Imperial College London
Earth Science and Engineering Department

Modelling the vertical mixing associated with bathymetry and internal waves in modern and ancient coastal oceans

Olivia Bentley

Supervised by:

Peter Allison and Matthew Piggott

Submitted in part fulfillment of the requirements for the degree of Doctor of Philosophy
and the Diploma of Imperial College London

January 21, 2018

DECLARATION OF ORIGINALITY

I hereby declare that I have properly acknowledged all material within this thesis which is not my own work. The work in this thesis has not been submitted in support of any other degree or qualification at this or any other university or institute of learning.

Olivia Louise Bentley

DECLARATION OF COPYRIGHT

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

ABSTRACT

Interactions of tidally-induced flows with complex bathymetry often result in vertical mixing and locally enhanced turbulence, in addition to the formation of internal waves which can propagate away and break elsewhere resulting in further vertical mixing. Vertical mixing is important both for ocean dynamics and also as it can increase biodiversity and productivity by facilitating nutrient exchange between the deep and surface ocean.

More generally, turbulent mixing in the ocean is an important, challenging and widely investigated matter. Understanding and modelling of the processes by which energy is converted from larger scales into small scale mixing, and how this impacts upon ocean circulation, present major challenges within the field.

In order to improve the understanding of the influence of bathymetry in shallow coastal oceans on mixing, two dimensional vertical slice numerical simulations are presented conducted with both Fluidity and MITgcm. These simulations analyse the effect of the interaction of tidal currents with complex bathymetry such as sills and flooded river valleys. Simulating vertical mixing processes explicitly generally requires very high numerical resolution to properly capture the dynamics and to diagnose correct mixing rates. A potential approach to mitigate this is considered through the application of mesh adaptivity for Loch Etive, which is shown to

reproduce dynamics well while using between one and two orders of magnitude fewer degrees of freedom than a fixed uniform mesh.

Idealised and more realistic simulations of the modern day real world fjord systems of Loch Etive and Idefjorden are first used to further validate model results and to better understand the processes which occur in these systems, including the influence of the second sill and the inter-sill spacing, before idealised simulations of the geological ancient are conducted. Targeted simulations of the geological ancient explore major stages in a transgression and regression during an icehouse and greenhouse world due to changing coastline geomorphology and the resulting nutrient heterogeneity.

ACKNOWLEDGMENTS

I would like to thank my supervisors Peter Allison and Matthew Piggott for their continued support and guidance. Mattias Green of Bangor University was also a great support hosting me at the university for collaboration on Idefjord for chapter 4. I would also like to thank the Grantham Institute for Climate Change and the Environment for their sponsorship, without which this work could not have been undertaken. In addition I wish to thank my fellow members of the Applied Modelling and Computation Group (AMCG) especially Alexandros Avdis who has provided support many times. And finally thank you to my husband, parents and friends.

CONTENTS

Declaration of originality i

Declaration of copyright ii

Abstract iii

Acknowledgments v

1 Introduction and Background 1

1.1 Introduction 1

1.2 Thermohaline circulation 2

1.3 Vertical mixing 4

1.4 Internal waves 8

1.5 Ocean modelling 10

1.5.1 Parameterisation 10

1.5.2 Ocean modelling using fixed and adaptive meshes 12

2 Numerical ocean models: Fluidity and MITgcm 15

2.1 Introduction 17

2.2 Governing Equations 18
2.2.1 Formation of the Navier-Stokes Equations 18
2.2.2 The Advection-Diffusion Equation 25
2.2.3 The incompressible, non-hydrostatic, Boussinesq, rotating Navier-Stokes equations 26
2.3 Meshes 27
2.3.1 Fixed Meshes - Vertical Co-ordinate Systems 27
2.3.2 Unstructured Meshes 28
2.4 The Fluidity model 29
2.4.1 Spatial Discretisation 29
2.4.2 Temporal Discretisation 36
2.4.3 Linear Solvers 38
2.4.4 Mesh Adaptivity 41
2.5 The MITgcm model 52
2.5.1 Sponge regions 52
2.5.2 Temporal discretisation 53
2.5.3 Spatial discretisation 53

3 Simulating mixing due to internal wave generation, propagation and breaking with adaptive mesh methods 57

3.1 Synopsis 59
3.2 Introduction 60
3.2.1 Sea Lochs and Fjords 60
3.3 Numerical Set-up 61
3.3.1 Fluidity discretisation 62
3.3.2 Mesh Adaptivity 63
3.4 Mixing Diagnostics 63
3.4.1 Critical Richardson area 63

3.4.2 Background potential energy and turbulent mean tracer variance decay rate . . . 65
3.4.3 Background Potential Energy Method . . . 66
3.4.4 Turbulent Mean Tracer Variance Decay Rate Method . . . 73
3.5 Simple test case: Internal seiche . . . 77
3.5.1 Set-up . . . 77
3.5.2 Results . . . 78
3.6 Lab validation: tilting-tank . . . 83
3.6.1 Set-up . . . 84
3.6.2 Results . . . 85
3.7 Pseudo real-world application: Loch Etive . . . 89
3.7.1 Set-up . . . 91
3.7.2 Initial benchmarking against Berntsen et al. (2009) . . . 94
3.7.3 Results . . . 100
3.8 Conclusions . . . 118

4 The influence of a second sill: a study of Idefjorden a two sill fjord 121
4.1 Synopsis . . . 123
4.2 Introduction . . . 124
4.2.1 Idefjord . . . 125
4.3 Methods . . . 127
4.3.1 Diagnostics . . . 130
4.4 Comparison with previous idealised studies . . . 132
4.5 Effect of Realistic bathymetry and stratification based on Idefjorden 134
4.6 Validation against real world data . . . 137
4.7 Tidal dynamics between two closely spaced sills in a narrow fjord: the development of a standing internal wave . . . 140
4.7.1 Realistic stratification and bathymetry for Idefjord . . . 140

4.7.2 Increasing the depth of the Svinesund sill top 143
4.7.3 Increasing the depth of the surface layer 144
4.7.4 Varied sill spacing and a deeper surface layer 146
4.8 Svinesund sill height change and bed shear stress 153
4.9 Discussion 154
4.10 Conclusions 159

5 Vertical mixing generated by interaction with bathymetry in Ancient Epicontinental Seas 163

5.1 Synopsis 164
5.2 Introduction 165
5.2.1 Epicontinental seas 166
5.2.2 Parameters to define complex bathymetry 168
5.3 Physical set-up 173
5.3.1 Diagnostics 176
5.4 Influence of the number of bumps 177
5.5 Transgression and Regression scenarios 183
5.6 Discussion and conclusions 188
5.7 Further Discussion – application to coral reef refugia 191

6 Conclusions 195

7 Bibliography 201

A The influence of the inclusion of the Coriolis force for Loch Etive 227

LIST OF FIGURES

1.2.1 Highly simplified illustration to show the thermohaline circulation from Rahmstorf (2002) 3
1.3.1 Schematic from Munk and Wunsch (1998) showing an impressionistic tidal energy flux budget 6
1.3.2 Illustration showing the basic stages of the biological pump based on Herndl and Reinthaler (2013) 8
1.4.1 Dissipation of M_2 tidal energy shown spatially, from Egbert and Ray (2003) achieved by satellite altimetry. 10
1.5.1 A diagram to show the difference between z-coordinate and σ-coordinates. 13

2.2.1 Figure showing the phenomena associated with various scales in the ocean and whether they are by nature hydrostatic or non-hydrostatic, extracted from Marshall et al. (1997b). 24
2.4.1 One dimensional schematic of piecewise linear (a) and piecewise quadratic (b) discontinuous shape functions and the two dimensional equivalents (c) and (d) respectively 32
2.4.2 The main stages of a non-linear iteration loop (AMCG, 2015) . . . 37

2.4.3 Mesh optimisation processes for two dimensional simulations. 45
2.4.4 An example of a super mesh formed from a donor mesh and a target mesh 48
2.4.5 Comparison of different types of mesh adaptivity used in this thesis: fully unstructured adaptivity, vertically structured adaptivity and vertical only adaptivity 49

3.4.1 The set-up for the simple test case used to explain the mixing diagnostics, showing the temperature field. 66
3.4.2 The isoscalar coordinates representing the reference state for the simple test case used to explain the mixing parameters. 69
3.4.3 The background potential energy, rate of change of background potential energy and the reference state for the simple test case used to explain the mixing parameters. 72
3.4.4 The background potential energy, rate of change of background potential energy, integral of D^{num} and D^{num} for the simple test case used to explain the mixing parameters. 76
3.5.1 The effect of aspect ratio on wave speed for fixed (MITgcm and Fluidity) and adaptive mesh results (Fluidity), compared with linearised theory and results using SUNTANS (Fringer et al., 2006). . 79
3.5.3 The differences in temperature fields for discontinuous Galerkin, control volume and continuous Galerkin discretisation methods at $t = 157.5$ s and $t = 1252.5$ s for the internal seiche simulations. . . 82
3.5.4 Mixing diagnosed in the seiche problem using both the method of background potential energy and the turbulent mean tracer decay method. 83
3.6.1 Number of nodes throughout the adaptive simulations of the tilting tank 86

3.6.2 A comparison of results for the tilting-tank using fixed and adaptive meshes with laboratory results from Horn et al. (2001) 87
3.6.3 A comparison of the mixing occurring in simulations of the tilting tank using different meshes. 88
3.6.4 Comparison of interface height and mixing for a vertical only adaptivity simulation and a coarse fixed structured mesh simulation . . . 90
3.7.1 Schematic of Loch Etive displaying bathymetry and initial linear temperature stratification as in Xing and Davies (2006) 92
3.7.2 Schematic showing the domain geometry for model runs for comparison with Berntsen et al. (2009) 96
3.7.3 Comparison of Fluidity simulations with Berntsen et al. (2009) for the temperature field at T/4 for a resolution of $\Delta x = 1.5625$ m . . . 97
3.7.4 A comparison of w velocities at the detector with Berntsen et al. (2009) for various resolutions 99
3.7.5 Snapshots of temperature and velocity fields at $T/4$ for a range of horizontal resolutions for the viscosity ν_c 101
3.7.6 Snapshots of temperature and velocity fields at $T/4$ for a range of horizontal resolutions for the viscosity ν_b 102
3.7.7 Snapshots of temperature and velocity fields at $T/4$ for a range of horizontal resolutions for the viscosity ν_a 103
3.7.8 Snapshots of temperature and velocity fields at $T/4$ for a range of vertical resolutions at a horizontal resolution of $\Delta x = 0.625$ m . . . 105
3.7.9 Convergence of $w_{\max}(\mathrm{T}/4)$ and $w_{\min}(\mathrm{T}/4)$ in terms of mesh resolution at all three viscosities used also showing an adaptive mesh result for each viscosity 107
3.7.10 Convergence of $\overline{A}_{\mathrm{Ri}}$ in terms of mesh resolution at all three viscosities used also showing an adaptive mesh result for each viscosity . . 108

3.7.11 Snapshots of the temperature field at each viscosity for a low resolution mesh, a high resolution (converged) mesh and an unstructured adaptive mesh at T/4 109

3.7.12 Snapshots of the critical Richardson area field at each viscosity for a low resolution mesh, a high resolution (converged) mesh and an unstructured adaptive mesh at T/4 110

3.7.13 Snapshots, stretched by a factor of ten vertically, of the temperature field for ν_a at T/4 for a converged high resolution mesh and adaptive mesh simulation. This also shows the adaptive mesh for a section of the domain 111

3.7.14 The relationship between time-averaged critical Richardson area and the number of nodes 113

3.7.15 Comparison of the mixing for lower and higher resolution fixed and adaptive mesh simulations at ν_a over the diagnostics region 116

4.3.1 Map of Idefjord (Johnsson et al., 2007) showing the location of the Svinesund and Bjällvarp sills, real world bathymetry for the two sills and two non-linear stratifications used for studies here 131

4.4.1 Comparison of snapshot results of Xing and Davies (2009a) and results of our study also using MITgcm for the same set-up 134

4.5.1 Qualitative comparison of velocity snapsots and temperature contours for Idefjorden set-ups with incresing complexity. 136

4.5.2 Qualitative comparison of tidally averaged temperature snapshots for Idefjorden set-ups with increasing complexity. 138

4.6.1 Comparison of field data acquired by Johnsson et al. (2007) with simulation data for sill top velocities at Svinesund 139

4.7.1 Contours of 0.023 kg m^{-3} density at four time levels at intervals of $T/16$ between 0 and $T/2$ inclusive for the stratification based on Johnsson et al. (2007) for the three velocities (lower, control and higher) 141

4.7.2 The power spectral density for depth averaged vertical velocity normalised by the M_2 frequency for the "Non-linear 1" stratification for the three velocities (lower, control and higher) 142

4.7.3 Density contours of 0.021 kg m^{-3} shown for every $T/16$. Results shown are for the stratification based on Johnsson et al. (2007) for the three depths of the Bjällvarp sill 145

4.7.4 The power spectral density for depth averaged vertical velocity normalised by the M_2 frequency for the stratification based on Johnsson et al. (2007) for the three depths of the Bjällvarp sill 146

4.7.5 Density contours of 0.021 kg m^{-3} shown for every $T/16$. Results shown are for the "Non-linear 1" stratification for the control velocity and the higher velocity 147

4.7.6 The power spectral density for depth averaged vertical velocity normalised by the M_2 frequency for the "Non-linear 1" stratification for the control velocity (a) and the higher velocity 148

4.7.7 Density contours of 0.021 kg m^{-3} shown at every $T/16$. Results shown are for the "Non-linear 2" stratification for the control velocity with a range of inter-sill separations 149

4.7.8 The power spectral density for depth averaged vertical velocity normalised by the M_2 frequency for the "Non-linear 1" stratification for the control velocity with a range of inter-sill separations 150

4.7.9 Density contours of 0.021 kg m^{-3} (black) shown for every $T/16$. Results shown are for the "Non-linear 1" stratification for the higher velocity with a range of inter-sill separations 152

4.7.10 The power spectral density for depth averaged vertical velocity normalised by the M_2 frequency for the "Non-linear 1" stratification for the higher velocity with a range of inter-sill separations 153

4.8.1 Approximate bed shear stress plotted against Svinesund sill depth . 155

4.8.2 Snapshots of the horizontal velocity field for the region of the sills at 17/6 T for various Svinesund sill depths 156

5.2.1 A schematic illustrating the different idealized coastal plane topographies during transgression in an icehouse versus a greenhouse 169

5.2.2 Regimes for the parameter space governed by the tidal excursion parameter, topographic steepness and Froude number based on Garrett and Kunze (2007) 172

5.3.1 Location in Wales, UK used for idealised bathymetries and temperature stratification based on July data for the Adriatic Sea 176

5.4.1 Snapshots of tidally-averaged (over three tidal cycles) change in buoyancy frequency with contours of time-averaged temperature for simulations using one bump to ten bumps 178

5.4.2 Snapshots of Richardson number at 23/8 T with contours of time-averaged temperature, for simulations using one bump to ten bumps 179

5.4.3 The critical Richardson area versus the number of sills (bumps) after three tidal cycles 181

5.4.4 The temperature profile at 250 m after the first bump crest for each icehouse simulation with between one and ten bumps after three tidal cycles. 182

5.5.1 Snapshots of time averaged temperature, Richardson number and passive tracer for icehouse world simulations 184
5.5.2 Vertical time-averaged temperature profile between the fifth and sixth bump for icehouse and greenhouse scenarios 185
5.5.3 Vertical distribution of a passive tracer for icehouse and greenhouse scenarios at $3T$ 187
5.5.4 Snapshots of time-averaged temperature, Richardson number and passive tracer for greenhouse world simulations 188

A.0.1 Comparison of Fludity results for Loch Etive without and with the application of the Coriolis force 228

LIST OF TABLES

3.7.1 List of adaptive mesh simulations and parameters for Loch Etive . . 95

4.3.1 Set-up parameters for idealised and more realistic set-ups of Idefjord and simulations to validate against physical studies by Johnsson et al. (2007) 130

5.3.1 Parameters for simulations of transgression and regression scenarios during an icehouse and greenhouse world 175

CHAPTER 1

INTRODUCTION AND BACKGROUND

Contents

1.1 Introduction 1

1.2 Thermohaline circulation 2

1.3 Vertical mixing 4

1.4 Internal waves 8

1.5 Ocean modelling 10

1.5.1 Parameterisation 10

1.5.2 Ocean modelling using fixed and adaptive meshes 12

1.1. Introduction

Oceans are complex systems and understanding their circulation is a branch of fluid dynamics where the range of scales present in the problems can be vast. Advances in understanding are made through a combination of theoretical, physical and numerical studies. However, physical studies for oceanography are more difficult to achieve than the laboratory experiments associated with other branches of fluid dynamics due to the scales involved, observational limitations and natural variation of the systems (Wunsch, 1996; Stewart, 2008b). Through the use of numerical models which solve the Navier-Stokes equations we can hope to better understand the dynamics of the oceans. Models such as Fluidity, which uses mesh adaptivity, can also reduce the computational cost of these large scale studies (AMCG, 2015; Piggott et al., 2009).

Understanding ocean circulation is of significant importance for a range of reasons, but of particular interest in this thesis is the importance in understanding the climate of the past and future, in addition to biological productivity and thus carbon draw down. The oceans transport large amounts of heat around the globe and the circulation of this heat is in part controlled by the intensity of vertical (or diapycnal) mixing, which is also responsible for nutrient influx into surface layers that control biological productivity (and results in the storage of fifty times more carbon in the ocean than in the atmosphere) (Ganachaud and Wunsch, 2000; Stewart, 2008a; McDougall, 1988; Wunsch and Ferrari, 2004; Toggweiler and Key, 2003; Rahmstorf, 2002). Changes in the rate of diapycnal mixing, which closes the thermohaline circulation and thus controls its intensity, influences how heat is distributed across the globe. Diapycnal mixing is a complex turbulent problem and is discussed further in section 1.3. Mixing itself is composed of the combination of stirring and molecular diffusion. Fluids are stirred by turbulent eddies which

increase the amount of surface area for molecular diffusion to occur across (and thus increasing the rate at which it occurs). The process of molecular diffusion then irreversibly changes the properties of the fluid.

1.2. Thermohaline circulation

The circulation of the world's oceans is an intrinsic part of the planet's climate. Ocean circulation is driven by the wind as well as differences in density imparted by subtle variations in salinity and temperature. In its simplest (and arguably overly simplistic) form the thermohaline circulation can be thought of as a giant conveyor belt that transports heat from near the equator towards the poles at the ocean's surface. In turn, cold water returns equatorward at depth. This process transports heat and salinity around the world and thus impacts upon the global climate (Ganachaud and Wunsch, 2000; Stewart, 2008a; McDougall, 1988). One of the ways that this circulation impacts on climate is a function of the ocean's heat capacity and the significant volume of surface water on the planet. The oceans can store vast amounts of heat, which can be transported in the ocean's circulation before being transferred to and from the atmosphere (Rahmstorf, 2002). In general, the thermohaline circulation transports 10^{15}W of heat poleward (Ganachaud and Wunsch, 2000; Rahmstorf, 2002; Toggweiler and Key, 2003). In the Atlantic, this transfer of heat has the effect of moderating the NW European climate as well as the amount of Arctic sea ice (Vellinga and Wood, 2002). Without the thermohaline circulation providing a regulating effect on the climate, less heat from low latitudes would be transported towards higher latitudes, meaning that the countries surrounding the North Atlantic could be up to 10–20^oC colder during winter (Rahmstorf, 2003). Understanding this circulation properly is therefore vital for investigating today's climate and the perturbations

that might affect it (Ganachaud and Wunsch, 2000).

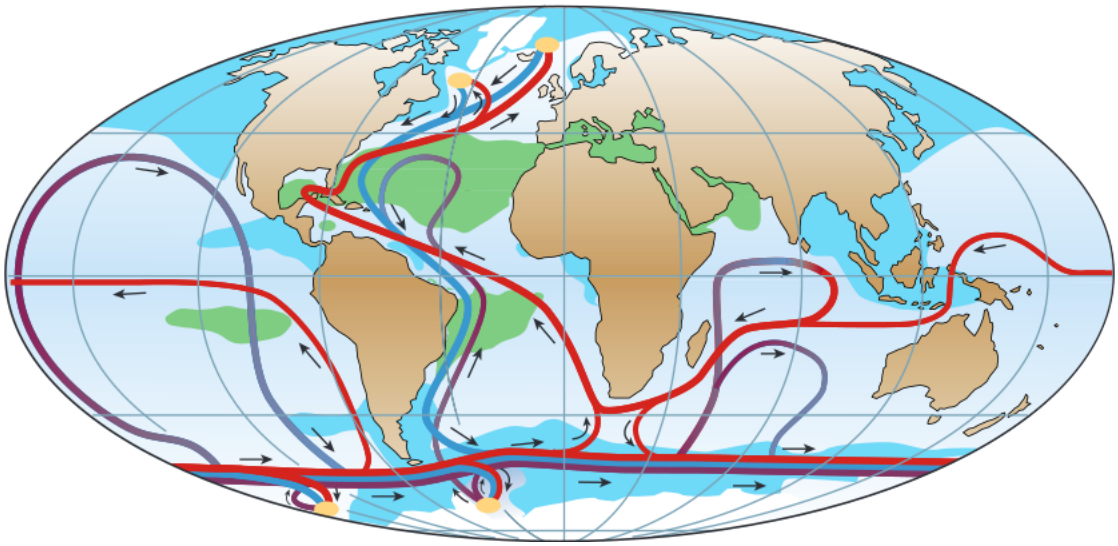


Figure 1.2.1.: Highly simplified illustration to show the thermohaline circulation transporting water around the globe, it also shows the main areas of downwelling in yellow – these are in the Nordic Sea and the Labrador Sea in the Atlantic, and the Ross Sea and Weddel Sea around Antarctica. The red lines depict the surface flow of warm water and the blue lines indicate the return flows of cold water at depth (Image from Rahmstorf (2002)).

Generally oceans are stably stratified by density (controlled by temperature and salinity), with the warmest, least saline water at the surface and the coldest, most saline waters at depth (Stewart, 2008a; McDougall, 1988). Although it should be noted that warm salty water may be more or less dense than cold fresh water, depending on the exact temperatures and salinities. Thermohaline circulation drives less dense warmer waters at the equator to higher latitudes where they consequently cool, mainly by evaporation (Gordon, 1986; Toggweiler and Key, 2003). Once cooled the waters become denser still (further losing buoyancy) due to an increase in the proportion of salt after the formation of sea ice and therefore they sink (or downwell) (Gordon, 1986; Toggweiler and Key, 2003). These downwellings occur in the Nordic and Labrador Seas and around the entire coastline of Antarc-

tica (Webb and Suginohara, 2001; Gordon, 1986; Ganachaud and Wunsch, 2000; Rahmstorf, 2002). As water downwells, more surface water is drawn poleward to replace it, and in turn becomes denser and itself downwells becoming bottom water. In this way, the movement of warm surface waters is controlled by the rates of downwelling (and in a similar manner upwelling) (Gordon, 1986; Ganachaud and Wunsch, 2000). Bottom water in the Atlantic travels southwards and eventually entrains some Antarctic Bottom Water from below; these deep waters then flow into the Pacific and Indian oceans (Ganachaud and Wunsch, 2000). The process of thermohaline circulation is illustrated in figure 1.2.1.

A reasonable amount is therefore known about the more obvious and localised downwellings of cold dense water, however the passage of the flow and the processes controlling its upwelling, which occurs globally, are less well understood (Webb and Suginohara, 2001; Toggweiler and Key, 2003).

1.3. Vertical mixing

Vertical mixing is responsible in large part for the upwelling which closes the thermohaline circulation. Due to the stable stratification of the oceans, mixing must occur between layers of different density (diapycnal mixing) in order to enable water properties to upwell from depth. However, to achieve this strong buoyancy forces must be overcome (Stewart, 2008a; McDougall, 1988). The breaking of internal waves, over seabed bathymetry, powered by the tide and wind forcings, provides a source of energy to overcome these forces (Wunsch and Ferrari, 2004).

Diapycnal mixing is far less prevalent than mixing along layers of equal density (isopycnal mixing) as more energy is required to overcome the restoring effect of stratification and buoyancy. However, it is of greater importance, as this small scale process controls the intensity of circulation by controlling the rate at which

ocean deep waters re-surface, and also re-supplies energy lost in the interior of the ocean by overturning and eddy-generation (Wunsch and Ferrari, 2004; Stewart, 2008a; McDougall, 1988). It also regulates sea surface temperatures (SST) and alters the depth over which surface heat is uniformly distributed, controlling the amount of heat mixed to and from depth where it can influence weather and climate (Chen et al., 1994; Moum and Smyth, 2001; Wunsch and Ferrari, 2004; Goosse et al., 1999).

Although we know less about upwellings than downwellings, field measurements (Polzin et al., 1997) have revealed that vertical mixing does not occur uniformly over the oceans as previously assumed, which would result in a different circulation to that observed. Strong mixing is in fact highly localised and enhanced near rough topographies due to the generation of turbulence, with a lower background level elsewhere (Webb and Suginohara, 2001; Stewart, 2008a; Ganachaud and Wunsch, 2000; Wunsch and Ferrari, 2004). This lower background level is characterised by a diffusivity of $\sim 10^{-5} \mathrm{m}^2\mathrm{s}^{-1}$ over the ocean interior (Ledwell et al., 1998; Gregg, 1989) while at rough topographies this can be up to three orders of magnitude larger (Polzin et al., 1997). Vertical mixing near complex topography is responsible for dissipating roughly 1TW of energy in the ocean (half of the amount needed for thermohaline circulation) through breaking internal waves in the deep ocean and shallow seas as shown in Figure 1.3.1 (Egbert and Ray, 2000; Munk and Wunsch, 1998). Turbulence at complex bathymetry is caused directly by the interaction of barotropic tides with bathymetric features. In addition to this, internal tides (internal waves at a tidal frequency) are also formed from this interaction which may break locally causing further localised turbulence or travel away from this area as internal waves and break further afield at complex bathymetry elsewhere (Munk and Wunsch, 1998).

In addition to closing the thermohaline circulation, influencing its intensity and

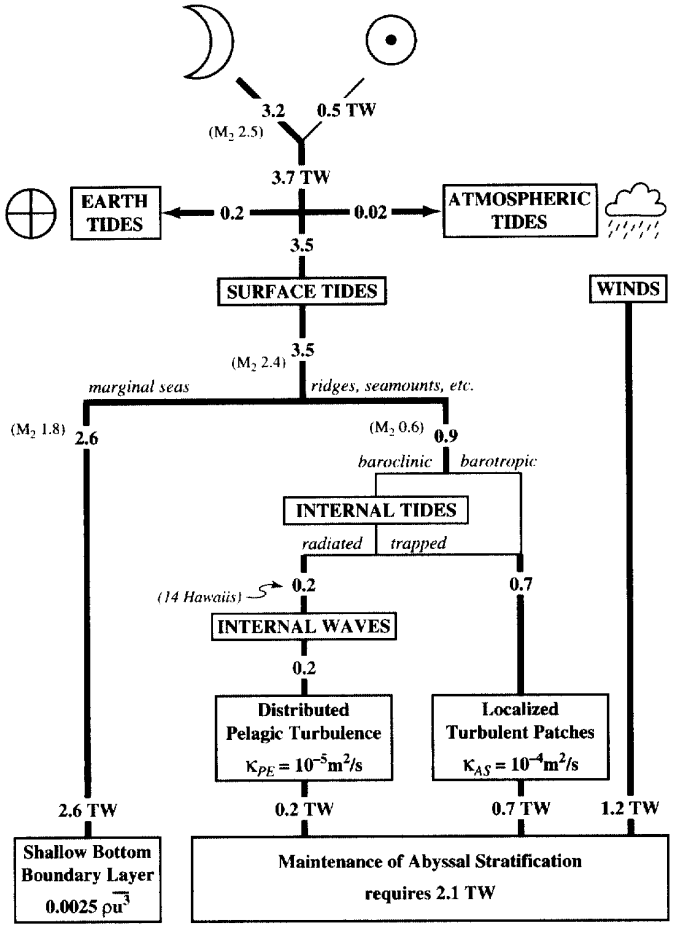


Figure 1.3.1.: Schematic from Munk and Wunsch (1998) showing an impressionistic tidal energy flux budget. Based on TOPEX/POSEIDON altimetric data results from (Egbert, 1997). 0.9TW of tidal energy is dissipated at ocean ridges and seamounts due to the generation of internal tides. Light lines represent speculation with observational support. This diagram shows that roughly half of the amount of energy required for the maitenance of abyssal stratification and thus thermohaline circulation is provided by the interaction of tides with complex bathymetry.

affecting global heat transport, vertical mixing also influences climate via carbon dioxide (CO_2) and nutrient exchange between the deep and surface ocean (Wunsch and Ferrari, 2004; Toggweiler and Key, 2003; Xing and Davies, 2006; Goosse

et al., 1999; Stewart, 2008a; McDougall, 1988). Nutrients such as phosphorus are typically depleted in surface waters due to demand. Phosphorus has been found to be the ultimate limiting nutrient on total ocean biological productivity as it is required to synthesis genetic material (DNA and RNA), as well as ATP (which is how energy is stored in living organisms' cells) (Tyrrell, 1999). Providing phosphorus to a system, for example through vertical mixing bringing it up from depth, is the most effective way of increasing the draw-down of carbon; providing one mole of phosphorus will draw down 106 moles of carbon according to the Redfield ratio (Redfield, 1963). The Redfield ratio is the ratio of chemical components of particulate matter in the ocean and was calculated via the composition of plankton which was found to be P:N:C:O_2 = 1:16:106:138 (Redfield, 1963).

Vertical mixing means waters from depth, rich in nutrients such as phosphorus and nitrogen, are ventilated by the atmosphere (increasing the content of carbon and oxygen). This enhances organism productivity and activates the "biological pump", which increases the dissolved greenhouse gas content of the ocean during biogeochemical cycles as shown in figure 1.3.2 (Toggweiler and Key, 2003; Xing and Davies, 2006; Goosse et al., 1999; Rahmstorf, 2002). This ventilated water is then taken deep into the ocean, meaning CO_2 from the atmosphere can be stored in the deep ocean reducing levels in the atmosphere (Xing and Davies, 2006; Goosse et al., 1999). Due to this process, the ocean contains around 50 times as much carbon as the atmosphere. Fluctuations in this level alter the pH of the ocean and could affect ecosystems and the climate in the future due to anthropogenic emissions (Rahmstorf, 2002; Caldeira and Wickett, 2003).

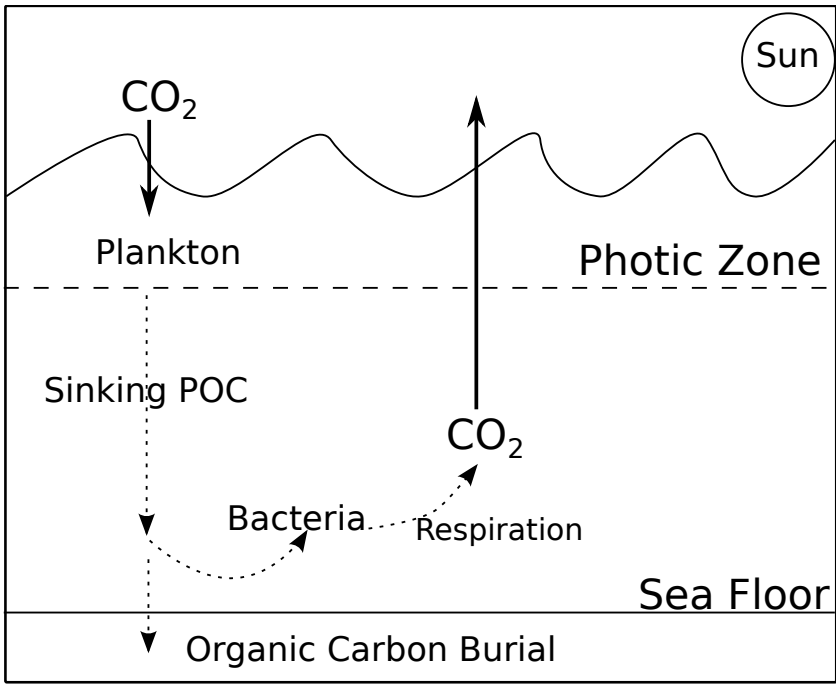


Figure 1.3.2.: Illustration to show the biological pump based on Herndl and Reinthaler (2013). This image shows the capture of carbon dioxide by phytoplankton during photosynethesis in the surface layers of the ocean where there is light, and then its transformation to particulate organic carbon (POC). 1 to 40 percent of which descends below the photic zone where it is fed on by bacteria resulting in the release of some carbon dioxide by respiration. Finally around 1 percent is eventually buried in the sea floor.

1.4. Internal waves

Internal waves and internal tides (internal waves generated at the tidal frequency (Gerkema and Zimmerman, 2008)) are intrinsically linked to rates of vertical mixing due to them leading to both distributed pelagic turbulence and localised turbulent patches (figure 1.3.1). Internal waves are generated when water particles that are disturbed from their equilibrium depth are moved back to their original locations by restoring buoyancy forces and as such are a form of gravity wave. These waves propagate along density interfaces within stably stratified liquids;

liquids are defined as stable if $N^2 > 0$ where $N^2 = -\frac{g}{\rho_0}\frac{\partial \rho}{\partial z}$, g is the gravitational acceleration, ρ is density, z is depth, N is termed the buoyancy frequency defined as the frequency that a displaced water particle would be expected to vibrate at due to buoyancy.

As outlined in the previous section and figure 1.3.1, internal waves are formed by two main processes (both at low frequencies): wind causing perturbations to the base of the mixed layer, and barotropic tidal forcing of water over complex bathymetry. As tides (a predominantly horizontal movement) interact with bathymetry, water is forced over the obstruction forming a vertical velocity component (like a wave maker) which oscillates with a tidal frequency and therefore produces an internal tide.

Roughly a third of tidal energy is dissipated at complex bathymetry in the deep ocean. To verify this the M_2 tidal energy dissipation can be viewed spatially (fig. 1.4.1). Peaks in energy dissipation can then be seen which correspond with major bathymetric features such as the Mid-Atlantic ridge and Hawaiin ridge. Other peaks in energy dissipation can be observed in shallow seas such as the South China Sea. In shallow seas it is known that flow interaction with bathymetry is enhanced and typically between a half and two thirds of energy is dissipated in these regions (Egbert and Ray, 2003; Miller, 1966; Garrett and Munk, 1979).

Internal waves transport energy (in diagonal wave beams in the vertical plane, which are reflected at the ocean floor and surface) over vast distances both in the vertical and horizontal with wave amplitudes of tens of metres (far greater than those of surface waves due to the weaker restoring forces), wavelengths of around a kilometre and speeds of up to tens of cm s^{-1} (Gerkema and Zimmerman, 2008; Garrett and Munk, 1979). This energy is eventually dissipated, due to the superposition of internal waves (causing steepening and breaking), by distinctive Kelvin-Helmholtz instabilities. Kelvin-Helmholtz instabilities occur at the sheared

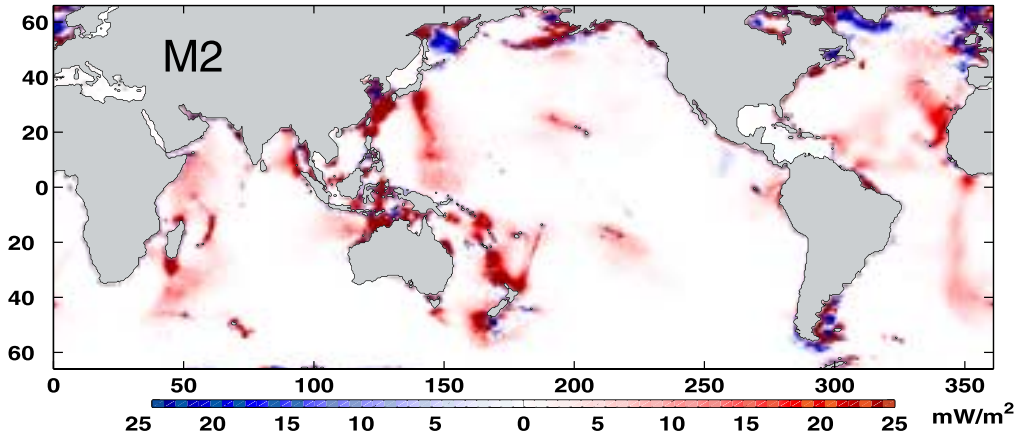


Figure 1.4.1.: Dissipation of M_2 tidal energy shown spatially across the globe extracted from Egbert and Ray (2003) where results are found by satellite altimetry. Areas are shown in red in the deep ocean where dissipation of this componenet of tidal energy occurs and can be seen to correlate with areas of complex bathymetry such as the Mid-Atlantic ridge and Hawaiin ridge. Results in shallow marine regions are however not as accurate and in some areas represented by anomalous blue colours

interface between two homogeneous fluid bodies of different densities and display an overturning billowing nature which increases the surface area for mixing to occur across.

1.5. Ocean modelling

1.5.1. Parameterisation

Mixing involves the transfer of energy from large scales to sub-centimetre scales. Therefore to explicitly capture it in numerical models, extremely fine mesh resolution is required (Siedler et al., 2001; Xing and Davies, 2009a). Large scale

general circulation models (GCMs) are not yet capable of the scale needed to resolve these complex processes, these scales are dictated by the Kolmogorov scales (the scale where turbulent kinetic energy is dissipated into heat due to viscosity) and Batchelor scales (the size of the smallest length scale of scalar concentration fluctuation before the dominance of molecular diffusion). Models can only represent phenomena of spatial scales at a minimum of twice the grid size, therefore these processes are generally parameterised via enhanced vertical diffusivity that can vary with location, time and Richardson number ($R_i = N^2/(\partial U/\partial z)$ a ratio between water column stability (represented by the Buoyancy frequency, N) and vertical velocity shear, which allow us to mimic a fine-resolution model (Haidvogel and Beckmann, 1999; Moum and Smyth, 2001; Wunsch and Ferrari, 2004; McDougall, 1988; Ganachaud and Wunsch, 2000; Gent and McWilliams, 1990). Parameterisation works by trying to produce the most realistic circulation, however they must incorporate incorrect energetics with no link to the sources of turbulence to achieve this (Wunsch and Ferrari, 2004; McDougall, 1988).

Some past parametrisations of vertical mixing by vertical diffusivity in GCMs, led to inaccuracies in the poleward heat transfer and the thermocline depth, which affected atmospheric heat gradients and the simulated climate. This highlights the need for ongoing improvements to modelling vertical mixing so that GCMs can be used to accurately monitor climate change (Pacanowski and Philander, 1981; Bryan, 1987; Ganachaud and Wunsch, 2000). Efforts to improve the parameterisation of mixing in ocean models have involved embedding 1D mixed layer models and turbulence closure schemes into GCMs (Gaspar et al., 1990). Turbulence closure schemes are a compromise, not being as fast as the simplest methods and not as detailed as higher-order schemes (Gaspar et al., 1990).

Parameterisations are generally assessed by comparing model results against ocean weather station data as well as using chemical tracers to test the circulation

and water mass formation produced by the model (England and Maier-Reimer, 2001; Gaspar et al., 1990). These methods show that even modern parameterisations display significant model errors (Gaspar et al., 1990). To further improve parametrisations it is vital to better understand the complex mixing processes themselves and their mechanisms, which occur at the boundary and interior (McDougall, 1988; Armi, 1978). It is particularly important that parameterisations are trustworthy and firmly physically based when using GCMs to study paleoclimates where a number of factors differ (continents configuration, atmospheric forcing and sea levels) (Large et al., 1994; Wunsch and Ferrari, 2004; McDougall, 1988).

In order to improve the parameterisation of vertical mixing, mesoscale (10–10,000 m sized) models are required that can be finer in resolution allowing the modelling and study of vertical mixing at smaller scales without parameterisation (Haidvogel et al., 2000). Numerical models such as Fluidity and MITgcm are both capable of these types of simulations and are used in this study.

1.5.2. Ocean modelling using fixed and adaptive meshes

Fixed structured meshes have historically been widely used for ocean modelling purposes, often with finite difference discretisation of the governing equations. These fixed structured meshes can be categorised through their use of three different types of vertical coordinate scheme, these are: z-coordinates, σ-coordinates and isopycnal coordinates. Typically large-scale ocean models use either z or isopycnal co-ordinates, while regional models typically use σ-coordinates (Griffies et al., 2000b; Vitousek, 2013). Since small-scale models are used in this study either z-coordinates or σ-coordinates (figure 1.5.1) are most appropriate, as isopycnal coordinates are most useful in large-scale models (Griffies et al., 2000a). In models of mixing, isopycnal coordinates are also inappropriate as they generally rely on the hydrostatic approximation and therefore cannot resolve overturning

(Vitousek, 2013).

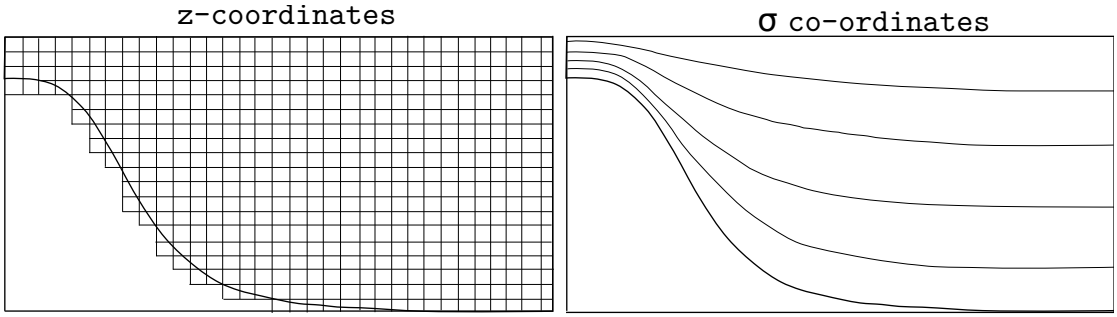


Figure 1.5.1.: A diagram to show the difference between z-coordinate and σ-coordinates. Thin lines show the grid lines whereas thick lines show the bathymetry and edge of the domain. σ-coordinates (right) are shown to follow the bathymetry in smooth lines whereas z-coordinates (left) form rectangular cells and show the "stair stepping" bathymetry representation. Based upon a figure by Vitousek (2013).

In models with complex bathymetry, σ-coordinates, in which the vertical coordinate is a percentage of depth, provide a smooth representation of bathymetry, which improves representation of the bottom boundary layer. However, their non-horizontal bathymetry-following grid cell orientations introduce errors in the horizontal baroclinic pressure gradient that can lead to spurious density flows (Kliem and Pietrzak, 1999; Vitousek, 2013). Although z-coordinates avoid these pressure gradient errors with their horizontal rows of grid cells, they instead have problems due to a staircase representation of bathymetries, which do not provide an accurate representation of terrain, although these can be partially solved by implementation of shaved cells (Barad et al., 2009; Adcroft et al., 1997; Vitousek, 2013). While neither z-cordinates nor σ-coordinates are perfect, the disadvantages (pressure gradient errors and stepped bathymetries) of both systems are undergoing significant investigation.

Due to the complex turbulent dynamics associated with vertical mixing, high mesh resolutions are necessary. When using fixed structured meshes, where this

high resolution is not desired over the whole mesh, grid nesting may be used to reduce the computational cost. This involves areas of higher resolution being embedded to resolve complex small scale dynamics in certain parts of the domain. However, these areas must be decided upon a priori and also result in abrupt changes in mesh resolution, which potentially result in corrections being necessary (Laugier et al., 1996). The recent use of unstructured meshes (typically using finite element or finite volume based discretisations) for ocean applications represents a potential solution to this problem. These meshes are able to smoothly vary resolution to resolve both small and large scale processes without requiring nesting, which is beneficial to represent coastlines and complex bathymetry (Pain et al., 2005). They do this usually by using triangular (two-dimensions) or wedges/tetrahedral (three-dimensions) elements, which tessellate enabling smoothly varying resolution. However, unstructured meshes in the vertical can also develop errors similar to those of σ-coordinate meshes and this must be combated for accurate resolution of the dynamics. A thorough review and history of the methods of unstructured mesh ocean modelling are provided by Pain et al. (2005) and the referenced papers.

Fixed meshes, although capable of yielding areas of high resolution, require the locations of this high resolution to be pre-defined. This is not necessary when using adaptive meshes, since adaptive meshes enable the locations of high mesh resolution to be varied according to the dynamics of the solution. This enables the capture of developing flow features where the location and extent cannot always be known in advance. Adaptive meshes involve coarsening or refining the mesh according to the complexity of the dynamics. Mesh adaptivity can therefore capture complex flow dynamics with the use of fewer nodes. This can produce a computational saving even when the overheads associated with the adaptivity process (and with the necessary finite element discretisation) are taken into account.

CHAPTER 2

NUMERICAL OCEAN MODELS: FLUIDITY AND MITGCM

Contents

2.1 Introduction 17

2.2 Governing Equations 18

2.2.1 Formation of the Navier-Stokes Equations 18

2.2.2 The Advection-Diffusion Equation 25

2.2.3 The incompressible, non-hydrostatic, Boussinesq, rotating Navier-Stokes equations 26

2.3 Meshes 27

2.3.1 Fixed Meshes - Vertical Co-ordinate Systems 27

2.3.2 Unstructured Meshes 28

2.4 The Fluidity model 29

2.4.1 Spatial Discretisation 29

2.4.2 Temporal Discretisation 36

2.4.3 Linear Solvers 38

2.4.4 Mesh Adaptivity 41

2.5 The MITgcm model 52
2.5.1 Sponge regions 52
2.5.2 Temporal discretisation 53
2.5.3 Spatial discretisation 53

2.1. Introduction

This study uses both the computational fluid dynamics numerical model Fluidity and the general circulation model MITgcm. Fluidity is an open source, general purpose code that numerically solves the Navier-Stokes and accompanying equations over unstructured or structured meshes, using finite element or control volume options (AMCG, 2015). Fluidity's ability to use unstructured meshes discretised using the finite element method provides great geometric flexibility (Ford et al., 2004). MITgcm (Marshall et al., 1997a,c) is an open-source numerical model for the study of atmospheric, oceanic and climate problems using a common hydrodynamical kernel. MITgcm solves the Navier-Stokes and accompanying equations over fixed meshes with z-coordinates using finite volume methods (Adcroft et al., 2013). Both Fluidity and MITgcm have non-hydrostatic options which is vital for attempting to fully resolve overturning processes and diapycnal mixing.

In addition Fluidity is capable of adapting its computational mesh to various solution fields throughout a simulation in order to increase resolution in areas where flow complexity evolves, while maintaining low resolutions elsewhere. Fluidity employs a so-called mesh optimisation approach to update meshes, three forms of this are examined within this thesis: fully unstructured, vertically structured and vertical only adaptivity (AMCG, 2015).

This chapter will introduce Fluidity by covering the governing equations, spatial and temporal discretisation methods, linear solvers and adaptive meshing techniques. Focus is placed on the meshes used in this thesis in section 2.3 including fixed meshes with a range of vertical co-ordinate systems and a range of adaptive meshes both unstructured and partially structured. In addition, MITgcm is described in further detail in section 2.5 including spatial and temporal discretisation of the governing equations.

2.2. Governing Equations

The governing equations used are outlined here along with the common approximations made.

2.2.1. Formation of the Navier-Stokes Equations

Mass conservation and incompressibility

Conservation of mass implies that for a fixed parcel of fluid with volume Ω the rate of change of mass is balanced by the flow of fluid across the boundaries ($\partial\Omega$) of the parcel:

$$\frac{\partial}{\partial t}\int_{\Omega}\rho dV = -\int_{\partial\Omega}\rho\mathbf{u}.\mathbf{n}dS, \tag{2.2.1}$$

where $\mathbf{n}$ is the unit vector normal to the boundary facing outwards, $\mathbf{u}$ is the velocity vector and t the time. This can be written alternatively using the Reynold's Transport theorem as:

$$\frac{\partial}{\partial t}\int_{\Omega}\rho d = -\int_{\Omega}\nabla.(\rho\mathbf{u})dV. \tag{2.2.2}$$

Since the quantity of material in the parcel through time is constant $\frac{\partial}{\partial t}\int_{\Omega}\rho dV$ must therefore be equal to zero. This means that equation (2.2.2) may be simplified, loosing the integral as the domain boundary is fixed, to:

$$\frac{D\rho}{Dt} = -\rho\nabla.\mathbf{u}, \tag{2.2.3}$$

where

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + \mathbf{u}.\nabla, \tag{2.2.4}$$

is the material derivative.

Momentum conservation

The momentum of a unit volume of fluid is $\rho\mathbf{u}$, if we neglect the effect of viscosity then the rate of change of momentum can be expanded as:

$$\frac{\partial \rho\mathbf{u}}{\partial t} = \rho\frac{\partial \mathbf{u}}{\partial t} + \frac{\partial \rho}{\partial t}\mathbf{u}. \tag{2.2.5}$$

Together with the incompressible flow equation (2.2.15) and Euler's equation (Batchelor, 2000):

$$\frac{\partial \mathbf{u}}{\partial t} = -\mathbf{u}.\nabla\mathbf{u} - \frac{1}{\rho}\nabla p, \tag{2.2.6}$$

this results in:

$$\frac{\partial \rho\mathbf{u}}{\partial t} = -\nabla p - \nabla.(\rho\mathbf{u}\mathbf{u}), \tag{2.2.7}$$

where $\mathbf{uu}$ is a dyadic product of vectors and is a tensor. By taking account of the effects of viscosity on fluid movement additional terms are required resulting in the following (Cauchy momentum equation):

$$\rho\frac{D\mathbf{u}}{Dt} = -\nabla.\overline{\overline{\sigma}} + \rho F, \tag{2.2.8}$$

where F represents forces such as gravity and Coriolis (per unit mass) and $\overline{\overline{\sigma}} = -p\mathbf{I} + \overline{\overline{\tau}}$ is the stress tensor. The viscous stress tensor for a Newtonian fluid such as water is:

$$\overline{\overline{\tau}} = 2\mu(\frac{1}{2}(\nabla u + (\nabla u)^T) - (\nabla.\mathbf{u})\mathbf{I}/3)). \tag{2.2.9}$$

For a Newtonian fluid where the appropriate additional forces (Coriolis and buoyancy) for scenarios studied here are added this results in:

$$\rho\frac{D\mathbf{u}}{Dt} + 2\rho\Omega\times\mathbf{u} = -\rho g\mathbf{k} - \nabla p + \frac{1}{2}\nabla.(\overline{\overline{\mu}}(\nabla\mathbf{u} + \mathbf{u}^T)). \tag{2.2.10}$$

where $\mathbf{k}$ is a unit vector in the vertical direction.

Coriolis force: 2.5D versus 3D methods

The Coriolis term, $2\rho\Omega \times \mathbf{u}$, may be expanded, given that Ω is:

$$\Omega = \Omega \begin{pmatrix} 0 \\ \cos(\varphi) \\ \sin(\varphi) \end{pmatrix}, \tag{2.2.11}$$

as follows:

$$2\rho\Omega \times \mathbf{u} = 2\rho\Omega \begin{pmatrix} w\cos(\varphi) - u\sin(\varphi) \\ u\sin(\varphi) \\ -u\cos(\varphi) \end{pmatrix}. \tag{2.2.12}$$

Since the systems of interest in this thesis are small by ocean modelling standard, the f-plane approximation ($f = 2\Omega\sin(\varphi)$) is used. Therefore the Coriolis terms can be simplified to $f\rho\mathbf{k} \times \mathbf{u}$, where $f = 2\Omega\sin\varphi$ is the Coriolis parameter, φ is the latitude and $\Omega = 7.292 \times 10^{-5}$ s^{-1} is the Earth's rotation rate.

The Coriolis force is applied in different methods for two dimensional vertical slice systems using MITgcm and Fluidity. Two dimensional MITgcm domains are effectively very thin three dimensional domains made up of one cell in the third dimension and periodic boundary conditions. Due to this v velocities are possible and the full three dimensional Navier-Stokes are solved and hence the Coriolis force is included.

However, two dimensional Fluidity simulations do not include a v velocity, therefore Coriolis force cannot be included in the equation for u velocity in this way. Instead it can be realised that the Navier-Stokes equation for the v velocity resem-

bles an advection diffusion equation and may be simplified from:

$$\frac{\partial v}{\partial t} + u\frac{\partial v}{\partial x} + v\frac{\partial v}{\partial y} + w\frac{\partial v}{\partial z} = -\frac{1}{\rho}\frac{\partial p}{\partial y} - 2\Omega u\sin(\varphi) + F, \tag{2.2.13}$$

where F denotes terms involved with additional forces, to

$$\frac{\partial v}{\partial t} + u\frac{\partial v}{\partial x} + w\frac{\partial v}{\partial z} = -fu + F, \tag{2.2.14}$$

by assuming that the dynamics are uniform in the y direction, which allows all terms involving derivatives in this direction to be removed. This equation for the pseudo v velocity may then be solved after the coupled Navier-Stokes equations for u and w are solved. Therefore without explicitly using a third dimension, a v velocity due to inertial motion can still be calculated which may then be used as an fv term in the equation for the u velocity.

This method allows coupling, however in a weaker way than the full 3D Navier-Stokes equations. It is therefore referred to as a pseudo three dimensional (or 2.5 D) method. Within the Fluidity flml file this is accomplished by adding an additional Vvel field which calculates a velocity with the source term $-fu$. This field is then used as part of the source term fv for the 2D velocity field.

The Boussinesq approximation

If variations in density, ρ', about the reference state can be assumed to be minimal ($\rho' \ll \rho_0$), then inertial forces due to these variations may be neglected (Boussinesq, 1903; Spiegel and Veronis, 1960). This has the effect of replacing ρ with ρ_0 in all terms in the Navier-Stokes equations except the buoyancy term where density is multiplied by gravity and changes in density drive convection.

Incompressibility implies that $\frac{D\rho}{Dt} = 0$ for a particular fluid parcel as variations

in density are minimal within the parcel. This reduces equation (2.2.3) to:

$$\nabla.\mathbf{u} = 0, \tag{2.2.15}$$

which is the equation of continuity and implies the non-divergence of fluid flow.

