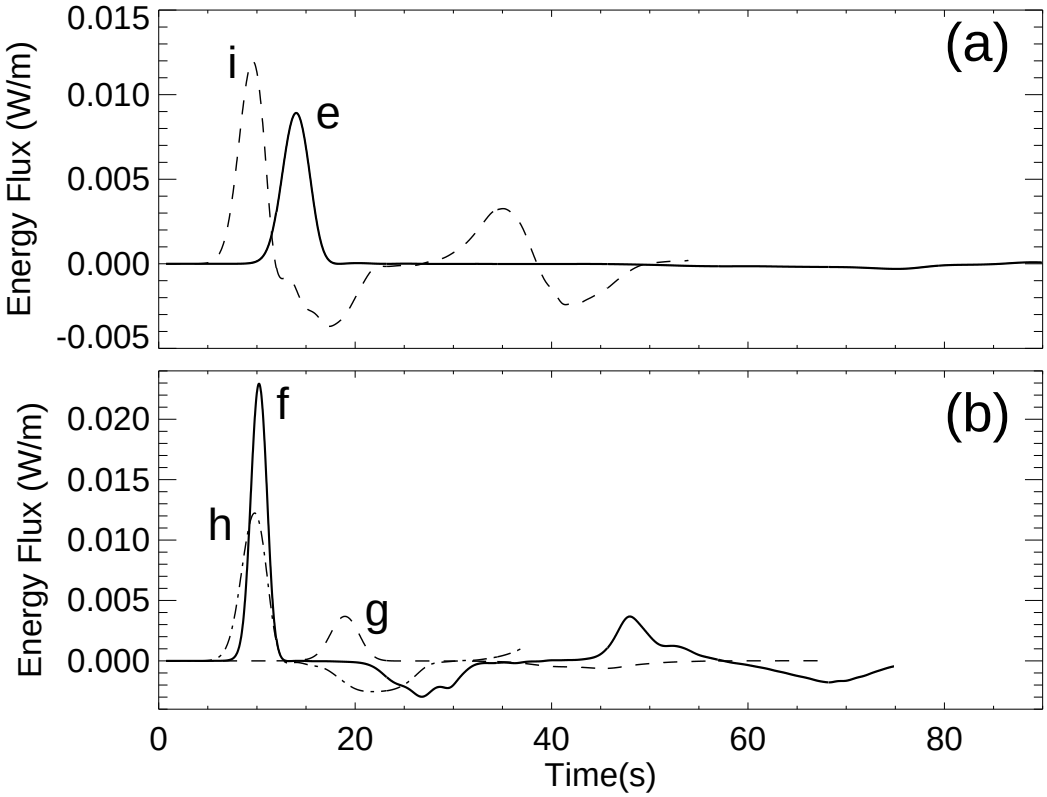

Fig. 15. Energy flux (W/m) against time (seconds) for model experiments (e) to (i) in Table 2.