The Boussinesq approximation is applicable in the ocean where changes in density around the background value of 1035 kg m^{-3} are less than 2% (Gill, 1982). Under the Boussinesq approximation the Navier-Stokes equations (2.2.10) are:

$$\rho_0 \frac{D\mathbf{u}}{Dt} + f\rho\mathbf{k} \times \mathbf{u} = -\rho g\mathbf{k} - \nabla p + \frac{1}{2}\nabla.(\overline{\overline{\mu}}(\nabla\mathbf{u} + \mathbf{u}^T)). \tag{2.2.16}$$

Since the density of water in the ocean typically only varies slightly around a reference state, ρ_0, it can therefore be given by a linear equation of state, which when the Boussinesq approximation is applied is:

$$\rho = \rho_0(1 - \alpha(T - T_0) + \beta(S - S_0)), \tag{2.2.17}$$

where $\alpha = -\frac{1}{\rho}\frac{\partial\rho}{\partial T}$ is the thermal expansion coefficient and $\beta = \frac{1}{\rho}\frac{\partial\rho}{\partial S}$ is the haline contraction coefficient. Reference values of temperature T and salinity S are represented by T_0 and S_0 respectively. In order to further simplify the complex dynamics occurring in simulations in this thesis all changes in density are assumed to be due to temperature alone, and for the reference temperature to be zero, further simplifying the equation of state to:

$$\rho = \rho_0(1 - \alpha T). \tag{2.2.18}$$

The hydrostatic approximation

The hydrostatic approximation is commonly used in (large-scale) ocean modelling, however it is inappropriate at the scales of interest in this thesis. According to Marshall et al. (1997b) the hydrostatic approximation is applicable when:

$$n = \frac{\alpha}{Ri} \ll 1, \tag{2.2.19}$$

where α is the aspect ratio of the system of motion (h/L) and $Ri = N^2h^2/U^2$ is the Richardson number. For systems of interest here typical values are $h \sim 10$ m, $U \sim 1$ ms^{-1}, $L \sim 100$–1000 m and $N \sim 0.01$ s^{-1} (e.g. Loch Etive), these values result in $n \sim 1$ and therefore show that the approximation would not be valid. Typical ocean phenomena, and whether or not they are hydrostatic, are shown in figure 2.2.1, where it can be seen that internal waves, a form of gravity waves are typically non-hydrostatic.

Previous studies of fjord systems, such as Knight Inlet, that use the approximation have been shown to resolve dynamics poorly when compared with non-hydrostatic studies and physical data (Stacey and Zedel, 1986; Cummins, 2000; Klymak and Gregg, 2003). This is because in simulations where vertical mixing and turbulent overturning are required to be fully resolved the hydrostatic approximation is not appropriate. Therefore simulations undertaken in this thesis use the non-hydrostatic capabilities of Fluidity and MITgcm. The necessity to use non-hydrostatic models for these types of simulations was also further studied by Xing and Davies (2007) for Loch Etive. Therefore the following explanation of the hydrostatic approximation is included purely for completeness.

The hydrostatic approximation assumes that pressure at a point is purely due to the weight of water above the point plus any pressure at the water's surface due

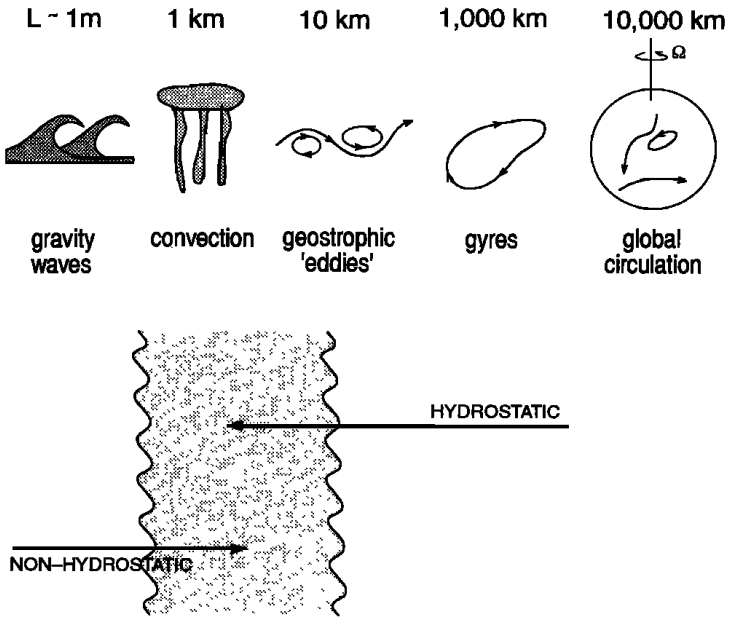


Figure 2.2.1.: Figure showing the phenomena associated with various scales in the ocean and whether they are by nature hydrostatic or non-hydrostatic, extracted from Marshall et al. (1997b).

to air pressure. This is known as hydrostatic balance:

$$\frac{\partial p}{\partial z} = \rho g. \tag{2.2.20}$$

Upon integration this gives:

$$p(z) = p_s + \int_z^{z_{top}} \rho g dz, \tag{2.2.21}$$

which essentially decomposes pressure into background value p_s (at the surface) plus a perturbation p' (depth dependent). This simple equation for pressure means that the vertical momentum equation can be simplified. Ignoring viscosity this

gives the following equation:

$$\frac{Dw}{Dt} = -\frac{\partial p}{\partial z} + \rho g = 0. \tag{2.2.22}$$

Therefore this approximation is only applicable when vertical accelerations are negligible as it assumes that vertical accelerations are unimportant related to the gravitational acceleration and vertical pressures. This approximation works well for large-scale simulations of ocean circulation and gyres, but begins to falter when horizontal scales of motion are smaller (1-10 km) and more comparable to the vertical scales. Smaller scale processes such as internal waves where these scales of motion are similar are inherently non-hydrostatic in nature (Marshall et al., 1997b). Large-scale ocean simulations that can successfully use the hydrostatic assumption for large scale processes must use parameterisations in order to account for the vertical mixing processes.

2.2.2. The Advection-Diffusion Equation

The equation governing the advection-diffusion of a scalar tracer θ (e.g. Temperature or Salinity) is given by:

$$\frac{\partial \theta}{\partial t} + \nabla.(\mathbf{u}\theta) = \nabla.(\kappa \nabla \theta) - \sigma\theta + F, \tag{2.2.23}$$

where we use the notation: $\mathbf{u}$ the velocity vector, t time, κ a tensor of (e.g. thermal or haline) diffusivities (in this thesis an anisotropic diffusivity is used with a horizontal and vertical component (κ_h and κ_v respectively)), σ an absorption coefficient (used in sponge regions for MITgcm simulations in chapters 4 and 5 and explained further in section 2.5) and F represents sources and sinks for the tracer.

For an incompressible fluid where $\nabla.\mathbf{u} = 0$ the advection term $\nabla.(\mathbf{u}\theta) =$

$\mathbf{u}.\nabla\theta + (\nabla.\mathbf{u})\theta$ becomes $\mathbf{u}.\nabla\theta$. And the diffusion term, responsible for mixing via molecular level processes, is $\nabla.(\kappa\nabla\theta)$. This results in the advection-diffusion equation below:

$$\frac{D\theta}{Dt} = \nabla.(\kappa\nabla\theta) - \sigma\theta + F, \tag{2.2.24}$$

where:

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + \mathbf{u}.\nabla. \tag{2.2.25}$$

2.2.3. The incompressible, non-hydrostatic, Boussinesq, rotating Navier-Stokes equations

Given the conservation equations and approximations outlined in the the previous sections, a full list of the equations solved for simulations in this thesis is reiterated below. These include the incompressible, Boussinesq, non-hydrostatic Navier-Stokes equations, the continuity equation, a linear equation of state and the advection-diffusion equation for temperature (we omit salinity in this thesis as it is not required to provide the stratification profiles used here) and passive tracers (c) (passive tracers are used in chapter 5 in order to model the mixing of nutrients, such as phosphorus, in the water column):

$$\frac{D\mathbf{u}}{Dt} + f\rho\mathbf{k}\times\mathbf{u} = -\frac{\rho g}{\rho_0}\mathbf{k} - \nabla\left(\frac{p}{\rho_0}\right) + \nabla.(\nu\nabla\mathbf{u}), \tag{2.2.26}$$

$$\nabla.\mathbf{u} = 0, \tag{2.2.27}$$

$$\rho = \rho_0(1-\alpha T), \tag{2.2.28}$$

$$\frac{\partial T}{\partial t} + \mathbf{u}.\nabla T = \nabla.(\kappa_T\nabla T), \tag{2.2.29}$$

$$\frac{\partial c}{\partial t} + \mathbf{u}.\nabla c = \nabla.(\kappa_c\nabla c), \tag{2.2.30}$$

where $\nu = \mu/\rho_0$ is the kinematic viscosity and is used as the equation is divided through by ρ_0 for simplicity.

2.3. Meshes

Due to the finite element and control volume discretisation methods used in Fluidity as well as the availability of a mesh optimisation capability, comparisons between results run with fixed and adaptive meshes may be made. This is because these discretisation methods allow the same set-ups to be used on both types of meshes. All meshes used in Fluidity are comprised of triangles in two dimensions (as opposed to the cuboids used in MITgcm simulations), and for this thesis are generated using Gmsh (Geuzaine and Remacle, 2009). Structured meshes have regular connectivity where each node has the same connectivity as neighbouring nodes, whereas unstructured meshes have irregular connectivity. The advantages and disadvantages of these two types of meshes are discussed in the following sections.

2.3.1. Fixed Meshes - Vertical Co-ordinate Systems

When modelling vertical mixing, the vertical co-ordinate system or mesh structure used can be extremely important and influences numerical results. Three main co-ordinate systems exist in ocean models, z-coordinates, σ-coordinates and isopycnal coordinates, each with their own advantages and disadvantages as presented in chapter 1.5.2. For small-scale models z-coordinates or σ-coordinates are most appropriate.

While neither z-coordinates or σ-coordinates are perfect, the disadvantages (pressure gradient errors and stepping bathymetries) of both systems are undergoing significant investigation (Vitousek, 2013). Fluidity is capable of using both

σ-coordinates and z-coordinates for fixed meshes (however adaptive meshes provide even greater flexibility) and both are used throughout this thesis, while MITgcm is only capable of z-coordinates.

In a step towards adaptive meshing, telescoping meshes are also used in chapter 3 for Loch Etive, where a fixed structured mesh is used that has a fine mesh (e.g. 10 m resolution) over the region of interest and a coarsening mesh away from this area until it reaches a low resolution (e.g. 100 m resolution). This allows the mixing to be captured more fully and the domain to be large enough to reduce the interactions of internal waves with the boundaries, while lowering the computational expense of the simulation without the need to make use of mesh adaptivity.

2.3.2. Unstructured Meshes

Unstructured meshes, which are possible within Fluidity due to the use of an underlying finite element discretisation approach, are useful due to their geometric flexibility. Their geometric flexibility stems from the fact that each grid cell can have a varying number of neighbouring cells (Cushman-Roisin and Beckers, 2011). These meshes are capable of capturing varied bathymetries and coastlines with far greater ease than fixed meshes, due to their ability to use varied sizes of triangular or tetrahedral elements to tessellate or represent interesting parts of the domain. They are also capable of placing more resolution tangential to coastlines and bathymetry in order to capture boundary currents. However, unstructured meshes are also susceptible to pressure gradient errors in the same way as sigma–coordinate models.

Unstructured meshes are also capable of smooth variations in resolution that aid the capture of processes at large and small scales (Pain et al., 2005). These properties mean that unstructured meshes are advantageous for modelling geo-

physical fluid dynamics problems, however in the past they have been underused in these areas due to challenges in achieving stable discretisation. The ability to use smooth variations in resolution also mean that unstructured meshes, especially when combined with adaptive mesh techniques, make capturing the cascade of energy associated with breaking internal waves more tractable than with fixed meshes (Fringer et al., 2006).

Powerful adaptivity algorithms may also be used with unstructured meshes; adaptivity is a major potential benefit of using Fluidity to study waves and mixing and will be discussed in the next section.

2.4. The Fluidity model

2.4.1. Spatial Discretisation

This section provides an overview of how the equations described in the previous sections are spatially discretised within Fluidity. Fluidity uses finite element methods (in addition to the control volume methods which are used here for the solution of the scalar advection-diffusion equations only for comparison in terms of the mixing generated in chapter 3). The boundary of Ω is defined as $\partial\Omega$ which may be split into different sections for different boundary condition types: Neumann, Dirichlet or Robin:

$$\partial\Omega = \partial\Omega^N \cup \partial\Omega^D \cup \partial\Omega^R. \tag{2.4.1}$$

Advection–diffusion equation

The advection-diffusion equation in conservative form for a scalar tracer θ can be written as:

$$\frac{\partial\theta}{\partial t} + \nabla.(\mathbf{u}\theta) - \kappa\nabla^2\theta = 0, \tag{2.4.2}$$

where the boundary can then be divided using (2.4.1) and boundary conditions applied for θ giving:

$$\mathbf{n}.(\kappa\nabla\theta) = g_N \text{ on } \partial\Omega^N, \tag{2.4.3}$$

$$-\mathbf{n}.(\kappa\nabla\theta) = h(\theta - \theta_a) \text{ on } \partial\Omega^R, \tag{2.4.4}$$

$$\theta = g_D \text{ on } \partial\Omega^D. \tag{2.4.5}$$

For continuous Galerkin methods (CG) we write the equations in their weak form by pre-multiplying by a test function φ and integrating over the domain:

$$\int_\Omega \varphi\Big(\frac{\partial\theta}{\partial t} + \nabla.(\mathbf{u}\theta) - \kappa\nabla^2\theta\Big) = 0. \tag{2.4.6}$$

Integrating the spatial derivative terms by parts using the first Green's identity yields:

$$\int_\Omega \varphi\nabla.(\mathbf{u}\theta)\Bigg[= \int_\Omega udv = \int_{\partial\Omega} uv - \int_\Omega vdu\Bigg] = \int_{\partial\Omega} \varphi(\mathbf{u}\theta).\mathbf{n} - \int_\Omega (\mathbf{u}\theta).\nabla\varphi, \tag{2.4.7}$$

$$\int_\Omega \varphi\kappa\nabla^2\theta\Bigg[= \int_\Omega udv = \int_{\partial\Omega} uv - \int_\Omega vdu\Bigg] = \int_{\partial\Omega} \varphi\bar{\bar{\kappa}}.\nabla\theta.\mathbf{n} - \int_\Omega \bar{\bar{\kappa}}.\nabla\theta.\nabla\varphi. \tag{2.4.8}$$

Together this means that (2.4.6) may be rewritten as:

$$\int_\Omega \varphi\frac{\partial\theta}{\partial t} - \nabla\varphi.\mathbf{u}\theta + \nabla\varphi.(\kappa\nabla\theta) + \int_{\partial\Omega} \varphi(\mathbf{n}.\mathbf{u}\theta - \mathbf{n}.(\kappa\nabla\theta)) = 0. \tag{2.4.9}$$

Assuming closed boundaries ($\mathbf{u}.\mathbf{n} = 0$) and applying homogeneous Neumann boundary conditions of $\partial\theta/\partial n = 0$ yields:

$$\int_\Omega \varphi\frac{\partial\theta}{\partial t} - \nabla\varphi.\mathbf{u}\theta + \nabla\varphi.\bar{\bar{\kappa}}.\nabla\theta = 0. \tag{2.4.10}$$

A solution θ is termed a weak solution of the advection diffusion equation if (2.4.10) is true for all φ in a space of test functions V. For a weak solution it is only a requirement that the first derivatives give finite values when integrated with the test functions, here equation 2.4.9 only contains first derivatives of θ. Weak solutions are not smooth enough to be classical solutions to the problem where second derivatives would be required.

Finite Element Discretisation The finite element method solves this weak form of the equation by looking for a solution within a finite-dimensional subspace, instead of finding the continuous solution from the whole function (Sobolev) space. Solutions may therefore be given as a linear combination of a finite number of functions (trial functions, φ_i) that form a trial space:

$$\theta(\boldsymbol{x}) = \sum_{i=1}^{N} \theta_i \varphi_i(\boldsymbol{x}). \tag{2.4.11}$$

where θ_i are coefficients that may be written as vectors, whose dimension is that of the trial space. This finite set of trial functions, means an infinite number of test functions are no longer required to find unique solutions to the weak form of the equation. For Galerkin methods the choice is made for the test space and the trial space to be the same, effectively a Galerkin projection of the weak equation to a finite dimensional subspace.

The simplest continuous Galerkin methods involve a set of basis functions φ_i, found by choosing piecewise linear functions φ_i that satisfy:

$$\varphi_i(x_i) = 1 \quad \forall i$$

$$\varphi_i(x_{j\neq i}) = 0 \quad \forall i, j,$$

where x_i are the locations of "degrees of freedom" of the discretisation. Since in the Galerkin method test and trial space are the same for P_N (where P_N are arbitrary degree polynomials), test functions can be linear combinations of the same set of basis functions. In order to define the trial and test spaces, polynomials defined over an element are selected, of degree $n \leq N$, these are called P_N discretisations. For example P_1 is a piecewise linear function, where the value of the function needs to be known only at the vertices of an element, this can be seen in figure 2.4.1. Continuous Galerkin methods restrict (2.4.11) so that θ is continuous between elements, whereas discontinuous Galerkin methods allow θ to be discontinuous between elements. They use careful integration by parts within the equations to achieve this.

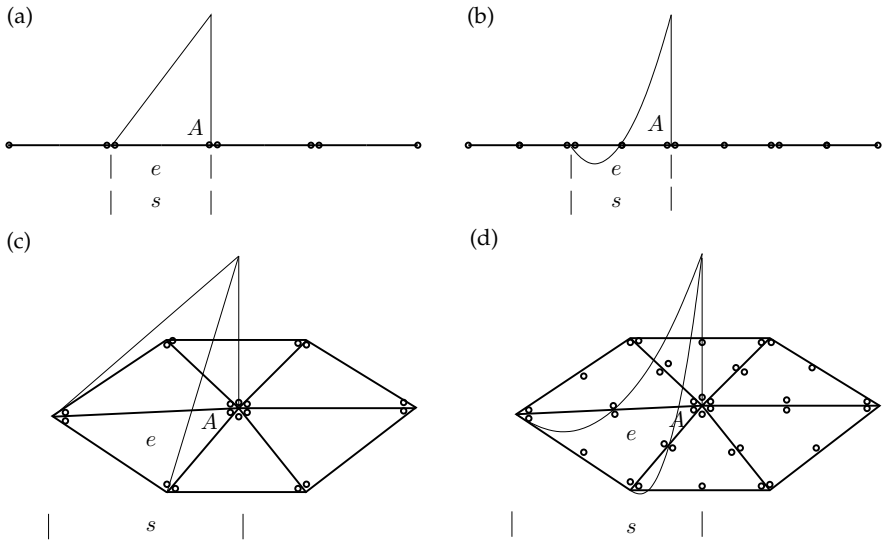


Figure 2.4.1.: One dimensional schematic of piecewise linear (a) and piecewise quadratic (b) discontinuous shape functions and the two dimensional equivalents (c) and (d) respectively. At A the value of the shape function is 1 while it takes the value 0 at surrounding nodes. The number of nodes per element e depends on the order of the polynomial (piecewise linear has one per vertex, while quadratic has more), the support s covers the same area as the element (extracted from AMCG (2015)).

To ensure that the equation being solved integrate to zero with all test functions, we must ensure that it integrates to zero with the basis function. This means that the discretised version of (2.4.10) becomes:

$$\sum_j \left\{ \int_\Omega \varphi_i \varphi_j \frac{d\theta_j}{dt} - \nabla\varphi_i . \mathbf{u} \varphi_j \theta_j + \nabla\varphi_i . \bar{\bar{\kappa}} . \nabla\varphi_j \theta_j \right\} = 0, \text{ for all } \varphi_i. \tag{2.4.12}$$

where $\theta = \sum_j \varphi_j \theta_j$ has been substituted in. This may be be written in matrix form as:

$$M\frac{d\theta}{dt} + A(\mathbf{u})\underline{\theta} + K\underline{\theta} = 0. \tag{2.4.13}$$

where M, A and K are termed discretisation matrices.

However, this thesis primarily uses discontinuous Galerkin methods, which will now be explained. The discontinuous test and trial functions associated with discontinuous Galerkin (DG) methods means that some of the terms in equation (2.4.9) are only well-defined across the elements and not between elements and integration by parts is necessitated when taking derivatives of discontinuous functions. Since over an element functions are continuous and well defined, we can define:

$$\int_e \varphi \frac{\partial\theta}{\partial t} - (\nabla . \varphi \mathbf{u})\theta + \nabla\varphi . \bar{\bar{\kappa}} . \nabla\theta + \int_{\partial e} \varphi \widehat{\mathbf{n}.\mathbf{u}\theta} - \varphi \widehat{\mathbf{n}.\bar{\bar{\kappa}}.\nabla\theta} = 0, \tag{2.4.14}$$

where the fourth and fifth terms demonstrate fluxes across the element boundaries, between elements. These terms do not have a unique value, however a consistent choice should be made, this highlights the importance of flux schemes for DG methods. Boundary conditions for DG are used in the same manner as CG, therefore by summing over all of the elements, e, and assuming weak boundary conditions

are applied, we have the following equation:

$$\Sigma_e \Bigg\{ \int_e \varphi \frac{\partial \theta}{\partial t} - (\nabla . \varphi \mathbf{u})\theta + \nabla \varphi . \bar{\bar{\kappa}} . \nabla \theta + \int_{\partial e \cap \partial \Omega_-^{D_c}} \varphi \mathbf{n}.\mathbf{u} g_D + \int_{\partial e \cap \partial \Omega_+^{D_c}} \varphi \mathbf{n}.\mathbf{u}\theta - \int_{\partial e \cap \partial \Omega^{D_c}} \varphi \mathbf{n}.\bar{\bar{\kappa}}.\nabla\theta + \int_{\partial e \cap \partial \Omega^{N_c}} \varphi \mathbf{n}.\mathbf{u}\theta - \varphi \mathbf{n}.\bar{\bar{\kappa}}.\nabla g_N + \int_{\partial e \ \partial \Omega} \widehat{\varphi \mathbf{n}.\mathbf{u}\theta} - \widehat{\varphi \mathbf{n}.\bar{\bar{\kappa}}.\nabla\theta} \Bigg\} = 0, \quad (2.4.15)$$

where + and − signs denote opposite sides of the element edge, $\partial\Omega^{D_c}$ represents the part of the boundary where Dirichlet boundary conditions are applied and $\partial\Omega^{N_c}$ is where Neumann boundary conditions are applied.

In this thesis simulations are run using the standard upwind advective flux. Upwind values are then used instead of θ in the advection terms of equation (2.4.15) and this equation is integrated by parts again. More information on upwind advective flux may be found in the Fluidity manual (AMCG, 2015). where ∂e_- is the inflow part of the element boundary, $\hat{\mathbf{u}}$ is the value of velocity in the interior of the element due to the averaging method, θ_{int} is the value of c inside the element and θ_{upw} is the upwind value. At boundaries θ_{int} cancels with θ_{upw}, but not at boundaries. Integration by parts twice within Fluidity may be more accurate for the numerical evaluation of the integrals, as the internal outflow boundary integrals are cancelled exactly.

DG methods can produce overshoots and undershoots for the scalar advection equation, therefore vertex-based limiters (Kuzmin, 2010) are used here to mitigate this. For linear elements, θ may be written as:

$$\theta = \bar{\theta} + \theta', \text{ where } \bar{\theta} = \frac{\int_e \theta dV}{\text{Vol}(e)}, \quad (2.4.16)$$

with the effect of the limiter meaning:

$$\theta = \bar{\theta} + \alpha\theta' \tag{2.4.17}$$

where a maximum value for α is found such that at each vertex θ is bounded by the maximum and minimum values of θ on every element at that vertex.

Discretisation of the diffusion term within Fluidity requires the Poisson equation to be solved $\nabla^2\theta = f$. This is difficult with DG methods since discontinuities in the trial and test function mean if one tries to find the Laplacian by integration parts of the integrals are not defined. Therefore this can instead be calculated for each element separately where functions are continuous, or a new vector may be introduced (and later eliminated locally) such that the Poisson equation may be re-written as a series of first-order equations. Here this is achieved using the so-called compact discontinuous Galerkin methods (Peraire and Persson, 2008).

Momentum equation

Discretisation of the momentum equation is very similar in form to that for the advection–diffusion equation and may be written in matrix form as equation (2.4.13) where $\mathbf{u}$ is substituted for θ and with the addition of a pressure gradient matrix C. Along with the advection-diffusion equation weak boundary conditions are also applied. Fluidity allows a separate set of test and trial functions to be used for the pressure term which occurs in the momentum equation (mixed-element discretisation).

The Boussinesq approximation is then dealt with by removing the density in the mass and advection terms. The discretised equation may be written as:

$$M\frac{d\mathbf{u}}{dt} + A(\mathbf{u})\underline{u} + Cor\underline{u} - K\underline{u} + C\underline{p} = \underline{b}(\rho') + \underline{\mathbf{F}}. \tag{2.4.18}$$

The continuity equation is also written in matrix form as $B^T\underline{u} = M_D\underline{g_D}$ and for CG pressure/DG velocity simulations the integration by parts of this means only continuous functions are integrated and extra face integrals are not required.

2.4.2. Temporal Discretisation

The advection-diffusion equation and momentum equation must also be discretised temporally, this process is similar for both equations. At the end of each timestep (n) the values of each field over the mesh are known we then need to calculate these values at the end of the next timestep $(n+1)$. Picard iteration (a non-linear iteration scheme) is used in Fluidity to give a number of tentative solutions ($\tilde{c}^{n+1}$ for variable c) for the value of the field and after the final Picard iteration they are accepted as the final values. The final iteration occurs either when a limit has been reached (here two iterations are used) or there is convergence. Picard iteration is required for the momentum equation due to the nonlinear term $\mathbf{u}.\nabla\mathbf{u}$. The process of Picard iteration is displayed in figure 2.4.2.

In order for the values at the end of the next time step to be calculated, a timestepping scheme is required. Here we use a Crank-Nicolson scheme for this, which uses a linear combination of values of the field (c) at the current time step (n) and unknown values at the next time step $(n+1)$:

$$c^{n+\theta} = \theta c^{n+1} + (1-\theta)c^n. \tag{2.4.19}$$

For a Crank-Nicolson scheme, which is a second-order semi-implicit, θ is set as $1/2$, but for other schemes must be in the range $[0, 1]$. Crank-Nicolson schemes are one of many time stepping schemes but are useful as they enable the use of larger timesteps as stability is maintained for larger Courant numbers compared with explicit schemes, where the Courant number must be below 1 (Crank and

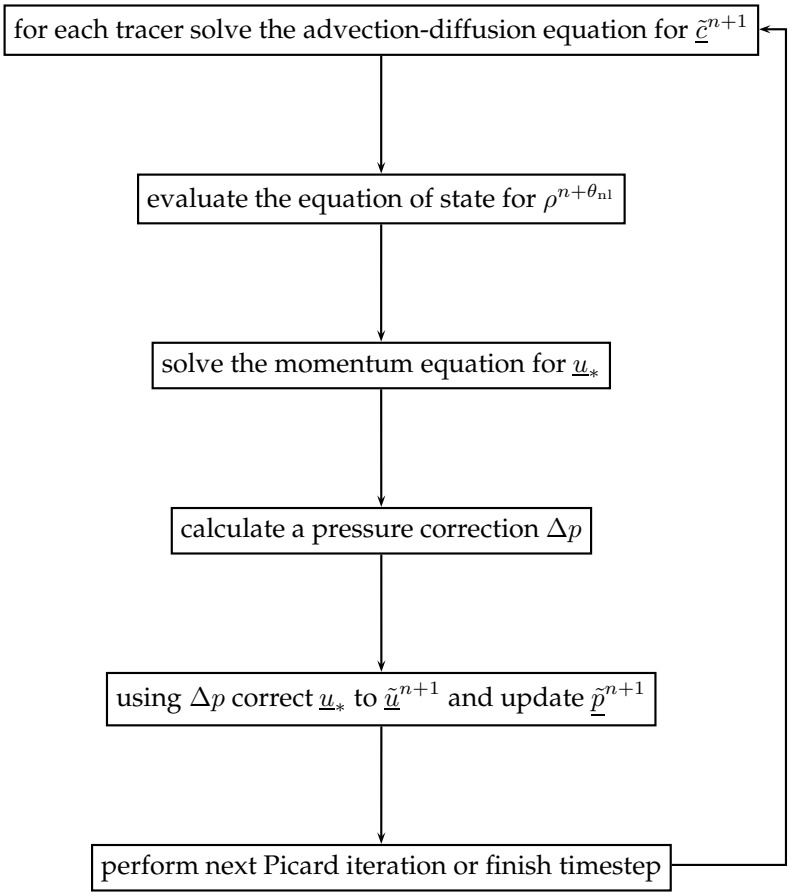


Figure 2.4.2.: The main stages of a non-linear iteration loop (AMCG, 2015)

Nicolson, 1947). The Crank-Nicholson scheme is used for both the advection–diffusion equation and the momentum equation in this thesis.

Non-linear relaxation is required where one variable is also used in the solution of another, for example **u** is required to solve the advection–diffusion equation. For these variables there are two possible values that may be used: those at n and those at $n+1$. The nonlinear relaxation paramater θ_{nl}, whose value is in the interval $[0, 1]$, is used to choose between these values. The value of the field to be

used is then given by:

$$\mathbf{u}^{n+\theta_{nl}} = \theta_{nl}\tilde{\mathbf{u}}^{n+1} + (1-\theta_{nl})\mathbf{u}^{n}. \tag{2.4.20}$$

At the first iteration however we usually use $\tilde{\mathbf{u}}^{n+1} = \mathbf{u}^{n}$ as we don't yet have a guess at $\mathbf{u}^{n+1}$. However, when a guess for pressure is used in the momentum equation the continuity equation is not satisfied, therefore a pressure correction step is used in order to maintain divergence free velocity by updating the values of pressure and temperature.

Since discontinuous Galerkin formulation is widely used in this thesis, we must consider the use of slope limiters (as described in section 2.4.1 (Finite Element Discretisation)) while temporally discretising the advection–diffusion equation, slope limiters only ensure a bounded solution in conjunction with an explicit advection scheme. However instead of treating the whole equation explicitly, certain terms may be treated explicitly while others implicitly. This is accomplished by first advecting the tracer and then allowing diffusion to occur, this two stage method splits the boundary condition components into Dirichlet and Neumann components. Due to the explicit treatment of advection, the timestep is split into subtimesteps to satisfy Courant number restrictions from the user, after each subtimestep the slope limiter is applied.

2.4.3. Linear Solvers

Once the equations required have been discretised they form a linear system of equations, these may be written in the form:

$$A\underline{x} = \underline{b}, \tag{2.4.21}$$

where A is a matrix, $\underline{x}$ is a vector of values that will be solved for and $\underline{b}$ contains terms that do not depend on x, both A and $\underline{b}$ are known. The values that we solve for are usually the discretised values of the field at the nodes. Since finite element discretisations with finite support (figure 2.4.1) are used the matrices are typically sparse (with few non-zero elements). Systems of sparse matrices may be solved efficiently using iterative methods avoiding the high costs of constructing the inverse of matrix A. Iterative solvers solve the linear system by constructing a sequence of approximations $\underline{x}^k$ that converge to the exact solution $\underline{\hat{x}}$. Once the residual, which is calculated at each iteration as follows:

$$\underline{r}^k = b - A\underline{x}^k, \tag{2.4.22}$$

falls below the value set by the user the iterative process finishes and $\underline{r}^k = A(\underline{\hat{x}} - \underline{x}^k)$ as $\underline{r}^k \to 0$ and $\underline{x}^k \to \underline{\hat{x}}$. Krylov subspace methods are used with preconditioners from the PETSc library to make matrix A easier to solve (Balay et al., 2014). The generalised minimal residual method (GMRES) is the Krylov subspace method used to solve the discretised momentum equation while the congugate gradient (CG) method is used for the solution of the temperature and pressure equations in most simulations here, with succesive over-relaxation (SOR) being the preconditioner used for velocity and temperature and a multigrid preconditioner is used for the pressure field (AMCG, 2015; Balay et al., 2014; Kramer et al., 2010). The relative solver tolerance error is set at 10^{-6} for all fields.

Krylov subspace methods are iterative methods, consider the following simple iterative method

$$\underline{x}^{k+1} = \underline{x}^k + \alpha_k \underline{r}^k, \tag{2.4.23}$$

where α_k is a scalar coefficient, which may change each iteration. We can then

show that

$$\underline{r}^{k+1} = \underline{r}^{k} + \alpha_k A\underline{r}^{k}. \tag{2.4.24}$$

This means that the residual at each iteration is a linear combination of the previous iteration and A applied to this previous residual, therefore the residual may be given as a linear combination of vectors:

$$r^0, A\underline{r}^0, A^2\underline{r}^0, ..., A^k\underline{r}^0. \tag{2.4.25}$$

These vectors span a subspace called a Krylov subspace. GMRES is built from all of these vectors, which must be stored in memory, this becomes expensive when a large number of iterations are used. However, CG is more efficient at finding an optimal solution to the Krylov subspace without storing all of its vectors, and maybe used where the matrix of the field is symmetric positive definite (SPD), where:

$$\langle \underline{x}, A\underline{x} \rangle \geq 0 \quad \forall x. \tag{2.4.26}$$

Preconditioners, such as SOR, improve the conditioning of the system. However where multiple length scales are involved SOR is not as effective, for these scenarios multigrid methods are more effective solving the problem on a range of coarse to fine meshes (Trottenberg et al., 2000). The multigrid methods used in Fluidity are algebraic multigrid methods (AMG) (Stüben, 2001), which give good results for unstructured mesh discretisations, vertical lumping is also used in Fluidity which is useful for large aspect ratio simulations as it helps with the large separations between horizontal and vertical modes in the pressure equation (Kramer et al., 2010).

2.4.4. Mesh Adaptivity

The adaptive mesh capabilities of Fluidity are one of the main advantages of the model for the study of diapycnal mixing, as shown by Hiester et al. (2011). Mixing involves the transfer of energy, in the form of momentum and heat from large to small scales dictated by the Kolmogorov and Batchelor scales. By focusing high resolution in areas which contain the most complex, or important dynamics, and reducing resolution in areas of lower complexity, the range of scales may be captured fully without the need for high resolution throughout the entire domain. Thus allowing significant computational savings to be made.

Three distinct forms of adaptivity are possible and may be combined, h-adaptivity which changes the total number of nodes/elements and changes the mesh connectivity (Berger and Colella, 1989), r-adaptivity which maintains the same mesh connectivity while moving vertex positions (Budd et al., 2009) and p-adaptivity where the mesh size is fixed (unlike h and r adaptivity) but the degree of the piecewise polynomials is locally varied (Babuška and Suri, 1994).

Fluidity employs a particularly powerful method of hr-adaptivity with the greatest flexibility to adapt to features as it is not constrained by previous meshes. Mesh adaptivity within Fluidity involves three main stages: metric formation, mesh optimisation and mesh-to-mesh interpolation (AMCG, 2015; Piggott et al., 2008a; Hiester et al., 2014, 2011; Hill et al., 2012; Parkinson et al., 2014; Smith et al., 2016). The approach taken to parallise these steps is described in (Piggott et al., 2008b; Gorman et al., 2009). Metric formation determines how to adapt the mesh, metric optimisation makes changes to the mesh based on that metric, and mesh-to-mesh interpolation moves data between the old and new meshes (AMCG, 2015). These three stages will be examined in sections 2.4.4 to 2.4.4, however further detail may be found in the Fluidity manual (AMCG, 2015) and relevant referenced

papers therein.

Metric formation

This first stage of adaptivity in Fluidity involves the formation of a metric tensor, M from the Hessian of the fields the user has selected to adapt to. This metric is a positive definite tensor providing a spatially varying field of desired anisotropic element sizes (George and Borouchaki, 1998; AMCG, 2015). A tensor field is required to store this information as mesh size is a function of space as well as direction, and edge lengths can vary over the dimensions of the domain.

The metric is used during optimisation of the mesh to control the distance between points (or the edge lengths). The distance, d_M, between two points x and y is given by the norm of the vector $\mathbf{v} = x - y$ in metric space:

$$d_M(x, y) = \|\mathbf{v}\|_M = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle}, \tag{2.4.27}$$

where $\langle \mathbf{v}, \mathbf{v} \rangle = \mathbf{v}^T M \mathbf{v}$. This gives the curvature of the field with respect to which the Hessian was computed in the direction of $\mathbf{v}$. A mesh is deemed to have satisfied this metric when $\mathbf{v}^T M \mathbf{v} = 1$, so that all edges have unit length with respect to the metric. This means if we interpret x and y as two element vertices, a large M means we need these vertices to be close (i.e. a small edge length) and vice-versa when M is small we allow a large edge.

For each field, f, that the user has decided to adapt to a separate metric, M_f, will be formed from the weighting defined by the user (the interpolation error bound ε, which may vary across the domain if wished) and the Hessian of that field, H. The default metric in Fluidity is the absolute metric:

$$M_{\infty f} = \frac{| H |}{\varepsilon}, \tag{2.4.28}$$

where the Hessian of a field f is $H \equiv \nabla^T \nabla f$ which is diagonalised as $H = Q\Lambda Q^T$ using the diagonal matrix Λ of eigenvalues of H and Q the matrix of normalised eigenvectors. The absolute value of this modified Hessian ($|H| = Q\,|\Lambda|\,Q^T$) is then used, as a positive value is required for mesh adapting. This essentially means we are concerned with high versus low curvatures of fields, but not the sign of this curvature. The interpolation error bound is defined as, for the L_∞ norm for an element e in the mesh (Frey and Alauzet, 2005):

$$\varepsilon = \mid f - \Pi_h f \mid_\infty \leq C \max_{\mathbf{x}\in e} \max_{\mathbf{v}} (\mathbf{v}^T\, |H(\mathbf{x})|\, \mathbf{v}), \tag{2.4.29}$$

where $\Pi_h f$ is the P_1 interpolation of field f, f is the exact solution of that field and C is a constant independent of the mesh (Castro-Diaz et al., 1997; Frey and Alauzet, 2005).

Adaptive simulations in this thesis also use the p-metric instead of the absolute metric where p=2, this is referred to as the M_2 metric (Chen et al., 2007). This metric is investigated in Hiester et al. (2014) and also Hiester (2012) for turbulent mixing simulations, where it was found to perform favorably. The M_2 metric is constructed by considering the L_2 norm of the interpolation error defined in two dimensions as:

$$M_{2f} = \frac{(\det(|H|))^{-\frac{1}{6}}|H|}{\varepsilon} = \det(|H|))^{-\frac{1}{6}} M_{\infty f}. \tag{2.4.30}$$

For this metric the L_2 norm of the interpolation error is bounded by:

$$\mid f - \Pi_h f \mid_2 \leq C N^{-1} \| \sqrt{\det H} \|_{\frac{2}{3}(\Omega)}. \tag{2.4.31}$$

By using the M_2 metric the small-scales in the dynamics should be better represented through the scaling by the field's curvature which encourages additional

resolution at higher and lower curvatures (Loseille and Alauzet, 2011).

Following the formation of metrics for each field a final metric must be found to guide the adapt that combines the information from all field metrics, this is given by the intersection of the superposition of individual field metrics. For the first two field metrics this is given as M =intersection(M_{f1}, M_{f2}) and for all subsequent field metrics (illustrated for a third field) M =intersection(M, M_{f3}) (Castro-Diaz et al., 1997).

Mesh optimisation

Mesh optimisation involves updating the mesh through a series of local topological and geometrical operations (Piggott et al., 2009) resulting in a new mesh, which better adheres to the element size requirements encoded within the metric. Other methods of adaptive remeshing are global remeshing (Peraire et al., 1987), where a new mesh is formed for the whole domain and local remeshing (Hassan et al., 1998), where areas of the mesh are entirely removed and completely remeshed.

During mesh optimisation several geometrical and local topological operations are performed in order to change the previous mesh, without being constrained by this previous mesh. These operations include: edge-splitting, edge-collapsing, edge to edge swapping and node movement in 2D (shown in figure 2.4.3), as well as face to edge swapping and edge to face swapping in 3D. All simulations in this thesis are two dimensional. In two dimensions the algorithms used for these operations are those of Vasilevskii and Lipnikov (1999).

As shown in figure 2.4.3 edge-splitting involves the addition of a node along an edge, splitting the element into more than one, while edge-collapse involves a node being removed and the surrounding elements merging into one and also resulting in the deletion of the addtional nodes associated with the original multiple elements.

In order to decide where to perform these operations a functional is used, which

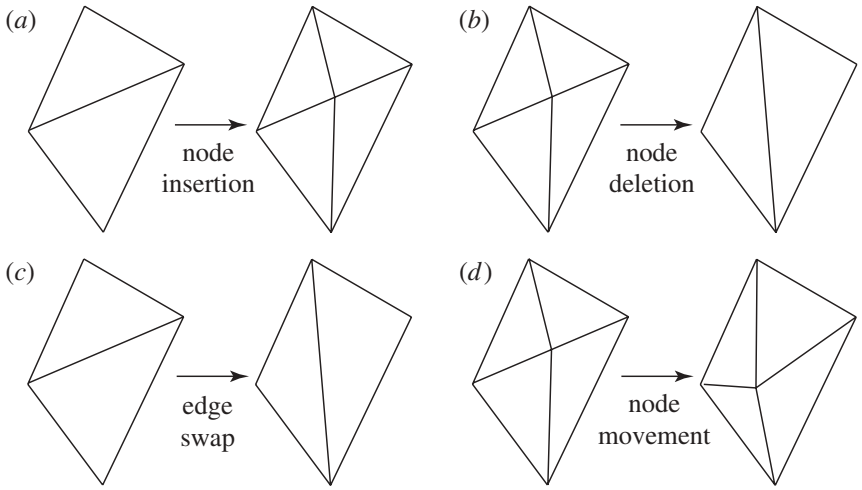


Figure 2.4.3.: Mesh optimisation processes for two dimensional simulations: (a) node insertion or edge splitting, (b) node deleteion or edge collapse, (c) edge swap and (d) node movement (taken from Piggott et al. (2009)).

determines the quality of an element and how this can be improved. The functional works to compare an element with the size and shape specified by the metric to decide how this element must be changed. For 2D simulations Fluidity uses the functional of Vasilevskii and Lipnikov (1999). The aim of mesh optimisation is to minimise this functional by moving through each element in turn.

Mesh interpolation

The final stage in the adaptive loop involves interpolating data, including solution fields, from the pre-adapted mesh onto the new post-adapt mesh. A number of options are available for this within Fluidity: no interpolation for pressure (where the field is recomputed from scratch using a pressure projection based approach (Ford et al., 2004)), consistent interpolation (Farrell et al., 2009) and Galerkin projection (Farrell et al., 2009), as well as geostrophic interpolation (Maddison et al., 2011) for velocity. The Fluidity manual ((AMCG, 2015)) encourages the

use of consistent interpolation unless for a discontinuous prognostic field (velocity and temperature in simulations presented in chapter 3), or if conservation of a particular field is crucial to the dynamics, in which case the use of Galerkin projection should be considered (Hiester, 2012).

Consistent interpolation involves taking the field values evaluated at the nodes of the old mesh as coefficients of the relevant shape functions. Mesh association (Manolopoulos et al., 2010) is used to identify which basis function to evaluate for a given node in the new mesh and thus to identify in which element (parent element) of the previous mesh each node of the new mesh lies. Finite element basis functions of the parent mesh are then used to perform the interpolation $f_j^T = \Sigma_{i=1} f_i^D \hat{\phi}_i^D$ (where the function f^T at node j is f_j^T, f^D is the function evaluated on the old (donor mesh) and f^T is the function evaluated on the new (target) mesh).

However, while this method is bounded and computationally cheaper than the other interpolation methods as it does not require the use of a super mesh, it does not necessarily preserve maxima, minima or the global integral of the field. In addition solutions must be well-defined at the mesh nodes in order for the use of the basis functions to work and hence it is unsuited to discontinuous Galerkin methods and is inherently non-conservative (AMCG, 2015).

Galerkin interpolation is conservative, non-dissipative but non-bounded and is the most accurate for minimising the L^2-norm of the residual of the difference between the function evaluated on the old and new mesh (Farrell et al., 2009). This residual is given as:

$$\|f_D - f_T\|_2^2 = \int_\Omega (f^D - f^T)^2 d\Omega. \tag{2.4.32}$$

In order to minimise this function, the minima is found by taking the differential with respect to target weights $\hat{f}_j^T$ and setting this to zero, and as the L_2 norm is

convex this gives:

$$\int_\Omega \Sigma_i \hat{f}_i^T \phi_i^T \phi_j^T d\Omega = \int_\Omega \Sigma_i \hat{f}_i^D \phi_i^D \phi_j^D d\Omega. \tag{2.4.33}$$

This may then be solved as a linear matrix equation, where $\int_\Omega \Sigma_i \phi_i^T \phi_j^T d\Omega$ is the mass matrix (M_T) and $\int_\Omega \Sigma_i \phi_i^D \phi_j^D d\Omega$ is the mixed matrix between the old mesh T and the new mesh D (M_{TD}) giving the system: $M_T f^T = M_{TD} f^D$. This system may be solved to find f^T after the mass and mixed matrices have been found. While the mass matrix may be easily calculated using the basis functions, the mixed matrix requires a super mesh (Farrell and Maddison, 2011). The use of the super mesh (an example of a super mesh is shown in 2.4.4) allows the preservation of the global integrals of the solution fields, and is formed as the union of the original mesh and the new mesh. Basis functions of the old and new meshes are continuous within a child element. Elements of the old mesh are therefore decomposed into children elements in the super mesh using mapping to calculate the relevant intersection integrals with the old mesh, in order to find the mixed matrix.

Galerkin interpolation introduces minimal diffusion compared to consistent interpolation, even though over and under shoots in the solution field are corrected by an additional diffusion step in the process (Farrell et al., 2009).

Other options for interpolation include Grandy interpolation which involves geometrically computing the intersection volume between the old and new mesh zones, pseudo-consistent which requires averaging for nodes on element boundaries and geostrophic which aims to preserve geostrophic balance in the interpolated velocity field.

Vertically structured, vertical only and fully unstructured adaptivity

Adaptive mesh capabilities in Fluidity are available for both fully and partially unstructured meshes (figure 2.4.5). *Fully unstructured adaptivity* involves adapting in multiple dimensions simultaneously, treating each identically and resulting in

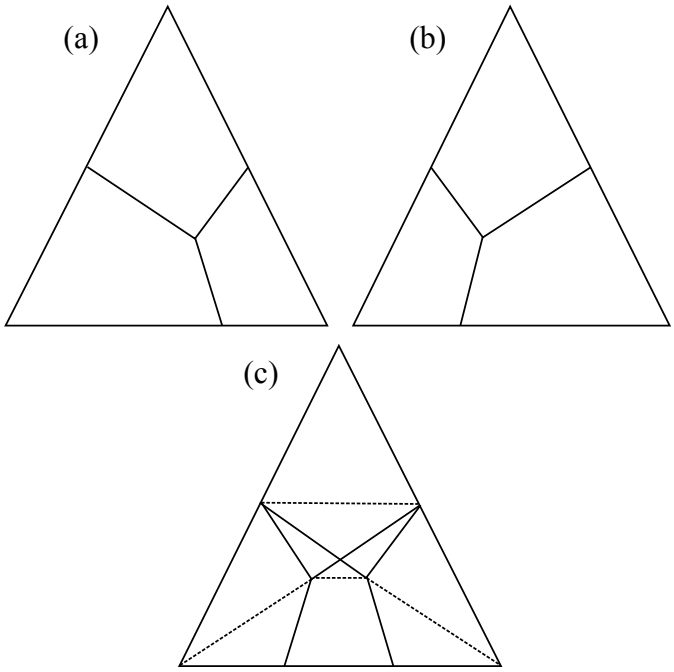


Figure 2.4.4.: An example of a super mesh (c) formed from a donor mesh (a) and a target mesh (c). The super mesh shows dashed lines where polyhedra have been divided into triangles for the calculation of integrals (adapted from Piggott et al. (2009)).

meshes which arbitrarily tessellate the domain. Partially structured/ unstructured adapted meshes seek to maintain some of the structure present in the mesh, typically the vertical alignment of the nodes. This is achieved with *vertically structured* adaptivity which involves adapting the horizontal and vertical separately. Adaptation is performed first in the horizontal, resulting in varying horizontal resolution and then a column of nodes is created under each surface node, with the spacing between nodes in each column being obtained through an independent 1D adaptive process. Once the vertical adapt is complete nodes are connected to form triangular (in 2D vertical slice set-ups) or tetrahedral (in 3D) elements (AMCG, 2015; Smith et al., 2016). It should be noted that, adaptation along only the horizontal or only the vertical, while keeping the other fixed, are also possible scenarios within this framework (figure 2.4.5). Vertically structured adaptivity is

advantageous in some situations, particularly for the stability of problems involving large-scale ocean aspect ratios, where spurious horizontal pressure gradients can be introduced (AMCG, 2015; Smith et al., 2016).

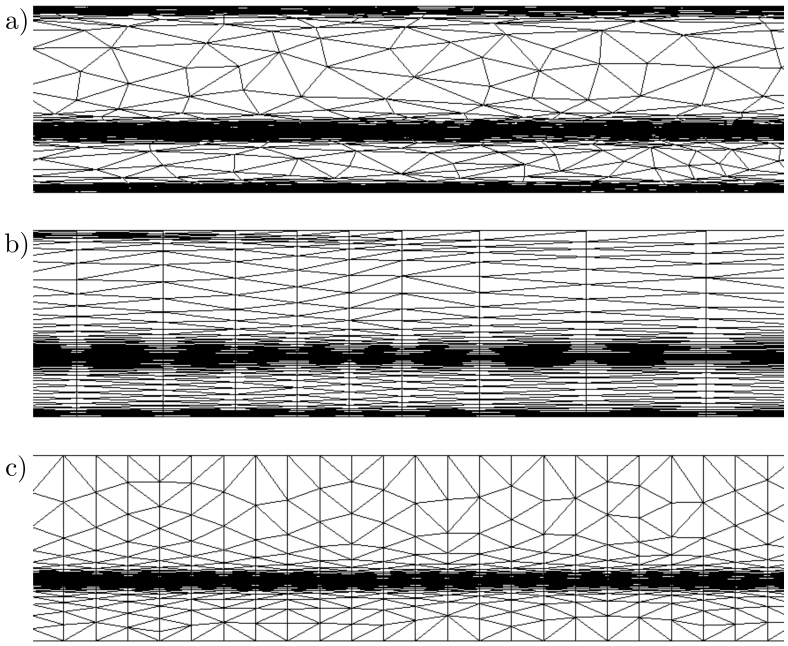


Figure 2.4.5.: Comparison of different types of mesh adaptivity used in this thesis; mesh segments shown are extracted from the same spatial and temporal location from meshes used for the study of the tilting tank in chapter 3. Mesh types are fully unstructured adaptivity (a), vertically structured adaptivity (b), and vertical only adaptivity (c). The unstructured mesh (a) shows how various element sizes and orientation are used to capture an interface, whereas meshes (b) and (c) also show various element sizes but also the vertical alignment of nodes dictating a columnar structure imposed on the mesh. The lack of a horizontal adaptive step for (c) results in uniform column widths, whereas the separate horizontal adaptivity step for (b) results in differing column widths according to the local dynamics. The final case of adapting the horizontal, but maintaining a prescribed (non-adapted) vertical mesh size is not displayed as it is not used in this paper where sharp vertical features need to be captured and tracked.

Metric advection

Metric advection is a process whereby the metric field (described section 2.4.4) is advected forward by the current flow velocities in the period of time between mesh adapts. As metric advection advects forward an estimate of the resolution needed per timestep before the next adapt is given (Wilson, 2009). The advection involves the movement of the components of the metric as scalars by control volume advection methods, these methods use a timestep which is determined by the Courant number restriction imposed by the user.

The effect of this process is that an area of interest should not as easily propagate away from an area of high resolution, as areas of high resolution are effectively "pushed forwards" with the dynamics. This process may increase the number of nodes in the simulation but could improve the representation of dynamics in some scenarios and has proved useful in simulations of complex dynamics by Hiester et al. (2011) to help maintain a sharp representation of interfaces. It also means that the mesh should not need to be adapted as often. Adapting the mesh too often, while meaning that the mesh is optimised for the state of the system more frequently, can be costly and can increase numerical diffusion due to the more frequent application of interpolation required.

Parallelisation and cost

In order to resolve the scales required to capture vertical mixing, high mesh resolutions are required resulting in large computational requirements. While adaptive meshes greatly decrease the number of nodes required to resolve a scenario, it should be acknowledged that there are also overheads associated with the processes of metric formation, remeshing and interpolation, however generally these are relatively small. During mesh interpolation, methods requiring the formation of a super mesh are slower than consistent interpolation. Metric-advection also

increases the cost of metric formation. Generally the cost of forming the Hessian is equivalent to the cost of forming any of the metrics referred to in section 2.4.4. A more thorough analysis of the costs can be seen for the lock-exchange problem in Hiester et al. (2011).

In order to meet the demands of high resolution simulations efficiently, parallel computation is required. Parallesiation of adaptivity methods must pay particular attention to dynamic load balancing, minimising communication overheads and the consistency of node and element information that is shared across sub domains boundaries. Fluidity has been shown to scale well to large numbers of processors as studied by Guo et al. (2010) for a wind-driven gyre.

During mesh optimisation resolution may be added to certain areas and reduced in others, this may have the affect of greatly increasing the number of nodes in one partition while decreasing them in another, this creates load imbalances. These possibly large load imbalances decrease computational efficiency, therefore Fluidity uses Zoltan (Devine et al., 2002) a library for parallel partitioning and distribution in order to re-partition the mesh amongst the processors being used. During intermediate adapt iterations edge-weighting is important, high edge-weighting is applied to low quality elements to ensure that they are not positioned on partition boundaries and therefore are not locked for the next iteration of the adapt, in order that this is prioritised the load is allowed to be slightly more imbalanced during these adapts. These processes, including changing partitions and migrating data, increases the computational cost. However relatively speaking these overheads are small when compared to the rest of the adaptive mesh procedure (Gorman et al., 2009).

2.5. The MITgcm model

MITgcm (Marshall et al., 1997a,c), a finite-volume based fixed structured mesh model, is a widely used model for ocean and atmospheric applications at a range of scales. MITgcm has been used extensively by (Xing and Davies, 2006, 2009a, 2010; Legg, 2004a,b; Legg and Huijts, 2006; Legg and Klymak, 2008a) amongst others for the study of flow over bumps using three dimensional or two dimensional vertical slice models at a range of scales. The use of MITgcm in this study allows us to benchmark Fluidity in chapter 3 and also to provide results quickly with a validated model for all other chapters. MITgcm is suitable for these studies as it also possesses non-hydrostatic options where pressure is obtained diagnostically and split into three components: surface (P_s), hydrostatic (P_{h}) and non-hydrostatic (P_{nh}). And although it uses z-coordinates in the vertical, the use of partial steps to better represent bathymetry minimises the disadvantages associated with this.

2.5.1. Sponge regions

Simulations conducted in chapters 4 and 5 use periodic boundary conditions where ten cell sponge regions are implemented at either end of the domain. These sponge regions are applied using the RBCS package to relax the temperature field back to the initial condition over the tidal period. This relaxation is implemented at each grid point by:

$$\frac{dT_{\mathrm{new}}}{dt} = \frac{dT_{\mathrm{old}}}{dt} - \frac{M_{rbc}}{\tau}(T_{\mathrm{old}} - T_{rbc}), \tag{2.5.1}$$

where M_{rbc} is the mask containing values of 1 within the sponge region and 0 elsewhere, $1/\tau$ is the relaxing timescale, T_{rbc} is the value the field is relaxed to, T_{old} is the temperature field before relaxation and T_{old} is the temperature field after relaxation.

2.5.2. Temporal discretisation

Time stepping methods employed in MITgcm are the quasi-second order Adams-Bashforth methods. Since problems studied here are well stratified, meaning the determination of a stable time-step may be limited by internal gravity waves, the staggered time step option is used which can be more efficient as it staggers in time the thermodynamic and flow variables. Thermodynamic variables are solved half a timestep later with the updated flow field and then pressure is computed to step forward, adding second order accuracy and stability, more information on this can be found within Adcroft et al. (2013). The problem however with this method is that the flow used to advect tracers is not located at the expected time, however this appears to cause no adverse affects (Adcroft et al., 2013).

2.5.3. Spatial discretisation

All simulations here use Cartesian coordinates. In order to maintain non-divergence of flow, components of velocity are staggered in space forming an "Arakawa C-grid" (Arakawa and Lamb, 1977). In the vertical z-coordinates are used with partial steps, which allows the bathymetry to take a piece-wise constant representation. The minimum physical thickness of these cells is determined by the user and a fraction of 0.2 is used in these studies, anything below this is rounded either to 0.2 times the cell depth or 0 whichever is closer.

Finite volume method

The finite volume method uses the integral forms of the governing equations and is a grid point method in the interior (second order centered finite difference). During the finite volume method the domain is divided into control volumes over which the equations are solved. Control volume discretisations are inherently conservative as

the flux into the control volume must equal the flux out of it. The finite volume method uses the *weak solution* as a solution for each finite volume, to the integral form of the conservation laws. The weak solution can be differentiable or made up of piece-wise constant elements.

To demonstrate this the 1D advection-diffusion equation may be written:

$$\frac{\partial c}{\partial t} + \frac{\partial}{\partial x}(uc - \kappa(\frac{\partial c}{\partial x})) = 0, \tag{2.5.2}$$

this is then discretised by integration over finite lengths (Δx_i)

$$\Delta x \frac{\partial c}{\partial t} + \delta_i F = 0, \tag{2.5.3}$$

which involves averaging the tracer c_i over the interval Δx_i and where F is the flux which is defined as:

$$F = u\bar{c} - \frac{\kappa}{\Delta x}\frac{\partial c}{\partial i}. \tag{2.5.4}$$

where i is the current index in the mesh relative to the x direction and $\bar{c}$ is the average of c over the interval Δx_i. The method for obtaining the flux can introduce truncation errors and in addition multiple methods are available to find $\bar{c}$, however all simulations here use second order flux limiters. If second order centered interpolation methods without flux limiters had been used however the result would be the same as those for finite difference for the interior and would only differ at boundaries.

Advection–diffusion equation

For the discretisation of the advection–diffusion equation in this thesis the second order flux limiter scheme is used, where flux is limited using the Superbee limiter. This non-linear method means that timestepping is switched to forwards in time.

Test cases shown in the manual show that flux limited schemes avoid the addition of extrema to model results at low and high Courant numbers and while they are slightly diffusive these affects are not significantly larger than for other schemes. Use of the flux limited method for passive tracers was also found to be vital to prevent values outside the 0 to 1 range specified as described in chapter 5, however second order methods were found be adequate to represent dynamics accurately.

Momentum equation

The momentum equations are discretised using the Eulerian flux form momentum equations as opposed to the "C-grid" discretisation used for the continuity equation. Advection of momentum is conducted similarly to tracers which conserves second moments and thus also kinetic energy.

CHAPTER 3

SIMULATING MIXING DUE TO INTERNAL WAVE GENERATION, PROPAGATION AND BREAKING WITH ADAPTIVE MESH METHODS

Contents

3.1 Synopsis 59

3.2 Introduction 60

3.2.1 Sea Lochs and Fjords 60

3.3 Numerical Set-up 61

3.3.1 Fluidity discretisation 62

3.3.2 Mesh Adaptivity 63

3.4 Mixing Diagnostics 63

3.4.1 Critical Richardson area 63

3.4.2 Background potential energy and turbulent mean tracer variance decay rate 65

3.4.3 Background Potential Energy Method 66

3.4.4 Turbulent Mean Tracer Variance Decay Rate Method 73
3.5 Simple test case: Internal seiche 77
3.5.1 Set-up 77
3.5.2 Results 78
3.6 Lab validation: tilting-tank 83
3.6.1 Set-up 84
3.6.2 Results 85
3.7 Pseudo real-world application: Loch Etive 89
3.7.1 Set-up 91
3.7.2 Initial benchmarking against Berntsen et al. (2009) 94
3.7.3 Results 100
3.8 Conclusions 118

3.1. Synopsis

Numerical simulations of an internal seiche, tilting tank and a pseudo real-world example of Loch Etive are performed using a combination of fixed structured and unstructured and partially structured adaptive meshes. These simulations allow the assessment of the influence of each of these types of mesh adaptivity on spurious numerical mixing and model dynamics. All models have some spurious numerical mixing, but since real values of physical mixing are very small, the study of the amount of spurious numerical mixing is important as we require models with very little numerical mixing in order that it does not mask the real mixing. In addition, these simulations allow for the comparison of two mixing diagnostics: background potential energy and turbulent mean tracer variance decay rate, which are found to be comparable. These diagnostics are compared as the turbulent mean tracer variance decay rate presents a simpler method of calculation for simulations with more complex bathymetry and open boundaries.

Fluidity successfully captures the complex dynamics associated with each of the three cases, with all of the mesh types used and compares well with physical studies of Loch Etive. The application of mesh adaptivity is assessed by its effect on a number of diagnostics: critical Richardson area, background potential energy and turbulent mean tracer variance decay rate. Overall adaptivity is found to be favourable, with fully unstructured adaptive simulations of Loch Etive in particular requiring between one and two orders of magnitude fewer degrees of freedom than a fixed uniform mesh. Adaptive mesh simulations are sensitive to the field weights and minimum edge length and so a metric is required in order to allow convergence to be observed, here the critical Richardson area is used as it was found to give a good representation of convergence.

3.2. Introduction

Numerical simulations of an internal seiche, tilting tank and a pseudo real-world example based upon Loch Etive are performed using a combination of fixed structured and unstructured, and partially structured adaptive meshes. These simulations predominantly use Fluidity but also use MITgcm for further benchmarking. The influence of adaptive mesh settings are investigated using Fluidity including several types of partially structured and unstructured mesh types, field weights and minimum edge lengths.

The simple scenario of an internal wave is compared against published model results from the SUNTANS model and linear theory using non-dimensionalised wave speeds, while the more complex internal seiche is compared against laboratory results for the interface height at the centre of the tank. The pseudo real-world example of Loch Etive includes a convergence study guided by the critical Richardson area and vertical velocity. Further adaptive simulations are guided by convergence of critical Richardson number which is shown to be effective at representing convergence. Additionally the effect of adaptive simulations is measured by its influence on mixing rates qualitatively for all scenarios using the background potential energy and turbulent mean tracer variance decay rate which are found to be comparable metrics with which to quantify this measure.

3.2.1. Sea Lochs and Fjords

Fjords and sea lochs (such as Loch Etive) have been proposed as analogous physical "laboratories" to help understand oceanic mixing processes more generally in the ocean due to their smaller dimensions (Berntsen et al., 2008; Vlasenko et al., 2002). Fjords are long, thin, deep, steep sided glacial valleys that have been flooded by the sea during a sea level rise and are therefore open to the sea at one end

and closed at the opposite end (making the dynamics easier to understand (and model) than those in the open ocean). Fjords also generally feature shallow sills at their seaward facing entrances, these bathymetric features are formed by glacial deposits and typically mark the connection to the sea (Thorpe, 2005). These sills trap dense, cold, saline water at depth within the fjord and as a result fjords are typically strongly stratified.

The oscillating barotropic tidal flow over the sill can cause a cascade of energy from large to small scales (Stigebrandt, 1999; Berntsen et al., 2008; Smethier Jr, 1981; Stashchuk et al., 2007). The exact response however according to Stigebrandt and Aure (1989) depends on the maximum barotropic flow (U) and the internal wave speed in the fjord (c_i). For "jet-type" basins such as Loch Etive during spring tides, where $U > c_i$, the response is a well-mixed bottom boundary layer and turbulent overturning vortices down-stream which may be accompanied by an internal wave only if the difference in speeds is small (Stigebrandt and Aure, 1989; Stigebrandt, 1999; Xing and Davies, 2006). Whereas for "wave-type" basins where $U < c_i$ solely internal waves are generated by the sill and propagate away, these basins are less energetic in the sense of smaller-scale turbulent structures than "jet-type" basins (Stigebrandt and Aure, 1989). However as velocity changes throughout the tidal cycle the system may alternate between states.

3.3. Numerical Set-up

Both Fluidity and MITgcm are used to solve the Navier-Stokes equations under the Boussinesq approximation, with the advection-diffusion equation for temperature as an active tracer and a linear equation of state in this chapter. Further equations and information on Fluidity and MITgcm may be found in chapter 2.5.3.

Fluidity is used here with several variants of fixed structured (using triangu-

lar elements in vertical-slice 2D domains) as well as unstructured and partially structured adaptive meshes (chapter 2.3) in order to examine the effects of different adaptive mesh strategies and parameters on the representation of vertical mixing. In addition, motivated by previous related work in the literature (Xing and Davies, 2006, 2009a, 2010), MITgcm (Marshall et al., 1997a,c), a finite-volume based fixed structured mesh model, has also been used to provide a further comparison and benchmark for the seiche. Note that all MITgcm simulations presented here are non-hydrostatic, use third-order direct space time schemes for advection of temperature, and a centered second-order scheme for momentum advection with staggered baroclinic Adams-Bashforth time stepping. Simulations that require a free-surface use the implicit linear free-surface option in MITgcm.

3.3.1. Fluidity discretisation

In order to discretise the Navier-Stokes equations using the finite-element method several choices are possible to represent velocity and pressure. All simulations presented here use a piecewise linear discontinuous Galerkin (unless otherwise stated for comparison) finite-element formulation for the velocity field with explicit sub-cycling of advection (with the number of sub cycles dictated by imposing a maximum Courant number of 0.1 per subcycle) and with the application of a slope limiter (Kuzmin, 2010) to help maintain stability (described further in section 2.4.2). Pressure is represented using a piecewise quadratic continuous Galerkin formulation. The resulting system is integrated in time using a Crank-Nicolson timestepping scheme and two Picard iterations are used per timestep to treat the nonlinear advection term in order to maintain temporal accuracy. A piecewise discontinuous Galerkin discretisation is also used for the temperature field (again with explicit sub-cycling of advection, with slope-limiting and with a maximum Courant number of 0.1 per subcycle).

3.3.2. Mesh Adaptivity

Adaptive mesh capabilities in Fluidity are available for both fully and partially unstructured meshes, both of which are used in this chapter to assess their influence on simulation dynamics and diapycnal mixing. The process by which mesh adaptivity is achieved is described fully in section 2.4.4 but essentially involves three main stages within Fluidity: metric formation, mesh optimisation and mesh-to-mesh interpolation (AMCG, 2015; Piggott et al., 2008a; Hiester et al., 2014, 2011; Hill et al., 2012; Parkinson et al., 2014; Smith et al., 2016).

3.4. Mixing Diagnostics

Diapycnal mixing is assessed here through the analysis of three different diagnostics: background potential energy, turbulent mean tracer variance decay rate, and the critical Richardson area. These diagnostics were chosen based on previous similar studies (Xing and Davies, 2006, 2009a, 2010; Berntsen et al., 2009, 2008; Tseng and Ferziger, 2001; Hiester, 2012; Hiester et al., 2014; Martin, 2014; Burchard and Rennau, 2008)

3.4.1. Critical Richardson area

The simplest diagnostic (calculated as Berntsen et al. (2008)) is the critical Richardson number area which is purely used to examine the extent of the domain that is dynamically unstable and could result in turbulent mixing. The Richardson number:

$$R_i = \frac{N^2}{\left(\partial u/\partial z\right)^2}, \tag{3.4.1}$$

where N is the buoyancy frequency and u is the horizontal velocity, is the ratio of the vertical density stratification to the vertical shear in velocity. Flow be-

comes unstable when the critical Richardson number ($R_i < 0.25$) is reached and significant mixing may occur. The area of the domain which achieves a critical Richardson number (A_{Ri}) is calculated at each timestep over the two-dimensional domain via

$$A_{\mathrm{Ri}} = \int_{0m}^{21500m} \int_{H(x)}^{\eta(x,t)} \delta(x,z)\, dz\, dx, \tag{3.4.2}$$

where

$$\delta(x,z) = \begin{cases} 1 & \text{if } R_i < 0.25,\ 6500 < x < 10000 \text{ and } -77/H(x) < z < -8/H(x) \\ 0 & \text{otherwise} \end{cases} \tag{3.4.3}$$

and a time-average is calculated ($\overline{A}_{\mathrm{Ri}}$)

$$\overline{A}_{\mathrm{Ri}} = \frac{1}{T} \int_0^T A_{\mathrm{Ri}}\, dt. \tag{3.4.4}$$

For this study the area calculated over is limited to a sub-region inland of the sill (figure 3.7.1), which is given by the limits in equation (3.4.3). This area prevents pollution by potential anomalous values close to the boundaries, limiting changes in $\overline{A}_{\mathrm{Ri}}$ to changes in the model dynamics. The vertical limits are chosen so that anomalous values close to the surface and bottom boundary are avoided while a fraction of the domain depth is chosen so that the box follows the edge of grid cells when σ-coordinates are used. The horizontal limits allow us to purely investigate the area of most interest for this study and represents the area covered by overturning rolls inland of the sill.

3.4.2. Background potential energy and turbulent mean tracer variance decay rate

Diapycnal mixing is energetically assessed by two main diagnostics within this thesis: background potential energy and turbulent mean tracer variance decay rate. These methods are outlined in this section including partial derivations and how to apply the methods through examples. Consider the advection diffusion equation:

$$\frac{\partial \theta}{\partial t} + \nabla.(\theta \mathbf{u}) - \nabla.(\kappa \nabla \theta) = 0, \tag{3.4.5}$$

where κ is the diffusivity, θ is temperature and $\mathbf{u}$ is velocity. There is an advective term $\nabla.(\theta \mathbf{u})$ which stirs the fluid and a diffusive term $\nabla.(\kappa \nabla \theta)$ which then mixes the fluid, changing the physical properties irreversibly. This process of stirring and subsequent mixing by molecular diffusion is referred to as physical mixing, however in addition to this the advection schemes (which are either diffusive or dispersive in nature) used in numerical models cause additional mixing, which is referred to as numerical mixing. The contribution of both the numerical and physical mixing gives the overall effective mixing during a numerical simulation. The tracer variance decay method allows us to differentiate the contribution from both numerical and physical mixing individually, whereas the background potential energy method will only quantify the sum of both.

A simple example is used to aid the comparison and explanation of these mixing diagnostics. This simulation consists of a simple two-dimensional vertical slice domain, 5 m by 5 m domain with temperature θ = 0.03 and a lower density circular region (θ= -0.03) released from rest in the center of this square (shown in figure 3.4.1), where the thermal expansion coefficient, α is set as -1.0. The lower density circle of fluid moves upwards before colliding with the top of the domain initiating enhanced physical mixing. In addition, a relatively coarse resolution

($\Delta x = \Delta z = 0.1$ m) fixed mesh is used to show the influence of higher levels of numerical mixing with the diagnostics.

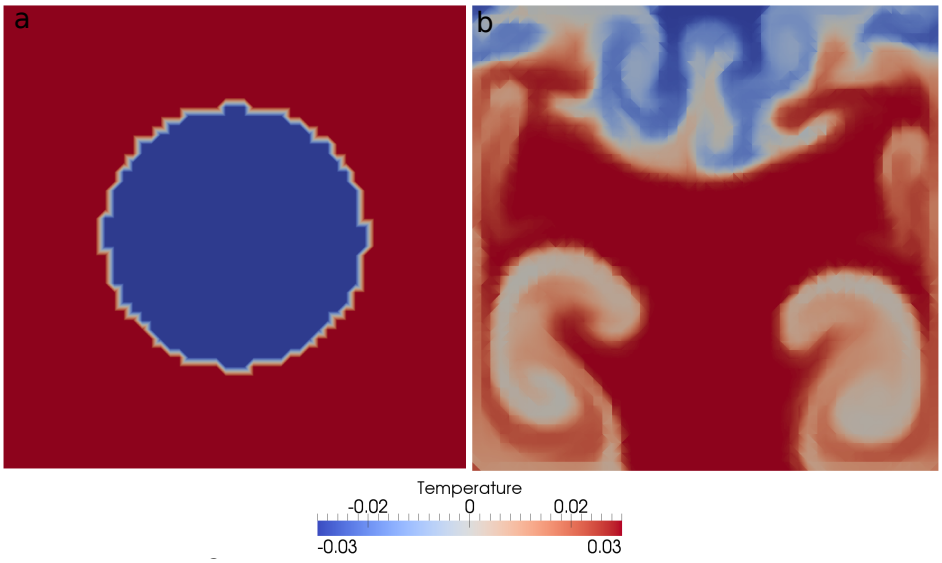


Figure 3.4.1.: The set-up for the simple test case used to explain the mixing diagnostics. The temperature field is shown here at $t = 0$ s (a) and $t = 15$ s (b) mixing can be seen to occur by eye as the area with intermediate temperatures (white) increases

3.4.3. Background Potential Energy Method

Background potential energy is the potential energy of the minimum energy state (or reference state) that can be achieved through adiabatic redistribution of the system (Winters and D'Asaro, 1996; Winters et al., 1995). For closed systems, changes in this reference state relate to changes in the background potential energy (Winters et al., 1995). Previous work which used this diagnostic includes (Ilıcak et al., 2009, 2012; Hiester, 2012), where Hiester (2012) quantified numerical mixing with Fluidity for a lock exchange, similar to the internal seiche and tilting-tank investigated here.

In this thesis background potential energy is calculated by the method of Tseng

and Ferziger (2001) using mixing bins to find the so-called isoscalar coordinate z_* defined below. Mixing bins are used to find the isoscalar coordinate due to their ease of application for simulations run in parallel, where volumes in each mixing bin can be summed across each processor, rather than by use of a sorting algorithm (Winters et al., 1995) which requires a high resolution mesh to sort the domain on before distributing the elements back onto the original mesh starting with most dense elements at the bottom.

Mixing bins contain information about the volume of the domain with a scalar value between certain preset bounds (or isoscalar surfaces). Isoscalar surfaces are surfaces of constant θ, where θ is the scalar of interest, the volume contained between these surfaces is well-defined and can be calculated from mixing bins. As mixing occurs the volume contained between isoscalar surfaces changes. The isoscalar coordinate (z_*) corresponds to these surfaces and therefore changes due to diffusive mixing, it has a single unique value anywhere on an isoscalar surface which obeys $z_*(\theta_2) < z_*(\theta_1)$ when $\theta_2 < \theta_1$. The isoscalar coordinate is defined as:

$$z_*(x,t) = \frac{1}{A}\int_\Omega H_S(\theta(x',t) - \theta(x,t))dV(x'), \tag{3.4.6}$$

according to Tseng and Ferziger (2001) where Ω is the domain and H_S is the Heaviside function. The difference between values of z_* is given by $|\triangle z_*| = \Delta V/A$ where ΔV is the total volume of fluid between two isoscalar surfaces and A is the cross-sectional area of the domain. The volume of fluid in the domain with value greater than or equal to θ is given by $Az_*(x,t)$. The isoscalar coordinate can be used to reorder the domain based on values of θ to give the reference state, which changes over time due to physical and numerical diffusive processes in closed systems. In this reference state there are no horizontal gradients of θ and the fluid is stably stratified. z_* represents the depth that a certain value for the scalar θ

would be found at in this reference state and is represented by a range between 0 and the height of the domain. If θ is a tracer that influences density then this reference state is the minimum potential energy state of the system.

The simple test case is used to show how mixing causes a spread of the isoscalar coordinates (here 101 bins are used) for a system as the tight interfaces between fluid of different θ are smeared and the temperature field evolves (figure 3.4.2).

Background potential energy (E_b) together with available potential energy (E_a) makes up the total potential energy (E_p) in the system, where according to Winters et al. (1995) potential energy is given by:

$$E_p = \int_\Omega \rho g z \, dV = E_b + E_a, \tag{3.4.7}$$

where ρ is density, g is acceleration due to gravity and z is depth. Available potential energy is the total amount of potential energy available for conversion to kinetic energy (due to buoyancy and dissipation) in order to reach a minimum energy state. Background potential energy is the energy of the minimum potential energy state and therefore as z_* is a vertical coordinate of the reference state it may be used in the calculation of E_b (Tseng and Ferziger, 2001):

$$E_b = \int_\Omega \rho g z_* \, dV = gA \int_0^H \rho z_* \, dz_*, \tag{3.4.8}$$

where this may be simplified as has been done on the right hand side of the above equation if the domain has constant height (H) and horizontal cross-sectional area (A).

For a scalar, θ, that affects density, mixing (which causes the only change to background potential energy) can be thought of as diapycnal mixing, therefore the rate of change of background potential energy gives the rate of diapycnal mixing. Diapycnal mixing will always cause an increase in the background potential energy.

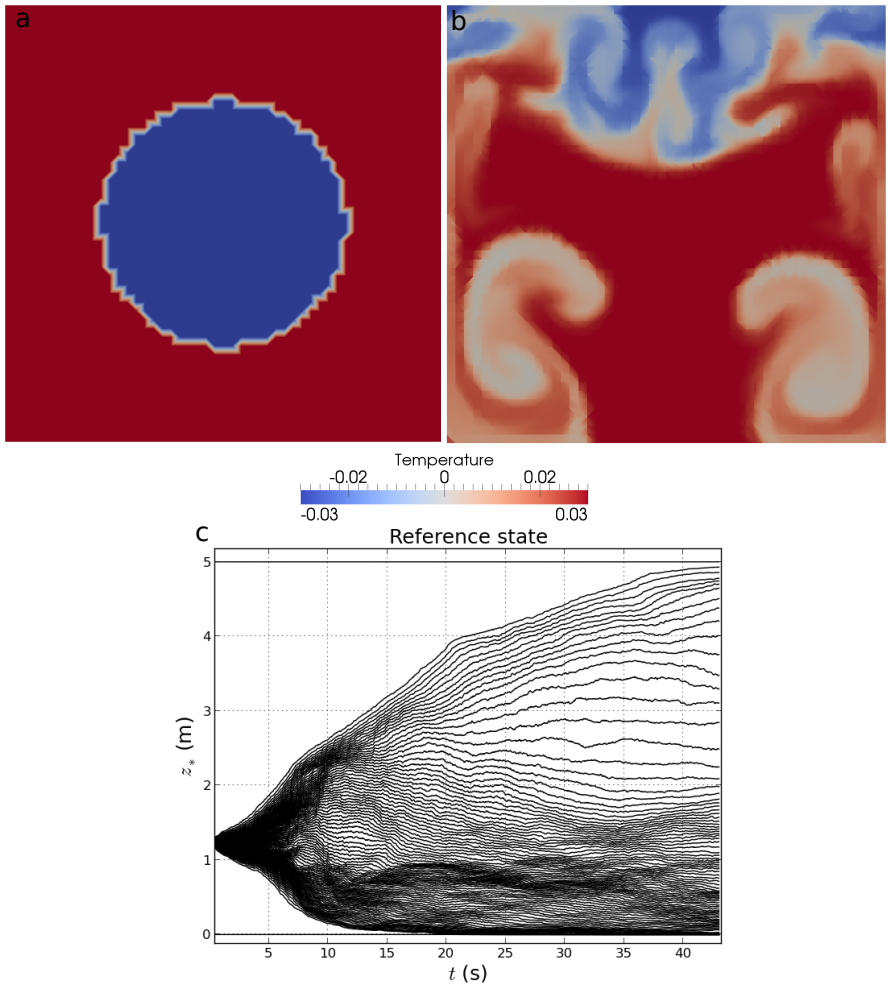


Figure 3.4.2.: The isoscalar coordinates representing the reference state. This shows the z_* values spreading throughout time, which represents larger volumes of fluid in the domain with intermediate values of θ than at the beginning where this area was minimal and due to the representation of the boundary of the circle. This is seen also as the increase in whiter areas of the domain beween the two timestages in figure 3.4.1.

Both background potential energy and its rate of change can be tracked over time for the simulation.

The isoscalar coordinate is required to calculate E_b using mixing bins. Mixing bins are given a set of bounds, where each bin returns the volume of the domain

between these bounds $\theta_i \leq \theta < \theta_{i+1}$ where $0 \leq i \leq N-1$ or $\theta_i = \theta_{min} + id\theta$ where $d\theta = (\theta_{max} - \theta_{min})/N$. This is calculated using the control-volume mesh at run time and each element is given a single value for θ, volumes for each bin can also be normalised by division by the volume of the domain (AMCG, 2015; Hiester, 2012). Following from this z_* is calculated after the simulation is complete, by establishing an equation to relate z_* to the probability density function, where the probability density function is:

$$P(\theta) = \frac{1}{V}\int_{\Omega} \delta(\theta - \theta')dV', \tag{3.4.9}$$

where V is the volume of the domain and $\delta(x)$ is the Dirac delta function. The isoscalar coordinate gives the thickness of fluid with a specific density in the reference state, so if all the layers have the same cross-sectional area, the volume in that layer is:

$$A\Delta z_* = VP(\theta)\Delta\theta, \tag{3.4.10}$$

which is integrated to give:

$$\int_0^{z_*} A(z)dz = V\int_{\theta}^{\theta_{max}} P(\theta)d\theta \tag{3.4.11}$$

or where horizontal area in the domain at different depths is constant:

$$z_* = H\int_{\theta}^{\theta_{max}} P(\theta)d\theta \tag{3.4.12}$$

where H is the height of the domain. As $P(\theta)$ defines the probability density function, therefore $P(\theta) = \int_{\theta_1}^{\theta_2} P(\theta)d\theta = 1$ is true. $P(\theta)$ can be related to the mixing bins by:

$$P(\theta) = b_i/A_b \tag{3.4.13}$$

where b_i is a mixing bin and A_b is a sum over the mixing bins:

$$A_b = \Sigma_{i=0}^{N-1}(\theta_{i+1} - \theta_i)b_i \tag{3.4.14}$$

z_* can then be calculated from:

$$z_* = \begin{cases} H\Sigma_i^{N-1}(\theta_{i+1} - \theta_i)P(\theta_i) \;\; i \leq N-1 \\ 0 \qquad i = N \end{cases} \tag{3.4.15}$$

which can be used to find E_b:

$$E_b = gA\int_0^H \theta z_* \, dz_* = \frac{1}{2}gA\Sigma_{i=1}^{N-1}(\theta_{i+1}z_{*i+1} + \theta_i z_{*i})(z_{*i+1} - z_{*i}) \tag{3.4.16}$$

As temperature is the scalar of interest here, this requires the use of the linear equation of state $\rho = \rho_0(1 - \alpha(T - T_0)) = \rho_0(1 + \bar{\rho})$, where α is the thermal expansion coefficient, T_0 is the reference temperature and $\bar{\rho} = -\alpha(T - T_0)$ is the perturbation density. Therefore E_b becomes:

$$\begin{aligned} E_b &= gA\rho_0\left(\int_0^H z_* \, dz_* + \int_0^H \bar{\rho}z_* \, dz_*\right) \\ &= gA\rho_0\left(\frac{H^2}{2} + \int_0^H \bar{\rho}z_* \, dz_*\right) \\ &= gA\rho_0\left((1 + \alpha\theta_0)\frac{H^2}{2} - \alpha\int_0^H \theta z_* \, dz_*\right), \end{aligned} \tag{3.4.17}$$

where the term that changes is solely $\int_0^H \theta z_* \, dz_*$. This means equation (3.4.16) can be written:

$$E_b = gA\int_0^H \theta z_* dz_* = gA\rho_0\left((1 + \alpha\theta_0)\frac{H^2}{2} - \frac{\alpha}{2}\Sigma_{i=1}^{N-1}(\theta_{i+1}z_{*i+1} + \theta_i z_{*i})(z_{*i+1} - z_{*i})\right). \tag{3.4.18}$$

The calculation of background potential energy is sensitive to the number of mixing

bins, any potential over and undershoots in the computed field and mesh resolution as investigated by Hiester (2012), and therefore comparison with alternative mixing analysis methods is interesting.

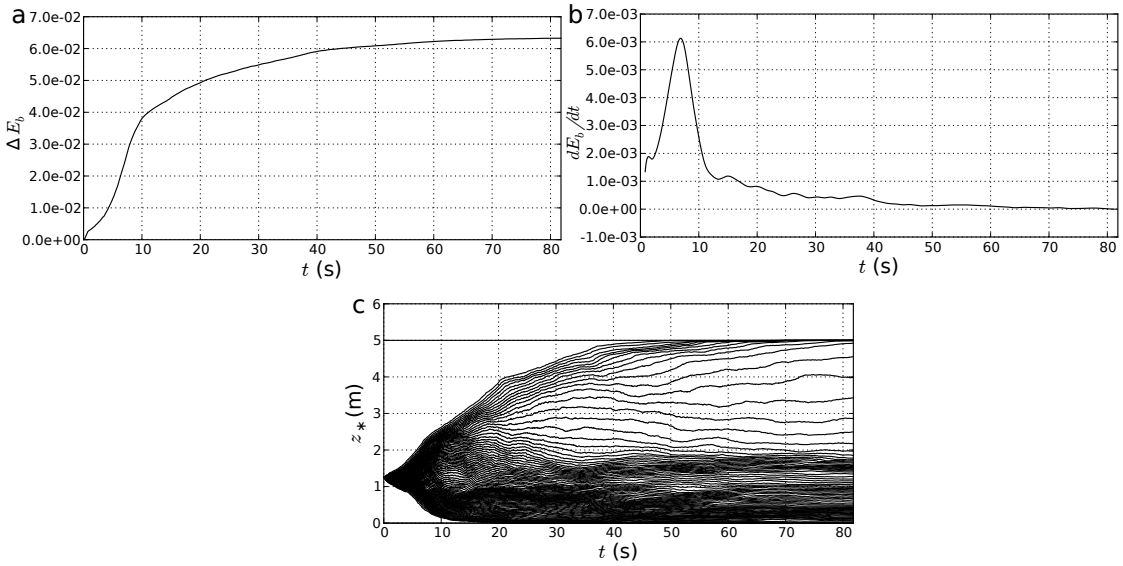


Figure 3.4.3.: The background potential energy (a), the rate of change of background potential energy (b) and the reference state (c) again for comparsion. As the isoscalar coordinates spread at different rates this can be observed as peaks and troughs in the graph of rate of change of background potential energy

The simple test case (figure 3.4.3) shows the relationship between background potential energy and the isoscalar coordinate in practice. Background potential energy rises throughout the simulation period as mixing is continuous throughout this time. The greatest rate of mixing occurs within the first 10 s and can be seen by the spike in the rate of change of background potential energy, this is due to the rapid mixing (smearing) of the sharp interface on this relatively coarse mesh. After this point mixing occurs at a lower rate however with smaller peaks in the mixing rate which correspond to changes seen in the z_* at $t = 20$ s and 38 s.

3.4.4. Turbulent Mean Tracer Variance Decay Rate Method

As suggested by Burchard and Rennau (2008) physical mixing is given as the turbulent mean tracer variance decay rate, while numerical mixing is given as the decay rate between the advected square of the tracer variance and the square of the advected tracer. Since micro-structure variance decay (which is the dissipation of tracer fluctuations by molecular diffusivity which occurs on the micro-scale) is the effective mixing process, tracer variance decay may be used as a measure of mixing. The numerical and physical mixing can then be combined to give the total mixing since they share the same units. This method is simpler than calculating the background potential energy and is computationally cheap (Burchard and Rennau, 2008).

The advection–diffusion equation for a conservative tracer θ without sources of sinks is:

$$\frac{\partial \theta}{\partial t} + \nabla.(\theta \mathbf{u}) - \nabla.(\nu^{\theta} \nabla \theta) = 0, \tag{3.4.19}$$

where t is time, $\mathbf{u}$ is the velocity field and the constant molecular tracer diffusivity is ν^{θ}. This can be split into mean and fluctuating parts for θ and $\mathbf{u}$, i.e. $\theta = \overline{\theta} + \theta'$ and $\mathbf{u} = \overline{\mathbf{u}} + \mathbf{u}'$ to give a conservation equation for the mean tracer:

$$\frac{\partial \overline{\theta}}{\partial t} + \nabla.(\overline{\theta}\overline{\mathbf{u}}) + \nabla.(\theta' \mathbf{u}') - \nabla.(\nu^{\theta} \nabla \theta) = 0, \tag{3.4.20}$$

and a dynamic equation for $\overline{\theta'^2}$ is:

$$\frac{\partial \overline{\theta'^2}}{\partial t} + \nabla.(\overline{\mathbf{u}}\overline{\theta'^2} + \overline{\mathbf{u}'\theta'^2} - \nu^{\theta} \nabla \overline{\theta'^2}) = -2\overline{\mathbf{u}'\theta'} \nabla \overline{\theta} - 2\nu^{\theta} \overline{(\nabla \theta')^2}. \tag{3.4.21}$$

For the square of the mean tracer this then gives:

$$\frac{\partial \overline{\theta}^2}{\partial t} + \nabla.(\overline{\mathbf{u}}\overline{\theta}^2 + 2\overline{\theta \mathbf{u}'\theta'} - \nu^{\theta} \nabla \overline{\theta}^2) = 2\overline{\mathbf{u}'\theta'} \nabla \overline{\theta} - 2\nu^{\theta} (\nabla \overline{\theta})^2. \tag{3.4.22}$$

Using the ocean model parameterisation of down-gradient fluxes with different diffusivities in the horizontal and vertical direction (K_h and K_v), the equation for the mean tracer and the square of the mean tracer become:

$$\frac{\partial \overline{\theta}}{\partial t} + \nabla.(\overline{\theta}\overline{\mathbf{u}}) - \frac{\partial^2 K_h \overline{\theta}}{\partial x^2} - \frac{\partial^2 K_h \overline{\theta}}{\partial y^2} - \frac{\partial^2 K_v \overline{\theta}}{\partial z^2} = 0, \tag{3.4.23}$$

and

$$\begin{aligned}\frac{\partial \overline{\theta}^2}{\partial t} + \nabla.(\overline{\theta}^2\overline{\mathbf{u}}) - \frac{\partial^2 K_h \overline{\theta}^2}{\partial x^2} - \frac{\partial^2 K_h \overline{\theta}^2}{\partial y^2} - \frac{\partial^2 K_v \overline{\theta}^2}{\partial z^2} &= -2K_h(\frac{\partial \overline{\theta}}{\partial x})^2 - 2K_h(\frac{\partial \overline{\theta}}{\partial y})^2 - 2K_v(\frac{\partial \overline{\theta}}{\partial z})^2 \\ &= -D^{phys}(\overline{\theta}^2).\end{aligned} \tag{3.4.24}$$

The sink terms on the right hand side give the physical mixing (D^{phys}) or the turbulent mean tracer variance decay, which can be regarded as equivalent to the micro-structure variance dissipation. Here also $\overline{\theta^2}$ can be substituted for $\overline{(\theta - \theta_0)^2}$ where θ_0 is any constant, which can also be the case for the mean tracer concentration equation for a closed domain with no boundary fluxes meaning that the volume mean of the sink terms is the decay for the mean tracer variance, despite $\overline{\theta^2}$ being the tracer squared and not the tracer variance. These sink terms also give D^{phys} for a domain with open boundaries and boundary fluxes, which is an advantage over the background potential energy method, which cannot accommodate these scenarios (Burchard and Rennau, 2008).

In addition to the physical mixing within a system due to diffusive dynamics, the numerical representation of physical processes introduces a separate non-physical diffusion. This diffusion is termed the numerical mixing (D^{num}), and can be calculated separately, following (Burchard and Rennau, 2008). This involves calculating the decay rate between the advected square of the tracer variance and the square of the advected tracer. In order to define this we examine the one-dimensional

case. The one-dimensional advection equation is given as below, when a spatial interval L is split into N spatial intervals and the time divided into timesteps, Δt

$$\frac{\theta_i^{n+1} - \theta_i^n}{\Delta t} + \frac{u}{\Delta x}(\theta_i^n - \theta_{i-1}^n) = 0, \tag{3.4.25}$$

$$A\left\{\theta^n\right\}_i = \theta_i^{n+1} = (1-c)\theta_i^n + c\theta_{i-1}^n, \tag{3.4.26}$$

where θ_i^{n+1} represents the tracer value at the new timestep, or the advected value of the last timestep and $c = \frac{u\Delta t}{\Delta x}$ is the Courant number. When equation (3.4.25) is multiplied by $(\theta_i^{n+1} + \theta_i^n)$ this gives:

$$\frac{(\theta_i^{n+1})^2 - (\theta_i^n)^2}{\Delta t} + \frac{u}{\Delta x}(\theta_i^n - \theta_{i-1}^n)(\theta_i^{n+1} + \theta_i^n) = 0. \tag{3.4.27}$$

This can be rearranged to give:

$$\frac{A\left\{(\theta^n)^2\right\}_i - (\theta_i^{n+1})^2}{\Delta t} = u\Delta x(1-c)\frac{(\theta_i^n - \theta_{i-1}^n)^2}{(\Delta x)^2}, \tag{3.4.28}$$

where the right hand side is always positive and $u\Delta x(1-c)$ is due to numerical diffusivity meaning the right hand side is a combination of diffusivity and discretisation of the tracer gradient squared. Therefore a measure of numerical mixing may be defined as below (calculated node-wise) for a one-dimensional model:

$$D^{num}\left\{((\theta^n)^2\right\} = \frac{A\left\{(\theta^n)^2\right\} - (\theta^{n+1})^2}{\Delta t} = \frac{A\left\{(\theta^n)^2\right\} - (A\left\{\theta^n\right\})^2}{\Delta t}. \tag{3.4.29}$$

In three dimensions with the addition of the general advection-diffusion equation for a three dimensional model:

$$\frac{\partial \theta}{\partial t} + Adv(\theta) + Mix(\theta) = Q, \tag{3.4.30}$$

where we have a general advection term Adv, a general mixing term Mix and a source or sink term Q. The numerical mixing (calculated node-wise) becomes:

$$D^{\text{num}} = \frac{A\{(T^n)^2\} - (A\{T^n\})^2}{\Delta t}. \tag{3.4.31}$$

Effective mixing D^{eff} is then simply the sum of D^{num} and D^{phys}, and can be directly compared with the rate of change of background potential energy.

Using the simple test case the results from the turbulent mean tracer variance method may be compared with the background potential energy method introduced in the previous subsection (figure 3.4.4). These results show a good agreement with a peak in mixing observable 8 s ater the start of the simulation as a sharp increase in both the rate of change of background potential energy and D^{num}. In addition smaller mixing events are also present for both diagnostics at 15, 20, 27, 30 and 38 s.

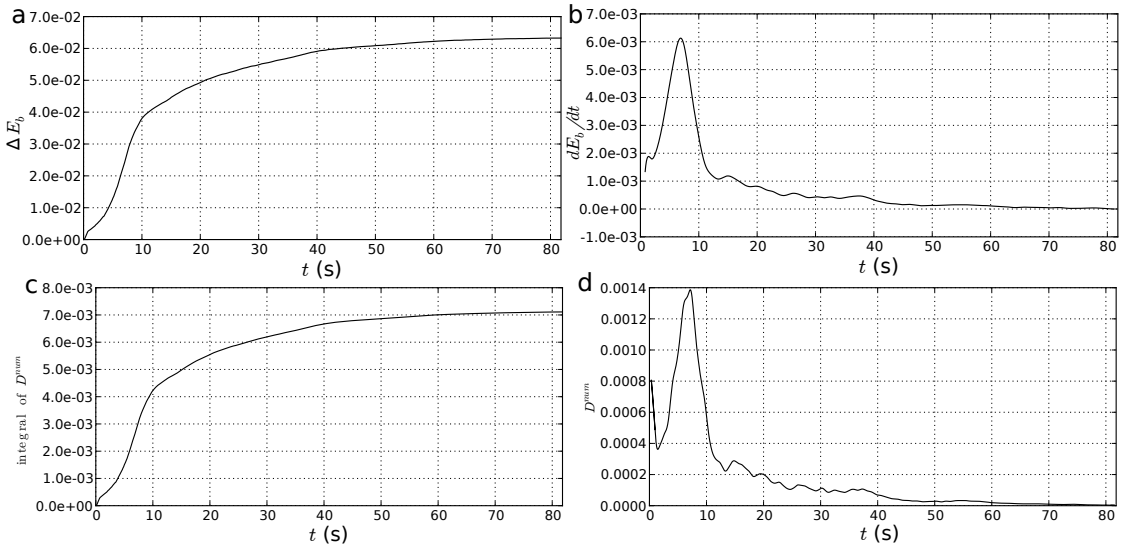


Figure 3.4.4.: The background potential energy (a), the rate of change of background potential energy (b), the integral of D^{num} (c) and D^{num} (d). There is good quantitative agreement between (a) and (c) and between (b) and (d), allowing periods of high and low mixing to be clearly identified.

3.5. Simple test case: Internal seiche

Simulations of an internal seiche are used as an initial benchmark of Fluidity including its different adaptive mesh options. In addition they are also used to compare the influence of discretisation methods and to provide a more complex test case to compare the background potential energy and turbulent mean tracer variance decay rate mixing diagnostics. No explicit thermal diffusivity is used for simulations of the seiche in order to focus on the influence of numerical mixing. Propagation of an internal seiche (a standing internal wave between a fluid of higher and lower density), has been studied by Fringer et al. (2006) using SUNTANS (a non-hydrostatic, unstructured finite-volume model).

3.5.1. Set-up

Simulations, following Fringer et al. (2006), consist of 2D vertical slice domains, with dimensions $L = 100$ m and $D = 10 - 160$ m resulting in varying aspect ratios ($\varepsilon = D/L = 0.1 - 1.6$). These are run for a period of 250 s with a timestep of 0.025 s, kinematic viscosity of $\nu = 10^{-6} m^2 s^{-1}$, and top to bottom density difference of $\Delta\rho/\rho_0 = 0.06$. A fixed structured triangular mesh is used with a resolution of $\Delta x = 1$ m and $\Delta z = 0.5$ m as in Fringer et al. (2006). Minimum and maximum edge length constraints for adaptive meshes were set to $\Delta x_{\min} = 1$ m and $\Delta z_{\min} = 0.5$ m and $\Delta x_{\max} = 5$ m and $\Delta z_{\max} = 2.5$ m, with a maximum of 100,000 nodes allowed. For vertical only adaptive meshes horizontal resolutions were held constant at $\Delta x = 1$ m throughout the simulation. The mesh was adapted every 5 timesteps to enable it to "keep up" with flow dynamics. Simulations predominatently use discontinuous Galerkin discretisation methods for temperature, however for comparison continuous Galerkin and control volume discretisations were also used to compare their influence on mixing.

The simulations were conducted with a free-surface and using free-slip boundary conditions on the walls. They were started from rest using the initial stratification for density given below:

$$\rho = -\frac{\Delta\rho}{2\rho_0}\tanh(\frac{2\tanh^{-1}\alpha_s}{\delta}(z+\frac{D}{2}-a\cos(kx))) \tag{3.5.1}$$

where $\Delta\rho/\rho_0$ is the density difference from top to bottom, D is depth, a is the seiche amplitude, k is the horizontal wavenumber, δ is the interface thickness and $\alpha_s = 0.99$.

3.5.2. Results

Effect of aspect ratio on wave speed

Following Fringer et al. (2006), to establish agreement of models and with the linearised dispersion relation, which is defined as:

$$\left(\left(\frac{u}{u_{DW}}\right)^2 = \tanh\left(\frac{\pi\epsilon}{2}\right)\right), \tag{3.5.2}$$

the square of the non-dimensionalised internal wave velocity is examined $\left(\left(\frac{u}{u_{DW}}\right)^2\right)$, where:

$$u_{DW}^2 = \frac{g'}{2k}\Big(1+\frac{k\delta}{2}\Big)^{-1}, \tag{3.5.3}$$

is the velocity in the deep water limit and $g' = 0.588\ ms^{-2}$ is the reduced gravity. Figure 3.5.1 shows a reasonable agreement between the results of MITgcm, Fluidity and those taken from Fringer et al. (2006) (SUNTANS) for fixed and adaptive meshes, especially at lower aspect ratios D/L which are important for large-scale ocean applications.

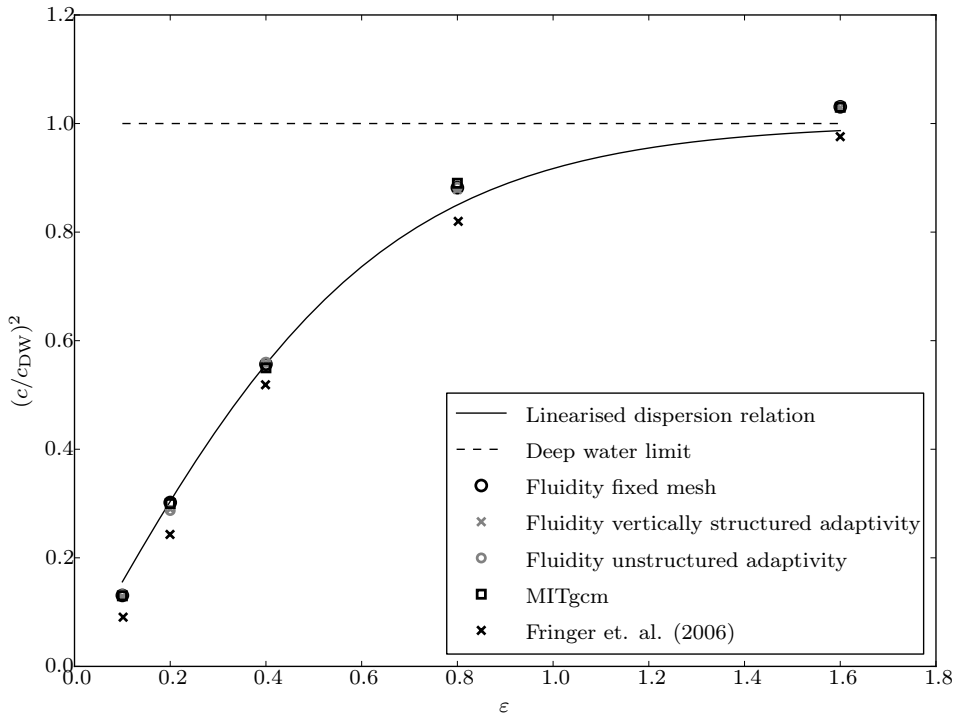


Figure 3.5.1.: The effect of aspect ratio on wave speed for fixed (MITgcm and Fluidity) and adaptive mesh results (Fluidity), compared with linearised theory and results using SUNTANS (Fringer et al., 2006). Results for Fluidity using fixed and adaptive meshes cluster around a similar value at each aspect ratio. MITgcm was also run here to provide additional benchmarking.

Mixing due to discretisation method

The internal seiche was run with continuous Galerkin, discontinuous Galerkin and control volume discretisations in order to compare mixing rates (figure 3.5.2). Results for all discretisations demonstrate a periodic nature in mixing diagnostics, which is clearest in the D^{num} results (but is also visible in the smoothed results for dE_b/dt), with peaks of high mixing rate and troughs of lower mixing rates. The periodicity is linked to the speed of the standing internal wave, where peaks are

observed when the interface is flat (with a higher vertical velocity) and troughs when the interface is sinusoidal (and vertical velocities are lower). For both control volume and discontinuous Galerkin methods the periodicity remains fairly regular as does the speed of the internal wave. However, for the more diffusive continuous Galerkin method mixing becomes less intense over time and the period between the episodes of more intense mixing increases, indicative of a slowing internal wave speed. This slowing is a direct response to the increasing width of the interface (δ) over time, due the relationship between internal wave speed and interface thickness:

$$c^2 = \tanh\left(\frac{\pi\varepsilon}{2}\right)\frac{g'}{2k}\left(1+\frac{k\delta}{2}\right)^{-1}, \tag{3.5.4}$$

where c is the speed of the standing internal wave, ε is the aspect ratio of the tank, g' is the reduced gravity and k is the horizontal wave number.

Continuous Galerkin discretisation results also display additional smaller peaks and troughs on top of the larger trend especially for D^{num} before t = 400s, this is due to instabilities developing on the interface, which overturn and break periodically, leading to increased numerical mixing. Control volume and discontinuous Galerkin discretisations cause less numerical mixing than continuous Galerkin, with final ΔE_b and integrated D^{num} values an order of magnitude smaller for control volume and roughly two orders of magnitude smaller for discontinuous Galerkin than continuous Galerkin (figure 3.5.2). This conclusions may also be drawn qualitatively from the interface diffusion visible in the temperature fields (figure 3.5.3). Discontinous Galerkin is therefore the least numerically diffusive discretisation method investigated and is therefore solely used for further studies. However, it should be noted for fairness that when comparing discretisation methods for the same number of nodes, discontinuous Galerkin methods result in a greater number of degrees of freedom, for a further discussion and similar

comparison of methods see (Hiester, 2012). This study has also shown very close agreement between results found for mixing using both diagnostics.

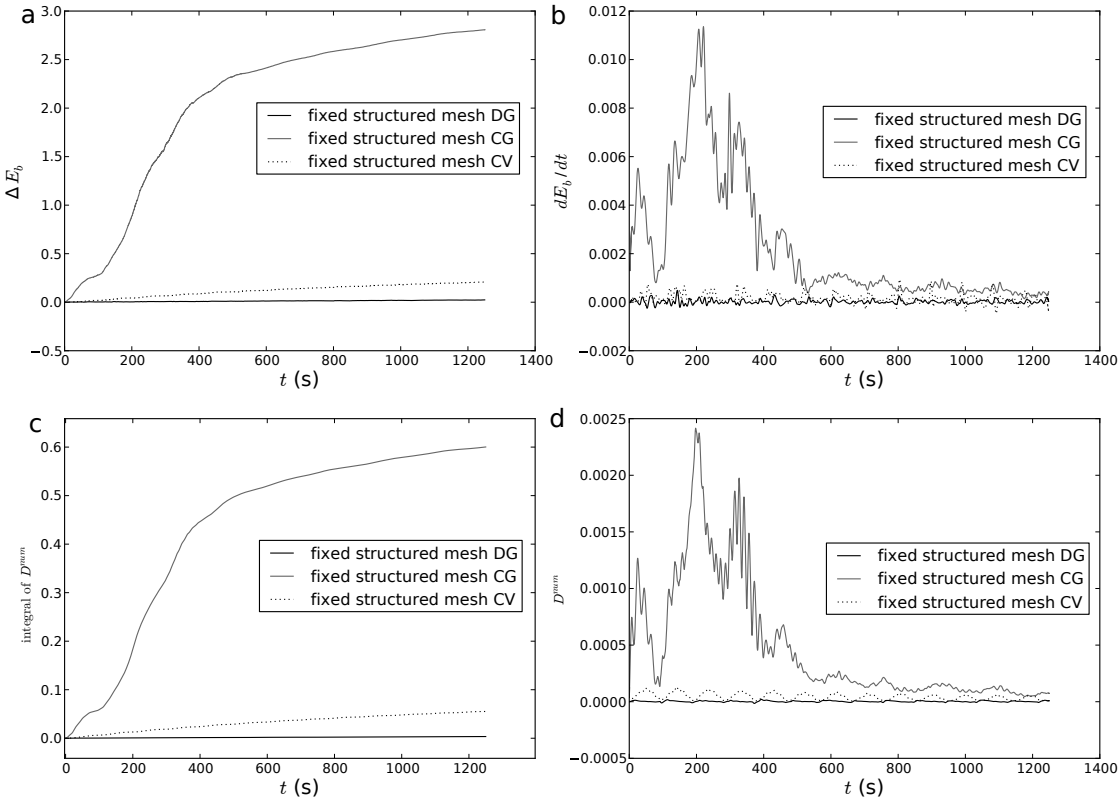


Figure 3.5.2.: Mixing observed using both the method of background potential energy and the turbulent mean tracer decay method due to the discretisation method used for the internal seiche simulations.

Mixing due to type of adaptive mesh

The internal seiche is also used to compare the mixing resulting from the choice of the type of mesh adaptivity. Results from both mixing diagnostics are consistent in terms of the relative amount of mixing present with each adaptive mesh approach, with oscillations relating to the propagation of the seiche back and forth across the tank visible in both (figure 3.5.4). Mesh adaptivity does not greatly affect the numerical mixing occurring in the simulation, with fixed mesh results comparing

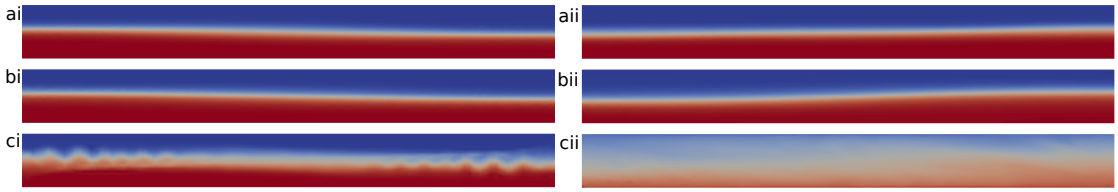


Figure 3.5.3.: The differences in temperature fields for (a) discontinuous Galerkin (b) control volume and (c) continuous Galerkin discretisation methods at (i) $t = 157.5$ s and (ii) $t = 1252.5$ s. This figure shows how the numerical mixing due to these discretisation methods affects the interface thickness visually and the impact this may have on the dynamics involved. In particular (ci) shows the instabilties developing on the interface and how this adds to the mixing visible in (cii).

well with unstructured and partially structured adaptivity based results using a similar average number of nodes. Vertically structured and vertical only adaptivity cause the most mixing in this case, while fixed meshes and unstructured adaptive meshes cause the least. Depending on the method of assessment the least diffusive adaptive mesh type causes 2–4 times less mixing than the most diffusive, however all methods used here cause minimal spurious mixing. Results comparing the discretisation methods (section 3.5.2) show that using a control volume method for temperature instead of the discontinuous Galerkin method used here causes an order of magnitude more spurious mixing although it still represents the dynamics well. Therefore the effect of mesh type is concluded to be minimal for this seiche test case. While the fixed mesh does give very slightly lower numerical mixing, this simulation is very simple and the main aim here is to check that adaptive meshes do not introduce a significant amount of numerical mixing, which given the comparable results is the case. This allows us to use adaptive meshes in the following more complex scenarios, where the benefits of adaptivity will be greater.

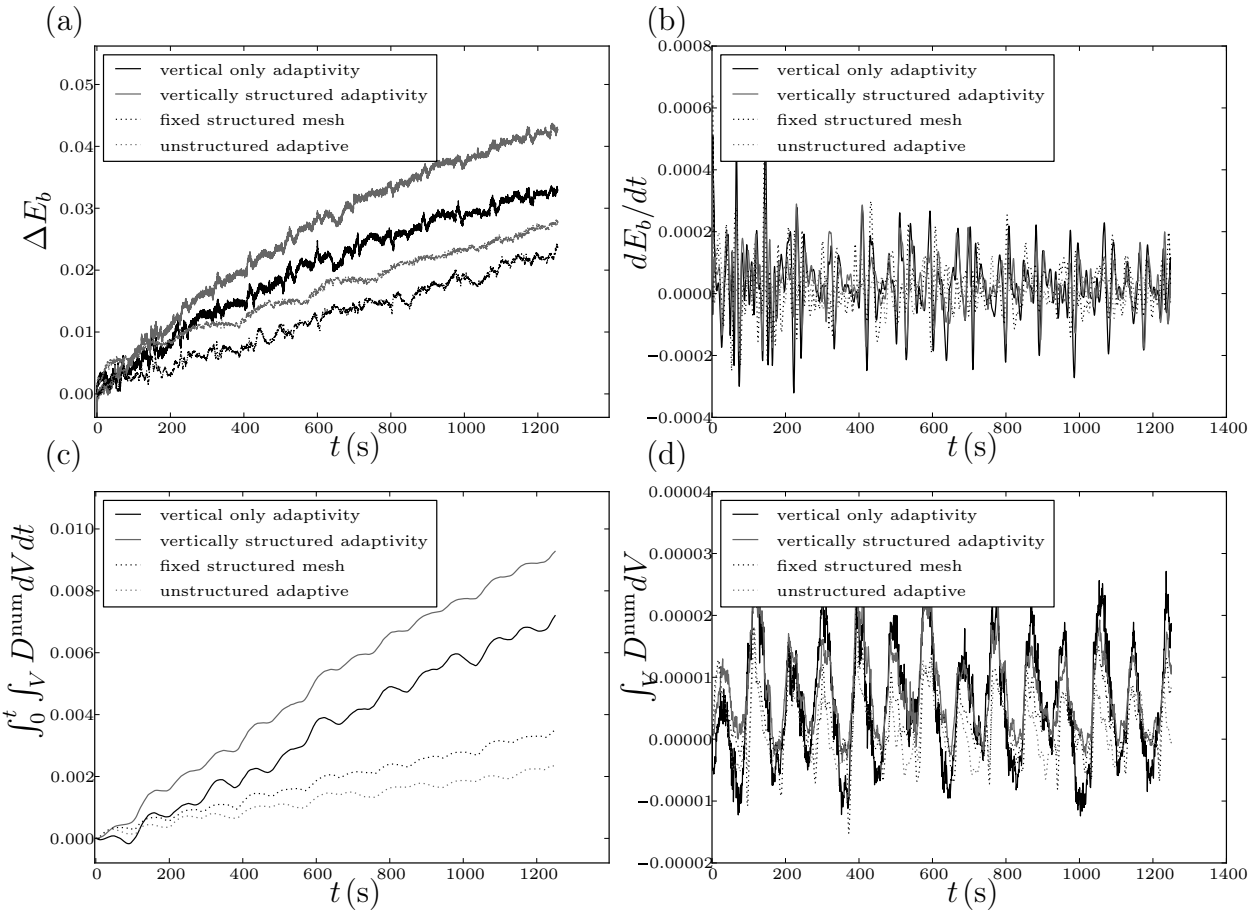


Figure 3.5.4.: Mixing diagnosed in the seiche problem using both the method of background potential energy (a and b) and the turbulent mean tracer decay method (c and d). The periodicity observed appears directly linked to the speed of the internal wave, where peaks are observed when the interface is flat (and vertical velocities highest) and troughs when the interface is sloped (as the original state or its mirror image).

3.6. Lab validation: tilting-tank

The tilting tank simulation shows the evolution of solitary internal waves along the thermocline, which are induced by the initial angle of the tank (Martin, 2014; Horn et al., 2001). This simulation, which is more complex than the internal seiche, is well-constrained experimentally and enables benchmarking against the laboratory

experiment of Horn et al. (2001). It also provides a good test case in order to explore the subtle differences invoked by the choice of mesh type and to ensure a good representation of dynamics using mesh adaptivity. Again as for the seiche no explicit thermal diffusivity is used in order to focus on the influence of mesh types on numerical mixing. This builds upon from work by Martin (2014) to model the laboratory experiment of Horn et al. (2001) using Fluidity.

3.6.1. Set-up

The domain set-up was 6 m long and 0.29 m deep, with a rigid lid, no slip boundaries and an initially sloped thermocline to mimic the tilting of the laboratory tank at the simulation start:

$$\rho(x,z) = -0.5\frac{\Delta\rho}{2\rho_0}\tanh\Big(\frac{z-\zeta(x)}{\alpha}\Big), \tag{3.6.1}$$

where the sloping interface is defined by $\zeta(x) = x\tan\theta + h$, where θ is the angle at which the tank is tilted and h is the height of the interface above the bottom of the flat tank (Horn et al., 2001; Martin, 2014). The interface thickness is given by $\delta_p = 2\alpha\tanh^{-1}(0.99)$ and initially takes the value of 0.01 m. A kinematic viscosity of 10^{-6} m^{2}s^{-1} was prescribed and a quadratic body force term $-\frac{0.036|\mathbf{u}|\mathbf{u}}{B}$ (where $B = 0.3$ m is the width of the tank), established by trial and error by Martin (2014), was used throughout to mimic friction with the sides of the tank which is not explicitly represented in this 2D vertical slice set-up. If a 3D set-up had been used this body force term would not be required as instead drag would be applied on the lateral boundaries.

Fixed structured mesh results had a resolution of $\Delta x = \Delta z = 0.001$ m (1,746,291 nodes) with coarser mesh results (shown in Figure 3.6.4) using $\Delta x = \Delta z = 0.03$ m (resulting in only 2211 nodes). In adaptive mesh simulations the mesh was

updated every 5 timesteps with minimum and maximum edge length constraints for fully unstructured adaptive meshes and vertically structured adaptive meshes set to $\Delta x_{\min} = \Delta z_{\min} = 0.001$ m and $\Delta x_{\max} = \Delta z_{\max} = 0.5$ m. The resultant number of nodes throughout the adaptive simulations are shown in figure 3.6.1, with a maximum of 19,600 and an average of 12,110 nodes for fully unstructured adaptivity, a maximum of 4,640 nodes and average of 3,380 nodes for vertical only adaptivity and a higher maximum of 213,540 nodes and average of 65,782 nodes for vertically structured adaptivity. Vertically structured adaptivity simulations require more nodes due to the imposed vertical structures meaning that when increased horizontal resolution is required it must be added for the entire column. For vertical only adaptive meshes horizontal resolutions were held constant at $\Delta x = 0.005$ m throughout the simulation.

3.6.2. Results

Tilting tank results are benchmarked against laboratory data (Horn et al., 2001) by tracking the interface (defined as the $\Delta\rho/2$ contour) height at 3.1 m along the tank throughout the simulation. The slight deviation from the centre of the tank accounts for the time to tilt the tank in the physical experiment (and the physical probe location).

All meshes provide a good representation of the dynamics of the internal wave (shown in figure 3.6.2), where the interface shows the movement of a train of waves past the centre of the tank. The first of the train is largest and is followed by smaller and smaller waves, this wave train travels backwards and forwards across the tank while gradually decaying in amplitude. The first wavetrain however is made up of fewer waves in the physical results due to transient effects causing their collapse, the longer wavetrain shown for our numerical results here is also demonstrated in other numerical studies (Maderich et al., 2012), increasing confidence in our

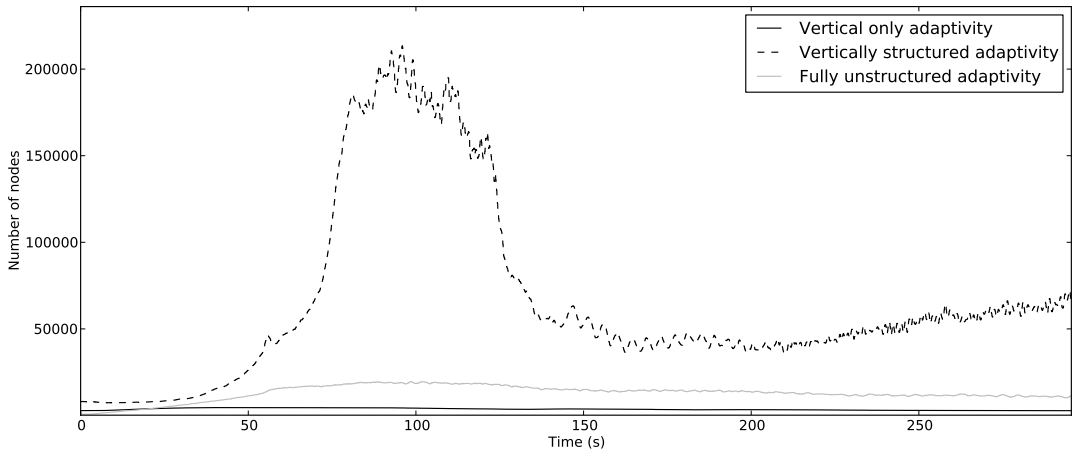


Figure 3.6.1.: Number of nodes throughout the adaptive simulations. An initial large increase in nodes is observed as the dynamics become progressively more complex as the density interface adjusts resulting in the adaptive simulations trying to resolve instabilities on the interface. The number of nodes then stabilises and begins to decline as the simulation becomes less complex as a simple internal wavetrain is initiated and propagates backwards and forth across the domain, although peaks are present in the number of nodes when the wavetrain interacts with the ends of the tank.

results. However, due to this difference comparison is limited to after this point. The numerical results are all very similar for the different meshes in their ability to capture the evolution of this wavetrain. All simulations conducted here capture the correct number of waves in the wavetrain except for vertical only adaptivity which misses the last (eighth and sixth respectively) wave at 175 s and 226 s. This is potentially because this method of adaptivity has resulted in slightly more mixing than the others (shown in 3.6.3c), as described previously mixing reduces the g' across the interface and as a result means the wave travels more slowly. The high resolution fixed structured mesh appears to represent the dynamics best while the unstructured adaptive mesh results produce the best adaptive mesh representation.

An investigation of the mixing occurring for the tilting tank using different

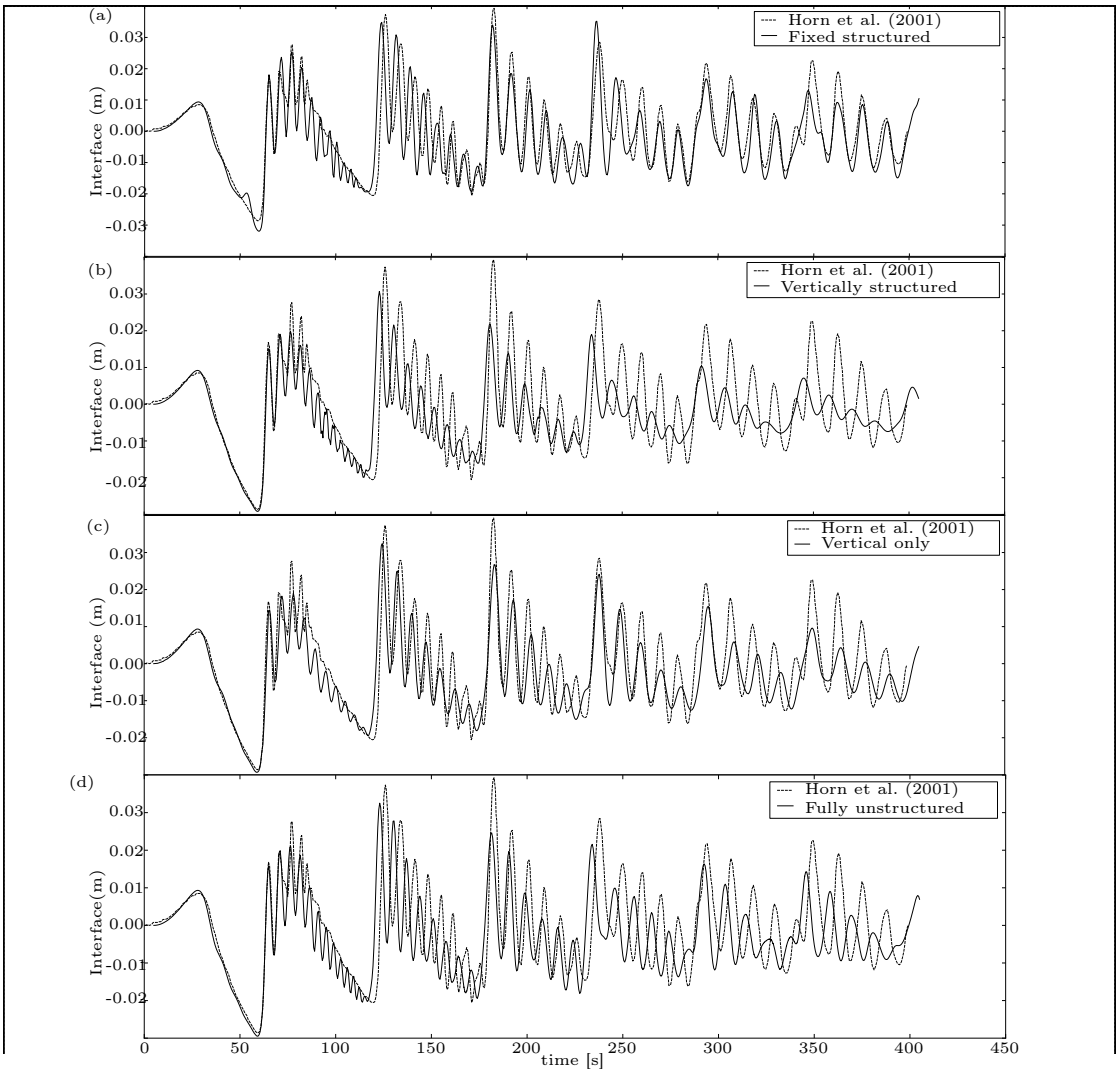


Figure 3.6.2.: A comparison of results for the tilting-tank with laboratory results from Horn et al. (2001) for (a) a high resolution fixed structured mesh, (b) vertically structured adaptivity (c) vertical only adaptivity and (d) fully unstructured adaptivity. For comparison the coarse mesh results are shown in a later figure.

mesh types reveals that the fixed and vertical only adaptive mesh results produce the most mixing while unstructured adaptivity results in the least mixing and is slightly anti-diffusive meaning density gradients are increased (which could help explain the wavetrain being slightly faster than the laboratory wave). Although

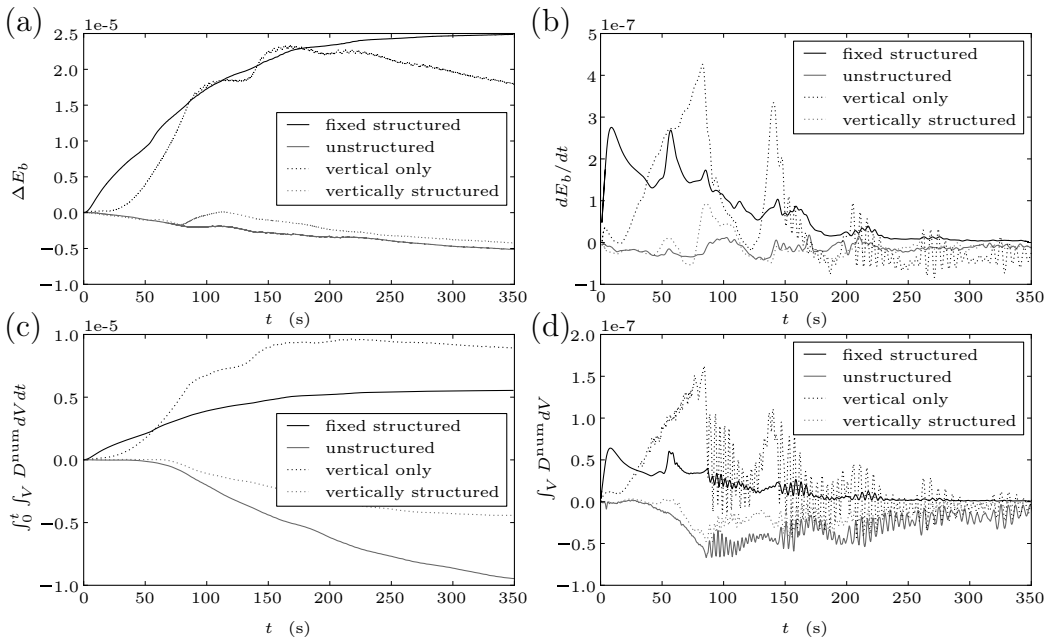


Figure 3.6.3.: A comparison of the mixing occurring in simulations of the tilting tank using different meshes. Fully unstructured, vertically structured and vertical only adaptive mesh simulations are compared to a high resolution fixed structured mesh (coarse fixed mesh results are shown in a later figure). The upper graphs show the background potential energy method, while the lower graphs show the results of the turbulent mean tracer variance method. The methods agree reasonably well in terms of trend but not as well for amplitudes.

it should be noted that vertically structured simulations use the most nodes here. These results are slightly different to those of the seiche due to the differences in the number of nodes used for each scenario and for each mesh type and due to the coarser mesh used for the seiche. However all meshes used for the tilting tank result in very little mixing (significantly lower than for the seiche) and very small values of ΔE_b between 2.5 $\times 10^{-5}$ and -5 $\times 10^{-6}$ kg m^2 s^{-2}, meaning that as well as accurately capturing the dynamics of the problem, adaptive mesh simulations of the tilting tank do not spuriously affect the mixing occurring greatly. Both methods of characterising the mixing again agree reasonably well, capturing the

peaks and troughs in mixing as the wavetrain travels across the tank in the D^{num} and $\Delta E_b/dt$ plots. Although there are slight differences between the two methods in terms of the amplitudes, when comparing the two methods the general trends are the same.

In order to demonstrate the advantages of mesh adaptivity for the capture of the tilting tank dynamics, vertical only adaptivity results (with the least number of nodes: a maximum of 4630 nodes and an average of 3380 nodes) are shown compared with a fixed mesh of similar number of degrees of freedom with $\Delta x = \Delta z = 0.03$ m. In figure 3.6.4(a) a comparison of the interface height evolution through time and the amount of mixing is shown using both mixing diagnostics. Results show that the coarse fixed mesh is incapable of capturing the full dynamics of the internal wave. While the general motion of the wavefront is captured, it propagates too slowly with the coarse fixed mesh, further proof of the numerical diffusion occurring at the interface. The analytical formula for the initial condition also remains the same for the interface but the grid size is too coarse to represent this which results in an initial spike in mixing shown in Figure 3.6.4b. Using a similar number of degrees of freedom the vertical only adaptivity however provides a good representation of the internal wave dynamics, capturing the correct speed of the wavefront throughout the simulation and a good representation of the wavetrain. The improvements in efficiency by using vertical only adaptivity versus the coarse fixed mesh are also seen in the mixing diagnostics: the change in ΔE_b and integral of D^{num} is at least an order of magnitude higher for the coarse fixed mesh.

3.7. Pseudo real-world application: Loch Etive

Loch Etive is a 30 km long, 145 m deep sea loch in Western Scotland with six sills present at 15 m average depth. It consists of several basins although previous

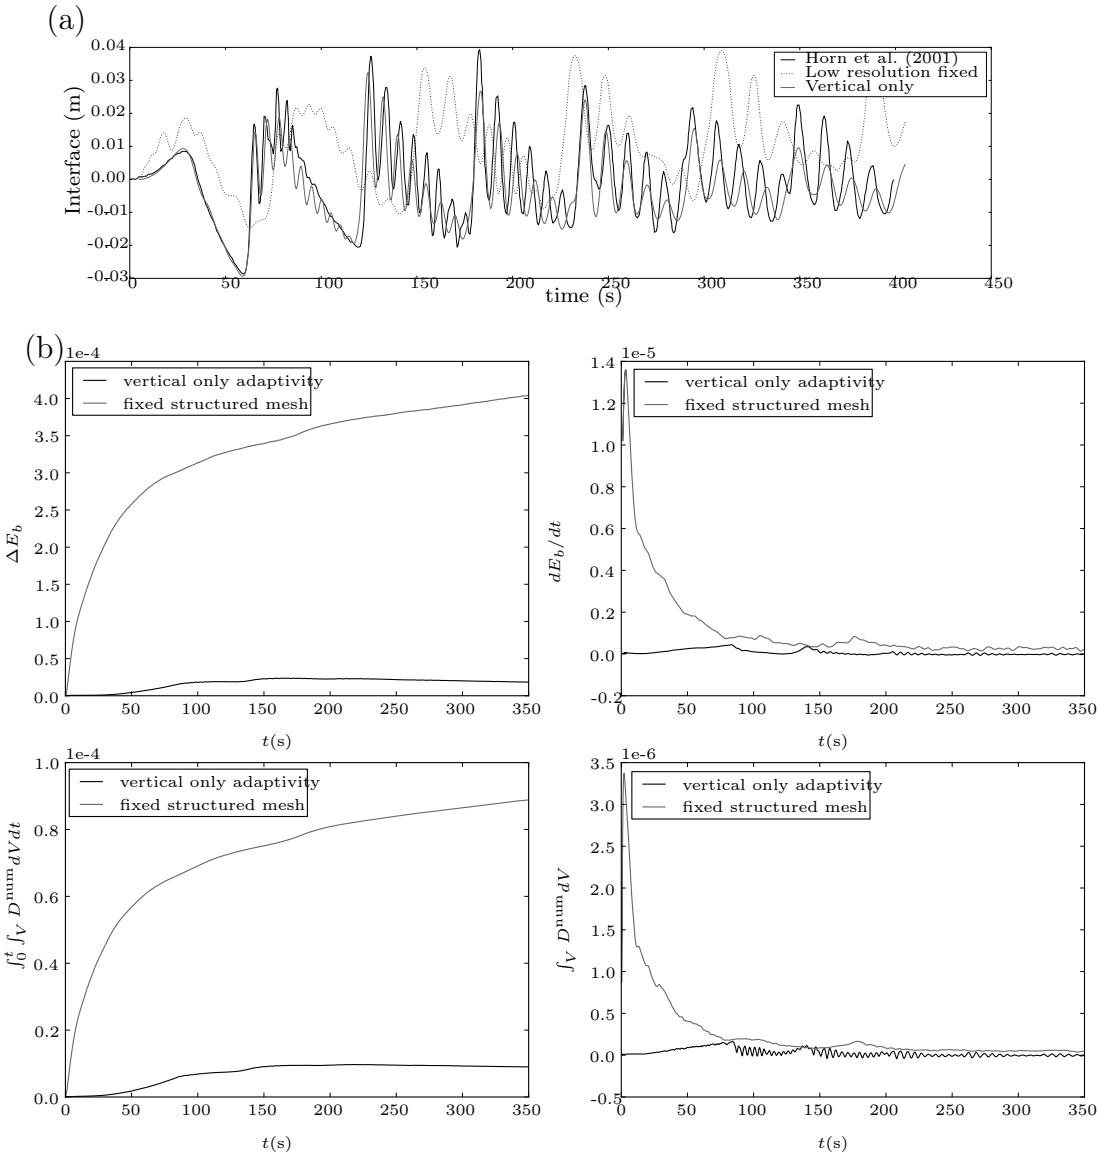


Figure 3.6.4.: Figure (a) displays the interface evolution throughout time for the vertical only adaptivity simulation and a coarse fixed structured mesh simulation (using a similar number of degrees of freedom) compared with Horn et al. (2001). Figure (b) demonstrates the differences in mixing caused by the two meshes, the upper two graphs show the background potential energy method which broadly agrees with the results of the turbulent mean tracer variance method below. The trends seen in both are very similar and we can see roughly an order of magnitude difference in mixing between the two simulations.

modelling studies have generally focussed on the largest upper basin (Inall et al., 2004). Loch Etive has been extensively studied both by field measurements and 2D vertical slice numerical models with approximate idealised geometries (Inall et al., 2004, 2005; Inall, 2009; Xing and Davies, 2006, 2009a, 2010; Berntsen et al., 2009, 2008). This makes it a good real-world test problem with which to validate the Fluidity model with mesh adaptivity. In earlier work (Xing and Davies, 2006, 2009a, 2010; Berntsen et al., 2009, 2008) conducted with the MITgcm and Bergen ocean model (BOM), simulations were shown to be highly sensitive to mesh resolution, with finer meshes showing small-scale non-hydrostatic processes superimposed on the large-scale flow. Numerical models show that as tidal currents flow over the topography, flow separation occurs at the sill crest with turbulent overturning downstream resulting in vertical mixing. Lee waves are generated on the sill slope when the Froude number $F_r = U/C_i$ (where U is the cross sill velocity and C_i is the first mode internal wave speed) in the region of the sill is greater than unity (a so-called "jet type" fjord system in (Stigebrandt and Aure, 1989)). These waves propagate back towards the sill once the tide reverses leading to further enhanced mixing of the region (Xing and Davies, 2006, 2009a, 2010; Berntsen et al., 2009, 2008).

3.7.1. Set-up

A 2D cross-sectional model with highly idealised bathymetry motivated by the largest upper basin of Loch Etive was thus formulated and simulated with Fluidity based upon the set-up presented in Xing and Davies (2006, 2007) which used MITgcm. Telescoping horizontal resolution is used for fixed mesh simulations, whereby horizontal resolution is finest (we use a range of Δx values from 80 m down to 0.625 m) over the sill, and coarsens either side up to $\Delta x = 100$ m. The

bathymetry of the domain is given by:

$$H(x) = \begin{cases} 50 - \frac{35}{1+(\frac{x-6500}{500})^4}, & \frac{x-5500}{500} \leq 2 \\ 100 - \frac{85}{1+(\frac{x-6500}{500})^4}, & \frac{x-5500}{500} > 2. \end{cases} \tag{3.7.1}$$

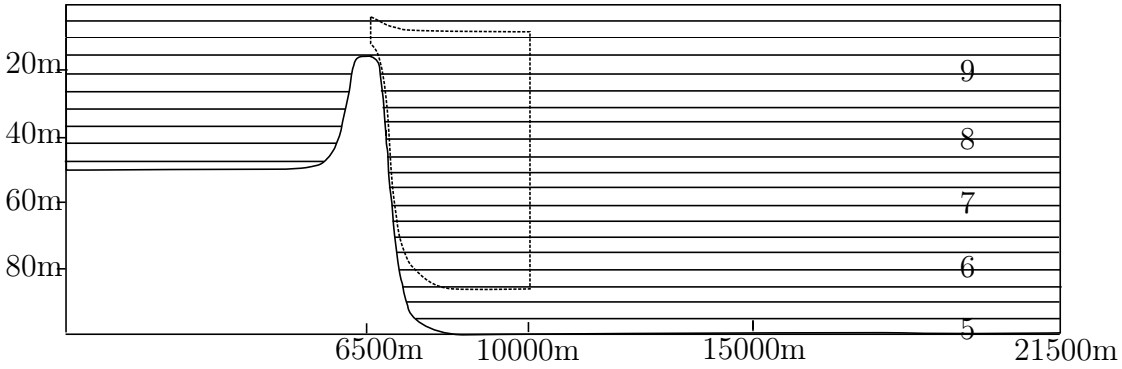


Figure 3.7.1.: Schematic of Loch Etive displaying bathymetry and initial linear temperature stratification with 0.5^oC contours between 10^oC at the surface and 5^oC at depth as in Xing and Davies (2006). The domain is 21,500m long with a sill located at 6,500m with a 15m minimum depth. A relatively long domain is used to avoid interactions with the end of the domain influencing results. A dashed region indicates the area that certain diagnostic quantities are computed over.

An M_2 tidal forcing of $U = U_0 \sin(\frac{2\pi}{T}t)$, where $T = 44714$ s is the M_2 tidal period and $U_0 = 0.3\text{ms}^{-1}$ is the peak tidal velocity, is applied at the shallower ($H = 50$ m) Western (left) end of the domain, while the deeper (H= 100 m) Eastern end is closed. Simulations were initiated from rest with an initial buoyancy frequency of $N = 0.01\ s^{-1}$ and Coriolis frequency of $f = 1.2$ x 10^{-4}s^{-1}, a free surface is also used and small constant values of diffusivity $K_h = K_v = 10^{-7}$ m^2s^{-1} as in (Berntsen et al., 2009). No slip, no normal flow velocity boundary conditions were used at the seabed. Berntsen et al. (2009) used viscosities of $A_h = 10^{-1}\text{m}^2\text{s}^{-1}$, $A_v = 10^{-3}\text{m}^2\text{s}^{-1}$, however studies in section 3.7.2 demonstrated that results become too turbulent at these viscosities to allow mixing processes

to be investigated without recourse to parameterisation or extremely high mesh resolutions and small timesteps for reasonably converged results. Section 3.7.2 also explains the lack of agreement between Berntsen et al. (2009) and Fluidity in that section. Therefore simulations using higher viscosities with consequently simpler dynamics are used, these viscosities are 1000, 100 and 10 times (ν_c, ν_b, ν_a) the Berntsen et al. (2009) values, with the same ratio between horizontal and vertical values maintained. Results are obtained over a quarter of a tidal cycle to focus on lee wave generation and breaking. Fixed mesh Fluidity simulations used 100 sigma layers by default (in addition to a range of 50–300 layers also being investigated) and a range of horizontal resolutions in the region of the sill from $\Delta x = 80$ m to $\Delta x = 0.625$ m. A range of timesteps were used between $T/7200$ s and $T/28800$ s. Adaptive simulations used a timestep of $T/28800$ s for the lowest viscosity choice for ν_a (since smaller mesh sizes are required and for consistency across results) and $T/7200$ s for all other viscosities, and the meshes are updated every 20 timesteps. In addition, for further comparison high resolution MITgcm simulations were also conducted here at each viscosity value.

In order to examine the benefits of adaptivity in terms of computational savings, a brief convergence study was also necessary. This was conducted for the range of viscosities described in the previous paragraph. This enabled the differences in mesh resolution required to capture the dynamics in systems of increasing complexity and increasing levels of multi-scale behaviour to become apparent. Simulations using higher viscosities (ν_c and ν_b) with simpler dynamics were used in order to examine the effect of a range of adaptivity settings such as: minimum and maximum edge length, interpolation error bounds (ε, used to control the sensitivity of the adaptive metric), maximum aspect ratio, use of metric advection, and frequency of adapts, before proceeding to the more complex lower viscosity simulations. Due to the length of the domain and processes of interest only occurring close to the

sill, interpolation error bounds were increased away from this area, so that even lower resolutions were used in the outer, less important extents of the domain. This telescoping uses an exponentially increasing function for ε with the length of the telescoped region investigated. This mimics the telescoping used for the fixed mesh resolutions which was based upon (Berntsen et al., 2009).

Adaptive mesh simulations here solely use unstructured adaptivity as previous idealised test cases show that it gives a good representation of dynamics while limiting spurious mixing. In addition, the aspect ratio for this problem is also not sufficiently high to require a structured adaptive mesh (i.e. vertically structured or vertical only adaptivity). The parameters used for the adaptive simulations conducted (where both temperature T and vertical velocity, w, were adapted to, but not horizontal velocities, u, due to the dominance of the tidal signal) are listed in table 3.7.1.

3.7.2. Initial benchmarking against Berntsen et al. (2009)

Set-up

This brief study uses a similar set-up to that outlined in section 3.7.1. However for direct comparison with Berntsen et al. (2009) a shorter (3,000 m) domain is considered with 100 σ-layers for direct comparison with the σ-layer using Bergen Ocean Model (BOM). In addition viscosities of $A_h = 10^{-1}\mathrm{m}^2\mathrm{s}^{-1}$ and $A_v = 10^{-3}\mathrm{m}^2\mathrm{s}^{-1}$ are used. A range of horizontal resolutions were investigated between $\Delta = 100$ m and $\Delta = 1.5625$ m. Due to the shorter domain the bathymetry is represented by:

$$H(x) = \begin{cases} (50 - \frac{35}{1+(\frac{x}{500}-2)^4}) & \frac{x}{500} \leq 2, \\ (100 - \frac{85}{1+(\frac{x}{500}-2)^4}) & \frac{x}{500} > 2. \end{cases} \tag{3.7.2}$$

Table 3.7.1.: List of adaptive mesh simulations for Loch Etive including parameters used, which are as follows: user-defined weight for temperature field (ϵ_T), user defined weight for vertical velocity (ϵ_w), minimum horizontal resolution ($\Delta x_{\min}$), minimum vertical resolution ($\Delta z_{\min}$), type of metric used (either absolute (∞) or p-metric, where p is 2) this is discussed in further detail in (Hiester et al., 2011), whether metric advection is used and the length of the region of ϵ telescoping (m).

Reference code	ϵ_T	ϵ_w	$\Delta x_{\min}$	$\Delta z_{\min}$	Metric Metric	Metric advection	Telescope length (m)
Aa	0.05	0.0017	2.5	0.1	p	no	500
Ab	0.05	0.0017	2.5	0.1	p	no	50
Ac	0.05	0.0017	2.5	0.1	p	no	1000
Al	0.05	0.0017	10	0.4	p	no	500
Ah	0.05	0.0017	0.625	0.025	p	no	500
A2	0.5	0.017	2.5	0.1	p	no	500
A3	0.025	0.00085	2.5	0.1	p	no	500
Ax	0.005	0.00017	2.5	0.1	∞	no	500
Ay	0.05	0.0017	2.5	0.1	p	yes	500
Az	0.05	0.0017	2.5	0.1	p	no	500
Av3	0.25	0.008	10	0.5	p	no	500
Av2	0.05	0.0005	5.0	0.2	p	no	500

a schematic of the domain is presented in figure 3.7.2. The simulation was run for a quarter of the tidal period.

Results

Simulations conducted with Fluidity were directly compared with results presented in Berntsen et al. (2009) using the Bergen Ocean Model (BOM) both qualitatively and quantitatively. BOM is a non-hydrostatic model which uses σ-coordinates and allows comparison with fine mesh solutions. A smaller domain simulation allowed a brief study of the effect of resolution on model results and an opportunity to benchmark Fluidity against another numerical model before conducting more rigorous comparisons using the longer domain set-up used by Xing and Davies

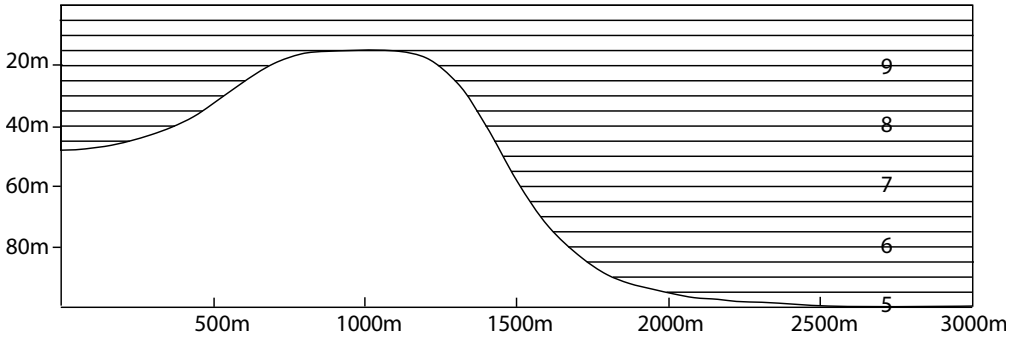


Figure 3.7.2.: Schematic showing the domain geometry for model runs for comparison with Berntsen et al. (2009). The domain is 3000m long and 100m deep at its deepest on the Eastern side of the domain inside the loch, while on the seaward side the depth is just under 50m, the idealised sill geometry can also be seen giving a sill depth of 15m. The diagram also shows the initial temperature stratification varying linearly from 10^oC at the surface to 5^oC at depth

(2006, 2007). This also allows us to highlight the differences between Fluidity and BOM for this example.

Qualitatively by comparing the temperature field at $T/4$ (figure 3.7.3) the results are observed to bear significant similarities to one another. Both results demonstrate that as tidal flow forces water over the sill a high temperature gradient with flat temperature contours and a short jet extending from the sill ($\sim$ 200 m long for Fluidity yet only $\sim$ 150 m for BOM) is produced. At the end of the jet, a number of overturning vortices are formed and propagate away from the sill (this was also observed in physical studies by Inall et al. (2004)), these vortices gradually mix the region and result in well-mixed Lochs in both numerical studies. Both results capture a similar number of overturning vortices, and a similar depth to which these vortices extend, at $\Delta x = 1.5625$ m resolution, vortices in Fluidity extend only $\sim 5-10$ m deeper than Berntsen et al. (2009). However the form of the eddies is quite different, results presented in Berntsen et al. (2009) involve higher temperature gradients in the horizontal direction and eddies extend a roughly constant

depth across the domain whereas the comparable Fluidity result shows the base of the eddies becoming deeper with distance from the sill. Bottom boundary layer separation occurs in comparable positions: Fluidity shows separation at $\sim$ 65 m depth, $\sim$ 500 m horizontally from the sill top while Berntsen et al. (2009) demonstrates the same lateral position but a depth of $\sim$ 68 m. However, this does not agree well with physical studies of the system (Inall et al., 2004), which indicate separation at $\sim$ 40 m depth.

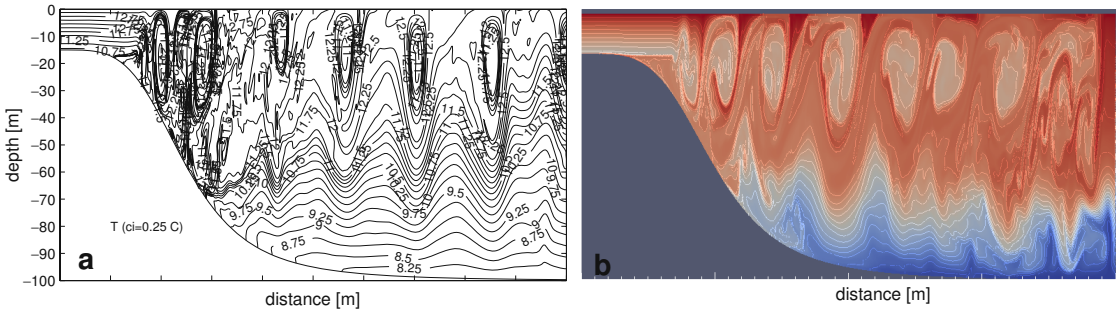


Figure 3.7.3.: A comparison of the Fluidity simulation (right) with Berntsen et al. (2009) (left) for the temperature field at T/4 for a resolution of $\Delta x = 1.5625$ m. Comparisons are made in the region behind the sill where dynamics are most complex at this point in the tidal cycle. Temperature fields are shown with contours at 0.25^oC intervals.

Qualitatively the internal wave produced by both sets of results (figure 3.7.3) is also similar with a length scale of $\sim 200m$ with smaller scale oscillations ($\sim$ 50 m) in addition close to the sill. Further quantitative comparison of the internal waves was undertaken using the vertical velocity at a detector located at $x = 1500$ m, $z = -20$ m for a range of resolutions (figure 3.7.4) allowing comparison of the frequency and amplitude of the wave. Results from this study show that as resolution is increased the amplitude and frequency of these internal waves increases to a maximum at $\Delta x = 6.25$ m, this agrees with the trend of Berntsen et al. (2009). In addition the time associated with the onset of the internal wave agrees well except for the very coarse $\Delta x = 100$ m resolution, which may be influenced by the

resolution required for the different numerical models. As resolution is increased the simple sinusoidal signal becomes more complex, with the addition of higher frequency components. The period of the internal waves agrees fairly well, with $\sim$8 mins for Fluidity and $\sim$7 mins for Berntsen et al. (2009) simulations with a resolution of $\Delta x = 12.5$ m.

However, despite general similarities between the results presented, there are two significant differences: firstly the amplitude of our results do not increase throughout the simulations as shown in Berntsen et al. (2009) and secondly the amplitudes of the internal waves studied here are generally lower than those of Berntsen et al. (2009). However, a final comparison with Berntsen et al. (2009) reveals that average sill top velocities of $\sim$ 0.9 m s^{-1} are observed in both sets of high resolution results, which is slightly larger than those measured in physical studies by Inall et al. (2004).

This study reveals that Fluidity is too heavily influenced by interaction with the end walls which are located too close to the area of interest to produce comparable results to previous studies for this shorter domain set-up. It also reveals that Fluidity responds slightly differently to BOM, as resolution is increased the amplitude of the internal wave does not increase as much. However, similarly to Berntsen et al. (2009), increasing model resolution superimposes small scale phenomena on the large scale dynamics resulting in a higher frequency component in the internal wave signal observed at the detector and changing the turbulent eddy structures. In order for full agreement between the models at the viscosity used in this section it is possible that, without parameterisation, much higher mesh resolutions may be required and thus further studies from this point on use higher viscosities to more fully resolve dynamics at manageable mesh resolutions.

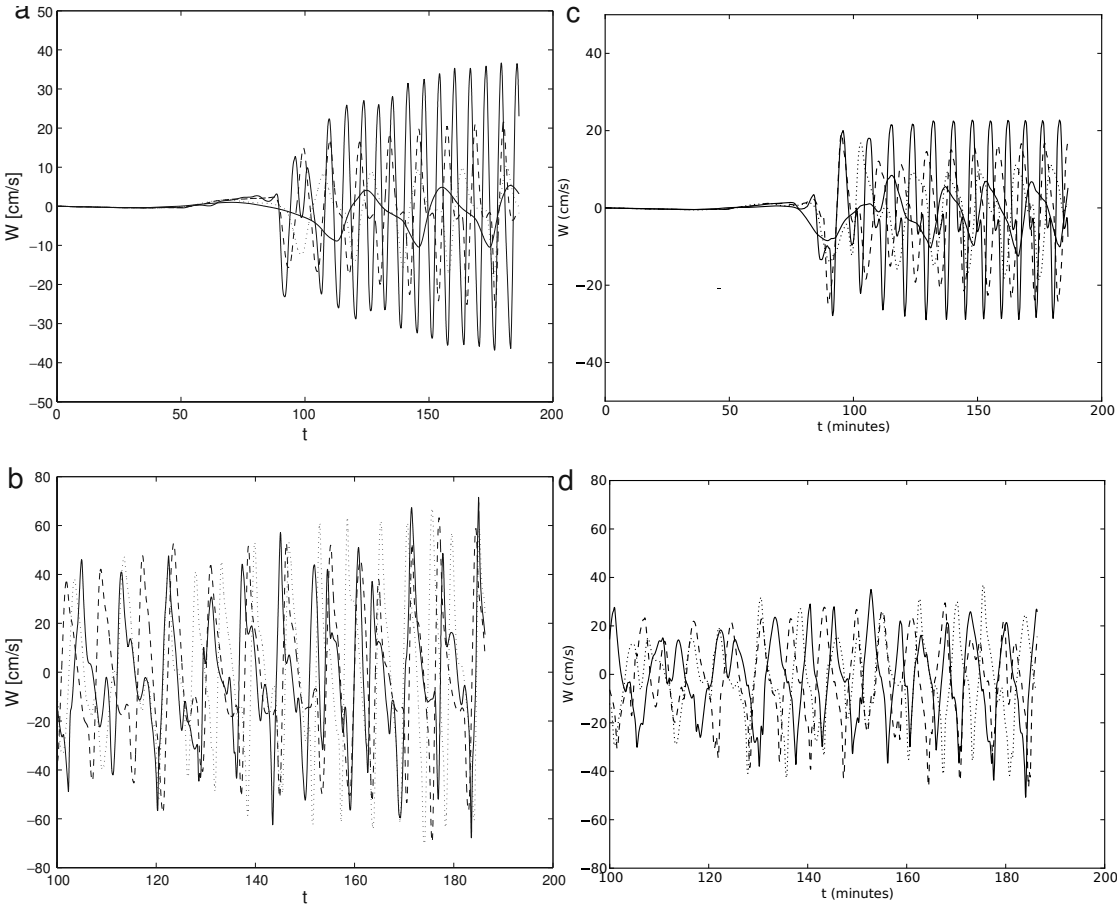


Figure 3.7.4.: A comparison of w velocities at the detector with Berntsen et al. (2009) for various resolutions. Coarser resolution results are presented in (a) for Berntsen et al. (2009) and (c) for results from this study for the following resolutions: $\Delta x = 100m$ (lower amplitude solid line), $\Delta x = 50m$ (dotted line), $\Delta x = 25m$ (dashed line), $\Delta x = 12.5m$ (larger amplitude solid line). Higher resolution results are presented in (b) for Berntsen et al. (2009) and (d) for results from this study for the following resolutions: $\Delta x = 6.25m$ (dotted line), $\Delta x = 3.125m$ (dashed line) and $\Delta x = 1.5625m$ (solid line).

3.7.3. Results

Convergence with fixed quasi-uniform meshes

In order to demonstrate the benefits of mesh adaptivity, it is essential that comparisons be made with largely converged results. Convergence is examined both qualitatively (figures 3.7.5 – 3.7.7) and quantitatively (figures 3.7.9 and 3.7.10) at three different viscosities each an order of magnitude different to one another. Dynamics are compared at a range of horizontal resolutions via snapshots of the temperature and velocity fields.

Using the highest viscosity ν_c (figure 3.7.5), the resulting dynamics are the most simple. Tidal velocity forces water over the sill top, resulting in a peak in horizontal velocity at the sill top, a peak in upward vertical velocity before the sill and a maximum downward velocity after the sill, down the inland side. The return of the overturning limb to the surface results in a much weaker upward vertical velocity which increases in amplitude beyond the initial $T/4$ period. Forcing over the sill creates interaction with the bottom boundary forming an overturned limb on the inland side, however no other more complex overturning is present for this high viscosity case. The overturned limb, which is visible in the temperature field extends $\sim$ 1750 m from the sill top and separation from the bottom boundary occurs just below the sill top at $\sim$ 30 m depth. At this viscosity diapycnal mixing is very limited with instability in the water column only within the overturned limb, which is verified by the Richardson number. As resolution is increased for this example, no additional dynamics are captured, however the representation of the overturned limb becomes sharper and less diffuse. Limited changes are observed as resolution is increased to be finer than $\Delta x = 5 - 10$ m.

Changes in dynamics for ν_b as horizontal resolution is increased however are more notable. Dynamics here show that as the fluid is forced over the sill by

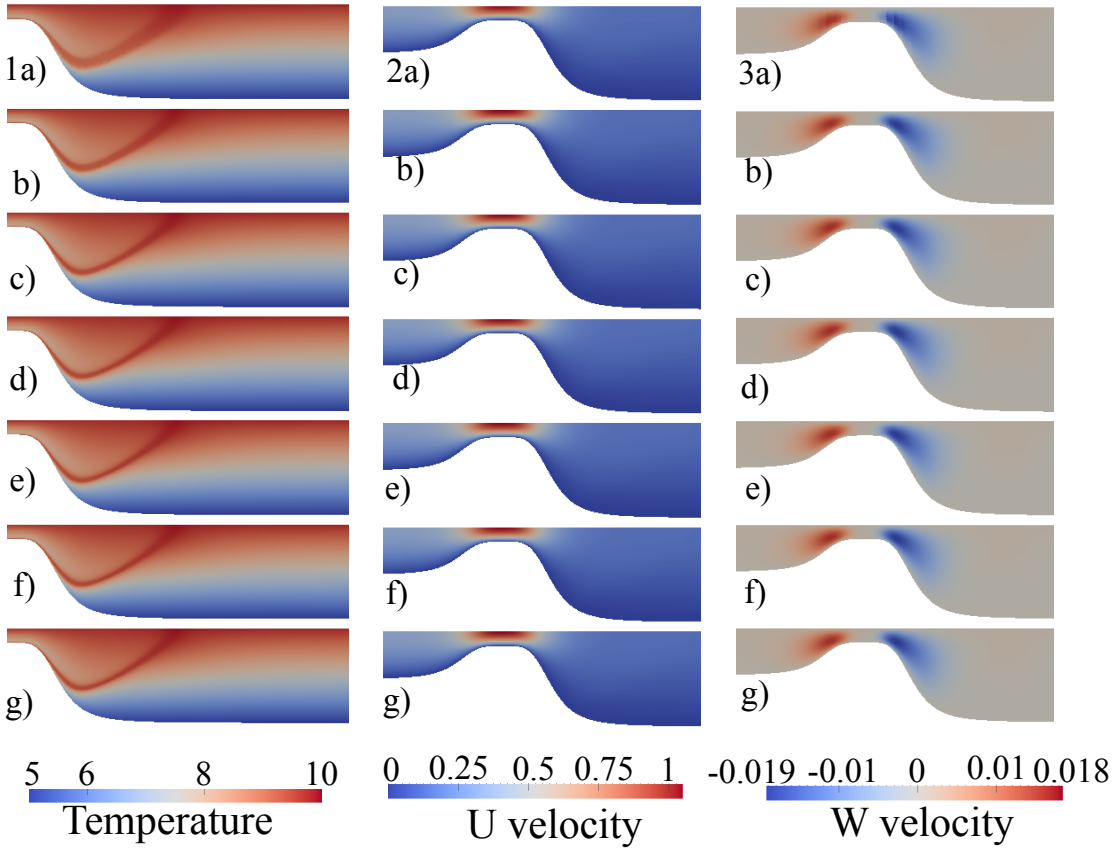


Figure 3.7.5.: Snapshots (stretched by a factor of ten vertically) of temperature (1), horizontal velocity (2) and vertical velocity (3) fields at $T/4$ for the highest viscosity ν_c ($A_h = 10^{-2}\mathrm{m}^2\mathrm{s}^{-1}$, $A_v = 1\mathrm{m}^2\mathrm{s}^{-1}$). Results are presented for a range of horizontal resolutions between $\Delta x = 40$ (a) and $\Delta x = 0.625$ m (g). All simulations at this viscosity use 100 vertical layers and a timestep of $T/7200$ s.

the tide an overturned limb is initially formed, sharing significant similarities with the dynamics of ν_c until $T/8$. Although similarities are apparent the limb for the ν_b simulations extends further from the sill, is thinner and is of a lower angle (figure 3.7.6a). After $T/8$ for ν_b, as viscous forces are not as strong as those of ν_c instabilities develop along the limb ~ 500 m from the sill top which begin to develop into overturning rolls. Overturning rolls are shown most clearly in the vertical

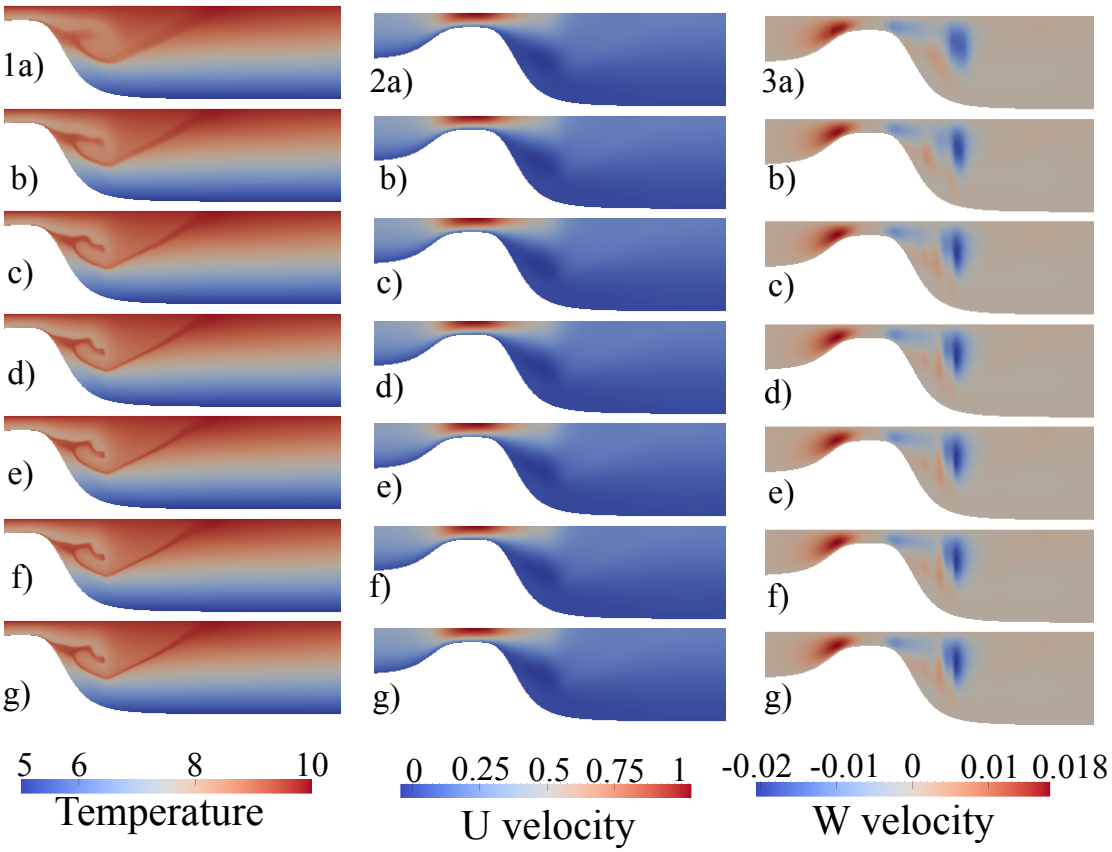


Figure 3.7.6.: Snapshots (stretched by a factor of ten vertically) of temperature (1), horizontal velocity (2) and vertical velocity (3) fields at $T/4$ for the viscosity ν_b ($A_h = 10^1 \mathrm{m^2 s^{-1}}$, $A_v = 10^{-1} \mathrm{m^2 s^{-1}}$). Results are presented for a range of horizontal resolutions between $\Delta x = 40$ (a) and $\Delta x = 0.625$ m (g). All simulations at this viscosity use 100 vertical layers and a timestep of $T/7200$ s.

velocity snapshots (figure 3.7.6(3)) where positive and negative velocities indicate the sides of the roll. These developing overturning rolls include smaller length scales than those of a simple overturned limb and therefore require higher mesh resolution to capture. When considering the temperature and Richardson number fields, visual convergence appears achieved by $\Delta x = 5 - 10$ m. However, due to overturning rolls in the velocity fields, convergence is not visible until $\Delta x = 2.5 - 5$ m as otherwise rolls appear merged.

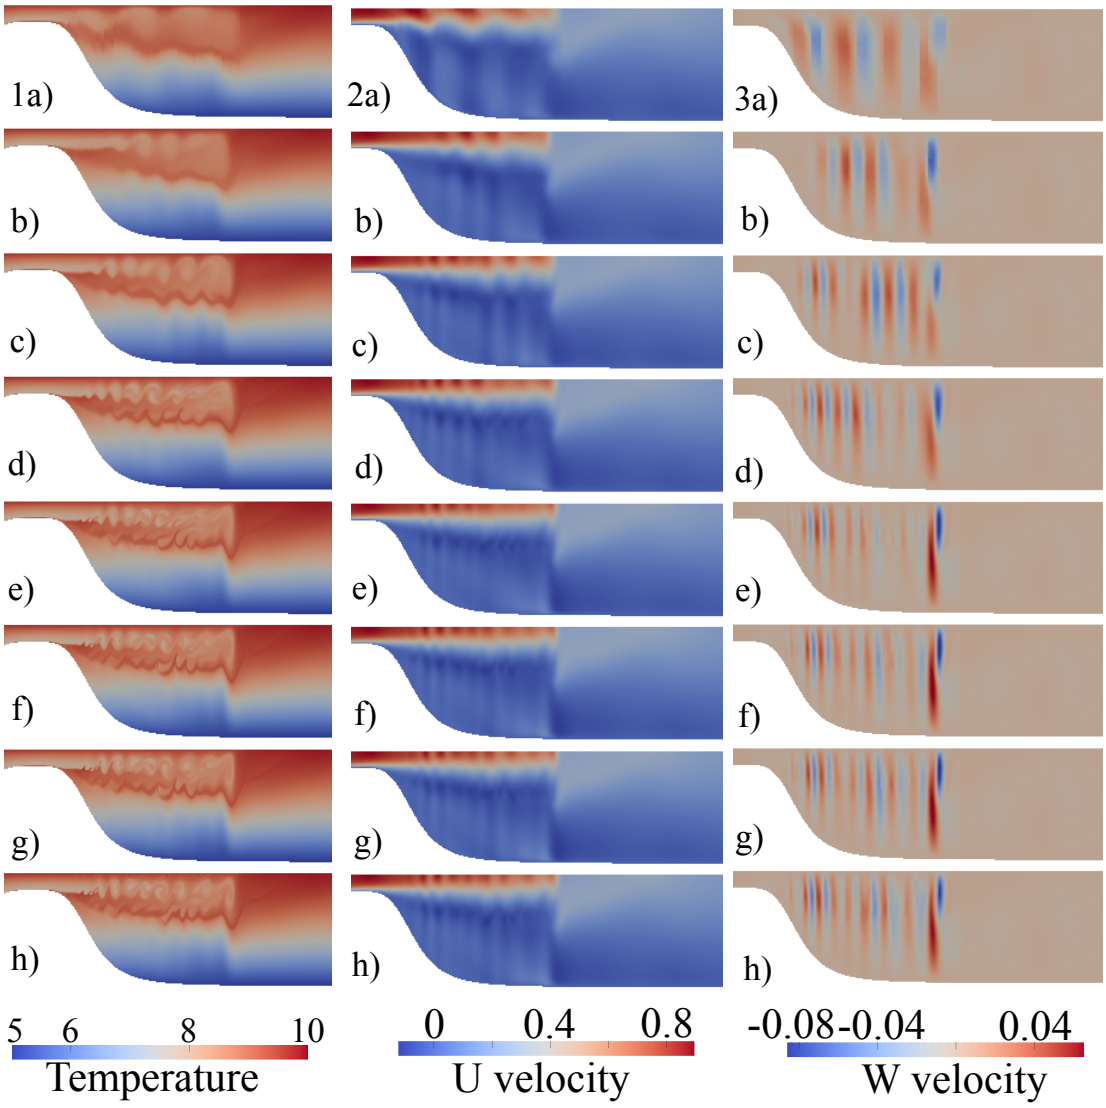


Figure 3.7.7.: Snapshots (stretched by a factor of ten vertically) of temperature (1), horizontal velocity (2) and vertical velocity (3) fields at $T/4$ for the viscosity ν_a ($A_h = 1\mathrm{m}^2\mathrm{s}^{-1}$, $A_v = 10^{-2}\mathrm{m}^2\mathrm{s}^{-1}$). Results are presented for a range of horizontal resolutions between $\Delta x = 40$ (a) and $\Delta x = 0.3125$ m (h). All simulations presented here use 100 vertical layers and a timestep of $T/7200$ s.

The most complex dynamics are those for ν_a, where after initial similarities with ν_b and ν_c again for the first $T/8$ of the simulation, the dynamics rapidly become

more turbulent due to having the lowest viscous forces. Instead of forming an overturned limb as for ν_b and ν_c, the lower viscosity leads to the formation of a short jet after $T/8$ which extends from the sill $\sim$ 300 m and at this distance then begins to form turbulent overturning rolls which propagate away from the sill (figure 3.7.7). These turbulent overturning rolls extend $\sim$ 1500 m into the domain at $T/4$ and to a depth of $\sim$ 45 m. By $T/4$ seven distinct rolls are present and are also clearly identified by the vertical and horizontal velocities (figures 3.7.7(2) and (3)) which clearly depict the side of each roll. Due to the more complex turbulent dynamics with smaller length scales present than ν_b and ν_c, even higher mesh resolution than previously used is required, and the consequence of coarse meshes are shown. For example a mesh resolution of $\Delta x = 40$ m results in only four much wider overturning rolls being captured, and therefore the wrong aspect ratio of these features and a poor representation of the model dynamics. Comparing these results qualitatively the number of overturning rolls increases as the horizontal resolution is increased until a resolution of $\Delta x = 1.25$ m when the number of rolls (and also their shape) convergences on seven. As resolution is increased the vertical velocities (3.7.7(3)) also show a change in intensity, at resolutions finer than $\Delta x = 2.5$ m the roll furthest from the sill is more energetic than the others. For $\Delta x = 2.5$ m this energetic roll is preceded by two weaker rolls however at $\Delta x = 1.25$ m and higher resolutions all of the rolls preceding this roll are of similar intensity. Therefore, visually, convergence is reached at $\Delta x = 1.25$ m.

Qualitatively as the vertical resolution is increased (figure 3.7.8) there is not a very significant change in the temperature and velocity fields, except that rolls are more sharply captured. However while the number and basic structure of the rolls is similar, on close observation the rolls at middle distance from the sill become more elongate horizontally as resolution is increased. When more than 200 layers are used, multiple layers of turbulent features are present. In addition the second

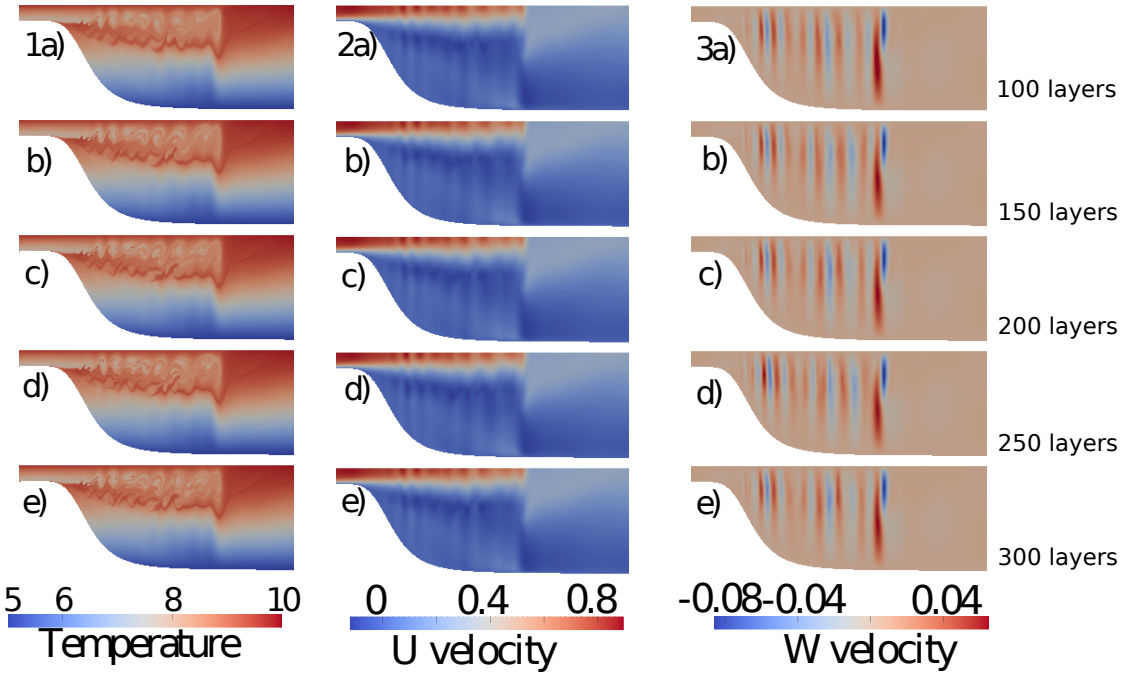


Figure 3.7.8.: Snapshots (stretched by a factor of ten vertically) of temperature (1), horizontal velocity (2) and vertical velocity (3) fields at $T/4$ for the highest viscosity ν_a. Results are presented for meshes with a range of number of vertical layers between 100 (a) and 300 (e) at a horizontal resolution of $\Delta x = 0.625$ m.

to last roll and roll furthest from the sill start to merge as vertical resolution is increased. These slight changes in features indicate qualitative convergence at approximately 200 layers.

Quantitative convergence analysis is performed in figures 3.7.9 and 3.7.10, which show convergence analyses in terms of $w_{\max}(\mathrm{T}/4)$, $w_{\min}(\mathrm{T}/4)$ and the time-averaged critical Richardson area, $\overline{A}_{\mathrm{Ri}}$, for each viscosity value considered. These parameters are considered here as vertical velocities are representative of the overturning dynamics and are more sensitive to resolution than horizontal velocities, and $\overline{A}_{\mathrm{Ri}}$ indicates changes in the area of possible mixing occurring. Convergence is examined at T/4 as instabilities that develop later in the simulations make model convergence far harder to demonstrate even at higher viscosities. Convergence

graphs indicate that resolutions of approximately $\Delta x = 5$ m and 50 sigma layers are required for reasonable convergence at ν_c, $\Delta x = 2.5$ m and 100 layers for ν_b and $\Delta x = 1.25$ m and 250 layers for ν_a. A study of the influence of Δt is also undertaken for ν_a with convergence at $\Delta t = T/28800$ s. In addition, for each viscosity value, the adaptive results (Av3, Av2 and Aa also shown in figures 3.7.9 and 3.7.10 and with parameters presented in table 3.7.1) agree well with the zone of convergence for velocity and $\overline{A}_{\mathrm{Ri}}$. The results at the lowest viscosity used here (ν_a) use a higher resolution than that of previous investigations in the literature despite using a viscosity ten times larger, which may suggest these previous studies have not completely converged in terms of the overturning dynamics.

The dynamics at T/4 for each of these viscosities can be seen in figure 3.7.11 for temperature and figure 3.7.12 for $\overline{A}_{\mathrm{Ri}}$ at the converged resolutions specified above, compared with a fixed low resolution of $\Delta x = 80$ m and 100 layers. These figures show that by decreasing viscosity the amount of complex turbulent overturning dynamics increases and also these complex dynamics begin earlier within the simulation, requiring higher mesh resolutions. ν_c shows no turbulent overturning, while overturning cells are just starting to form at the end of the jet at T/4 for ν_b, however ν_a exhibits strong overturning features from T/8. Converged simulations of ν_a also agree well qualitatively with physical studies by (Inall et al., 2004), in terms of the extent of the overturning rolls from the sill and the separation point location.

Dynamics appear fairly well captured even in the low resolution snapshots for ν_c (Figures 3.7.11 (1a) and 3.7.12 (1a)), with the only noticeable differences being a slightly wider more diffuse front. However, as viscosity is reduced the difference between these low resolution snapshots and the converged snapshots increases. Once overturning characteristics develop in ν_a the low resolution snapshots show fewer diffuse overturning rolls and a longer wavelength internal wave, resulting in

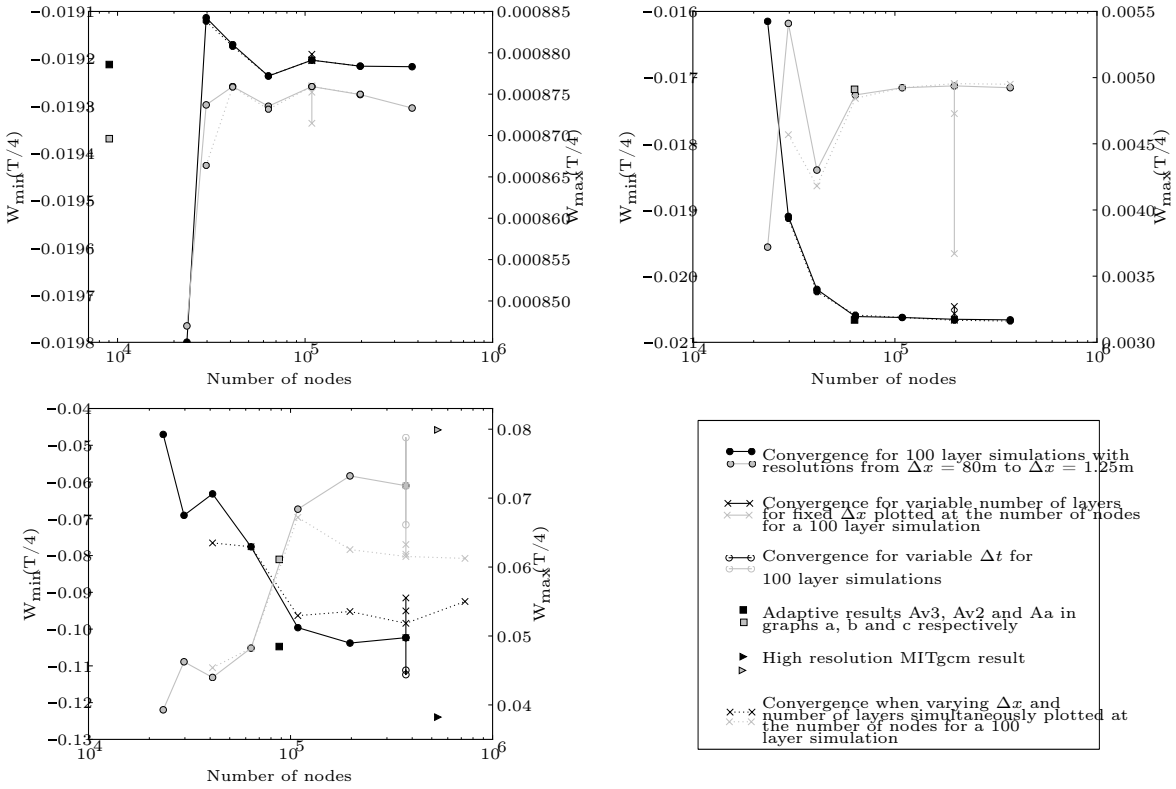


Figure 3.7.9.: Convergence of $w_{\max}(\mathrm{T}/4)$ (grey) and $w_{\min}(\mathrm{T}/4)$ (black) for (a) ν_c, (b) ν_b and (c) ν_a. A range of horizontal resolutions, layers (25 – 125 for (a), 25 – 150 for (b) and 100 – 300 for (c) and Δt ($\frac{T}{7200}$ – $\frac{T}{57600}$) are explored as defined in the legend. Results using mesh adaptivity (Av3, Av2 and Aa) and high resolution mesh MITgcm are also presented. Studies of Δt and the number of layers are conducted for a fixed horizontal resolution and therefore in these cases the total number of nodes is a multiplier of the x value. For simulations where the horizontal resolution is varied simultaneously the total number of nodes is again a multiplier of the x value in order for ease of comparison.

a large difference to the area that turbulent mixing may occur over.

As viscosity is decreased and resolution is increased $\overline{A}_{\mathrm{Ri}}$ and $w_{\max}(\mathrm{T}/4)$ increase which indicates overturning occurs over a larger area and at a higher intensity. At lower resolutions the area of critical Richardson number is a block with little structure whereas at high resolutions the intricacies of where overturning and

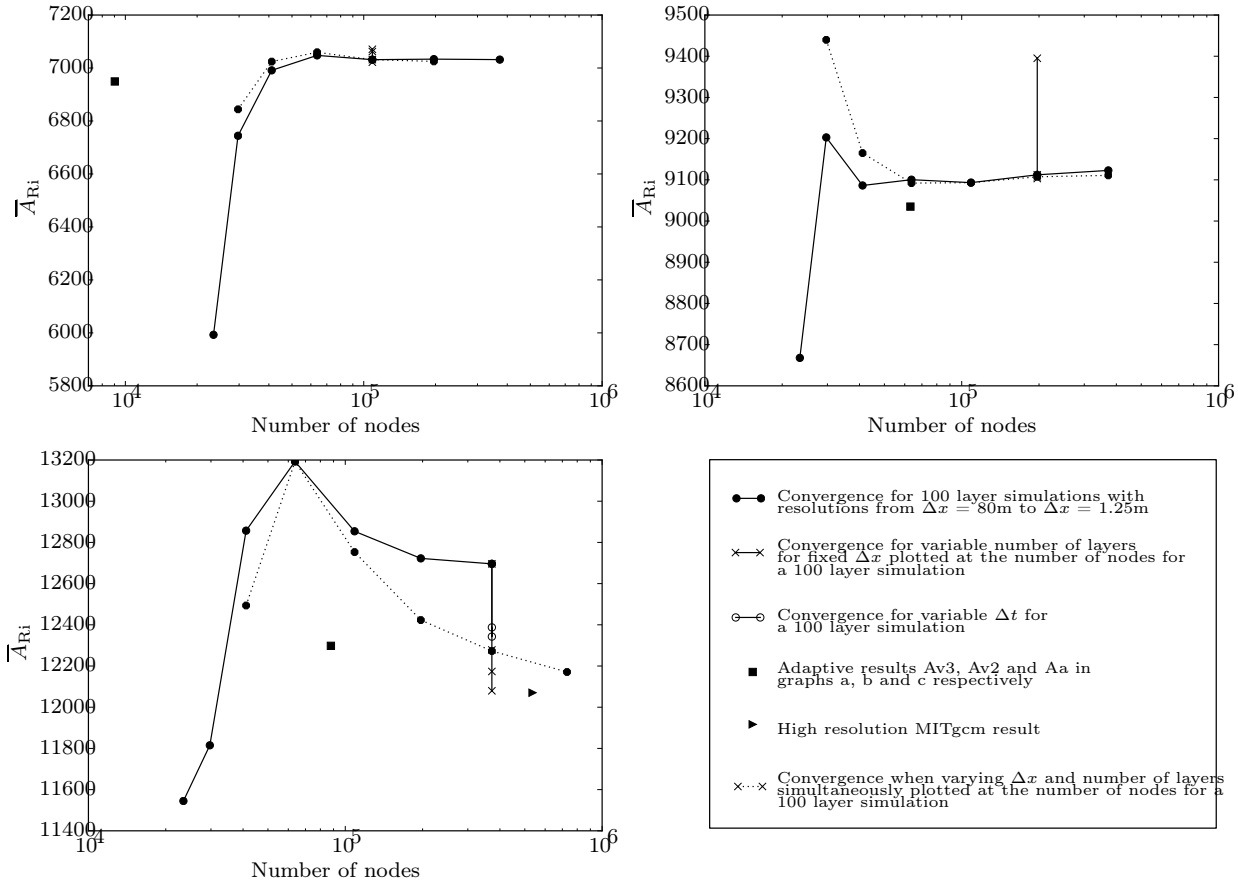


Figure 3.7.10.: Convergence of $\overline{A}_{\mathrm{Ri}}$ for a) ν_c, (b) ν_b and (c) ν_a. As figure 3.7.9 a range of horizontal resolutions and layers are shown for each converged simulation as well as a range of Δt. Results using mesh adaptivity (Av3, Av2 and Aa) and high resolution mesh MITgcm are also presented. Studies of Δt and the number of layers are conducted for a fixed horizontal resolution and therefore in these cases the total number of nodes is a multiplier of the x value. For simulations where the horizontal resolution is varied simultaneously the total number of nodes is again a multiplier of the x value in order for ease of comparison.

mixing occurs are visible (figure 3.7.12), which over time leads to the more intricate dynamics of the temperature fields.

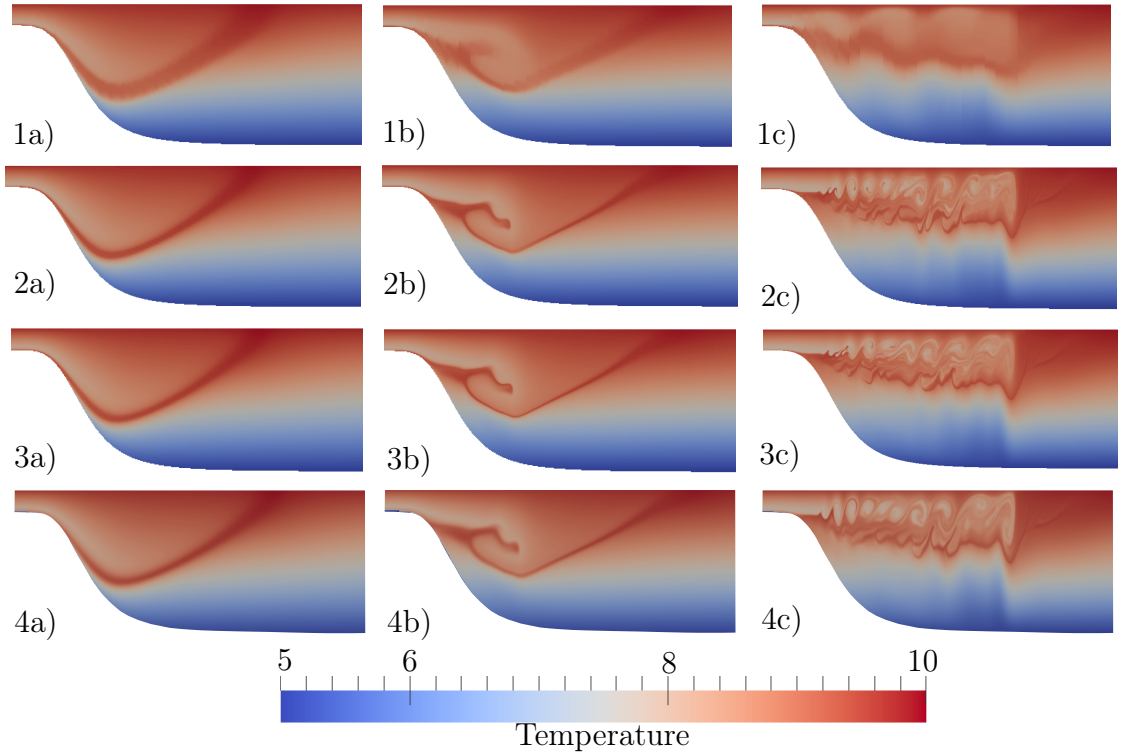


Figure 3.7.11.: Snapshots (stretched by a factor of ten vertically) of the temperature field for each viscosity (a) ν_c, (b) ν_b and (c) ν_a at T/4. Results are shown for a low resolution mesh ($\Delta x = 80$ m with 100 layers) (1), a high resolution mesh at the resolution as stated in the main text as being required for convergence (2), an unstructured adaptive mesh (3) and a high resolution MITgcm result (4). Resolutions shown for the high resolution Fluidity meshes at each viscosity are: (a) $\Delta x = 5$ m and 50 layers, (b) $\Delta x = 2.5$ m and 100 layers, (c) $\Delta x = 1.25$ m and 250 layers, while MITgcm uses (a) $\Delta x = 10$ m and $\Delta z = 1$ m, (b) $\Delta x = 2.5$ m and $\Delta z = 1$ m, (c) $\Delta x = 1.25$ m and $\Delta z = 0.5$ m (where these resolutions where determined using a brief convergence study.).

Comparisons with MITgcm

High resolution MITgcm simulations show very good agreement (figure 3.7.11 and 3.7.12) with both converged fixed and adaptive mesh Fluidity simulations, which further verifies model convergence. At ν_c and ν_b, temperature fields display well

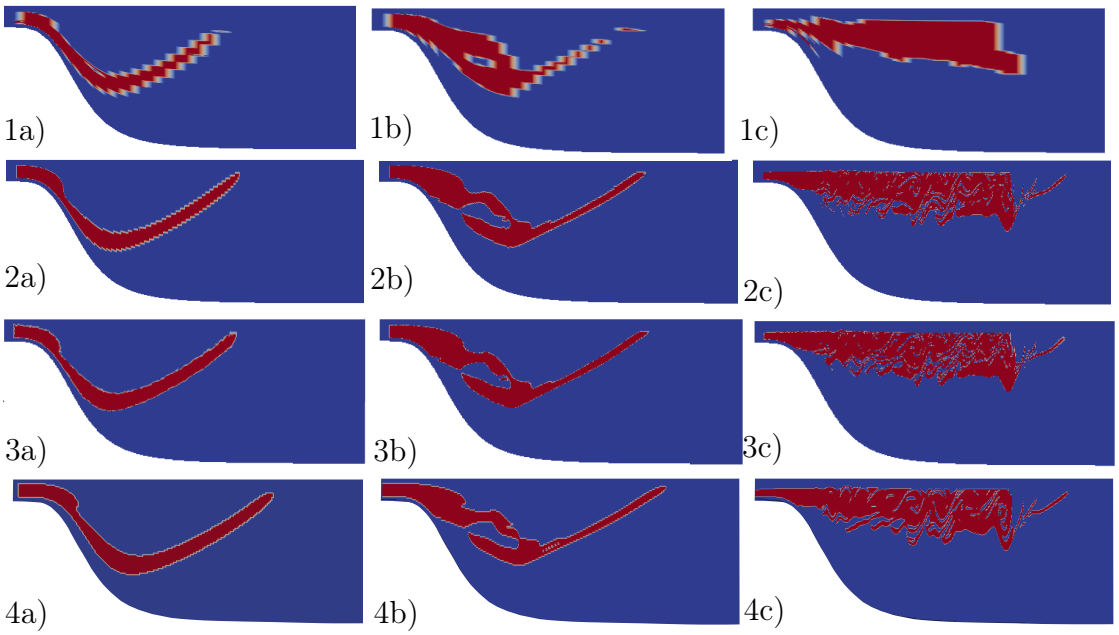


Figure 3.7.12.: Snapshots (stretched by a factor of ten vertically) of the critical Richardson area field for each viscosity (a) ν_c, (b) ν_b and (c) ν_a at T/4. Results are shown for a low resolution mesh ($\Delta x = 80$ m with 100 layers) (1), a high resolution fixed mesh at the resolution (as previous figure) required for convergence (2), an unstructured adaptive mesh (3) and a high resolution MITgcm result (4). Resolutions shown for the high resolution Fluidity meshes at each viscosity are: (a) $\Delta x = 5$ m and 50 layers, (b) $\Delta x = 2.5$ m and 100 layers, (c) $\Delta x = 1.25$ m and 250 layers, while MITgcm uses (a) $\Delta x = 10$ m and $\Delta z = 1$ m, (b) $\Delta x = 2.5$ m and $\Delta z = 1$ m, (c) $\Delta x = 1.25$ m and $\Delta z = 0.5$ m (where resolutions were determined by a separate convergence study).

defined and similar dynamics while only slight differences are present for $\overline{A}_{\mathrm{Ri}}$, where a small additional unstable area is present at the sill top. These small differences are due to slight differences in the velocity shear and stratification in this area and indicate that $\overline{A}_{\mathrm{Ri}}$ is a robust diagnostic for convergence analysis. Simulations at ν_a agree well in terms of the major features of the dynamics in addition to the extent and depth of overturning rolls; slight differences however in the very small scale detail and shape of the rolls are present. These differences are influenced in

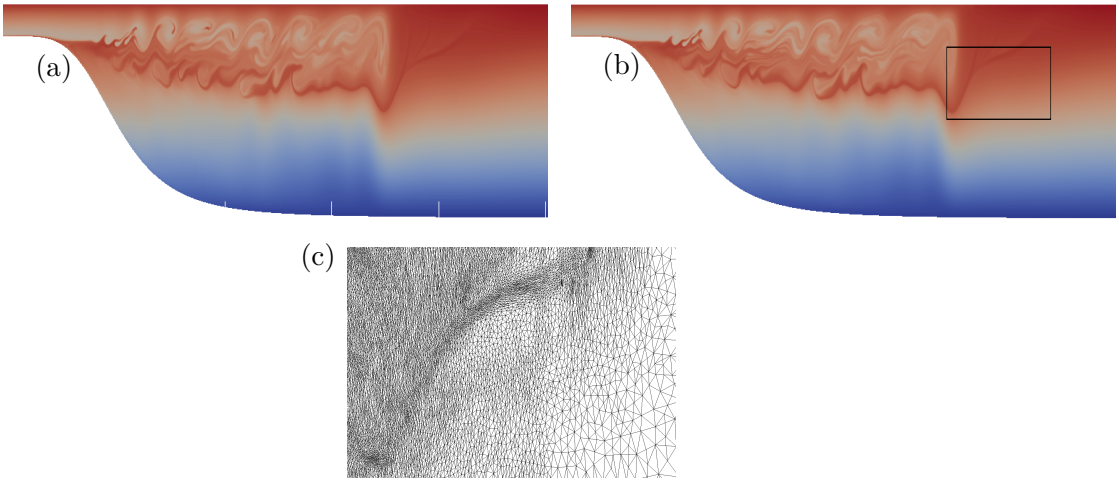


Figure 3.7.13.: Snapshots, stretched by a factor of ten vertically, of the temperature field for ν_a at T/4. Results are shown for the high resolution fixed mesh resolution (a), compared with the adaptive mesh simulation (Aa) (b), where the adapted mesh in the box is shown in (c). The mesh follows the features of the dynamics as resolution has been enhanced here and elements are aligned with the structures.

part by the differences in vertical co-ordinates used. Fluidity uses σ- coordinates which results in non-uniform element shapes at sharp changes in bathymetry but also enhances vertical resolution on the sill top. Fluidity simulations therefore have 100 layers at any point in the domain, whereas MITgcm with a constant Δz results in the equivalent resolution of only 15 layers at the sill. Despite these small differences figure 3.7.10 shows measurements of $\overline{A}_{\mathrm{Ri}}$ for MITgcm agree within 5% of converged Fluidity values. Values for $w_{\mathrm{max/min}}(\mathrm{T}/4)$ show reasonable agreement but do not agree as well as $\overline{A}_{\mathrm{Ri}}$ as they are instantaneous measures of extremes and more easily influenced by small differences between models than averaged diagnostics.

Mesh adaptivity – the capture and representation of dynamics

Unstructured adaptive mesh simulations (figures 3.7.9 – 3.7.12) capture the dynamics with little difference observable compared with those of the high-resolution fixed converged mesh at all viscosities. As the dynamics become more complex at lower viscosities, mesh adaptivity provides a larger computational saving, using on average 17% of the number of nodes of the converged fixed mesh at ν_c and only 9% of the number of nodes at ν_a. At ν_c the front is as narrow and well-defined for both the adaptive and high resolution fixed, converged Fluidity results. For the slightly more complex ν_b the front is also well-defined and sharp, in addition the flow separation occurs at the same location and the end of the jet (which is on the verge of producing overturning rolls) is the same shape and is possibly better defined for the adaptive simulation. For the ν_a simulation the dynamics are far more complex, flow separation occurs at the crest of the sill and is followed by a short jet that produces a number of overturning rolls that extend away from the sill and increase the area for mixing to occur across. For both adaptive and high resolution Fluidity simulations, the location of separation and the length of the jet are the same in addition to the depth to the bottom of the area of turbulence and the number of overturning rolls. The structure of the overturning rolls and the distance they propagate from the sill is also comparable for both the fixed high resolution and adaptive simulation. This is achieved with a minimum edge length (2.5 m) imposed within the mesh adaptivity step, which is larger than the Δx used for the converged fixed mesh (1.25 m), possibly due to the versatility of an unstructured adaptive mesh, where mesh elements are capable of aligning with flow features such as overturning rolls unlike a fixed structured mesh which forces the use of vertical and horizontally (or bathymetry) aligned elements to model these elliptical features. This mesh alignment can be seen in figure 3.7.13, where in (c) the shapes of the overturning rolls are apparent within the mesh.

Effect of mesh adaptivity settings

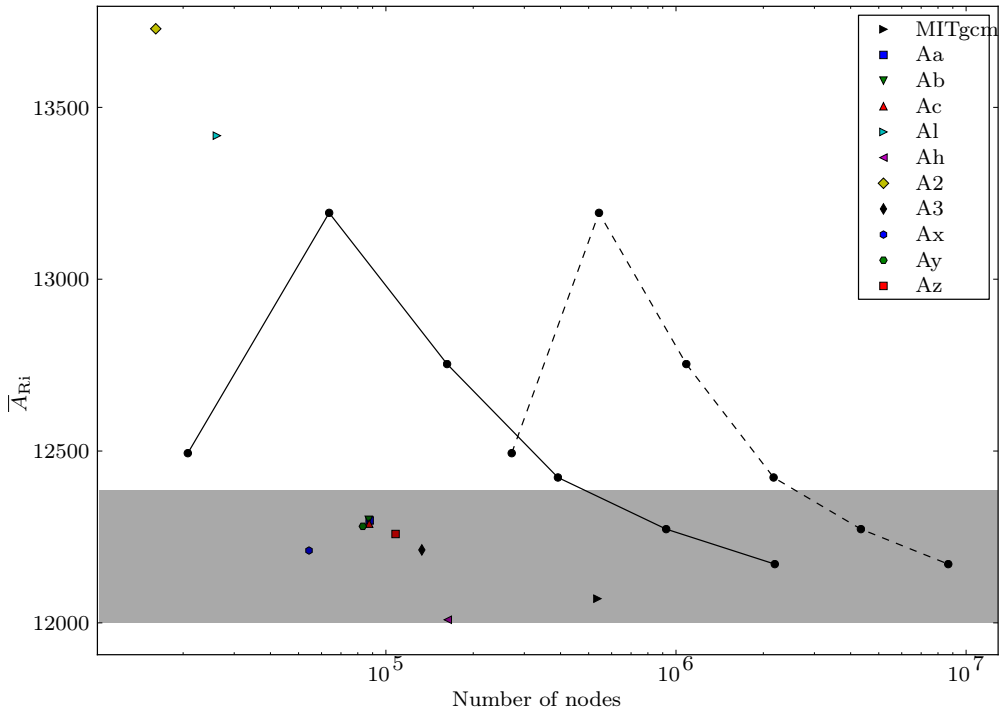


Figure 3.7.14.: The relationship between time-averaged critical Richardson area and the number of nodes, where the number of nodes for adaptive simulations is the average over the simulation time. Adaptive results are presented (defined in table 3.7.1) along with fixed mesh results (solid line) for simulations varying Δx (circles) and number of layers simultaneously and a high resolution MITgcm result (black triangle) also shown in figure 3.7.10. The dashed line indicates the number of nodes that would have been used by the fixed mesh without mesh telescoping. Fixed mesh results suggest a range of 12000 – 12400 m^2 for convergence based on results investigating Δt and number of layers, this is shown by the grey block on the graph and agrees with most adaptive results and results from MITgcm.

Figure 3.7.14 shows results from ten different fully unstructured adaptive mesh simulations for ν_a (with adaptivity parameters given in table 3.7.1) compared with the convergence plot for the fixed structured mesh simulations in terms of time-

averaged critical Richardson area where convergence is seen on the range 12000 – 12400 m^2). Different parameter settings were used in order to highlight their importance for convergence. Time-averaged critical Richardson area was used to demonstrate convergence here as it proved to be the most reliable diagnostic for this purpose at it is an averaged metric representative of mixing and was not as likely to be influenced by spurious affects which may influence maximum and minimum values.

Interpolation error bounds (or field weights) are used to control the number of nodes in the mesh, lower field weights mean that the result should provide a better representation of the dynamics. Simulations Aa, A2 and A3 show that the highest field weights result in a reduced number of nodes but also a solution outside the desired range for $\overline{A}_{\mathrm{Ri}}$. Simulation Aa uses weightings which equate to approximately 1% of the range of the relevant fields, which compared with A2 (which corresponds to 10%), results in a larger number of nodes and a solution within the desired range. Simulation A3 further reduces these weightings to $\sim$ 0.5%, however only a small increase in nodes is observed and a small change in $\overline{A}_{\mathrm{Ri}}$ indicating convergence upon the desired range. Since the dynamics of interest are located close to the sill, in order to further save computational time this field weighting is specified as a function of space and telescoped away from the sill using an exponential function across a set length, changing field weightings to $\sim 10\%$ at distance from the sill. Results however indicate that the length of this telescoped zone is relatively unimportant as restrictions on the changes in edge length from element to element are likely to have prevented any disturbance to the dynamics.

Minimum edge lengths also influence the total number of nodes and the accuracy of the results. By increasing the minimum edge length (Al) the resultant mesh has fewer nodes but with a solution outside the desired range as we force the model to use edges longer than the value used for the field weight may wish. This is a

similar effect to using an under resolved fixed mesh as shown earlier where excess diffusion resulted. By decreasing the minimum edge length (Ah) dynamics are represented slightly more sharply, however there is only a small change in $\overline{A}_{\mathrm{Ri}}$ and both results Aa and Ah are in the desired range for this parameter, indicating convergence at Aa (which is referred to as the best adaptive simulation due to the larger computational saving over Ah). The use of an adaptive timestep (Az) can also be beneficial in simulations such as this as the timestep necessary for stability and to reduce numerical diffusivity can vary throughout the simulation based on the changing flow velocities of a barotropic tide. For this simulation however there was little difference observed in terms of $\overline{A}_{\mathrm{Ri}}$ or the dynamics therefore we can be confident in the fixed timestep used for Aa.

Most simulations here use the M_2 metric which has been shown to improve the representation of small scale dynamics (Hiester et al., 2014; Loseille and Alauzet, 2011). However, for comparison the M_∞ norm has also been used (Ax) and in fact shows a very similar result to Aa with a slight computational saving, this shows that either metric may be used for these simulations. Further simulations however use the M_2 norm as there is little difference and it has been successfully used previously for modelling highly complex turbulent-like dynamics (Hiester et al., 2011). While metric-advection was previously found to be useful for simulations of a lock-exchange (Hiester et al., 2011), this application shows that a limited difference is made by this option despite the relatively long time between mesh adapts, where the method is usually most beneficial. Therefore due to the slight computational overheads further simulations have been conducted without this feature. However, it would be interesting to consider advecting the metric with the wave speed rather than the velocity field. It is worth noting that the previous study (Hiester et al., 2011) involved a front moving with the velocity field, which is not the case here.

The adaptive simulation that agrees with the dynamics of the high resolution fixed meshes and seems most stable while agreeing with the parameter choice of previous work to produce a good compromise between expense and solution accuracy is Aa. This simulation has an average of 9×10^4 nodes while the high resolution converged telescoping fixed mesh uses 1.1×10^6. It should be noted however that the telescoping mesh used in this fixed mesh run is already saving mesh resolution and without this it would need approximately 4×10^6 nodes. The use of an adaptive mesh here therefore saves approximately 1.5 orders of magnitude in terms of degrees of freedom used.

Assessment of mixing diagnostics

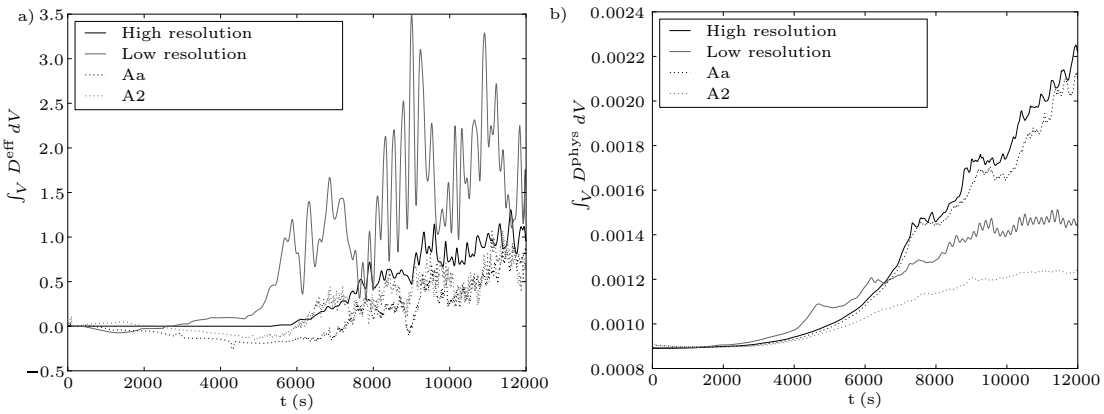


Figure 3.7.15.: Comparison of the mixing for lower and higher resolution fixed and adaptive mesh simulations at ν_a over the diagnostics region. Subfigure (a) shows that D^{eff} is highest for the low resolution fixed mesh and lowest for adaptive mesh simulations, however D^{phys} alone (b) is highest for the high resolution fixed mesh.

High resolution fixed meshes appear to produce $\sim 40\%$ less effective mixing than coarse resolution fixed meshes due to lower levels of spurious numerical mixing. For the high resolution fixed mesh the distinct lower frequency peaks in D^{phys} with a period of $\sim$2000 s can each be attributed to the formation of well-defined

energetic overturning rolls, which increase the area for mixing to occur across. These overturning features begin to form at $\sim$ 5000 s at the end of the jet and propagate away. Later in the simulation at $\sim$ 10000 s a larger increase occurs in the physical mixing rate for high resolution results probably due to the merge and partial collapse of rolls further from the sill, which is poorly captured in coarser results and A2. In addition to more physical mixing occurring for high resolution fixed meshes, we know due to $\overline{A}_{\mathrm{Ri}}$ that the area that this mixing occurs over is only 10% larger, therefore mixing is more intense and wider spread.

In contrast, the low resolution simulation displays a lower overall rate of physical mixing, still with distinct peaks due to overturning features, however these features are poorly captured and there are fewer than for higher resolution meshes. The effective mixing shown in figure 3.7.15(a) shows that lower resolution fixed meshes with less distinguishable features due to the dynamics are thus "overpowered" by spurious mixing due to coarse mesh resolution in the effective mixing plot.

In order to demonstrate the effectiveness of the adaptive meshes, results from Aa and A2 are also considered. Aa produces almost identical mixing characteristics, in terms of D^{phys}, to the high resolution simulation. These characteristics indicate that rolls form, propagate, merge and collapse in a very consistent manner. Simulation A2 (which uses the fewest number of nodes) however, represents similar features to the low resolution fixed mesh simulation for D^{phys}, although it still demonstrates the features associated with the most turbulent overturning rolls. Simulation A2 therefore demonstrates the large influence small changes in adaptivity parameters (and hence mesh) may cause. Both adaptive results generate less effective mixing than the fixed meshes due to less spurious mixing, this is due to the high mesh resolution that has developed to fully resolve the dynamics where and when needed.

These simulations are however only run for T/4 and so have not yet reached a

quasi-steady state, therefore no additional conclusions can be drawn on mixing. This also may explain the difference in amplitude of physical and effective mixing between the high resolution fixed mesh amplitude and Aa, at this stage the results display sufficient similarity to indicate a good representation of the dynamics has been achieved.

3.8. Conclusions

Simulations of a simple internal seiche and tilting tank have been used to compare vertically structured and unstructured approaches to mesh adaptivity in terms of their representation of dynamics and the levels of numerical mixing induced. Fully unstructured adaptivity produces the best representation of the tilting tank dynamics while minimising spurious mixing for both the tilting tank and seiche simulations. Vertical only adaptivity provides an efficient representation of the tilting tank, producing the largest computational saving whilst maintaining a good representation of the dynamics. Compared with a fixed mesh of a similar number of nodes it produces roughly an order of magnitude less spurious mixing. In addition these simple cases showed that the choice of the type of structured or partially structured adaptivity made little impact on the numerical mixing (two to four times the amount) especially when compared with the effect of discretisation type which can cause one or two orders of magnitude difference. This led to the use of discontinuous Galerkin discretisation methods for all simulations due to these being the least diffusive.

The simple internal seiche and tilting tank are also used to compare the two methods of quantifying mixing. They show that method of Burchard and Rennau (2008) provides a robust and comparable method that is computationally cheap and simple to apply for more complex scenarios such as Loch Etive. Results

obtained using both methods display the same characteristics for when mixing rates are high with only slight disparities in magnitude in some cases.

Studies of the semi-idealised case of Loch Etive demonstrate the importance of mesh resolution and viscosity on model results. By increasing the resolution, the dynamics become increasingly complex and turbulent, with an increase in the influence of overturning rolls. By studying three viscositites we observe convergence at $\Delta x = 5$ m and 50 layers for ν_c, $\Delta x = 2.5$ m and 100 layers for ν_b and $\Delta x = 1.25$ m and 250 layers for ν_a (with a timestep of $\frac{T}{28800}$ s). This shows the additional resolution required to fully resolve the more complex dynamics at lower viscosities. For each of these viscosities an adaptive mesh simulation was presented which was capable of resolving the dynamics and providing a computational saving. This also illustrated that as viscosity is reduced the savings made by adaptive mesh simulations increase using, on average, 17% of the number of nodes of the converged fixed mesh at ν_c and only 9% of the number of nodes at ν_a.

By using fully unstructured adaptive meshes, elements may align with features of the dynamics such as overturning rolls to help maintain sharp interfaces and reduce the numerical diffusion while using slightly lower resolution in these areas than fixed meshes. For convergence with ν_a, 250 layers were required resulting in $\Delta z = 0.06$ m over the sill and $\Delta z = 0.4$ m within the fjord, this heterogeneity may not have been necessary to resolve the local dynamics everywhere and the problem of this varied vertical resolution is also improved by use of adaptive meshes.

By further studying adaptivity parameters with ν_a, a good compromise in terms of number of nodes and accuracy of solution was found in simulation Aa. Solutions are sensitive to field weight and minimum edge length, and appear converged for Aa. However, features such as metric advection were not found to significantly assist the solution. The simulation Aa shows sufficient visual similarity to a high resolution fixed mesh simulation and represented an appropriate value of $\overline{A}_{\mathrm{Ri}}$. In

addition, it also displays very similar D^{phys} characteristics indicating the same processes of overturning roll formation, merging and breaking, and low spurious numerical mixing. When compared with a fixed mesh with no telescoping of mesh resolution that has converged, Aa saves 1.5 orders of magnitude in terms of the number of nodes. This indicates a considerable computational saving and should enable the modelling of more complex scenarios with greater ease. While computational time would not be further saved using an adaptive timestep this feature may increase the stability of the solution in some circumstances.

CHAPTER 4

THE INFLUENCE OF A SECOND SILL: A STUDY OF IDEFJORDEN A TWO SILL FJORD

Contents

4.1 Synopsis 123

4.2 Introduction 124

4.2.1 Idefjord 125

4.3 Methods 127

4.3.1 Diagnostics 130

4.4 Comparison with previous idealised studies 132

4.5 Effect of Realistic bathymetry and stratification based on Idefjorden . . 134

4.6 Validation against real world data 137

4.7 Tidal dynamics between two closely spaced sills in a narrow fjord: the development of a standing internal wave 140

4.7.1 Realistic stratification and bathymetry for Idefjord 140

4.7.2 Increasing the depth of the Svinesund sill top 143

4.7.3 Increasing the depth of the surface layer 144

4.7.4 Varied sill spacing and a deeper surface layer 146
4.8 Svinesund sill height change and bed shear stress 153
4.9 Discussion 154
4.10 Conclusions 159

4.1. Synopsis

A suite of numerical simulations of the real-world fjord system Idefjord are conducted to aid understanding of the scale at which bathymetric features interact via the hydrodynamics they generate. MITgcm was used for two-dimensional vertical slice simulations, which were found to agree well with physical studies of Idefjord conducted by Johnsson et al. (2007). This work contrasts with more idealised simulations conducted by Xing and Davies (2011) to model the two sill system with smooth Gaussian bathymetry and linear stratification.

Simulations revealed that for the conditions present during previous studies of the fjord the system acts as a "jet-basin" due to the shallow surface layer. Reducing the height of the second sill was considered, motivated by plans being considered by the Norwegian Coastal Administration to allow an easier shipping passage and to better understand a two sill system. This reduces the frequency of the peak power spectral density of the system at the basin centre due to the influence on superposition effects. Upon increasing the depth of the surface layer, the proximity of the sills leads to the development of a standing internal wave in the inter-sill basin for sub-critical Froude number simulations. These dynamics persist until the sills are separated by a distance roughly equivalent to the wavelength of the internal wave. Super-critical Froude number simulations demonstrate a reduced influence of a second sill as the dominant dynamics are those of a hydraulic jump and high frequency internal waves. In addition changes in the height of the Svinesund sill could potentially disturb toxic sediments by changing the velocity field and therefore bed shear stress.

4.2. Introduction

In shelf seas and fjords, the dissipation of the internal tide is crucial for sustaining shelf sea primary production (Sharples et al., 2007) and fjord basin water mixing, enabling replenishment of often stagnant basins (Stigebrandt and Aure, 1989; Johnsson et al., 2007). Fjord systems can be classified as one of two types based on a densimetric Froude number, $F_d = U/c_i$, where U is the barotropic tidal velocity over the sill and c_i is the internal wave speed in the basin (Stigebrandt and Aure, 1989). If $F_d < 1$ the flow is sub-critical and we have a "wave basin", in which internal tides are generated at the slopes of the sill. In the super critical case, $F_d > 1$ and we have a "jet basin", where tidal jets are generated at the sill. The densimetric Froude number may be considered as a ratio of the barotropic current component over the sill to the phase speed of the fastest internal wave mode. For a two layer system the phase speed of the fastest internal wave mode may be estimated as Johnsson et al. (2007) by:

$$c_i = \sqrt{g\frac{\rho_2 - \rho_1}{\rho_0}\frac{H_1 H_2}{H_1 + H_2}} \tag{4.2.1}$$

where g is the gravitational acceleration, ρ_1 is the density of the surface layer, ρ_2 is the density of the lower layer, ρ_0 is the characteristic density and H_1 and H_2 are the corresponding heights of the two layers.

The conversion rates of tidal energy (into kinetic energy by buoyancy and dissipation) on large scales are well constrained and parameterised (Green and Nycander, 2013), but there are ambiguities at which scale topographic features interact. Using observations, Johnsson et al. (2007) suggest that the two closely separated sills in the Idefjord (Figure 4.3.1) do not interact with each other in terms of tidal conversion. The modelling results by Xing and Davies (2011), on the other hand,

claim that this is not the case universally because they see a change in the spectra from the top of the sills as they are moved further apart in their (idealised) simulations. In addition, due to the proximity of the two sills they suggest a standing internal tide is generated, similar to a seiche, which is formed by the superposition of the internal waves generated at the two sills. Standing waves have been observed previously in fjords due to the interaction of internal waves (generated either by a sill (Allen and Simpson, 1998) or by the interaction of the wind (Arneborg and Liljebladh, 2001)) with their reflection from the end of the fjord.

The two investigations of the Idefjord presented by Johnsson et al. (2007) and Xing and Davies (2011) are obviously inconsistent. Here, we therefore revisit the Idefjord with the aim to re-investigate if two closely spaced sills do interact, and at which separation distance this interaction ends or begins. We accomplish this by using MITgcm to conduct a range of simulations with realistic bathymetry. Firstly to capture the dynamics at the time of measurement by Johnsson et al. (2007) and then by investigating a range of inter-sill separations while monitoring changes in the power spectral density and density contours. In addition changes in the height of the inner sill are used to examine the extent to which a second sill influences dynamics.

4.2.1. Idefjord

The Idefjord system is a 25 km long fjord system along the border between Sweden and Norway (see Figure 4.3.1a). The fjord has a total water volume of $400 \times 10^6 \mathrm{m}^3$, and can be divided into two sections with a 90 degree change in direction; this study focuses on the outer section and in particular the basin between the two outermost sills. The further inland sill is the Svinesund sill and the seaward sill is actually made up of two sills as it is forked: the Seläter sill and the Bjällvarp sill. However, due to their proximity and small separating basin they are frequently

considered as one sill and referred to as Bjällvarp. Svinesund and Bjällvarp are separated by a 1 km long basin with an average depth of 22 m, although the distance between sill tops is $\sim$ 1.5 km (Johnsson et al., 2007; Brattström et al., 2015). Large freshwater inputs from two rivers and high rates of precipitation during Autumn and Winter into the system result in a freshwater surface layer (2 m deep at the time of measurement), below which there is a gradual density gradient (Johnsson et al., 2007). Previous studies of other fjord systems have shown that high levels of freshwater runoff like that observed in Idefjord, lead to higher rates of productivity but also increase the stratification and therefore produce more stagnant bottom waters leading to anoxic periods (Gustafsson and Nordberg, 2002).

The largest town along the fjord system is Halden with a population of 30,000, located at the change in direction of the fjord and $\sim$ 6 km further inland than Svinesund. Waste water from industries and treatment plants flows via the Tista river into the fjord (Brattström et al., 2015; Asteman et al., 2015). Since the 1860s industrial discharge (mainly due to saw mills for the paper industry) of heavy metals (such as mercury and lead etc.) and toxic and oxygen consuming organic matter, has further depleted oxygen in the bottom water of the fjord, which is already low due to the natural phenomena of restricted water exchange due to the shallow sills. These bottom waters are recovering due to reduced industrial discharges and filtering, however soft unconsolidated fine particle sediments which are not deeply buried are still rich in toxic pollutants, with peaks in heavy metals at between 10 and 20 cm depth (Brattström et al., 2015; Asteman et al., 2015). Plans being considered by the Norwegian Coastal Administration for Idefjord include broadening and deepening the Svinesund sills to allow an easier shipping passage. However, it is believed that this may change fjord dynamics and the base of erosion, which could "stir up" unconsolidated sediments and thus heavy metal pollutants.

Single sill simulations where the Svinesund sill has been flattened are used to investigate the potential impact of such proposed changes.

4.3. Methods

This chapter uses the Massachusetts Institute of Technology general circulation model (MITgcm) (Marshall et al., 1997a,c). This model, which was discussed further in section 2.5.3 was chosen here due to its successful use for previous studies of the interaction of tides with bathymetry in both the deep ocean (Legg, 2004a,b; Legg and Huijts, 2006; Legg and Klymak, 2008a) and on the shelf (Vlasenko et al., 2014), as well as with sills in fjords (Xing and Davies, 2006, 2009a, 2010) and to allow direct comparison with previous simulations of a two sill environment motivated by Idefjord (Xing and Davies, 2011). Simulations presented here are non-hydrostatic and use the implicit linear free-surface options in MITgcm. The second-order flux limited advection scheme is used for temperature as in Legg and Klymak (2008b), with a centered second-order scheme for momentum advection with staggered baroclinic Adams-Bashforth time stepping.

This study seeks to understand the physical results of Johnsson et al. (2007) by using real bathymetry from the Idefjord rather than the idealised Gaussian bumps used by Xing and Davies (2011).

As in Xing and Davies (2011), two-dimensional vertical slice (x-z) models were used with periodic boundaries in the x-direction and a ten element sponge region at either end of the domain. Long domains of 60 – 125 km were used to minimize unrealistic interaction of internal waves. Mesh resolutions of $\Delta x = 10$ m in the region of the sills telescoping to $\Delta x = 800$ m at distance, and $\Delta z = 1$ m were used (Berntsen et al., 2008; Xing and Davies, 2011). In addition to using small constant values of the horizontal and vertical diffusivity, $K_h = K_v = 10^{-7}$ m^2 s^{-1}

and viscosity: the horizontal viscosity was set to $A_h = 10^{-1}\mathrm{m}^2\mathrm{s}^{-1}$, whereas in the vertical we use $A_v = 10^{-3}\mathrm{m}^2\mathrm{s}^{-1}$ to maintain numerical stability. These values are low enough to allow an understanding of the processes that lead to mixing but not to quantify the mixing itself.

The model was forced by application of a body force of the form $U_0(2\pi/T)\cos(2\pi t/T)$ (as in Xing and Davies (2011)), where $1/T = 1.40526\mathrm{x}10^{-4}\mathrm{s}^{-1}$ is the dominating semi-diurnal (M2) tidal frequency and U_0 is the peak tidal velocity (see Johnsson et al. (2007) for details). For studies of comparison against Xing and Davies (2011) initially a peak forcing velocity of U_0 =0.1 ms^{-1} was used, following this for comparison against physical data from Johnsson et al. (2007) a range of velocities of U_0 =0.1–0.3 ms^{-1} were used (as shown in table 4.3.1). Finally three peak forcing velocities were used for the remaining studies: the control velocity, U_0 =0.2 ms^{-1} (which gives the best agreement with sill top velocities and density contours compared to the data in Johnsson et al. (2007)), a higher velocity of U_0 =0.5 ms^{-1} and a lower velocity of U_0 =0.08 ms^{-1}. However, unlike Xing and Davies (2011) the Coriolis parameter, f, is set to 0 here, as the width of the fjord ($\sim$ 250 m along the Svinesund to Bjällvarp section) is narrower than the Rossby radius. For systems narrower than the Rossby radius ($c_i/\mid f \mid$), the influence of Coriolis effects are limited when compared with buoyancy effects. In addition, further simulations for Loch Etive demonstrated limited changes in the dynamics between simulations conducted with and without Coriolis effects enabled (Appendix A).

Three different bathymetry profiles were used in this chapter: an idealised Gaussian bathymetry as used by Xing and Davies (2011) to enable a direct comparison, Gaussian bathymetry scaled to Idefjord dimensions (with a domain of 40 m depth and sill top depths of 10 m) and finally real bathymetry from the Idefjord (Figure

4.3.1c). The Gaussian bathymetry of Xing and Davies (2011) is defined as:

$$H(x) = \begin{cases} \frac{70.0}{(1+((x-27512.0/125.0)^4)-100.0}, & x \leq 28262.0 \\ \frac{70.0}{(1+((x-29012.0/125.0)^4)-100.0}, & x > 28262.0. \end{cases} \tag{4.3.1}$$

Parameters for simulations used to compare against Xing and Davies (2009a) and to show the influence of realistic stratification and bathymetry and to fit data to that of physical studies by Johnsson et al. (2007) are shown in table 4.3.1.

In simulations with linear stratification (for direct comparison with Xing and Davies (2011)) $N = 0.01$ s^{-1} is used, while all other simulations use one of two non-linear stratification profiles (Figure 4.3.1b): that of a polynomial fit to density data from Johnsson et al. (2007) (termed "Non-linear 1"):

$$\rho = \begin{cases} 1000.0, & z > -2.0 \\ 1000.0(1 + (1 + (37.764/z + 25.6224)/1000.0)), & z \leq -2.0, \end{cases} \tag{4.3.2}$$

where density is fixed as 1000 for the surface layer due to high freshwater fluxes into the basin. An alternative stratification is also considered with a deeper surface layer (termed "Non-linear 2"), in order to increase the effective depth of the surface layer:

$$\rho = \begin{cases} 1000.0, & z > -10.0 \\ 1000.0(1 + (1 + (37.764/(z+8) + 25.6224)/1000.0)), & z \leq -10.0. \end{cases} \tag{4.3.3}$$

This alternative stratification ensures the system remains a wave-basin throughout the tidal cycle when using the control velocity (as the sill top velocity does not exceed the internal wave velocity).

The influence of a range of inter-sill spacings on the internal wave dynamics was investigated from 1.5 km representing Idefjord to 50 km using the stratification

Table 4.3.1.: Set-up parameters used for the suite of idealised and more realistic set-ups for Idefjord. Set-ups are those based on Xing and Davies (2009a) with gradual changes to parameters to incorporate real world bathymetry and stratification from physical studies by Johnsson et al. (2007), in addition to an exploration of forcing velocities to fit data to that of Johnsson et al. (2007).

Run Name	Maximum depth (m)	Sill depth (m)	Sill spacing (m)	Realistic bathymetry	Forcing amplitude (U_0)	Stratification
Id1	100	30	1500	No	0.1	Linear
Id2	40	10	1500	No	0.1	Linear
Id3	39.3	11.1	1500	Yes	0.1	Linear
Id4	39.3	11.1	1500	Yes	0.1	Non-linear 1
Idv1-5	39.3	11.1	1500	Yes	0.15	Non-linear 1
Idv2	39.3	11.1	1500	Yes	0.2	Non-linear 1
Idv2-5	39.3	11.1	1500	Yes	0.25	Non-linear 1
Idv3	39.3	11.1	1500	Yes	0.3	Non-linear 1

"Non-linear 2". These simulations involved moving the sills further apart and increasing size of the inter-sill basin at a depth of 26.82 m. Four inter-sill spacings are presented here: 1.5 km (as in the real fjord), 5 km, 10 km and 50 km. These allow us to examine the relationship between internal wave wavelength and sill interaction. An additional method of testing the influence of the sills on one another is by reducing the height of the second sill, at Svinesund. This is completed in two stages using the stratification of Johnsson et al. (2007): a simulation with the Svinesund sill flattened at 21 m depth, and a simulation with the Svinesund sill flattened at 28 m depth, which effectively opens the inter-sill basin. In order to examine the effect of the proposed deepening of the second sill on dynamics and possible disturbance of toxic sediments, a further range of Svinesund sill depths (11–28 m depth) were simulated using "Non-linear 1" stratification.

4.3.1. Diagnostics

Tidally averaged change in buoyancy frequency

For comparison with Xing and Davies (2011) the tidally averaged change in buoyancy frequency is used to give an indication of mixing. This parameter: $DN =$

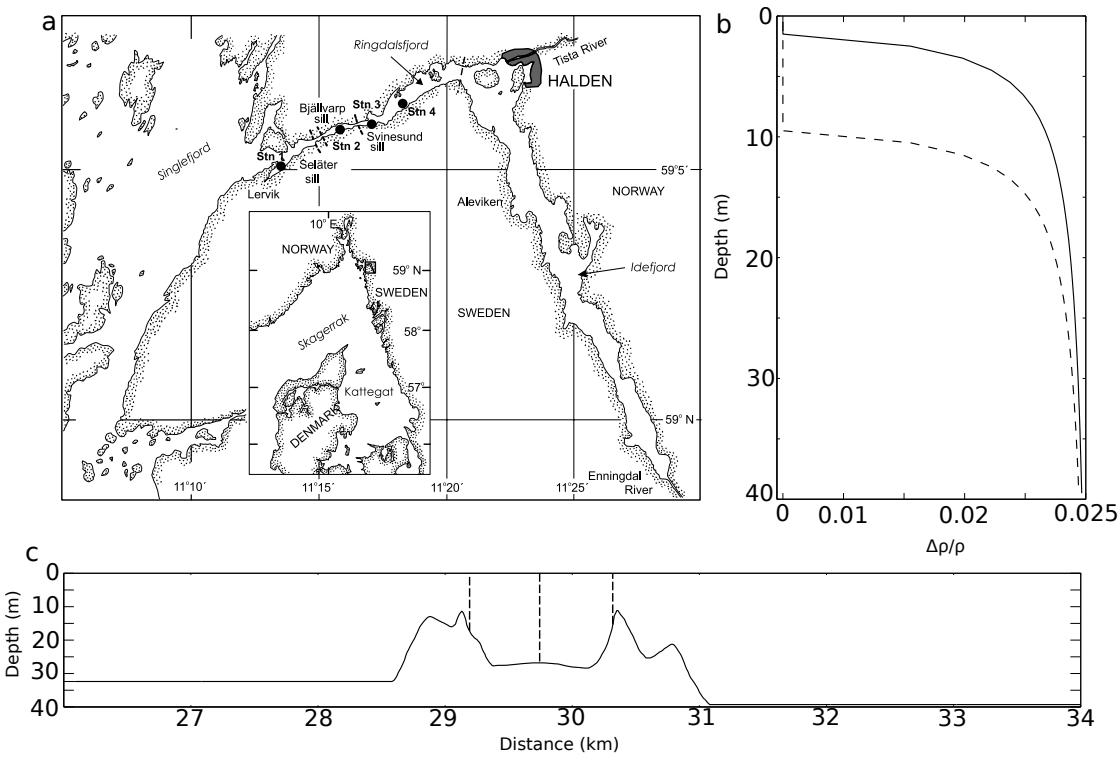


Figure 4.3.1.: Map of Idefjord (Johnsson et al., 2007) showing the location of the Svinesund and Bjällvarp sills in addition to their location with respect to the rest of the system (a). Schematic of the initial density stratification (b) showing "Non-linear 1" Johnsson et al. (2007) (solid line) and "Non-linear 2" (dashed line). Part (c) shows the profile for the real world scenario with a sill separation of 1.5km, in addition to locations of depth averaged vertical velocity PSD calculation are shown in dashed lines.

$\overline{N}^2 - N_0^2$, where $\overline{N}$ is the tidally averaged buoyancy frequency (at a single location) and N_0 is the initial buoyancy frequency. Since the parameter is averaged over the tidal cycle the impact of tidal advection is removed. Mixing is shown by negative values of DN, which implies the erosion of density stratification.

Bed shear stress

Estimates of bed shear stress (τ) in this study are calculated using the following equation:

$$\tau = \rho C_D |\overline{u}| \overline{u}, \tag{4.3.4}$$

where $\rho = 1023$ kgm^{-3} is the density of seawater (Pingree and Griffiths, 1981), $C_D = 0.0025$ is the drag coefficient with a suitable value chosen which is common in many oceanographic models, and $|\overline{u}| = \sqrt{u^2 + w^2}$ is the velocity magnitude (Wells et al., 2007). Estimates of bed shear stress may be used to indicate when sediments are moved by the force of the fluid's velocity. Sediments within Idefjorden are typically silt; medium to coarse silt has a critical bed shear stress of 0.0826 Nm^{-2} (Berenbrock and Tranmer, 2008) which equates to a velocity of 0.18 ms^{-1} using equation 4.3.4. However due to cohesion effects present in clay based sediments, higher bed shear stresses are often required to mobilise sediments, these equate to those of medium sand particles, which have a critical bed shear stress of 0.2 Nm^{-2} (Berenbrock and Tranmer, 2008) or a velocity of 0.28 ms^{-1}. Since average velocity of the water column is used for this diagnostic, it is likely to under-estimate due to the overturning nature of the dynamics (unlike for simpler river systems). In order to mitigate this snapshots of velocity are also shown to analyse the velocity near the bed.

4.4. Comparison with previous idealised studies

In order to validate our model set-up, Id1 results (defined in table 4.3.1) are compared with Xing and Davies (2009a, 2011). These comparisons are made in terms of the horizontal velocities, temperature, tidally averaged change in buoyancy frequency and tidally averaged temperature fields and are shown in figure

4.4.1. The results are shown for the main area of interest focused on the two sills and demonstrate good agreement between the two set-ups. At time 50/8 T an internal wave is generated at the first and second sill, the wavelength of these waves is comparable, as is the wave speeds and temperature distribution. However, our results do not capture the furthest wave from the sill in the wavetrain as shown in Xing and Davies (2009a, 2011). At 52/8 T a jet is shown to extend from the second sill in both results corresponding to a higher velocity and a depression in the temperature contours, however slight differences are observable in the exact distribution of velocities and temperature contours. A jet is also present at 56/8 T extending from the first sill corresponding with a negative velocity and positive velocities are present at mid depth between the sills. Comparison of tidally averaged change in buoyancy frequency and temperature show a region of negative change in buoyancy frequency at mid depth between the sills and the spreading of the temperature contours indicates mixing and the erosion of stratification in both results. The only difference noticeable in our results is a layer at the bottom of the domain showing positive change in buoyancy frequency indicating that temperature contours have been forced together here. These slight differences are observable between our results and those of Xing and Davies (2009a, 2011), may be due to the version of MITgcm and the discretisation methods used (which are not stated in Xing and Davies (2009a, 2011)). These differences are also only quite moderate, despite the lack of inclusion of the Coriolis force, which is included in results from Xing and Davies (2009a, 2011) and therefore shows further support for our neglect of this force in simulations here.

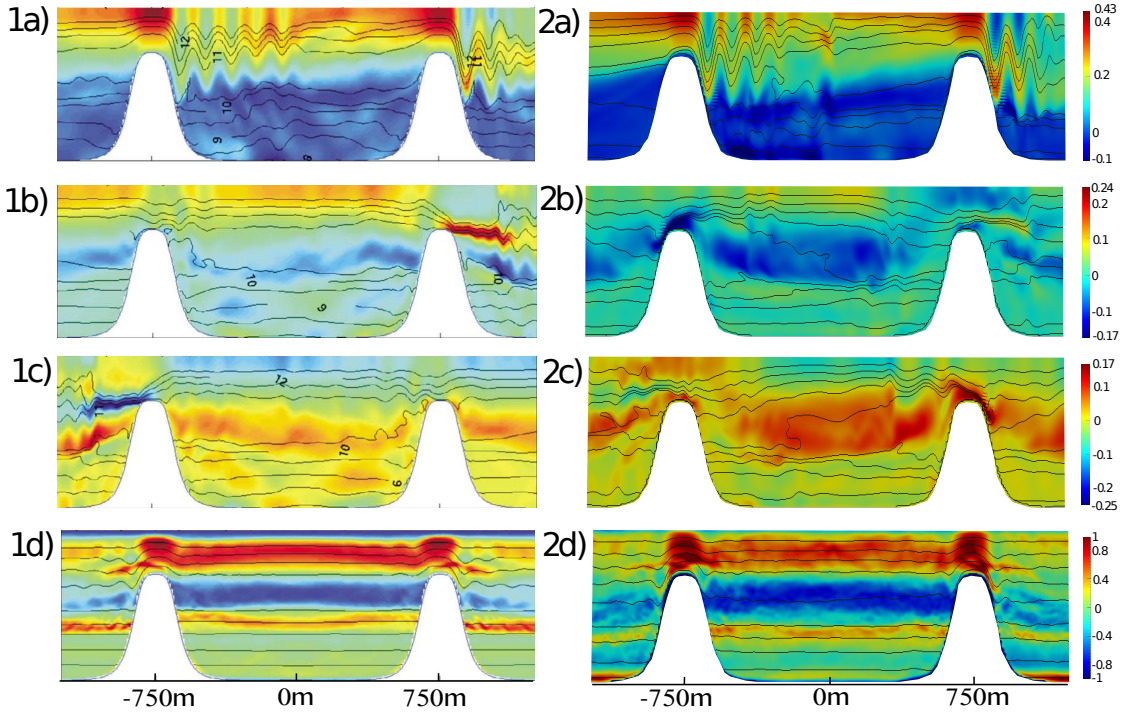


Figure 4.4.1.: Comparison of results of Xing and Davies (2009a) (1) and results of our study also using MITgcm for the same set-up (2). Results are compared for a sub-domain focussed on the two sills, spaced 1500 m apart. Results show the horizontal velocity with contours of temperature (spaced at 0.5^oC) at (a) 50/8 T, (b) 52/8 T and (c) 56/8 T for Id1 (and 56/8 T, 58/8 T and 62/8 T for (Xing and Davies, 2009a, 2011) with this shift in times due to a difference in forcing sinusoidal signal). Additionally (d) compares the tidally averaged change in buoyancy frequency with contours of tidally averaged temperature.

4.5. Effect of Realistic bathymetry and stratification based on Idefjorden

In order to examine the effect of realistic bathymetry and stratification a range of simulations were used which provide intermediate steps between the idealised set-up of Xing and Davies (2009a, 2011) (Id1) and the realistic sill depths, bathymetry and stratification (Id4) based on data for the area (Johnsson et al., 2007). Inter-

mediate simulations involve first reducing the maximum depth and sill depth while maintaining the stratification and idealised bathymetry (Id2), before replacing this with realistic bathymetry (Id3) and finally realistic stratification (Id4). Comparisons of these four simulations in terms of the horizontal velocity and temperature stratification are shown in figures 4.5.1 and 4.5.2.

Simulations show that consistent with theory, reducing the depth of the fjord and reducing the sill top depth while maintaining constant stratification and velocity (from simulation Id1 to Id2) results in more turbulent mixing with turbulent overturning rolls rather than simply internal waves with only limited overturning. This is expected as these changes result in an increase in the relative height of the topography and therefore the Froude number as water is forced through a smaller gap at the sill top for Id2 than Id1, these changes are more significant than the small reduction in topographic slope and steepness.

The greater turbulence for simulation Id2 is shown by a larger area with more negative values of the tidally averaged change in buoyancy frequency indicating that the stratification has been eroded by turbulent mixing between the sills (figure 4.5.2). This erosion of stratification for Id2 has also extended further from the sills than for Id1. The influence of Id3 on changes in the tidally averaged change in buoyancy frequency from Id2 is very small. However, the more complex bathymetry has resulted in some localised changes such as larger areas of negative tidally averaged change in buoyancy frequency above the sill tops and a slightly deeper surface layer of stratification intensification. Also on the left hand side of the seaward sill the area of stratification erosion does not extend as far from the sill which is also likely due to the shallower angle of the sill. The shallower angles of the sides of the sills also result in slightly lower velocities at the sill tops (figure 4.5.1) and slightly less turbulent overturning rolls which do not extend as far from the sills.

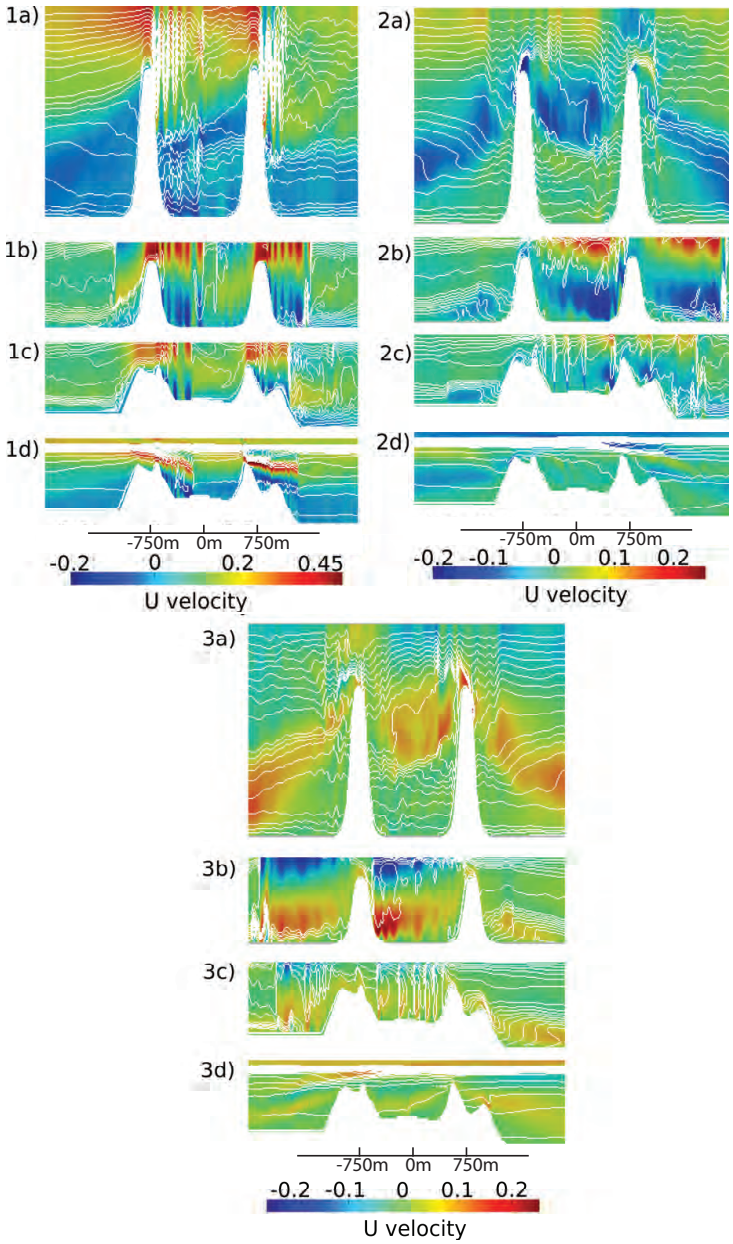


Figure 4.5.1.: Comparison of simulations Id1 (a), Id2 (b), Id3 (c) and Id4 (d), which display increasing complexity of model set-up. Results show the horizontal velocity with contours of temperature (spaced at 0.25^oC for Id1, 0.125^oC for Id2 and Id3 and 0.0005^oC below the thermocline for Id4) at (1) 50/8 T, (2) 52/8 T and (3) 56/8 T.

The changes between Id3 and Id4 are more pronounced due to the effect of a pronounced change in the stratification. Velocity snapshots (figure 4.5.1) show some similarities with Id1, with an internal wave signal clearly visible. However, there are still turbulent overturning rolls present (figure 4.5.1) these have a shorter wavelength and appear more energetic than Id3 at $50/8\ T$. At $52/8\ T$ the dynamics roughly mimic those of Id1 with a depression of the isotherms at the second sill top. Followed at $56/8\ T$ by a depression of the isotherms at the first sill and the initiation of an internal wave. The breaking of the higher frequency internal wave shows areas of intense mixing in figure 4.5.2, where values of the tidally averaged change in buoyancy frequency are much larger requiring a new scale showing intense mixing between the sills and extending away from them on both sides. Values of the tidally averaged change in buoyancy frequency for the layers nearer the surface could be influenced by the numerics trying to preserve the sharp boundary of the thermocline below the fresh surface layer.

4.6. Validation against real world data

In order to validate the model set-up with realistic bathymetry and stratification against physical studies by Johnsson et al. (2007) a series of five forcing velocities were used to compare results with Svinesund sill top velocities, three of these velocities are shown in figure 4.6.1 for ease of comparison. Velocities show a predominantly M_2 barotropic tidal signal, with two dominant peaks and troughs each day. Physical data (Johnsson et al., 2007) shows velocity peaks of ~ 0.75 ms^{-1} and troughs of ~ 1.25 ms^{-1} at 8.97 m depth although these values vary throughout the time of observation. Qualitatively these amplitudes are found to be best represented by a forcing velocity of U_0= 0.2 ms^{-1}. This qualitative agreement is supported by cross-correlation which results in values for $U_0 = 0.1, 0.15, 0.2, 0.25$

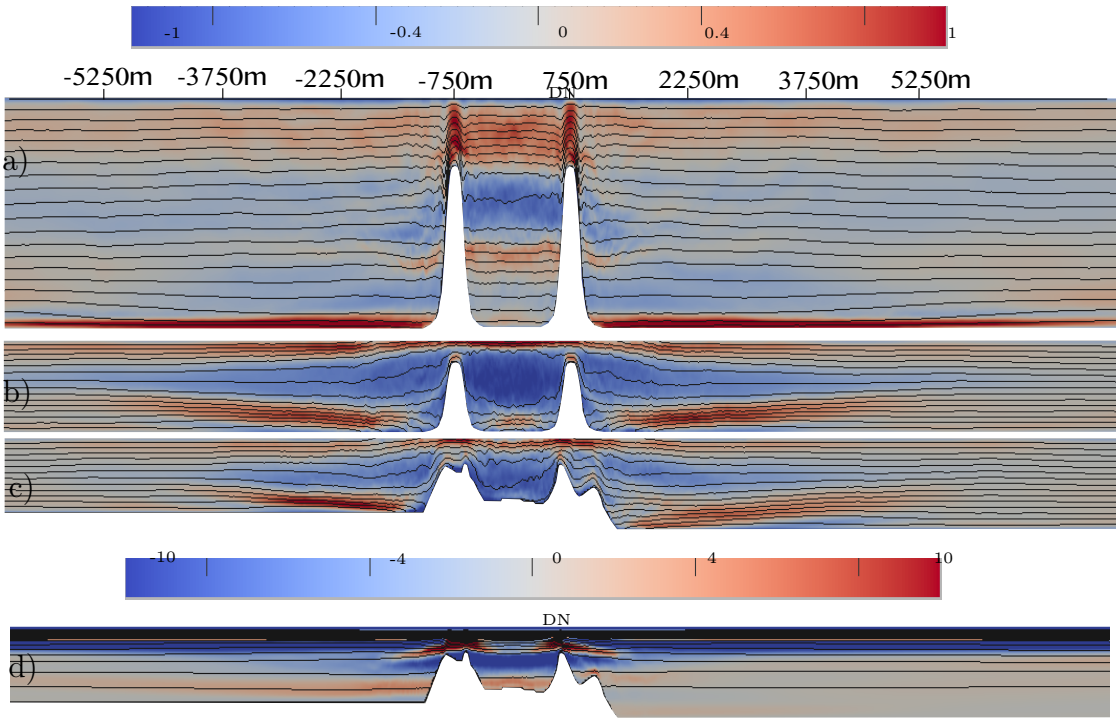


Figure 4.5.2.: Comparison of simulations Id1 (1), Id2 (2), Id3 (3) and Id4 (4) showing the tidally averaged change in buoyancy frequency with contours of tidally averaged temperature (spaced at 0.5^oC for Id1, 0.25^oC for Id2 and Id3 and 0.0005^oC below the thermocline for Id4) at $50/8\ T$.

and 0.3 ms^{-1} of 0.7479, 0.7283, 0.8181, 0.8169 and 0.8127 respectively, demonstrating greatest agreement at U_0=0.2 ms^{-1}. Agreement is not perfect due to the nature of comparing data with a real world system with external influences but also because data is not directly available at the depths of the physical survey from model data results. Modelled data must therefore be extrapolated for comparison with Johnsson et al. (2007).

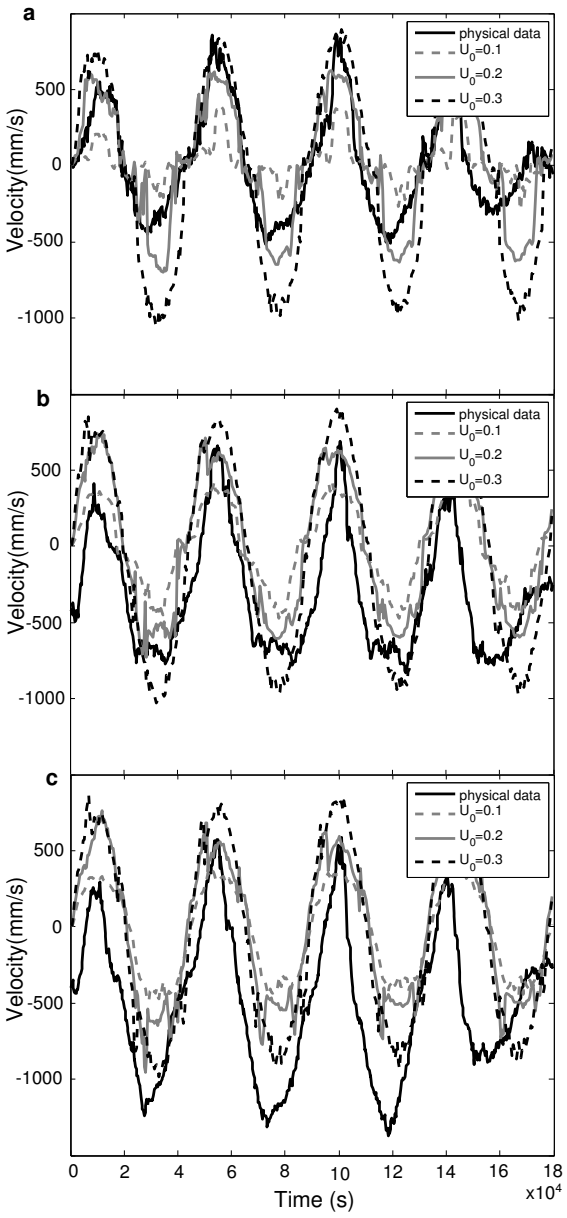


Figure 4.6.1.: Comparison of field data acquired by Johnsson et al. (2007) with simulation data for sill top velocities at Svinesund, the second sill in numerical studies conducted here. Results for E-W component of velocity are compared with u velocities in numerical studies at three depths (a) 3.97 m, (b) 5.97 m and (c) 8.97 m. Three different forcing velocities are shown $U_0 = 0.1$, 0.2, and 0.3 ms^{-1}

4.7. Tidal dynamics between two closely spaced sills in a narrow fjord: the development of a standing internal wave

4.7.1. Realistic stratification and bathymetry for Idefjord

Simulating the conditions in the fjord given in Johnsson et al. (2007) results in dynamics within the sill basin which are not entirely consistent with a wave-basin. Although initially an internal wave is set up, as the tidal velocity increases during the tidal cycle the dynamics change (fig. 4.7.1b). Once $T/24$ is reached a hydraulic jump develops and is followed by a high amplitude non-linear internal wave.

Power spectral density of depth averaged vertical velocity (PSD) (fig. 4.7.2b) at these locations shows a change at the edges of the basin (i.e. at the slope of the sills locations shown in Figure 4.3.1c) compared with the centre of the basin. At the slopes the signal is similar either side with a largest peak ($\sim 10^4$ $(\mathrm{m/s})^2$) in amplitude for M_2 and peaks clearly present for the M_2 frequency and its higher harmonics as there is still a propagating tide here. However, in the middle of the basin the signal has highest amplitude of power jointly for the M_4 and M_8 frequencies ($\sim 10^1$ $(\mathrm{m/s})^2$) and with peaks in power also for M_2. The weaker M_2 signal ($\sim 10^1$ $(\mathrm{m/s})^2$) is present due to depth averaging and is due to the contribution of the surface signal. The reduction in overall power in the signal for the middle of the basin demonstrates the losses which have resulted from the more turbulent jet-basin dynamics present and the higher power in higher frequencies shows the result of the hydraulic jump due to the break down of the M_2 tide (fig. 4.7.1b).

The predominant dynamics observed in these simulations are those of a jet-

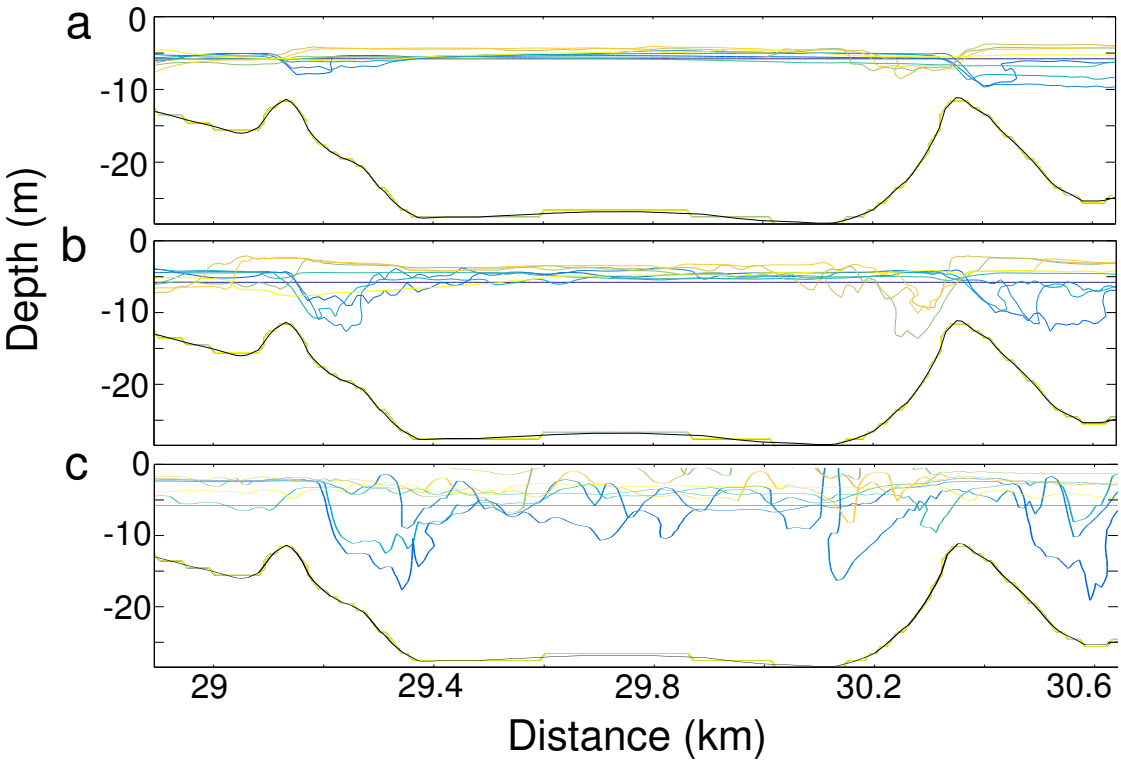


Figure 4.7.1.: Contours of 0.023 kg m^{-3} density at four temporal locations at intervals of $T/16$ between 0 and $T/2$ inclusive, shown in the parula colormap (where $T/8$ is dark blue and $T/2$ is yellow). Results shown are for the stratification based on Johnsson et al. (2007) for the three velocities: the lower velocity U_0 =0.08 ms^{-1} (a), the control velocity U_0 =0.2 ms^{-1} (b) and the higher velocity U_0 =0.5 ms^{-1} (c). Bathymetry is shown in black in each plot.

basin system rather than a wave-basin as discussed in Johnsson et al. (2007). This is due to the stratification present at the time of measurement, which may not be representative of average conditions in the fjord. Johnsson et al. (2007) assumed a surface layer that extended to the sill tops; however, as the system has a large contribution of freshwater at the time of the observations the system is effectively comprised of three layers. If modelled as a two layer system the surface layer should probably be much shallower than the sill depth and result in a predominantly super-critical system.

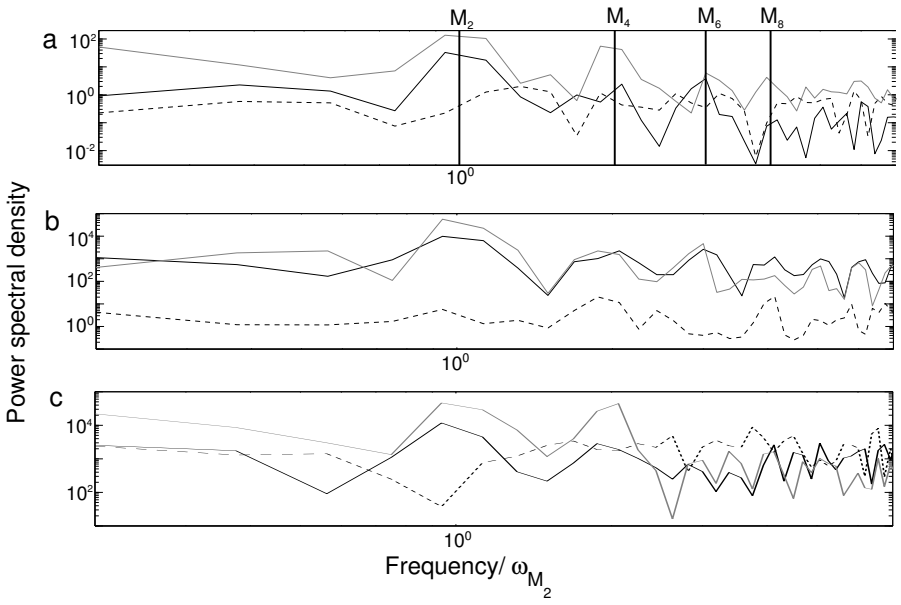


Figure 4.7.2.: The power spectral density for depth averaged vertical velocity normalised by the M_2 frequency for the "Non-linear 1" stratification for the three velocities: the lower velocity U_0 =0.08 ms^{-1} (a), the control velocity U_0 =0.2 ms^{-1} (b) and the higher velocity U_0 =0.5 ms^{-1} (c). Results are shown for three spatial locations in the basin between the two sills: 60 m from the first sill (black line), the midpoint between the sills (black dashed line) and 30 m from the second sill (grey). Tidal frequencies M_2, M_4, M_6 and M_8 have been marked in panel (a). Results are shown with different scales due to the different energetics of the systems in order that their power spectrums be presented in an easily interpretable form.

To further elucidate the basin dynamics we turn our attention to the simulations with an increased and decreased sill top velocity. Simulations using the higher velocity demonstrate that the internal hydraulic jump is formed earlier in the tidal cycle ($T/96$ as opposed to $T/32$) and almost no semi-diurnal waves exist (Fig. 4.7.1c). This is because the flow is super-critical and as the hydraulic jump breaks down we again see high amplitude, high-frequency internal waves being generated (fig. 4.7.1c). For this case, however, the amplitudes of the internal waves are much larger than in the control (up to $\sim$ 5–10 m as opposed to $\sim$2–3 m)

and they start to resemble Kelvin-Helmholtz billows (fig. 4.7.1c). This energetic system sees energy pumped into higher tidal frequencies with the highest power in the middle of the domain at the M_8 frequency ($\sim$ 20 $(\mathrm{m/s})^2$) and no peak in power at the M_2 frequency any longer as seen at the domain edges.

By using a lower forcing velocity the system remains a sub-critical wave-basin until $T/12$, longer than the control case. After this point the velocity becomes super-critical and a hydraulic jump develops, just as in the other simulations (fig. 4.7.1a). This hydraulic jump, however, remains fairly stable and only generates high frequency internal waves during a very small window ($3T/8$ – $5T/8$) of the tidal cycle (Fig. 4.7.1a). Consequently, the loss in power at the M_2 frequency (which is still the peak power at $\sim 10^3$ $(\mathrm{m/s})^2$) in the center of the basin is small (fig. 4.7.2a). Subsequent peaks at higher frequencies are similar to those observed at the sides of the basin.

4.7.2. Increasing the depth of the Svinesund sill top

Reducing the height of the Svinesund sill top has the effect of reducing its ability to convert energy and to separate water masses. The density contour plots (fig. 4.7.3) show that at a sill height of 21 m there is a significant change in the internal wave dynamics. There is only a very small hydraulic jump (of amplitude $\sim$ 2 m rather than $\sim$ 5–10 m), but on the right hand side of the sill a hydraulic jump still develops in a similar manner to the control simulations. When the Svinesund sill is removed the spectra at the location where the sill used to be is very similar to that at the (previous) center of the basin (even when sill height is only reduced to 21 m (Figure 4.7.4b)). That is hardly surprising, but peak power at the center of the basin also changes from that of independent peaks at M_4 and M_8 to a broader peak between M_4 and M_6 (10^1 $(\mathrm{m/s})^2$) and a loss of power at the M_8 frequency (of almost an order of magnitude). This is due to the lack of superposition of

internal waves formed at the two sills, and suggests that the main source of power at the M_8 frequency for the two full height sill simulations is wave-wave interaction. Simulations with a reduced height of the Svinesund sill also show a slight increase in the power of the M_2 frequency, both at the basin centre and at the sill itself (to 10^1 $(\mathrm{m/s})^2$). This is because an M_2 internal wave that forms at the slope and then remains on the right hand side of the Svinesund sill down to the depth of 40 m even when the sill is flattened to the inter-sill basin depth (shown at 30.7–31 km in the domain in fig. 4.3.1, on the far right of the Svinesund sill (at 30.6 km in Fig. 4.7.4c). The combination of this M_2 internal wave and the M_4 internal wave produced at the Bjällvarp sill could be the cause of the basin centre peaks at the M_6 frequency.

4.7.3. Increasing the depth of the surface layer

Increasing the depth of the surface layer by changing the stratification increases the speed of the internal wave that develops on the fluid interface and means that the system may remain a sub-critical wave-basin for a larger portion of the tidal cycle. Using the control velocity this results in a sub-critical velocity for the entire tidal cycle, the resulting wave-basin dynamics are those of a standing internal wave (or seiche) between the sills (fig. 4.7.5a). For this simulation the node is clearly present in the basin center and the standing internal wave therefore has a wavelength of 3 km. The power spectral density in the middle of the basin does not differ greatly from the edges of the basin for the high frequency components and the signal retains its largest peak at M_2 frequency (10^2 $(\mathrm{m/s})^2$), however the amplitude of the M_4 frequency is very similar.

Increasing the tidal forcing also initiates a seiche like behaviour (fig. 4.7.5b), but in this case it does not occur for the entire tidal cycle and at $5T/24$ the tidal velocities become super-critical a hydraulic jump forms, which again breaks

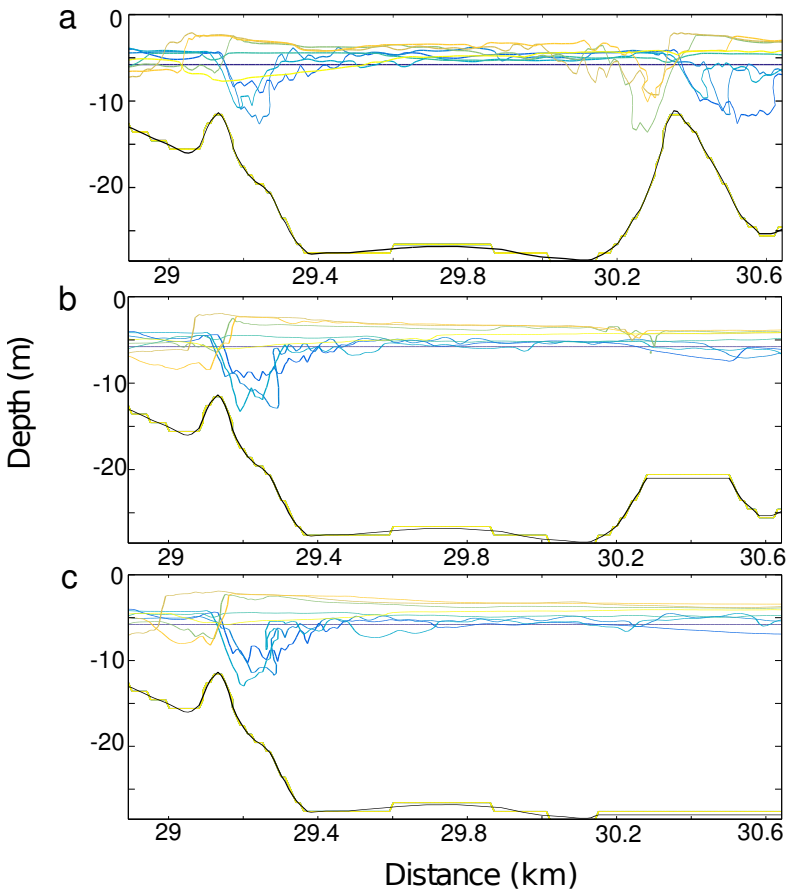


Figure 4.7.3.: Density contours of 0.021 kg m^{-3} shown for every $T/16$. Results shown are for the stratification based on Johnsson et al. (2007) for the three depths of the Bjällvarp sill (the real depth (a), 21 m (b) and 28 m (c)). Bathymetry profiles are shown in black for each scenario.

down into high frequency internal waves. These high frequency components then propagate into the interior until the tidal velocity is significantly reduced and are then initiated again once the tidal velocity increases. Due to periods with super-critical flow the signal has lost more power in the centre of the basin than the control velocity case. In addition the highest power values are at M_6 although peaks are still present at M_2 and M_4 due to the seiche present for part of the tidal cycle.

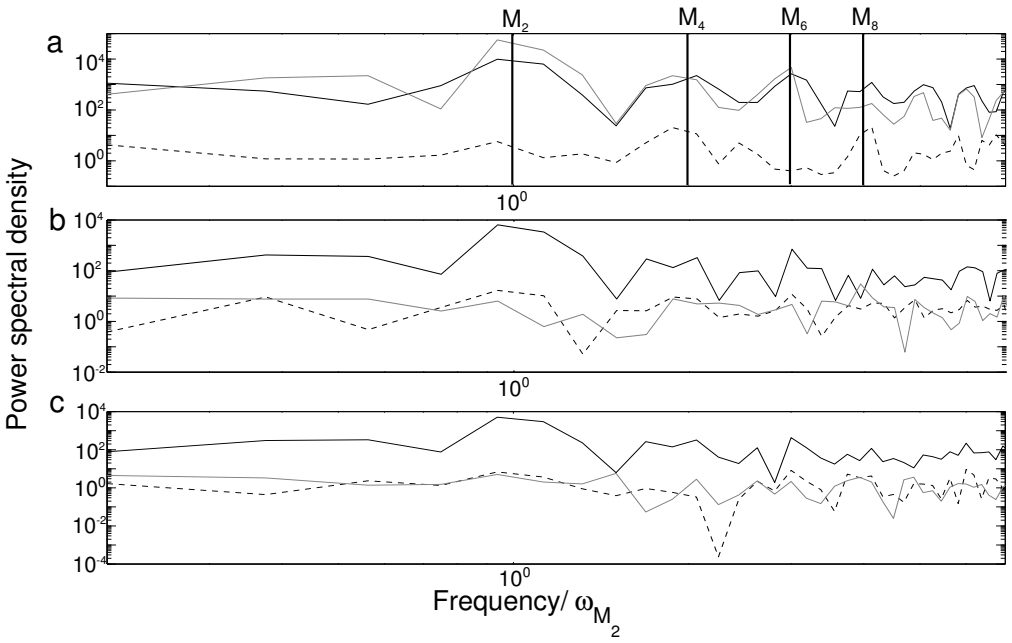


Figure 4.7.4.: The power spectral density for depth averaged vertical velocity normalised by the M_2 frequency for the stratification based on Johnsson et al. (2007) for the three depths of the Bjällvarp sill (the real depth (a), 21m (b) and 28 m (c)). Results are shown for three spatial locations in the basin between the two sills: 60 m for the first sill (black line), the mid-point between the sills (black dashed line) and 30 m from the second sill (grey). Tidal frequencies M_2, M_4, M_6 and M_8 have been marked in panel (a).

4.7.4. Varied sill spacing and a deeper surface layer

Control velocity

The previous section clearly shows that the two sills are interacting and that the product of this interaction is a standing internal wave in the basin. Therefore, to test the interaction of these sills we will now increase the separation between them. This results in the wave-basin behaviour beginning to break down at 5 km separation, which is seen by high frequency, low amplitude internal waves also propagating in the basin in addition to the seiche during peak flood tides (fig.

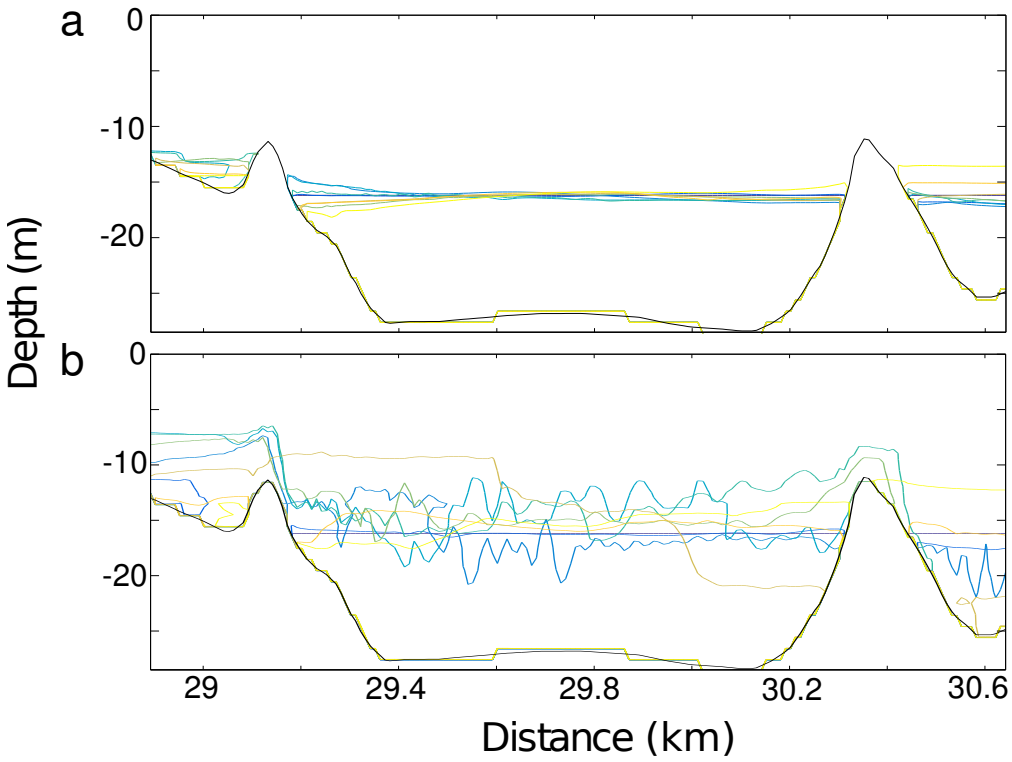


Figure 4.7.5.: Density contours of 0.021 kg m^{-3} shown for every $T/16$. Results shown are for the "Non-linear 1" stratification for the control velocity (a) and the higher velocity U_0 =0.5 ms^{-1} (b). Results show the development of a standing internal wave at both velocities. Bathymetries are shown in black for each scenario.

4.7.7b). In general, however, the basic seiche behaviour is retained with a node still present and a wave pattern of a propagating standing internal wave (Figure 4.7.7b). In terms of the power spectral density M_2 has been replaced with a peak at M_4 (10^{-1} (m/s)^2) as the highest power for the system (fig. 4.7.8b) at the basin centre. The barotropic M_2 frequency is still represented as a peak in power but this is the smallest peak of all of the barotropic frequencies (10^{-3} (m/s)^2) and shows energy has been started to be pumped into higher frequencies. Using a 5 km separation the sills are only slightly closer together than the ∼5.5 km internal wave wavelength estimated using equation 4.2.1, which may explain the slight

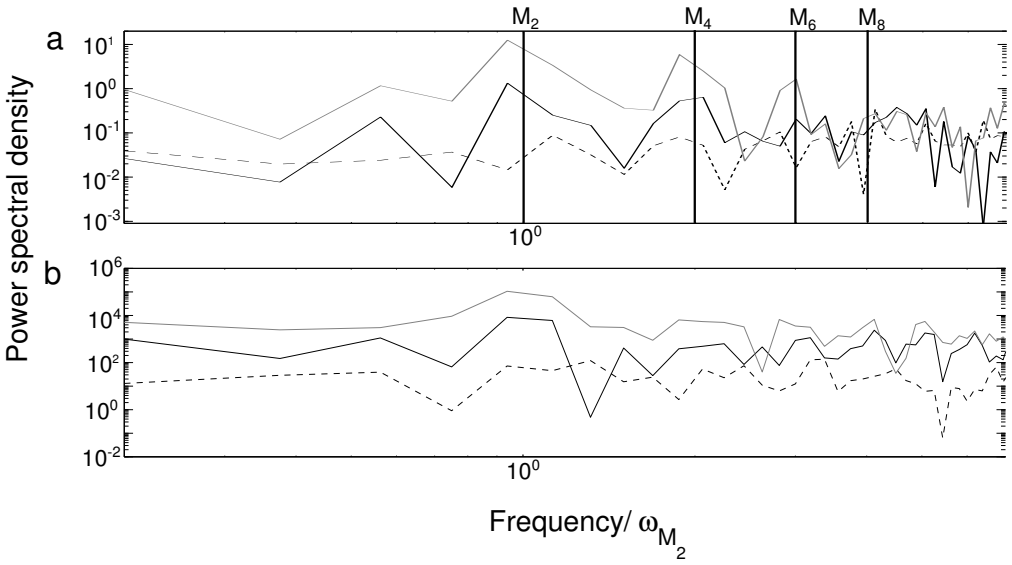


Figure 4.7.6.: The power spectral density for depth averaged vertical velocity normalised by the M_2 frequency for the "Non-linear 1" stratification for the control velocity (a) and the higher velocity U_0 =0.5 ms^{-1} (b). Results are shown for three spatial locations in the basin between the two sills: 60 m for the first sill (black line), the mid-point between the sills (black dashed line) and 30 m from the second sill (grey). Tidal frequencies M_2, M_4, M_6 and M_8 have been marked in panel (a).

breakdown in internal wave dynamics.

As the inter-sill distance is increased to 10 km the interaction between the sills is further reduced and internal waves can be clearly seen to be generated at both sills and allowed to freely propagate between the sills, with the distance too great to enable superposition to produce an effective standing wave in the basin (fig. 4.7.7c). Later in the simulations (after two tidal cycles), at peak flow, high frequency low amplitude internal waves which rapidly dissipate are formed. Although the wave field can no longer be described by a standing wave, a node can now be found about a third of the way into the domain from the first sill (at Bjällvarp). The spectrum at the centre of the domain shows the same features as for a 5 km separation: energy is being pumped into higher frequency components.

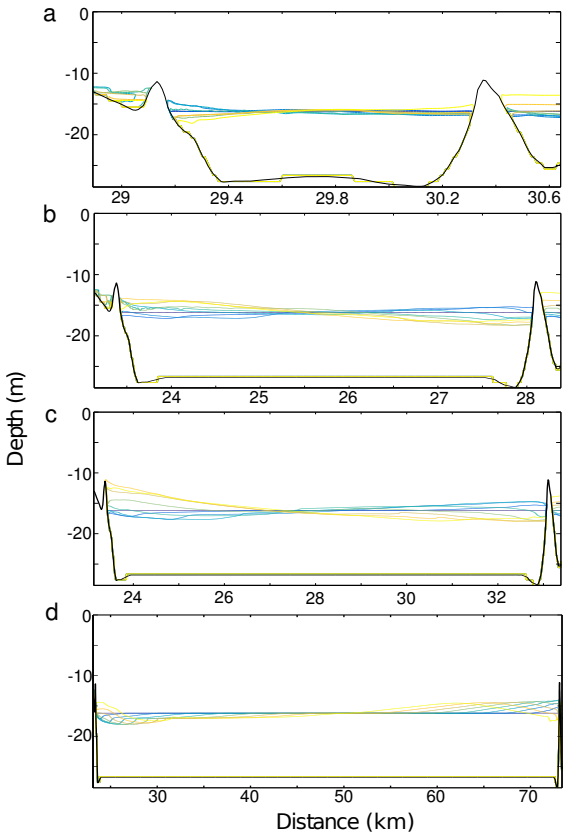


Figure 4.7.7.: Density contours of 0.021 kg m^{-3} shown at every $T/16$. Results shown are for the "Non-linear 1" stratification for the control velocity with a range of inter-sill separations: 1.5 km (a), 5 km (b), 10 km (c) and 50 km (d). Bathymetry profiles are shown in black for each scenario. Each result is plotted on a different x-axis scale in order to show the full inter-sill basins

For this separation the highest power signal is for the M_6 frequency ($\sim$0.2 $(m/s)^2$), closely followed by M_4 ($\sim$0.1 $(m/s)^2$) (Fig. 4.7.8c). There is also a slightly higher power than for a 5 km separation for all higher frequency components of the signal.

For the largest inter-sill separation investigated here (50 km) the standing internal wave is no longer present and the internal waves generated at the two sills interact less and can be observed as separate waves (fig. 4.7.7d). However, despite this, later in the simulation (from $2T$ onward) there is constructive superposition generating larger amplitude waves about a third of the way into the basin from

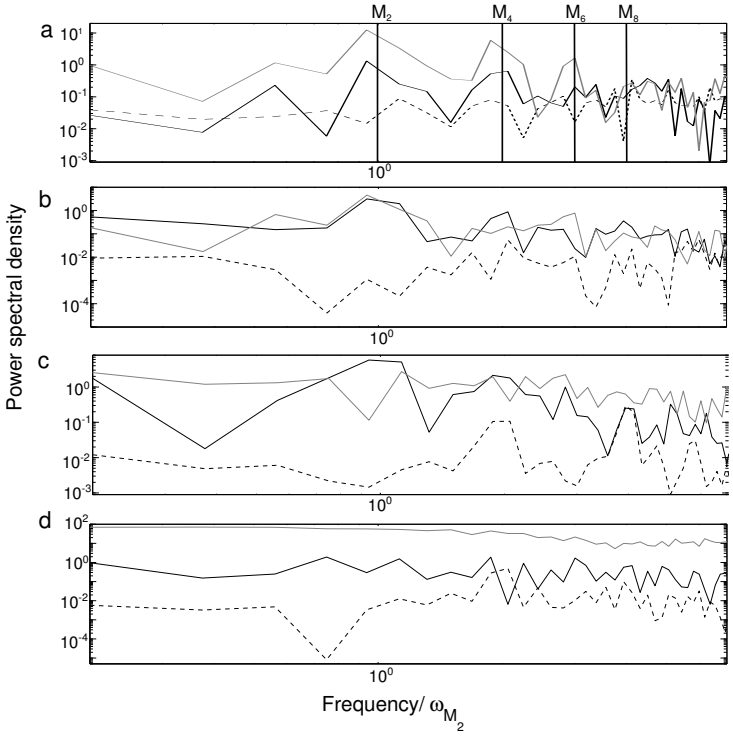


Figure 4.7.8.: The power spectral density depth averaged for vertical velocity normalised by the M_2 frequency for the "Non-linear 1" stratification for the control velocity with a range of inter-sill separations: 1.5 km (a), 5 km (b), 10 km (c) and 50 km (d). Results are shown for three spatial locations in the basin between the two sills: 60 m for the first sill (black line), the mid-point between the sills (black dashed line) and 30 m from the second sill (grey). Tidal frequencies M_2, M_4, M_6 and M_8 have been marked in panel (a).

both sills. Again there is a small signal from high frequency signals. However, the multiple higher harmonics found in some of the previous simulations are no longer present, and the peak power is found at the M_4 frequency ($\sim$0.4 (m/s)^2) (Fig. 4.7.8d). In addition, spectral peaks are present at non tidal frequencies (which are not present for smaller separations), and these must come from the interactions of the internal waves which are not synchronised.

Higher velocity

By forcing the model with a higher velocity we can test the influence of inter-sill distance for a super-critical system. As the inter-sill separation is increased the seiche that was present for part of the tidal cycle when sub-critical velocities were used is reduced (Fig. 4.7.9). The hydraulic jump is now fully developed, and when it breaks down it generates more overturning and high frequency internal waves which are present for more than half of the tidal cycle. The seiche now has a node located approximately one third of the way into the basin from the first sill (fig. 4.7.9b). The main peak in power at the centre of the basin is also shifted to a slightly higher frequency from the M_2 peak for the 1.5 km separation. In addition the power associated with higher frequencies such as M_8 and above is increased (to $\sim$10 $(\mathrm{m/s})^2$) (fig. 4.7.10b). By increasing the separation to 10 km the seiche has almost entirely disappeared and the interaction between the sills is reduced (Fig. 4.7.9c). However, as for the control velocity simulation there is still some interaction of the internal waves formed. In the control simulation this resulted in increased amplitudes, whereas here we see a steepening of the wave and overturning causing packets of high amplitude internal waves which propagate along the lower amplitude main component of the internal wave. The resultant power shows a further shift of the peak to the M_8 frequency (10^1 $(\mathrm{m/s})^2$), further demonstrating the loss of seiche features (fig. 4.7.10c).

Using the largest separation distance tested here we show that similar to a 10 km separation, it results in the generation of high frequency internal waves and even less interaction between the sills (Fig. 4.7.9d). Again we see the results of superposition of the internal waves but due to the greater distance between the sills there is some attenuation, and therefore the resulting steepening and higher frequency internal wave production is reduced. This reduction of high frequency internal wave production is also shown by the power spectrum at the centre of

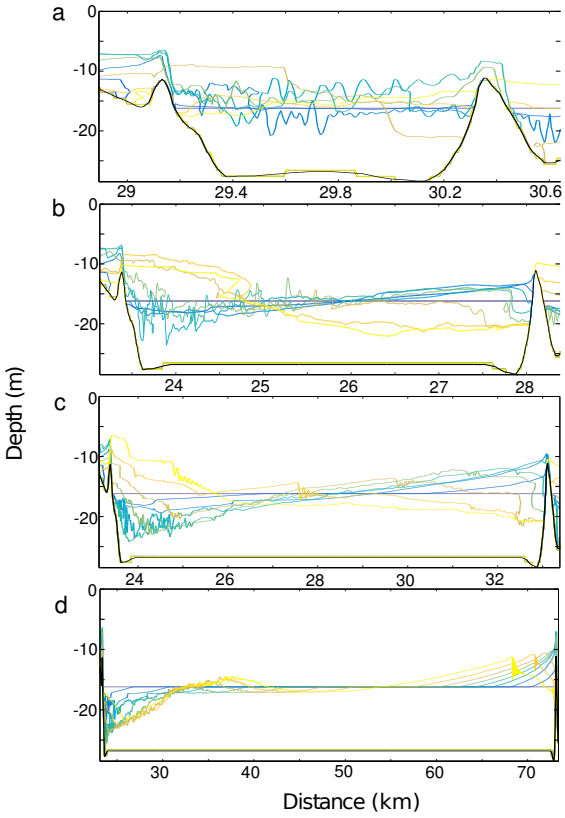


Figure 4.7.9.: Density contours of 0.021 kg m^{-3} (black) shown for every $T/16$. Results shown are for the "Non-linear 1" stratification for the higher velocity (U_0 =0.5 ms^{-1}) with a range of inter-sill separations: 1.5 km (a), 5 km (b), 10 km (c) and 50 km (d). Bathymetries are shown in black for each scenario.

the basin where the M_4 frequency dominates above higher frequencies again (∼40 $(m/s)^2$), however the power is not much higher than frequencies for M_{12} and higher (∼10–20 $(m/s)^2$) (fig. 4.7.10d).

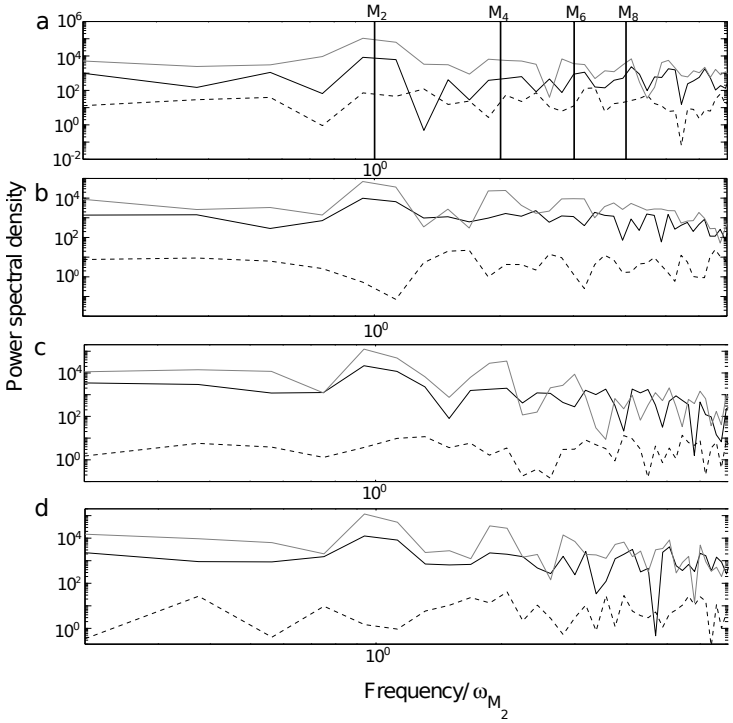


Figure 4.7.10.: The power spectral density for depth averaged vertical velocity normalised by the M_2 frequency for the "Non-linear 1" stratification for the higher velocity (U_0 =0.5 ms^{-1}) with a range of inter-sill separations: 1.5 km (a), 5 km (b), 10 km (c) and 50 km (d). Results are shown for three spatial locations in the basin between the two sills: 60 m for the first sill (black line), the mid-point between the sills (black dashed line) and 30 m from the second sill (grey). Tidal frequencies M_2, M_4, M_6 and M_8 have been marked in panel (a).

4.8. Svinesund sill height change and bed shear stress

In order to examine the effect of deepening the Svinesund sill on the disturbance of toxic loosely buried sediment a suite of simulations are used to analyse the influence on bed shear stress. By comparing the maximum bed shear stress (using the approximation given by equation 4.3.4) as the height of the Svinesund sill is reduced, it can be argued whether movement of the sediments is likely. Results

show (figure 4.8.1a) that initially as a small amount of the sill top is removed the approximate maximum bed shear stress reduces, this may be due to a reduction in the sill top velocities at this sill and a simplification of the dynamics. However, when the sill height is reduced further the bed shear stress begins to increase again, as less energy is removed from the system at the Svinesund sill and in addition higher velocities are present at the top and down the left hand side of Bjällvarp (figure 4.8.2). This increases the bed shear stress stress in these areas (figure 4.8.1b). By analysing the bed shear stress at Bjällvarp a linear increase is found as the Svinesund sill depth is increased (figure 4.8.1). When most of the Svinesund sill is removed and the sill top is at 28 m depth, the maximum bed shear stress across the profile is actually slightly increased, indicating that the removal of Svinesund could have significant negative effects on the environment in Idefjorden, depending on the exact location of the erosion caused and the most toxic sediments. The maximum bed shear stress of 1.25 Nm^{-2} for the sill top depth of 28 m is significantly larger than that required for the movement of the sediments found in Idefjorden.

4.9. Discussion

Comparisons of the simulations conducted here using the idealised two sill set-up used by Xing and Davies (2011) with published results reveals very little differences despite the Coriolis force being included in the published simulations. The results replicated the same general dynamics with only slight differences, which include the furthest wave in the wavetrain at 50 / 8 T not being present and minor differences in the velocity and temperature structures while the overall dynamics are consistent. These results provide confidence in the set-up used for further simulations.

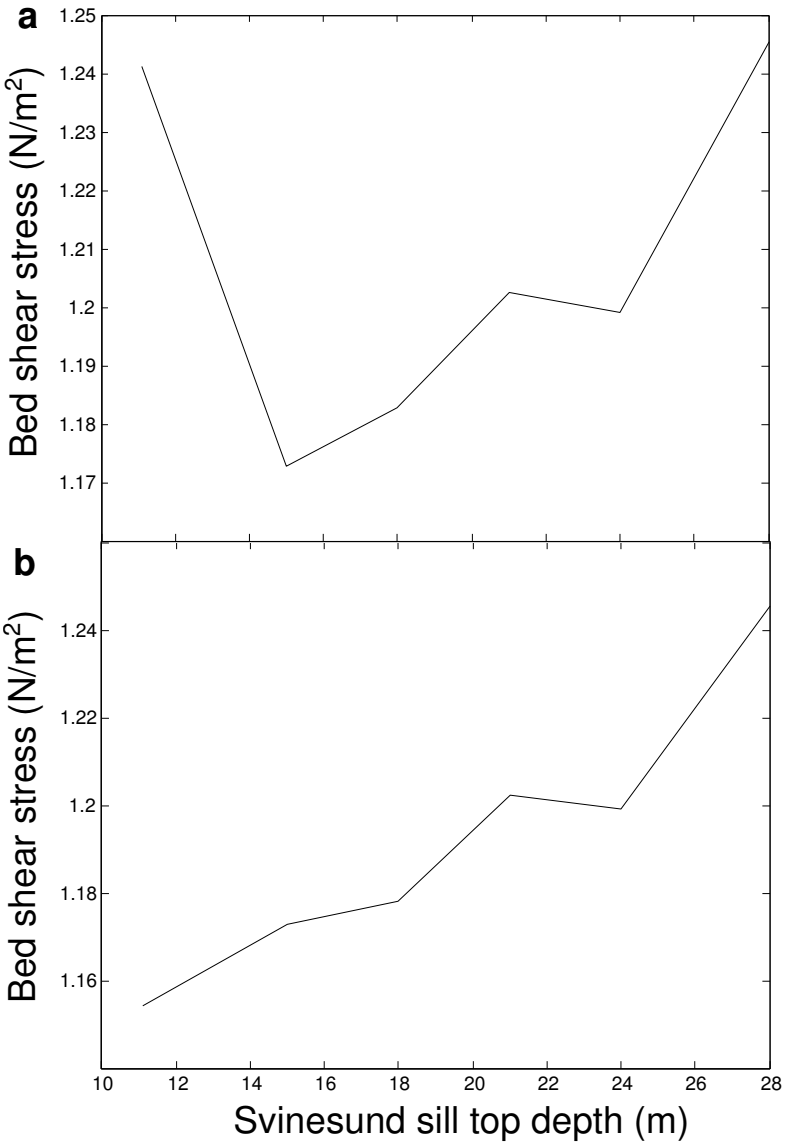


Figure 4.8.1.: Approximate bed shear stress plotted against Svinesund sill depth. Results are shown for the temporal and spatial maximum (a) and the temporal maximum bed shear stress at the Bjällvarp sill top (b) versus sill depth.

Changes in the dynamics are presented as the simulation is adjusted from the idealised simulation of Xing and Davies (2011) to that of a more realistic set-up based on survey results from Johnsson et al. (2007). The dynamics change

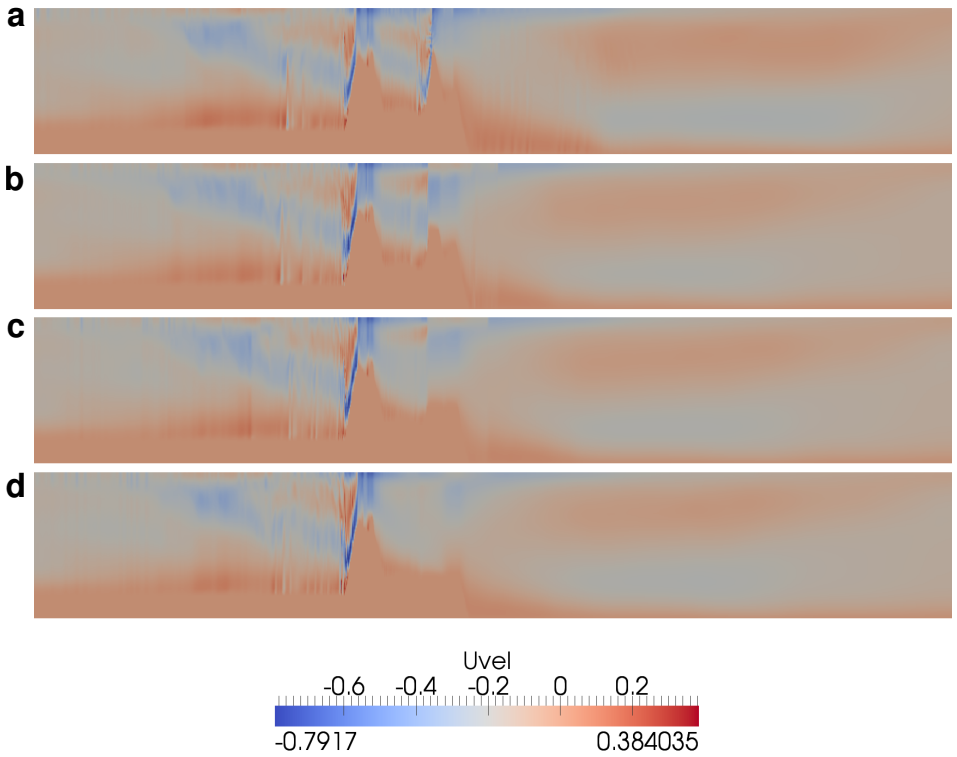


Figure 4.8.2.: Snapshots of the horizontal velocity field for the region of the sills at 17/6 T. Results are shown for (a)Id4 in addition to results using a Svinseund sill depth of 18 m (b), 24 m (c) and 28 m.

from those presented by Xing and Davies (2011) through the presence of higher frequency internal waves which break resulting in areas of more intense mixing between the sills, represented by more negative values of the tidally averaged change in buoyancy frequency. While intermediate simulations demonstrated the increase in mixing and turbulent dynamics which results from reducing the sill separation along with the sill top depth, as expected from theory. In addition to realistic bathymetry demonstrating that changes in the local slope angle may affect the distribution of mixing, with areas of gentle slope on the outside of the basin meaning that the plume of water with eroded stratification extending from the sills is reduced in size. These localised reductions in sill slope also reduce sill

top velocities and cause slightly less turbulent dynamics in general.

Validation of MITgcm model results using various forcing velocities against sill top velocities measured by Johnsson et al. (2007) shows that dynamics are largely M_2. The signal amplitudes of the real world data are largely well represented by the model results. However, differences are present especially at 8.97 m depth. These are due to variations in the physical signal caused by the nature of a real world system. The forcing velocity of $U_0 = 0.2$ m s^{-1} provides the best agreement with physical data qualitatively and quantitatively.

Simulations conducted using the control velocity with realistic stratification and bathymetry based on observations from the Idefjord result in disagreement with assumptions made by Johnsson et al. (2007). This is due to the hydraulic jumps observed in the numerical results and subsequent high amplitude non-linear internal waves which develop at the sill tops when the tidal velocities exceed those of the internal wave. These jet-basin dynamics break down the M_2 tide into high frequency waves. Internal wave velocities are exceeded at the sill tops resulting in a super-critical system in this study due to the speed of the internal wave; this speed is lower than that expected by Johnsson et al. (2007), who assumed a much deeper surface layer than is the case during the studied conditions due to the large volumes of surface fresh water. Super-critical jet-basin dynamics are also present for simulations with an increased and decreased sill top velocity, despite variation in the size and duration of the hydraulic jump due to the forcing velocity. Even using the lower forcing velocity gives results with sill top velocities greater than those of the internal wave speed.

The dynamics of the inter-sill basin were found to be fairly sensitive to the height of the Svinesund sill, with reductions in height reducing its ability to convert energy and separate water masses. Despite the obvious changes of the spectra at the Svinesund sill, which now mimics those of the previous basin centre, more

unexpected shifts in power spectral density at the basin centre were found with the addition of a peak at the M_6 frequency. This is due to the combination of an M_2 internal wave formed at the remaining slope of Svinesund sill with the M_4 internal wave formed at Bjällvarp sill.

By increasing the depth of the surface layer (from the "Non-linear 1" to the "Non-linear 2" stratification) the dynamics of a sub-critical wave-basin system with the control velocity were revealed. For this system a standing internal wave (or seiche) develops between the sills due to the superposition of the internal waves formed at each sill. While the higher forcing velocity resulted in seiche dynamics for only sub-critical parts of the tidal cycle and hydraulic jumps and high frequency internal waves for the remaining cycle.

Simulations conducted where the spacing of the sills was increased revealed for the sub-critical system that as spacing was increased above the internal wave wavelength that the seiche behaviour breaks down. This results in internal waves being generated separately at each sill and propagating away, these waves still interact however and this results in constructive superposition increasing wave amplitudes. In addition energy is pumped into higher frequency components.

The influence of sill spacing for a super-critical system is a more developed hydraulic jump, however partial seiche behaviour is still present at separation distances below the internal wave wavelength for sub-critical parts of the tidal cycle. At greater sill spacing the internal waves still interact resulting in steepening of the waves and overturning causing packets of high amplitude internal waves, which propagate along the lower amplitude main component of the internal wave. This break down in internal wave dynamics is also seen by a pumping of energy into higher frequency components as for the sub-critical case.

Reducing the height of the Svinesund sill in several stages results in changes in the velocity field. As Svinesund sill depth is increased the maximum bed shear

stress at the Bjällvarp sill increases due to an increase in the velocity on the slopes of the sill, this increased velocity occurs due to the reduction in the amount of energy removed from the system through interactions with the Svinesund sill. However, the spatial and temporal maximum bed shear stress initially decreases as the dynamics become simpler and velocities on the flanks of the Svinesund sill reduce, before increasing as the effect of increased sill flank velocities at the Bjällvarp sill dominate the signal. Results of maximum bed shear stress at the Bjällvarp sill are more significant as spatial temporal maximums may be influenced by spurious values and may not be representative of large scale processes. Depending on the relationship between locations of increases in bed shear stress and the most toxic sediments there may be a significant increase in the mobilisation of toxic sediments.

4.10. Conclusions

Using simulations with progressively less idealised, more realistic, set-ups than Xing and Davies (2009a, 2011) causes changes in the dynamics observed in a two sill system. These changes are due to the reduction in the depth of the fjord and the sill top depths and make the simulations more turbulent causing greater erosion of stratification. In addition, changes from a smooth to a realistic bathymetry results in some spatial heterogeneity in the stratification erosion, including the extent to which stratification is eroded in a jet before and after the sills. The Bjällvarp sill is less steep and therefore results in a smaller jet than the Svinesund sill. Finally the realistic non-linear Idefjord stratification results in stronger values of the tidally averaged change in buoyancy frequency and more turbulent overturning. These results show that using realistic simulations can make large changes to the dynamics. Therefore to understand a particular system the correct parameters should be

used, in particular stratification which is highly influential to the dynamics occurring. Realistic simulations conducted here with MITgcm capture the dynamics of Idefjord and agree well with physical data (Johnsson et al., 2007).

Using stratifications and bathymetries based on physical studies of Idefjord the dynamics show super critical jet-basin type dynamics for most of the tidal cycle at all but the lowest (U_0=0.08 ms^{-1}) forcing velocities used. At the control velocity and highest velocity this results in the formation of hydraulic jumps which in turn form high frequency internal waves in the sill basin. This result is contrary to assumptions made by Johnsson et al. (2007).

By reducing the height of the Svinesund sill the influence of the interaction between the sills is reduced, demonstrating that even a relatively small reduction in height can make a very large change in the dynamics. As the sill height is reduced the hydraulic jump at the Svinesund sill is no longer present although M_2 internal waves are still generated at the remaining slope of Svinesund. This results in a reduction of power at the Svinesund sill (from 10^3–10^5 (m/s)^2 for M_2–M_{12} to 10^0–10^1 (m/s)^2 for M_2–M_{12}) in addition to a shift in peak power from M_8 frequencies in the basin centre to M_4–M_6 due to the superposition of an M_2 and and M_6 signal from each sill.

For the control velocity using a stratification with a deeper surface layer (which may be representative of many fjord systems) at close distances the influence of two sills with sub-critical Froude numbers is to initialise a standing internal wave or seiche. This behaviour persists until the spacing of the sills is greater than 5 km, and shows that sill separation must be less than the wavelength of the internal wave that is generated for a standing wave to be generated. After this the standing wave is not generated and the degree to which the internal waves formed at the two sills interact is reduced. However, although a standing wave is not formed there is still evidence of constructive superposition increasing amplitudes at points

of the tidal cycle even with a 50 km inter-sill distance. As the inter-sill distance is increased the largest power amplitude of the signal at the basin centre is generally shifted from the M_2 to M_4 frequency.

In addition, simulations using a higher velocity also demonstrate that significantly less interaction occurs when the wavelength of the internal wave is exceeded. While separation is below this distance partial seiche-like dynamics are initialised with the addition of high frequency internal waves due to the break down of hydraulic jumps. However, as distance is increased above that of the internal wave wavelength the features of the seiche disappear although some superposition still occurs. The remaining dynamics consist of high frequency internal waves from the break down of hydraulic jumps and superposition effects which steepen the waves leading to overturning and breaking forming high frequency internal waves. As separation is increased however the violence of the interaction reduces resulting in less significant overturning due to superposition effects.

Additional simulations with a greater range of Svinesund sill heights show that the planned deepening of the Svinesund sill could disturb toxic sediments at the Bjällvarp sill. This is due to less energy being removed from the system by interaction with the Svinesund sill and higher velocities being present close to the sea bed at the top and sides of the Bjällvarp sill.

CHAPTER 5

VERTICAL MIXING GENERATED BY INTERACTION WITH BATHYMETRY IN ANCIENT EPICONTINENTAL SEAS

Contents

5.1 Synopsis 164
5.2 Introduction 165
5.2.1 Epicontinental seas 166
5.2.2 Parameters to define complex bathymetry 168
5.3 Physical set-up 173
5.3.1 Diagnostics 176
5.4 Influence of the number of bumps 177
5.5 Transgression and Regression scenarios 183
5.6 Discussion and conclusions 188
5.7 Further Discussion – application to coral reef refugia 191

5.1. Synopsis

Previous chapters have analysed the influence of mixing processes in fjord based systems in the present day as this is important for understanding the current influence of internal wave generation and mixing processes on ocean circulation. However, scenarios of tidally forced flow interaction with complex bathymetry may also be applied to the geological past. This may improve understanding of how variations in seabed rugosity brought about by eustatic changes in sea-level (changes due to the volume of water in the ocean) may have a profound impact on the localized efficiency of the oceanic carbon pump. These coastal geomorphological changes are dramatic during a sea-level cycle. The formation of incised river valleys during low-stands (low sea level) creates roughness whereas sediment blanketing during the transgression and high-stands (high sea level) forms smoother bathymetry.

In order to investigate the influence of seabed rugosity, idealised local scale ($10km$ long) 2D vertical slice simulations were set up whereby vertical mixing processes, temperature profiles and passive tracers were analysed. Passive tracers were used to represent the nutrient phosphorus which was initialised as a layer at depth and allowed the calculation of a rate of transfer to surface waters to be calculated.

These numerical simulations were performed using MITgcm and focused on a small section of an idealised continental shelf at various stages of regression and transgression during icehouse and greenhouse worlds. Seabed rugosity is likely to change during a transgression due to sediment blanketing, and also vary with the frequency and magnitude of eustatic change. These simulations allow us to understand the influence of this on vertical mixing and the associated nutrient (e.g. phosphorus) uplift to the surface ocean. Simulations showed that upwelling

produced by flow over rough bathymetry during the early part of an icehouse transgression is likely to be significantly higher than that observed from flow over a smoother seabed during a regression. Importantly during an early transgression a stretch of coastline just 20 km long by 5 km wide could result in an annual upward phosphorus flux of 10^5 tonnes, equivalent to the phosphorus input by the modern Mississippi.

5.2. Introduction

Vertical mixing in the oceans can help to increases productivity by supplying nutrients from the deep sea to the surface ocean where they can be utilized by photosynthetic plankton. This increases the productivity of surface waters and activates the "biological pump" which ultimately increases the dissolved greenhouse gas content of the ocean and impacts upon oceanic acidity (Toggweiler and Key, 2003; Goosse et al., 1999; Rahmstorf, 2002). As plankton dies and sinks to the seabed it is decomposed by bacteria with typically less than $\sim 1\%$ being buried beneath the sea-floor. However, where anoxic conditions exist beneath a stratified waterbody the percentage of carbon sequestered in sediment may be higher (Herndl and Reinthaler, 2013). According to Redfield stoichiometry (Redfield, 1934) the average ratio of phosphorus to carbon in marine organic matter is 1 : 106. However, the average phosphorus content of seawater is very low at only 71 $\mu g l^{-1}$ (Baturin, 2003). Phosphorus is required to synthesize genetic material (DNA and RNA) and to store energy in cells (as ATP) and is therefore a limiting nutrient for biological productivity (Tyrrell, 1999). Increasing the phosphorus content of marine surface waters is therefore an effective way of increasing the productivity and thus in turn increasing carbon drawdown. The productivity of the surface waters is spatially variable and dependent on nutrient supply, which may be supplied from terrestrial

weathering or vertical mixing (Demaison and Moore, 1980). Areas of complex bathymetry are known to increase the rate of vertical mixing, thereby increasing carbon draw-down at these sites. Carbon draw-down is increased due to the higher rates of productivity in the surface water and the subsequent decay of plankton at depth resulting in anoxia (Polzin et al., 1997; Demaison and Moore, 1980).

Areas of upwelling or vertical mixing are often sites where carbon rich, fine-grained sedimentary rocks such as black shales form (Demaison and Moore, 1980; Parrish, 1982; Moore et al., 1993) and such facies have sourced 90% of the worlds hydrocarbon resources (Klemme and Ulmishek, 1991; Tissot et al., 1979). An understanding of the spatial and temporal distribution of vertical mixing processes and the processes that supply nutrients to the sea surface is therefore vital to our understanding of how black shales and petroleum resources have formed (Otis and Schneidermann, 1997). Black shale formation is impacted by organic productivity, preservation of organic matter, and sedimentation rates (Katz, 2005; Tyson, 2005). Black shales commonly occur during transgressions but not during regressions when bioturbated facies dominate (Page et al., 2007). The reasons for this are relatively unknown and we suggest here may be due to variations in seabed rugosity.

5.2.1. Epicontinental seas

Modern shallow coastal oceans (essentially equivalent to the epicontinental seas of the geological past) comprise of relatively small areas and in total represent approximately $\sim 3\times10^6 km^2$ or 8% of the total marine surface area but are responsible for 30% of total marine productivity. This elevated productivity is at least partly due to vertical mixing and upwelling associated with flow/bathymetry interactions (Walsh et al., 1988, 1991; Longhurst et al., 1995). The association of upwelling with productivity in the modern marine realm is known to have extended into the

geological past. In the geological past epicontinental seas were extremely vast in extent (each covering up to $\sim 10^6 \text{km}^2$) and therefore it is likely that these regions would have been responsible for an even greater percentage of total global marine productivity due to increased flow/bathymetry interactions (Allison and Wells, 2006; Wells et al., 2005). Sediments deposited in these seas make up some of the oldest sedimentary rocks that one can find today and contain much of the archive of biodiversity and the vast majority of the planet's hydrocarbons in addition to information about past climatic changes (Scourse and Austin, 2002; Allison and Wells, 2006). A better understanding of the processes that occur in these seas therefore improves the understanding of the formation of these records (Scourse and Austin, 2002; Allison and Wells, 2006). It is already known that changes in tidal dynamics have had significant impacts on the climate due to the influence of changing rates of vertical mixing in epicontinental seas on primary production and the carbon cycle (Walsh, 1988; Scourse and Austin, 2002). However, the influence of bathymetry changes during trangressions and regressions in icehouses and greenhouses on vertical mixing, primary production and the carbon cycle has to our knowledge not been examined.

Transgressions, Regressions and their influences

Coastlines move laterally throughout time and essentially track the transgression and regression of sea-level. During sea-level fall the geomorphology of a coastline may change dramatically as a result of erosion and fluvial down-cutting and this can generate considerable topographic roughness. Conversely, the early stages of a transgression will submerge this roughened surface which is then slowly blanketed with sediment transported from the land. Milankovitch driven eustasy is a driver of such changes although local tectonics also have an effect (Hays et al., 1976; Chappell, 1974; Imbrie et al., 1984; Chappell and Shackleton, 1986; Shackleton,

1987; Collier et al., 1990). However, the magnitude of Milankovitch driven sea-level variation has changed throughout geological time and broadly follows two climatic modes (these are icehouse and greenhouse times). During icehouse times, ice sheets formed on land and the typical Milankovitch-driven 100 kyr climate cycle would have led to $\sim 100\ m$ of sea-level variation as the size of these large ice sheets varied over time. Conversely, greenhouse climates included smaller reservoirs of terrestrial ice which would melt and re-form over time and so sea level rise and fall during the 100 kyr cycle might have only been 10–20 m (Jacobs and Sahagian, 1995; Satterley, 1996; Pittet and Strasser, 1998). Climate also plays a role in erosion and hence the extent to which seabed rugosity can be blanketed by sedimentation. For example, sediment yields in non-glaciated rivers throughout a Milankovitch climate cycle are typically 20–50% lower during the coldest and driest part of the cycle (Van der Zwan, 2002; Blum and Hattier-Womack, 2009). It is not unreasonable to expect greenhouse climates in non-desert environments to produce the same effect.

Thus, all else being equal, we may assume that the coastal topography during an icehouse low stand would be rougher than that which formed during greenhouse times. The early stage of a transgression during icehouse times would therefore involve sea level rise over rougher topography (fig. 5.2.1 (1a,b)) than during greenhouse times (fig. 5.2.1 (2a,b)). Moreover, given that sediment supply is higher in warmer and wetter climates one might expect the greenhouse topography to be blanketed more quickly than in icehouse times.

5.2.2. Parameters to define complex bathymetry

As discussed previously complex bathymetry is widely known to cause turbulence and produce internal waves which may break locally or at a distance from this bathymetry causing enhanced vertical mixing. Vertical mixing and the production of internal waves by complex bathymetry is controlled by a number of parame-

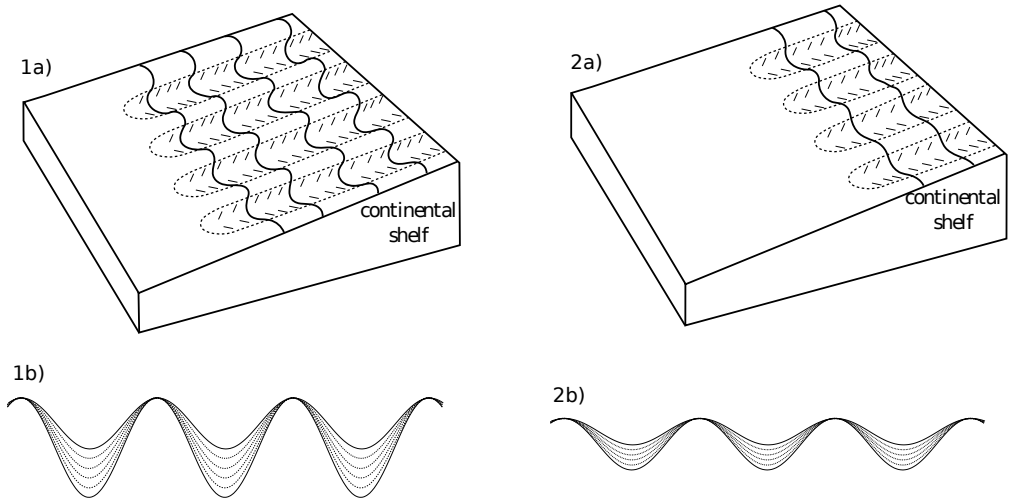


Figure 5.2.1.: A schematic illustrating the different idealized coastal plane topographies during an icehouse (1a) versus a greenhouse (2a). River valleys are shown to extend further along the coastal plane in an icehouse and to cut deeper. This results in the cross sectional topographies seen in (1b) and (2b). During higher sea levels these topographies are smoothed by sedimentation as shown by the dashed lines resulting in lower amplitude bathymetry shown by the lower amplitude solid black line.

ters: tidal velocity (U_0), stratification (N), barotropic frequency (ω_0), topographic height (h_0), Coriolis frequency (f), topographic length (L) and water depth (H). These parameters are used to define a selection of independent non-dimensional numbers which aid our understanding of how mixing is influenced by bathymetry. These are: the tidal excursion parameter ($R_L = U_0/(\omega_0 L)$), topographic slope ($h_0/L = dh/dx$), internal wave characteristic slope ($s = \sqrt{(\omega_0^2 f^2)/(N^2\omega_0^2)}$), relative height of topography (h_0/H), Froude number ($U_0/(Nh_0)$) and topographic steepness ($\gamma = (dh/dx)/s$).

Topographic steepness is the ratio of the slope of the topography to the slope of the wave rays that radiated gravity waves will propagate along. A large number of studies in the past focused on the effects of topographic steepness ascertaining that slopes are sub-critical for $\gamma < 1$ and super-critical when $\gamma > 1$. Initially sub-

critical slopes were examined (Bell, 1975) before critical slopes (Balmforth et al., 2002) and super-critical slopes were examined analytically (Laurent et al., 2003) and numerically (Khatiwala, 2003). However, most of these older studies assumed that $R_L < 1$ and that the ocean was infinitely deep (so internal waves that develop are not reflected at the surface) which is not the case for the continental shelf processes of interest here. Legg and Huijts (2006) however considered a range of R_L values while examining γ. In general increasing the steepness of the slope increases the amount of energy converted to higher harmonics ($n\omega_0$ harmonics are produced up to the buoyancy frequency) which break more easily and thus steeper slopes result in more mixing (Khatiwala, 2003; Legg, 2004b).

The tidal excursion parameter R_L, a measure of the non-linearity of waves, was explored by Bell (1975) for regimes with $\gamma < 1$ resulting in the understanding that for super-critical values ($R_L > 1$) baroclinic waves are produced at higher harmonic frequencies. This parameter, R_L, is related to the Froude number, and is equivalent to the Froude number if the horizontal scale of the obstacle and internal tide are the same. R_L describes the wavelength of the tide relative to the topography, and therefore how far the water particles are advected over the obstacle. When this is critical ($R_L \sim 1$) the flow interaction is most sensitive to the obstacle and nonlinearity becomes important, but energy conversion rates are smaller than when the parameter is super-critical (Bell, 1975; Llewellyn Smith and Young, 2002). However, as many analytical studies (Balmforth et al., 2002; Laurent et al., 2003) assumed $R_L << 1$, less is known about the regime where both R_L and γ are large (Nycander, 2005).

The Froude number tells us when flow is affected by bathymetry. When the Froude number is sub-critical (< 1) the flow is effectively blocked by the bathymetry and when super-critical the flow is relatively unaffected by the bathymetry and the flow response is linear (Nycander, 2005). For critical Froude numbers tran-

sient effects are expected and the Froude number may become super-critical due to acceleration over the topography (a response often present at fjord sill tops) (Legg, 2004b). In addition, for regimes with super-critical R_L and sub-critical Fr hydraulic control may be experienced (Legg, 2004b). Legg (2004b) also includes a study of linear and more realistic non-linear stratifications. Legg and Klymak (2008a) examined flow over isolated ridges resulting in two Froude number regimes for hydraulic jumps near the shelf break. These hydraulic jumps are similar to those seen in fjords however represent different regimes due to the greater depths. Studies by Qian et al. (2010) show that when the Froude number is small at super-critical ridges maximum energy flux is independent of γ and strong stratification around the top of the ridge is most important in controlling the amount of energy flux. When stratification around the ridge crest is weak, however, ridge height becomes a more important factor. For simulations in this study it is therefore important to consider the location of the thermocline in addition to the steepness of bathymetry.

Research by Garrett and Kunze (2007) divided the parameter space into regimes based on three of the non-dimensional numbers (R_L, γ and Fr) since certain terms are used in more than one non-dimensional number, these regimes can be seen in figure 5.2.2. Despite the range of research conducted (Bell, 1975; Balmforth et al., 2002; Laurent et al., 2003; Nycander, 2005; Garrett and Kunze, 2007; Nash et al., 2004; Gayen and Sarkar, 2011a; Legg and Klymak, 2008b; Qian et al., 2010) most is applicable to the deep ocean or continental slope where flow speeds are lower (by one to two orders of magnitude) and the water depth is larger, and thus characterise regimes that are unsuitable for understanding processes on the continental shelf. The earliest work by Bell (1975) focused on regimes one and two of fig. 5.2.2 which was furthered by Balmforth et al. (2002) in these regimes. In the deep ocean studies such as Legg and Klymak (2008b) examined the influence of the

Hawaiian ridge and explored dynamics in regime 5 of fig. 5.2.2. Further studies have examined the effects on the mixing, turbulence and conversion of barotropic to baroclinic energy at the continental slope, examining the criticality of the slope and its corrugations (Nash et al., 2004; Gayen and Sarkar, 2011b).

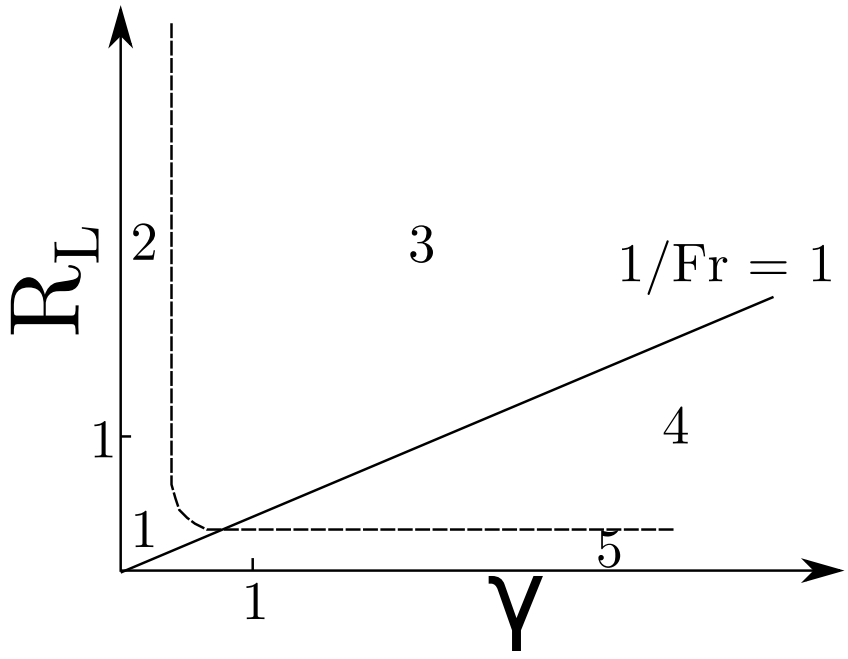


Figure 5.2.2.: Regimes to divide the parameter space are provided based on the tidal excursion parameter, topographic steepness and Froude number. Based on Garrett and Kunze (2007)

Limited studies have been undertaken in regimes suitable for understanding bathymetric features on the continental slope where depths are on the order of 100 m and flow velocities can reach a few metres per second. The parameters expected for these scenarios result in $R_L > 1$ and either $\gamma \sim 1$ or slightly super or sub-critical values for γ (weaker stratification results in more sub-critical values of γ and thus regime five) in addition to sub-critical values for Fr that become super-critical above bathymetric features. This combination of non-dimensional numbers shows that these scenarios correspond with regimes three, four and five. Previous studies within these regimes include a sub-set of Gaussian bathymetry

results from Legg and Huijts (2006) in addition to studies of tidal interaction with sills at the entrance to fjords such as those in chapters 3 and 4 and the references therein. One of these referenced studies (Xing and Davies, 2009a) showed clearly that increasing the depth of the sill top or increasing the sill width reduces the amount of energy converted to internal waves and also the amount of mixing. In addition Xing and Davies (2010) added small scale bathymetry to the side of the sill, showing that bumps of height $\sim$ 5m and wavelength $\sim$ 100 m can increase the amount and extent of mixing especially when positioned close to the region of flow separation. Studies by Xing and Davies (2009b) also considered sill aspect ratio and depth. Increasing the width or depth of the sill top reduced the intensity of the hydraulic transition (or resulted in no hydraulic transition) and reduced the amount of mixing and energy converted to baroclinic waves.

5.3. Physical set-up

This chapter uses the Massachusetts Institute of Technology general circulation model (MITgcm) (Marshall et al., 1997a,c). This model, which is discussed further in chapter 2.5.3 was chosen here due to its successful use for previous studies considering the effect of complex bathymetry on the generation of internal waves and mixing at larger scales (Legg, 2004a; Legg and Klymak, 2008b; Legg and Huijts, 2006) in addition to a number of studies on the continental shelf for fjord systems (Xing and Davies, 2006, 2009a, 2010). As in previous chapters simulations here are non-hydrostatic and use the implicit linear free-surface options in MITgcm. The second order flux limited advection scheme is used for temperature and additional tracers here as used by Legg and Klymak (2008b) with a centered second-order scheme for momentum advection with staggered baroclinic Adams-Bashforth time stepping.

Two dimensional vertical slice domains are used with no slip bottom boundary conditions and periodic lateral boundary conditions. Ten grid cell sponge regions at either "end" of the domain are used to relax the temperature and passive tracer fields back to their initial conditions over a tidal cycle. Mesh resolutions used here are based on those used by Xing and Davies (2009a) due to the similar scale and nature of the experiment and are $\Delta x = 10\ m$ (which telescopes away from the area of interest to $\Delta x = 800\ m$) and $\Delta z = 1\ m$. As in Xing and Davies (2009a) and previous chapters small constant values of diffusivity $K_h = K_v = 10^{-7}\ \mathrm{m^2 s^{-1}}$ were used and viscosities of $A_h = 10^{-1}\mathrm{m^2 s^{-1}}$, $A_v = 10^{-3}\mathrm{m^2 s^{-1}}$. In these simulation tidal velocity is applied as a body force, as in Legg (2004a), of the form:

$$\frac{\partial u}{\partial t} = U_0 \omega \cos(\omega t), \tag{5.3.1}$$

where U_0 is the forcing velocity amplitude (which is scaled linearly with depth in this study as shown in table 5.3.1) and $\omega = 1.41$ x $10^{-4}\ s^{-1}$ is the M_2 tidal frequency.

To understand the key stages of transgression and regression scenarios during an icehouse and greenhouse world six simulations were conducted and key parameters describing these are presented in table 5.3.1. Bathymetries used in these simulations are based on the coastal topography of South Wales (UK) (fig. 5.3.1). This topography has an approximate 20 m amplitude and 500 m wavelength, which is used in simulations I1 and I2 representing an icehouse bathymetry in the early and late stages of a transgression before sedimentation. Simulation I3 and I4 have a reduced topography that reflects the impact of sedimentary blanketing during the regression. Simulations for a greenhouse world use bathymetry which has valleys that are a quarter as deep because sea-level fluctuation is lower during these periods (fig. 5.2.1). Due to the limited sea level variation during a greenhouse only two

simulations are used (G1 and G2). These test the influence of sedimentation on bathymetry/flow interaction by using simulations before and after sedimentation. These simulations all use temperature stratification found in the present day shallow Adriatic Sea (which has a depth of $\sim 125\ m$) during July. This stratification (NOAA, 2016) (shown in fig. 5.3.1d) is approximated using the equation:

$$T(z) = 1.22995 \times 10^{-5} z^3 + 3.33974 \times 10^{-3} z^2 + 0.3026222 z + 23.38305, \quad (5.3.2)$$

where z is depth. As for simulations in previous chapters, the influence of salinity is neglected for simplicity. In addition the influence of Coriolis force is neglected in these simulations as it was found to have a limited influence on the overall dynamics at the scale of interest here (Appendix A).

Table 5.3.1.: List of simulation parameters for transgression and regression scenarios during an icehouse and greenhouse world.

Reference	Bump height (m)	Domain depth (m)	Peak velocity $(U_0\ (ms^{-1}))$	Stage
I1	40	50	0.1	Early transgression
I2	40	100	0.05	Late transgression
I3	20	100	0.05	Early regression
I4	20	50	0.1	Late regression
G1	10	30	0.17	Transgression
G2	5	30	0.17	Regression

Since scenarios of coastal roughness throughout periods of transgressions in the geological past would involve multiple bathymetric features due to multiple river valleys along the coastline, simulations here must use a set number of features to give a representative assessment of the affect of these bathymetries. As the amount of mixing which occurs due to interactions with bathymetric features is known to increase when additional features are added, as long as the separation distance is relatively small (Xing and Davies, 2011) a number of features must be chosen whereby the influence of the addition of further bumps is minimised. In

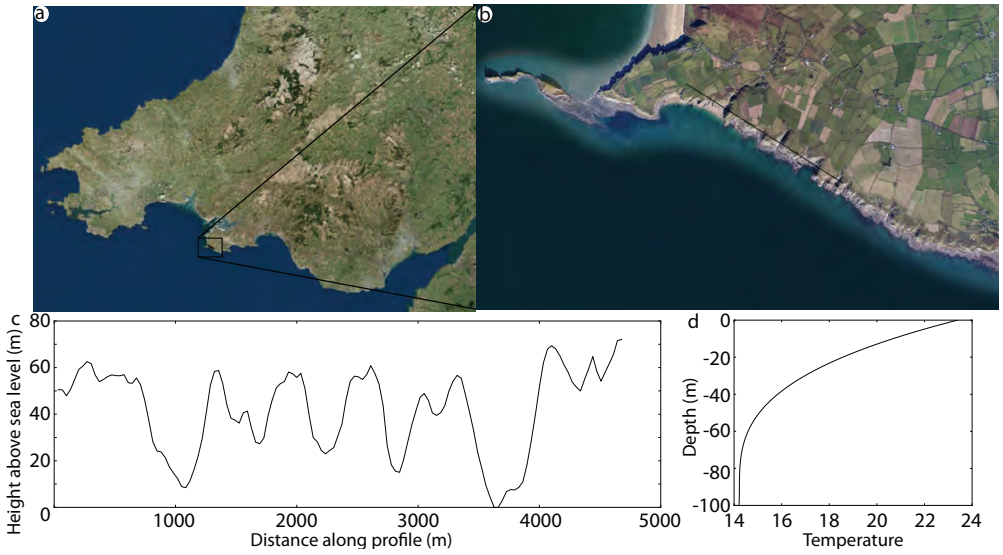


Figure 5.3.1.: Location (a) and zoom (b) of the south Coast of Wales to show the location used for idealised bathymetries (satellite images from Esri ArcGIS Basemaps). In figure b valleys can clearly be seen that form the complex bathymetry during a transgression. The bathymetry that results from the transect parallel to the coast shown in (a) and (b) is displayed graphically in (c). For completeness the vertical temperature profile suitable for this study is shown in (d), this was extracted from July data for the Adriatic Sea (a suitable shallow sea) (NOAA, 2016).

order to decide how many bathymetric features to include in this study a range of scenarios were investigated using between one and ten bathymetric bumps. These simulations used the same sinusoidal bathymetry with a 500 m wavelength as found approximately in fig. 5.3.1c and the non-linear stratification of equation (5.3.2). Simulations are representative of the domain depth and bump amplitude of simulation I1.

5.3.1. Diagnostics

For this study mixing is quantified by reference to the time-averaged temperature fields (where temperature is averaged over three tidal cycles), Richardson number

snapshots, critical Richardson area A_{Ri} (described in chapter 3.4), passive tracers and time-averaged temperature and passive tracer profiles (between the fifth and sixth bumps for icehouse and greenhouse simulations). Passive tracers are used to mimic the movement of phosphorus due to the generated turbulence. The passive tracer is initialized as a layer of constant concentration of $300\mu gl^{-1}$ (a value typical of stagnant basins such as the Black or Baltic Sea. However, the ocean average concentration is $71\mu gl^{-1}$ (Baturin, 2003) below the seasonal thermocline ($\sim 25\ m$ depth)). In addition, for studies of the influence of the number of bumps used, the change in buoyancy frequency (DN) is also used to give an indication of mixing as Xing and Davies (2009a, 2011). This parameter: $DN = \overline{N}^2 - N_0^2$, is the tidally averaged change in buoyancy frequency as presented in chapter 4.3.1.

5.4. Influence of the number of bumps

Simulations containing a range of numbers of bumps show that as additional bumps are added an area of negative DN (fig. 5.4.1) develops at bump crest depth. This indicates that additional mixing occurs as additional bumps are added as was shown for a two sill scenario by Xing and Davies (2011).

For a single bump scenario values of DN show negative values extending either side of the bumps crest at depths equivalent to around half the height of the bump. In addition, a positive change in DN is shown in the surface layer. This is due to the advection scheme attempting to preserve the rapid change in density at this depth due to the non-linear stratification used.

As an additional bump is added a large area of more negative DN is added, this area occurs between the two bumps at depths of 10–20 m (equivalent to bump crest depth). Additional areas of strong negative DN are added within each inter-bump basin created by the addition of bumps (at 10–20 m depth) until around

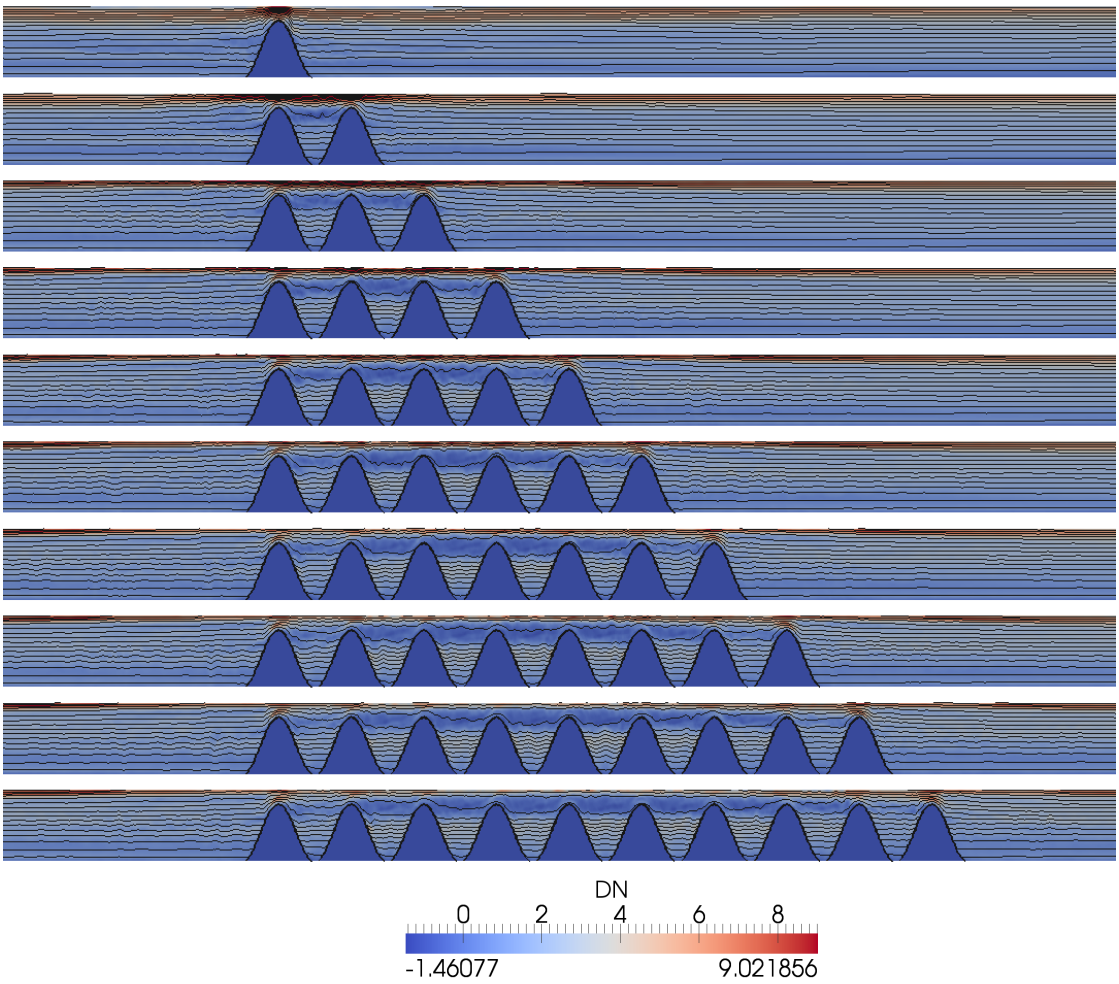


Figure 5.4.1.: Snapshots of tidally-averaged (over three tidal cycles) change in buoyancy frequency with contours of time-averaged temperature (for every 0.5^{o}C). Results are shown for one bump (top) to ten bumps (bottom). Dark blue areas show regions where buoyancy frequency changes indicate mixing has occured which coincide with areas where time-averaged temperature contours have spread indicating erosion of the stratification.

five bumps are present. When the sixth bump is added although another area of negative DN is added at similar depths in the new inter-bump basin, the intensity of the change in buoyancy frequency is reduced in this basin and also in the other outer basins (i.e. between bumps one and two). In addition to this the intensity

of the change in DN in the inner basins is slightly reduced when compared with simulations with fewer bumps. From this simulation and up until ten bumps are used, for every additional bump added the intensity of the negative change in DN in the outer inter-bump regions is reduced although the intensity of DN in the inner basins is maintained from six bumps onward.

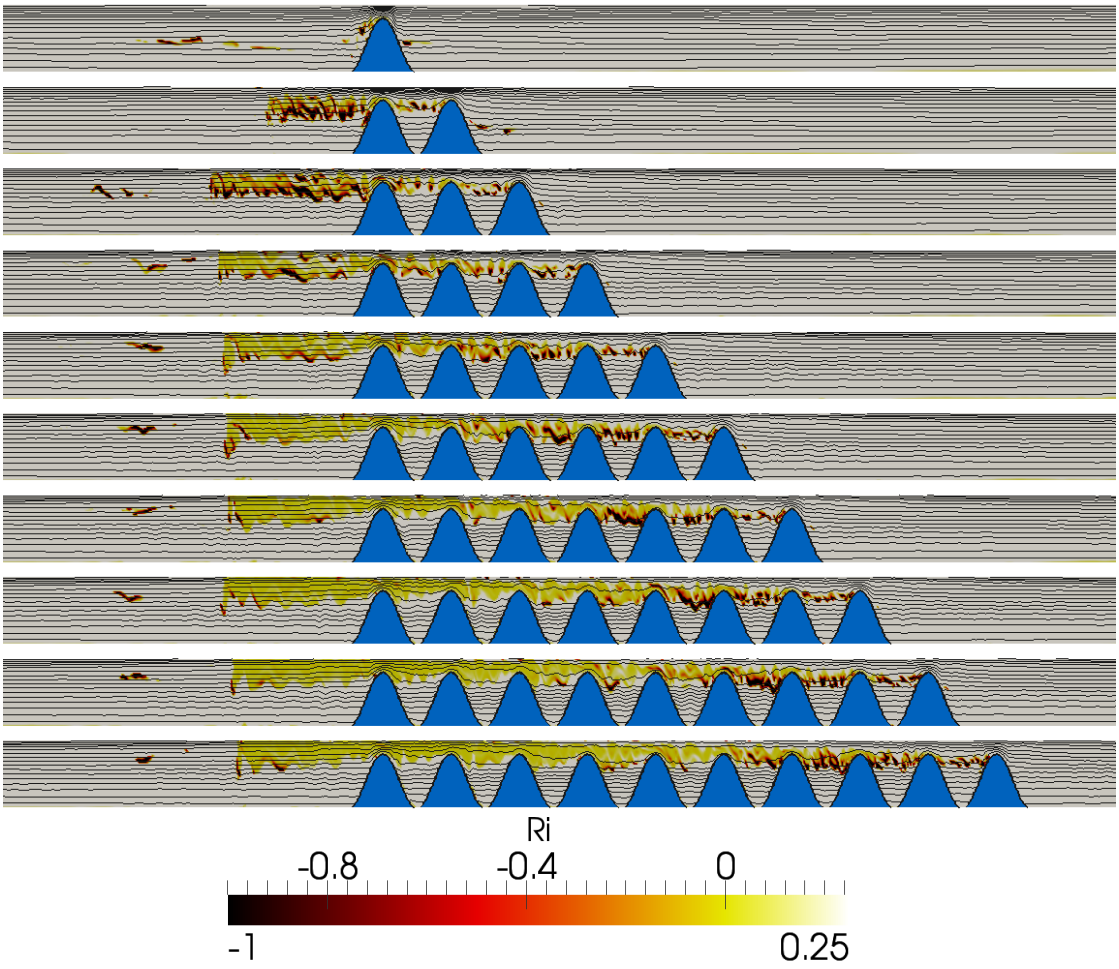


Figure 5.4.2.: Snapshots of Richardson number at 23/8 T with contours of time-averaged temperature (for every 0.5^{o}C), where flow is left to right. Results are shown for one bump (top) to ten bumps (bottom).

The changes in the Richardson number as additional bumps are added also show that the water column is unstable and that overturning and vertical mixing

will occur either side of a single bump at mid-depths and in between additional bumps at bump top depths (10–20 m depth). This agrees with the changes in buoyancy frequency (fig. 5.4.1) due to the mixing that has occurred. Simulations with up to three bumps demonstrate a dramatic increase in the area with critical Richardson numbers (at the time shown in fig. 5.4.2) as bumps are added while the values of the Richardson number remain extreme. The additional areas with critical Richardson number are: an increase in the length before the first bump and within the new inter-bump regions. As the number of bumps exceeds three the length of the unstable region before the first bump (on the left) remains constant and the only additional areas with critical Ri occur in the new inter-bump regions. In addition, when more than three bumps are used, the area with extreme values of Ri (dark patches in fig. 5.4.2) remains relatively constant at this point in the tidal cycle and is located towards the right hand side of the chain of bumps (the side from which tidal flow is propagating at this point). A potential reason for this area of extreme value not extending further into the inter-bump basins towards the left hand side is that energy has been removed from the flow by interaction with the first three bumps resulting in slightly less turbulent dynamics.

The additional areas with critical Richardson number however vary slightly throughout the tidal cycles and thus slightly different trends are present in fig. 5.4.3 to those discussed in the previous paragraph. Instead this graph which plots the critical Richardson area shows that approximate convergence of the critical Richardson area is not found until $\sim$8 bumps are used. This graph begins to converge as smaller critical areas are added with each additional bump as the number of bumps increases due to the extraction of energy by preceding bumps as explained previously.

The influence of vertical mixing due to bathymetry is to erode the stratification, fig. 5.4.4 shows that simulations using 2–4 bumps erode the most stratification.

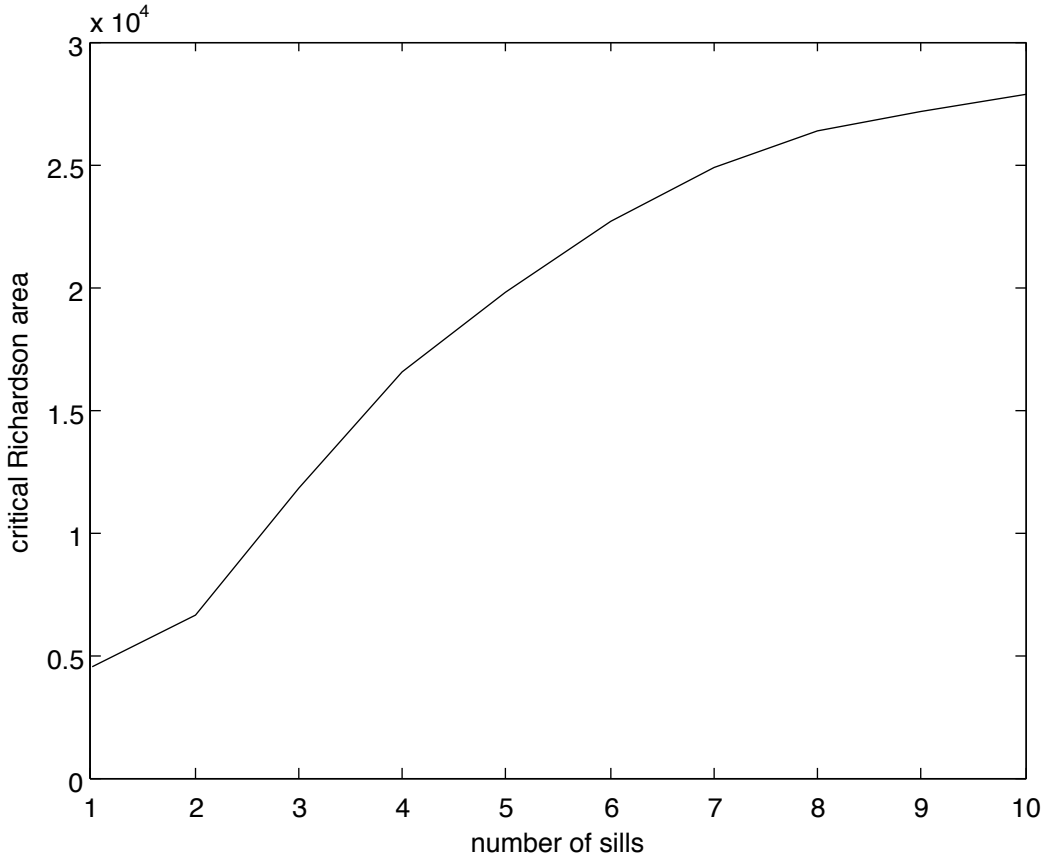


Figure 5.4.3.: The critical Richardson area versus the number of sills (bumps) after three tidal cycles. Simulations show that as more bumps are used the area that is liable to vertical mixing increases, however the extra amount of area with critical values begins to increase less quickly as more bumps are used.

This is evident by the "straightening" in the vertical in the plot of the temperature profile particularly between 8–20 m depth and is due to this location experiencing the most turbulent mixing at both directions of tidal flow. Simulations with more than four bumps result in slightly less mixing and an erosion of stratification between $\sim$ 5–30 m at this location due to the region only experiencing the most vigorous mixing when flow propagates from left to right, as when flow is reversed the mixing at bumps preceding this point extract energy from the system resulting

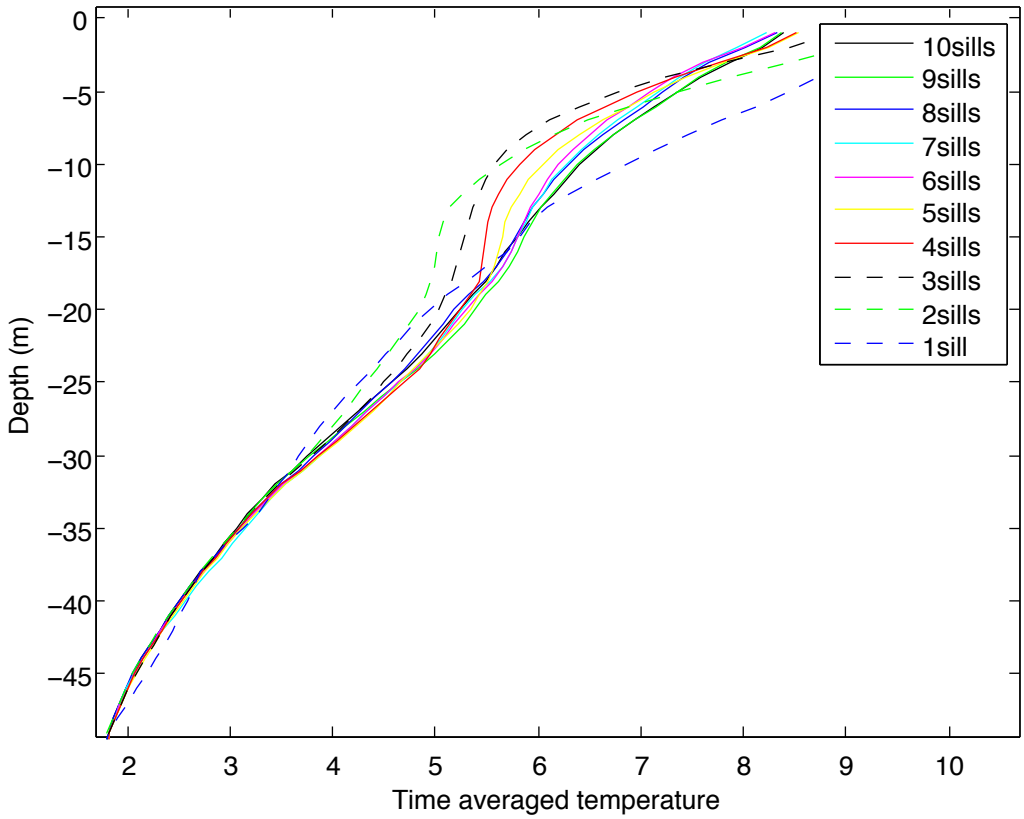


Figure 5.4.4.: The temperature profile at 250 m after the first bump crest for each icehouse simulation with between one and ten bumps after three tidal cycles.

in less turbulent mixing at this location. Graphically near convergence is shown from simulations using six sill and the results of nine and ten sill are almost identical.

All subsequent simulations therefore use ten sills as simulations with this number are fairly well converged in terms of the temperature profile and $\overline{A}_{\mathrm{Ri}}$, and in addition snapshots show the dynamics do not change considerably by adding additional bumps.

5.5. Transgression and Regression scenarios

In an icehouse world snapshots of time-averaged temperature (fig. 5.5.1a) show that during an early transgression (I1) the distinct thermocline is no longer present and instead a more gradual temperature gradient is left due to the erosion of stratification. This corresponds with a large area of critical Richardson number (fig. 5.5.1b) originating at the bump crests. Critical values of Ri demonstrate that turbulent mixing occurs for this simulation above ~ 25 m depth, this location indicates that mixing occurs across the thermocline. However, as the water column depth is increased to that used in simulation I2 of a late stage of a transgression, very little change is seen in the time-averaged temperature snapshot when compared with the original stratification. The thermocline remains largely straight and strong and only small perturbations are present due to internal wave beams. The area of the domain with critical Richardson numbers is also reduced and is mostly limited to depths at or below the peaks of the bumps ($\sim$60 m depth) and is therefore unlikely to be as influential on the erosion of the stratification at the thermocline or in moving nutrients above the thermocline.

As sedimentation develops and bump amplitude is halved (I3) the topographic slope and relative height of the bathymetry is also halved, this results in critical Richardson numbers only occuring deeper in the water column and with values that are slightly larger indicating a more stable water column than I1 and I2. Very little difference is visible between time-averaged temperature snapshots for I2 and I3 as any mixing occurring is below the thermocline and therefore less important for this study (and also more difficult to visualize due to the colour scheme). As sea level is reduced during a late stage of a regression (I4) once again mixing is visible in the time-averaged temperature snapshots by slight stratification erosion, and a more gradual temperature gradient is shown (fig. 5.5.2), however the change in

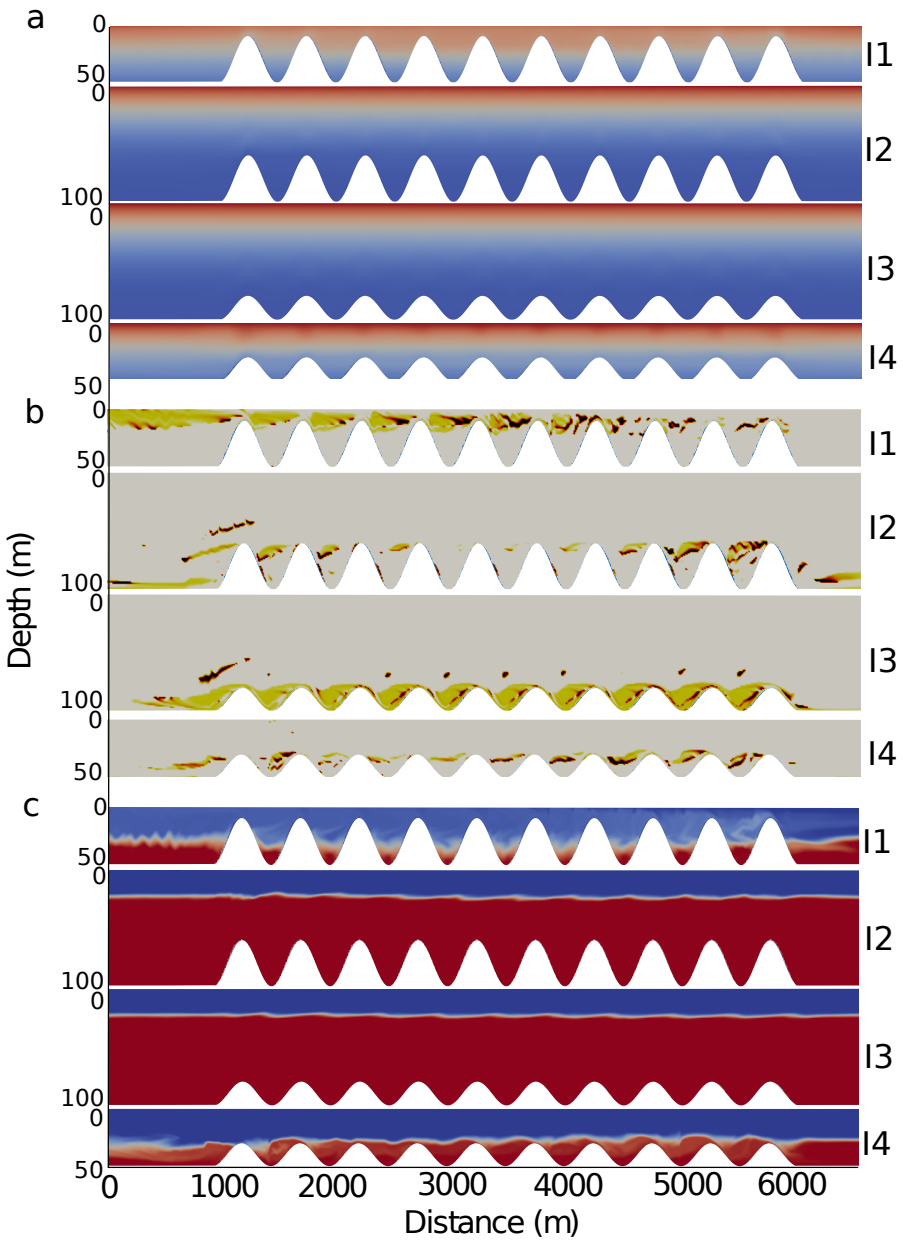


Figure 5.5.1.: Snapshots of (a) time averaged temperature, (b) Richardson number and (c) passive tracer for icehouse world simulations I1, I2, I3 and I4 as labelled. Snapshots are shown after three tidal cycles for (a) and at 23/8 T (where flow is right to left) for (b) and (c) for the centre of the domain.

stratification is smaller than for simulation I1. The effect of internal wave beams on time-averaged temperature is also observable for this scenario. Richardson number snapshots show smaller values of R_i than for I3. However due the lower topographic slope and relative height of the bathymetry the area with these values is smaller and deeper than for I1 and thus not as significant in terms of mixing across the thermocline.

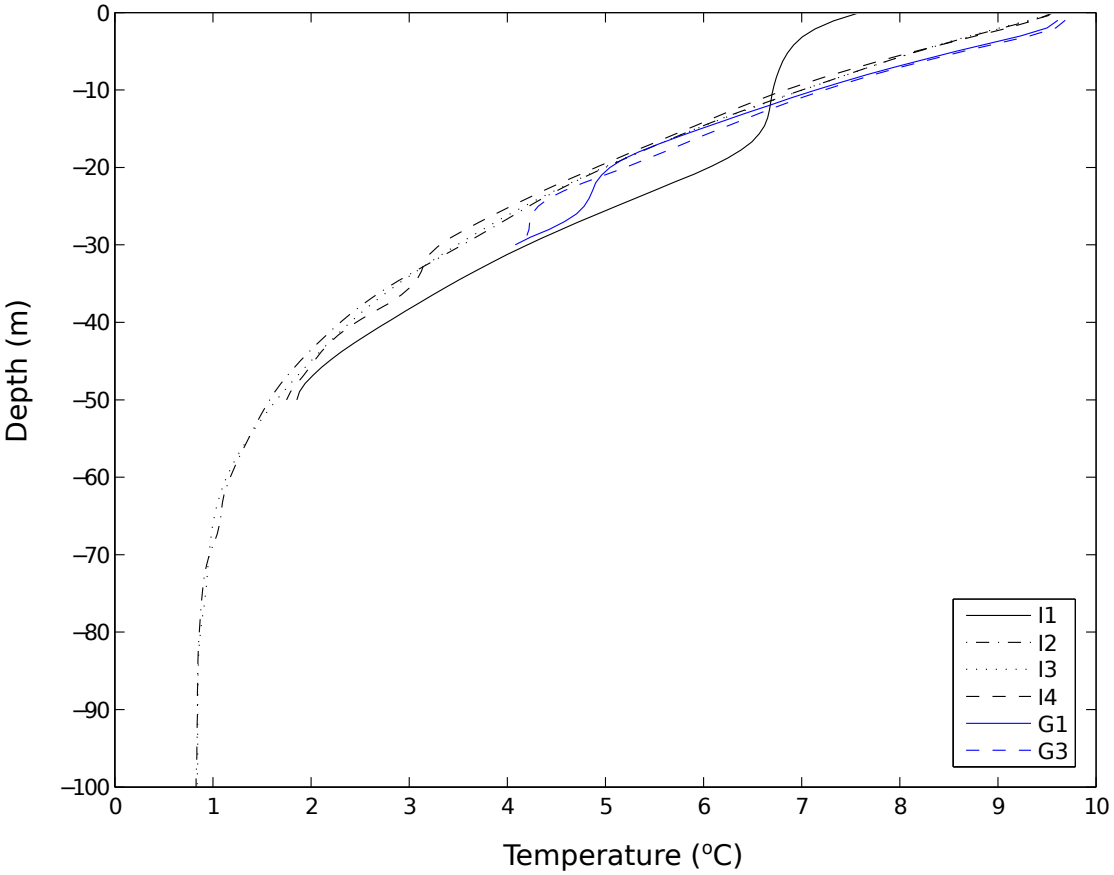


Figure 5.5.2.: Vertical time-averaged temperature profile between the fifth and sixth bump. Simulations for icehouse scenarios are shown in black and greenhouse scenarios are shown in blue.

The depth profiles of time averaged temperature (fig. 5.5.2) are further evidence that I1 has resulted in the most significant mixing. This mixing results in an erosion of stratification or "straightening" of the temperature profile between $\sim$

5–18 m depth for I1 (when compared with the original stratification shown in fig. 5.3.1d). For I2 despite the presence of some areas of unstable water (although these areas are smaller than I1) the location of the resultant mixing occurs below the thermocline and means there is very little change in the stratification, a slight "straightening" is however present between 70–100 m depth. After sedimentation (I3) there is very little mixing and therefore very little change in the thermocline until later in the regression when sea level is reduced (I4). At this stage there is again some significant mixing resulting in erosion of stratification above the bump crests at $\sim$ 35–38 m depth, although this is not as significant as for scenario I1.

The location of the passive tracer after three tidal cycles is presented both spatially (fig. 5.5.1c) and graphically between the fifth and sixth sills (fig. 5.5.3). These figures show that during the late transgression and early regression during an icehouse (I2 and I3) when sea levels are highest; there is very little vertical mixing of the tracer. This results in very little change to the original distribution, with only slight deflections present in the top surface (fig. 5.5.1c). Slightly more mixing of the passive tracer occurs for I4, resulting in the top of the interface being slightly above that of I2 and I3 (fig. 5.5.3) and the snapshots show there are slightly stronger perturbations to the interface of tracer than for I2 and I3. However, I1 includes far more tracer mixing than all other simulations (fig. 5.5.3), resulting in a concentration of 75 $\mu g l^{-1}$ in the upper 20 m of water. This concentration is present between every bump at this depth and extends in front of the first sill and beyond the last sill for $\sim$ 5 km and $\sim$ 3 km respectively, resulting in an elevated concentration over a greater area. Therefore during an icehouse world significant mixing in terms of tracer distribution would only occur during the early transgression.

By comparison during the greenhouse simulations there is less mixing. Sedimentation further reduces the topographic slope and relative height of the bathymetry

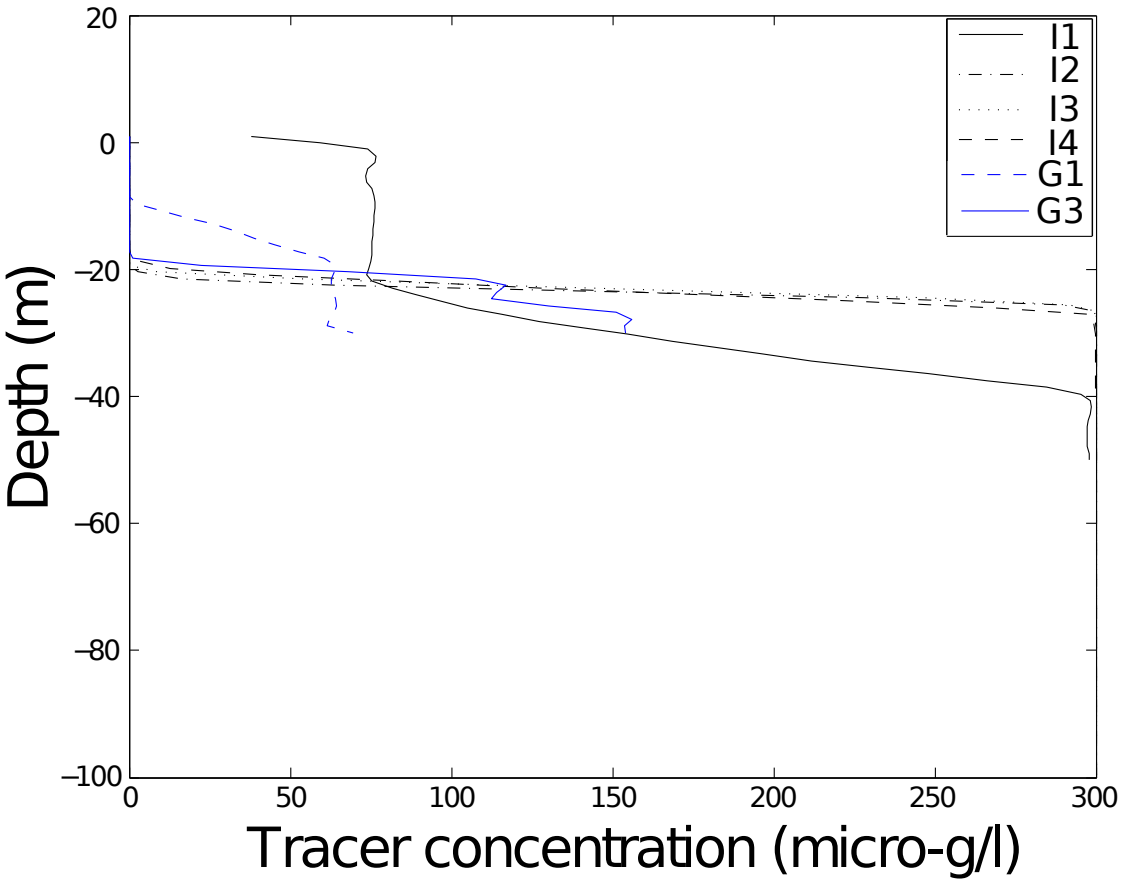


Figure 5.5.3.: Distribution of a passive tracer at $3T$, which was initialized as a layer below 25 m depth with constant concentration of 300 $\mu g l^{-1}$. Simulations for icehouse scenarios are shown in black and greenhouse scenarios are shown in blue.

and the time-averaged temperature shows spatially (fig. 5.5.4a) and graphically (fig. 5.5.2) less erosion of the stratification. Although there is some perturbation to the thermocline and slight overturning apparent (fig. 5.5.4a) for both G1 and G2. Richardson number snapshots (5.5.4b) show that G1 and G2 also result in only small areas of critical values which occur within 15 m of the seabed and thus this makes little impact on the temperature profile (mainly due to the lack of much denser colder water due to the limited depth of these scenarios). The distribution of tracer however is significantly affected for the G1 scenario (fig. 5.5.3), where a

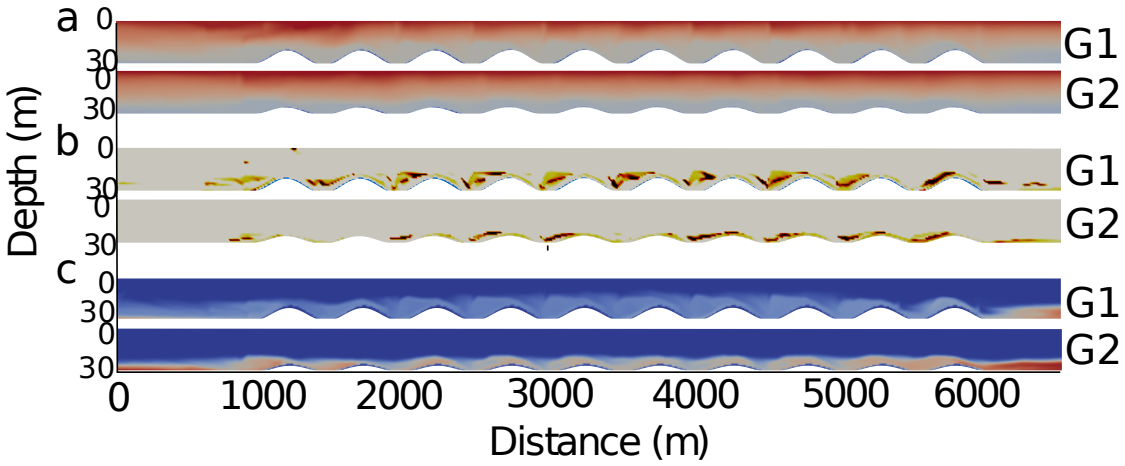


Figure 5.5.4.: Snapshots of (a) time-averaged temperature, (b) Richardson number and (c) passive tracer for greenhouse world simulations G1 and G2 as labelled. Snapshots are shown after three tidal cycles for (a) and at 23/8 T for (b) and (c) for the centre of the domain.

concentration of 25 μgl^{-1} is found ∼15m from the surface but again there is little influence for the regression after sedimentation scenario G2.

Therefore it is likely that significant mixing influencing the amount of nutrients available in the surface layer will only occur during the early stages of a transgression predominantly during an icehouse world but also to a lesser extent during a greenhouse world when rough bathymetry is present. It is unlikely to occur at any other period of transgression explored here. Mixing is significantly reduced to the lowest levels in any scenario explored here during the regression of a greenhouse world due to reduced fluvial down-cutting causing smoother bathymetry which is further blanketed by sediments.

5.6. Discussion and conclusions

The addition of further bumps to a simulation was shown to increase the amount of mixing occurring. This was shown by the increase in area with negative DN values

at 10–20 m depth in the new inter-bump basins, in addition to greater areas with critical Richardson numbers at similar bump-top depths, which demonstrate an unstable and overturning water column at these locations. However, the amount of mixing that occurs does not display a linear relationship with the number of bumps used. Snapshots of DN show that when more than five bumps are used the intensity of the negative DN values in outer basins is reduced, as with each additional bump the intensity of the outer basins is further reduced. In addition between five and six sills the intensity of DN in the inner basins is reduced and remains at this lower level when further bumps are added. A similar effect is shown for critical Richardson number whereby after three bumps are used the length of the unstable area extending from the first and last bumps remains fairly constant in addition to the size of areas with most extreme value of Ri. This is due to the energy being removed from the flow by interaction with the first few bumps resulting in slightly less turbulent dynamics for the remaining bumps. This is also demonstrated by the time-averaged temperature stratification between the first two bumps, which shows that initially as the number of bumps is increased there is greater removal of stratification. However, this converges on a slightly weaker removal of stratification when more than five bumps are used. This is also due to this region then only experiencing the most turbulent dynamics for half of the tidal cycle due to the direction of flow. Together the analysis of diagnostics and dynamics occurring suggest that the amount of mixing begins to converge after eight bumps are used and therefore ten bumps were used for further studies.

Transgression and regression simulations show that vertical mixing is most pronounced during the early inundation phase of topography and is proportional to topographic relief. The greater relief that might be expected in an icehouse world substantially elevates the level of vertical mixing by comparison with that which might reasonably be expected in a greenhouse climate. This is due to sediment

blanketing and also the frequency and magnitude of eustatic change, if reduced levels of eustatic variation in a greenhouse world are accompanied by lower levels of coastline rugosity (ignoring the effects of local tectonics) then coastal vertical mixing will be lower in a greenhouse world than it is in an icehouse world.

Productivity due to influx of phosphorus and nitrogen in the Mississippi delta results in "dead zones". These zones occur when planktonic productivity is elevated and the aerobic decay of sinking dead plankton depletes the dissolved O_2 in the water column. The Mississippi annually discharges 100,000 metric tonnes of phosphorus resulting in increased algal productivity and a "dead zone" of around 10,000 km^2 (Joyce, 2000). Rates of vertical advection of the passive tracer found from simulations here suggest that a stretch of coastline just 20 km long and 5 km wide (100 km^2) during an icehouse early transgression could result in a similar flux of 100,000 metric tonnes per year from depth to the upper 20 m of the water column. While nutrient flux is similar to that delivered by the modern Mississippi, the area in which this acts is two orders of magnitude smaller. One might therefore expect the effects of oxygen depletion and black shale deposition (nutrient loading) in surface waters to be more acute over areas of elevated topography. The role of topographic roughness on marine coastal productivity is clearly important and exerts an as yet poorly defined role on carbon cycling through time. Vertical mixing will be most intense during the earliest phase of a transgression but also much more intense during icehouse times than during greenhouse times.

The simulations used in our experiments are two-dimensional simplifications of the real world. In the real world coastal plains may be flattened surfaces or gently inclined ramps. Transgression over a gently inclined ramp-like surface roughened by incision would mean that there would always be a narrow zone of vertical mixing associated with topography but that this zone would be likely only to occur where the top of the topographic highs was close to the thermocline. If a horizontal but

incised coastal plain was inundated then the elevated mixing would be limited to a narrow period during the transgression when the thermocline was just above the topographic highs. Importantly mixing associated with seabed rugosity is highly localised and likely to create a coastal ocean with spatially variable productivity. Spatial heterogeneity in nutrient levels could positively benefit upon planktonic habitat variability and hence biodiversity.

This work focuses on comparing two highly idealized cases in different climatic modes and highlights the impact of submarine topography on the nutrient status of shallow coastal areas while also showing the differences that would be expected between icehouse and greenhouse times. It is also worth noting that the same effect would occur during the transgression of a block faulted basin where the rough topography was fault bounded controlled rather than being generated by erosion. It is important to note that nutrient advection and surface productivity does not always lead to benthic (sea-floor) anoxia as this is also dependent upon oxygen supply.

Modern nutrient upwelling associated with submarine topography can also impact positively upon marine diversity. It is thus likely that a rough coastal plain that was inundated during a transgression would produce a coastal ocean with a spatially diverse nutrient status and this could also impact upon the palaeobiodiversity of ancient epicontinental seas.

5.7. Further Discussion – application to coral reef refugia

Corals and in particular tropical scleractinian corals are sensitive to temperature changes as this affects the balance between the algal endosymbionts and the coral host leading to the expulsion of the algae and which in turn leads to coral bleaching

and mortality (Hoegh-Guldberg, 1999; Fitt et al., 2001; Glynn, 1991). In the present day the breaking of large amplitude internal waves at complex bathymetry is a known explanation for coral refugia. Studies of the Andaman Sea (Wall et al., 2015) have shown that tropical scleractinian corals, have persisted to be present despite rising sea surface temperatures. This is due to the vertical mixing and intermittent but regular advection of cold water from depth to the surface reducing the temperature of surface waters of the coastal ocean (Vlasenko and Stashchuk, 2007; Wall et al., 2015).

In periods of the geological past during coral crises after marine mass extinction events such as those of the end-Triassic, global warming and ocean acidification (Hautmann, 2012) meant that coral reefs disappeared globally. Remaining corals were typically found as solitary corals or as reefs only at very northerly paleolatitudes above 30^o (Lathuiliere and Marchal, 2009), however despite this reefs have been found at low latitudes. An example is the Early Hettangian (the earliest period postdating the Triassic) well developed coral reef (currently found in Southern France) at a paleolatitude of under 30^o, known as Elmi's reef (Kiessling et al., 2009) which has been considered an exception to the Early Jurassic reef decline, and has a continental shelf setting.

Possibly the survival of Elmi's reef may be attributed to similar reasons to the survival of the current day reefs of the Andaman Sea. Although studies of the Andaman Sea have examined the influence of breaking internal waves on temperature stratification to alleviate thermal stress on corals, it would not be unreasonable to suggest that other influences which affect surface temperature stratification could potentially also alleviate thermal stresses. The roughest areas of bathymetry examined in this chapter reduce temperatures in the surface layer by $\sim 2^oC$, which may be sufficient to provide refugia since corals are susceptible to thermal stress due to changes of 1–2^oC (Hoegh-Guldberg, 1999; Fitt et al., 2001). Since Elmi's

reef is situated on the continental shelf and during a period of sea level rise it is possible that its survival may be the result of such bathyemetry induced vertical mixing, however more detailed paleobathymetries would be necessary to comment further.

Although the work in this thesis primarily examines the formation of internal waves and the associated mixing a potential next stage of research would be to study internal wave breaking and implications this could have had in the geological past. Significant research has been conducted on breaking internal waves at sloping bathymetry including the influence of steepness and curvature, these studies include but are not limited to the following Cacchione and Wunsch (1974); Legg and Adcroft (2003); Venayagamoorthy and Fringer (2006); Vlasenko and Hutter (2002); Bourgault and Kelley (2007); Berntsen et al. (2006); Thiem and Berntsen (2009); Venayagamoorthy and Fringer (2007). However, simulations with realistic temperature stratification are more limited (although explored by Vlasenko and Stashchuk (2007) for the Andaman Sea), while stratification is known to have a significant effect on these processes (Thiem and Berntsen, 2009; Venayagamoorthy and Fringer, 2007). Studies of breaking internal waves over a select range of paleobathymetries with realistic temperature stratification could aid understanding of locations of potential refugia in periods of global warming of the past and is a potential continuation of this work.

CHAPTER 6

CONCLUSIONS

Contents

An investigation of the influence of bathymetric features in varying forms and numbers on vertical mixing and internal wave generation and propagation in shallow coastal seas and epicontinental seas has been undertaken. While the influence of bathymetric features has previously been examined in the deep ocean, there is much less research in shallow continental seas where flow-bathymetry interaction is of increased importance. Therefore the use of intermediate scale real world scenarios such as the fjord/ loch test cases of Loch Etive and Idefjord has helped to elucidate the processes at work. Following on from this the acquired knowledge of these systems has been used to model the influence of the complex bathymetric features as they change through transgressions in icehouse and greenhouse worlds, in order to aid understanding of vertical mixing, nutrient heterogeneity and the localized efficiency of the oceanic carbon pump. Simulations of increasing complexity have been used throughout this thesis, showing the influence less idealised simulations with more complex bathymetry and stratification have.

Adaptive mesh simulations of Loch Etive have shown that fully unstructured mesh adaptivity can be used successfully to capture the dynamics of these turbulent systems via the use of simple test cases of an internal seiche and tilting tank to aid the choice of adaptivity parameters. These simple test cases reveal that, of the three varieties of mesh adaptivity used, fully unstructured adaptivity produces the best representation of the tilting tank dynamics while minimising spurious mixing for both the tilting tank and seiche simulations. The use of fully unstructured adaptive meshes, enables elements to align with features of the dynamics such as overturning rolls to maintain sharp interfaces and reduce the numerical diffusion while using slightly lower resolution in these areas than fixed meshes. The simple test cases are also used to compare two methods of mixing quantification, revealing that the method of Burchard and Rennau (2008) provides a robust and comparable method that is computationally cheap and simple to apply. For Loch Etive

fully unstructured mesh adaptivity results in a 1.5 orders of magnitude saving in the number of nodes used when compared with a fixed mesh with no telescoping of mesh resolution that has converged and show similar values and characteristics of $\overline{A}_{\mathrm{Ri}}$ (which shows the extent of the domain that is dynamically unstable and could result in turbulent mixing) and D^{phys} (a measure of the resulting physical mixing). In addition, studies of Loch Etive also present the importance of viscosity and mesh resolution for the capture of the complex turbulent overturning dynamics. For each viscosity value studied fully unstructured mesh adaptivity resulted in a saving in the number of nodes used while capturing the dynamics well. This indicates a considerable computational saving and should enable the modelling of more complex scenarios with greater ease. This chapter also shows that fixed structured mesh and fully unstructured adaptive mesh simulations reproduce the dynamics of Loch Etive well when compared with MITgcm simulations and real world physical studies.

While dynamics for a single sill system have been extensively investigated at a range of scales, the influence of multiple bathymetric features on dynamics was less well understood with ambiguities in the literature as to at which scale these features interact. Fixed mesh simulations of Idefjord have been used to elucidate the dynamics of such a system. These simulations reveal that for sub-critical systems, with respect to the Froude number, bathymetric features lead to the formation of a standing internal wave (or seiche) in the basin between the features. This seiche is present for all scenarios tested here where the sill spacing is less than the generated internal wave's wavelength and it exists due to the superposition of internal waves generated at the two features. A seiche generated in such a way while suggested by Xing and Davies (2011) has not been discovered previously in the literature. In addition to this, realistic simulations show good agreement with physical survey data from Johnsson et al. (2007) and reveal the dynamics of the

system to be those of a jet-basin including hydraulic jumps and high frequency internal waves. This is contrary to the wave-basin assumptions of (Johnsson et al., 2007) due to the shallower surface layer (due to high inflow rates of freshwater) at the time of survey. Reducing the height of the Svinesund sill changes the dynamics (as expected) to those similar to a single sill system. However, some notable differences are presented including the replacement of M_8 with M_6 frequencies. This occurs due superposition of an M_2 internal wave which forms on the remaining slope of the Svinesund sill and the M_4 internal wave generated at the Bjällvarp sill. Additionally removal of the Svinesund sill causes changes in the bed shear stress in the system, which may result in disturbance of toxic sediments on the Bjällvarp sill due to higher velocities in this region as less energy is removed from the system by the Svinesund sill. Further studies of the system would help explain these risks by potentially modelling a larger number of fjord bathymetric features.

As these chapters have illustrated the influence of singular or multiple bathymetric features in the present day, the final stage of this work applies this to the geological past. Fixed mesh MITgcm simulations were used with ten sinusoidal bathymetric features, dimensions based on the Welsh coastline and temperature stratification based on the modern Adriatic Sea. These demonstrated the importance of various stages of transgression in Ice-house and Green-house worlds. Ten bathymetric features were used as studies showed that the extra mixing generated after adding further bumps at this point is reduced. This was determined by considering the parameters: DN (the tidally averaged change in buoyancy frequency which is used to give an indication of mixing), Ri (the ratio of the vertical density stratification to the vertical shear in velocity which shows where flow becomes unstable) and A_{Ri} (the extent of the domain that is dynamically unstable and could result in turbulent mixing). Simulations of key stages during the two climate modes showed that due to sediment blanketing and smoother bathymetry during

a regression that more mixing occurs during a transgression. More specifically more mixing occurs during an early transgression when sea-level is lower causing higher sill top velocities and therefore more turbulent systems. More significant mixing also occurs in an icehouse than a greenhouse and this is due to lower levels of coastline rugosity due to reduced levels of eustatic variation. The early stages of transgression in an Ice-house world therefore have the roughest bathymetry and shallowest sea-level and therefore result in the most mixing. For a section of coastline during this period just 20 km long and 5 km wide (100 km^2) this could result in a similar flux of phosphorus into the surface layers (shown by the vertical advection rates of a passive tracer) as is annually exported by the Mississippi, which for the Mississippi creates a "dead zone" of around 10,000 km^2 (Joyce, 2000). Therefore it may be expected that during the early transgression of an Ice-house that the effects of oxygen depletion and black shale deposition in surface waters may be more acute. The role of topographic roughness on marine coastal productivity is clearly important and exerts an as yet poorly defined role on carbon cycling through time. Similar effects may well be seen due to transgressions over block faulted basins where the rough topography was fault bounded controlled rather than being generated by erosion.

Further implications of this work could be the alleviation of thermal stress on coral during periods of coral crisis in the geological past due to vertical mixing eroding temperature stratification in the area. This could explain the survival of reefs at low latitudes after extinctions events.

In summary the work in this thesis presents the importance of understanding the interaction between tides and bathymetric features via a variety of systems of varied complexity, where the importance of such interaction is presented in terms of changes in the system dynamics, nutrient heterogeneity and stratification of the area.

CHAPTER 7

BIBLIOGRAPHY

Contents

BIBLIOGRAPHY

Adcroft, A., Campin, J., Dutkiewicz, S., Evangelinos, C., Ferreira, D., Forget, G., Fox-Kemper, B., Heimbach, P., Hill, C., Hill, E., et al., December 2013. MITgcm Manual. MIT Department of EAPS, Cambridge, MA, USA, 4th Edition.

Adcroft, A., Hill, C., Marshall, J., 1997. Representation of topography by shaved cells in a height coordinate ocean model. Monthly Weather Review 125 (9), 2293–2315.

Allen, G., Simpson, J., 1998. Reflection of the internal tide in upper Loch Linnhe, a Scottish fjord. Estuarine, Coastal and Shelf Science 46 (5), 683–701.

Allison, P. A., Wells, M. R., 2006. Circulation in large ancient epicontinental seas: What was different and why? Palaios 21 (6), 513–515.

AMCG, 4 2015. Fluidity manual v4.1.12. Imperial College London.
URL 10.6084/m9.figshare.1387713.v2

Arakawa, A., Lamb, V. R., 1977. Computational design of the basic dynamical processes of the UCLA general circulation model. Methods in computational physics 17, 173–265.

Armi, L., 1978. Some evidence for boundary mixing in the deep ocean. Journal of Geophysical Research: Oceans (1978–2012) 83 (C4), 1971–1979.

Arneborg, L., Liljebladh, B., 2001. The internal seiches in gullmar fjord. part i: dynamics. Journal of Physical Oceanography 31 (9), 2549–2566.

Asteman, I. P., Hanslik, D., Nordberg, K., 2015. An almost completed pollution-recovery cycle reflected by sediment geochemistry and benthic foraminiferal assemblages in a Swedish–Norwegian Skagerrak fjord. Marine pollution bulletin 95 (1), 126–140.

Babuška, I., Suri, M., 1994. The p and h-p versions of the finite element method, basic principles and properties. SIAM review 36 (4), 578–632.

Balay, S., Abhyankar, S., Adams, M., Brown, J., Brune, P., Buschelman, K., Eijkhout, V., Gropp, W., Kaushik, D., Knepley, M., et al., 2014. PETSc users manual revision 3.5. Tech. rep., Argonne National Laboratory (ANL).

Balmforth, N., Ierley, G., Young, W., 2002. Tidal conversion by subcritical topography. Journal of Physical Oceanography 32 (10), 2900–2914.

Barad, M. F., Colella, P., Schladow, S. G., 2009. An adaptive cut-cell method for environmental fluid mechanics. International journal for numerical methods in fluids 60 (5), 473–514.

Batchelor, G. K., 2000. An introduction to fluid dynamics. Cambridge university press.

Baturin, G., 2003. Phosphorus cycle in the ocean. Lithology and Mineral Resources 38 (2), 101–119.

Bell, T., 1975. Topographically generated internal waves in the open ocean. Journal of Geophysical Research 80 (3), 320–327.

Berenbrock, C., Tranmer, A. W., 2008. Simulation of flow, sediment transport, and sediment mobility of the lower Coeur d'Alene River, Idaho. US Geological Survey.

Berger, M. J., Colella, P., 1989. Local adaptive mesh refinement for shock hydrodynamics. Journal of computational Physics 82 (1), 64–84.

Berntsen, J., Xing, J., Alendal, G., 2006. Assessment of non-hydrostatic ocean models using laboratory scale problems. Continental Shelf Research 26 (12), 1433–1447.

Berntsen, J., Xing, J., Davies, A., 2009. Numerical studies of flow over a sill: sensitivity of the non-hydrostatic effects to the grid size. Ocean Dynamics 59 (6), 1043–1059.

Berntsen, J., Xing, J., Davies, A. M., 2008. Numerical studies of internal waves at a sill: Sensitivity to horizontal grid size and subgrid scale closure. Continental Shelf Research 28 (10), 1376–1393.

Blum, M. D., Hattier-Womack, J., 2009. Climate change, sea-level change, and fluvial sediment supply to deepwater depositional systems. SEPM Special Publication 92, 15–39.

Bourgault, D., Kelley, D. E., 2007. On the reflectance of uniform slopes for normally incident interfacial solitary waves. Journal of physical oceanography 37 (5), 1156–1162.

Boussinesq, J., 1903. Théorie analytique de la chaleur: mise en harmonie avec la thermodynamique et avec la théorie mécanique de la lumière. Vol. 2. Gauthier-Villars.

Brattström, T., Carlsson, T., Jusslin, P., Lindström, C., 2015. Idefjorden - a basis and support for an Environmental impact assessment, regarding an expansion of the svinesund sills. Tech. rep., Department of Earth Sciences, University of Gothenburg.

Bryan, F., 1987. Parameter sensitivity of primitive equation ocean general circulation models. Journal of Physical Oceanography 17 (7), 970–985.

Budd, C. J., Huang, W., Russell, R. D., 2009. Adaptivity with moving grids. Acta Numerica 18, 111–241.

Burchard, H., Rennau, H., 2008. Comparative quantification of physically and numerically induced mixing in ocean models. Ocean Modelling 20 (3), 293–311.

Cacchione, D., Wunsch, C., 1974. Experimental study of internal waves over a slope. Journal of Fluid Mechanics 66 (02), 223–239.

Caldeira, K., Wickett, M. E., 2003. Oceanography: anthropogenic carbon and ocean pH. Nature 425 (6956), 365–365.

Castro-Diaz, M., 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.

Chappell, J., 1974. Geology of coral terraces, huon peninsula, new guinea: a study of quaternary tectonic movements and sea-level changes. Geological Society of America Bulletin 85 (4), 553–570.

Chappell, J., Shackleton, N., 1986. Oxygen isotopes and sea level.

Chen, D., Busalacchi, A. J., Rothstein, L. M., 1994. The roles of vertical mixing, solar radiation, and wind stress in a model simulation of the sea surface temper-

ature seasonal cycle in the tropical pacific ocean. Journal of geophysical research 99 (C10), 20345–20.

Chen, L., Sun, P., Xu, J., 2007. Optimal anisotropic meshes for minimizing interpolation errors in l^p-norm. Mathematics of Computation 76 (257), 179–204.

Collier, R. L., Leeder, M., Maynard, J., 1990. Transgressions and regressions; a model for the influence of tectonic subsidence, deposition and eustasy, with application to quaternary and carboniferous examples. Geological Magazine 127 (2), 117–128.

Crank, J., Nicolson, P., 1947. A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type. In: Mathematical Proceedings of the Cambridge Philosophical Society. Vol. 43. Cambridge Univ Press, pp. 50–67.

Cummins, P. F., 2000. Stratified flow over topography: time-dependent comparisons between model solutions and observations. Dynamics of atmospheres and oceans 33 (1), 43–72.

Cushman-Roisin, B., Beckers, J.-M., 2011. Introduction to geophysical fluid dynamics: physical and numerical aspects. Vol. 101. Academic Press.

Demaison, G., Moore, G. T., 1980. Anoxic environments and oil source bed genesis. AAPG Bulletin 64 (8), 1179–1209.

Devine, K., Boman, E., Heaphy, R., Hendrickson, B., Vaughan, C., 2002. Zoltan data management services for parallel dynamic applications. Computing in Science & Engineering 4 (2), 90–96.

Egbert, G., Ray, R., 2000. Significant dissipation of tidal energy in the deep ocean inferred from satellite altimeter data. Nature 405 (6788), 775–778.

Egbert, G. D., 1997. Tidal data inversion: Interpolation and inference. Progress in Oceanography 40 (1), 53–80.

Egbert, G. D., Ray, R. D., 2003. Semi-diurnal and diurnal tidal dissipation from TOPEX/Poseidon altimetry. Geophysical Research Letters 30 (17).

England, M. H., Maier-Reimer, E., 2001. Using chemical tracers to assess ocean models. Reviews of Geophysics 39 (1), 29–70.

Farrell, P., Maddison, J., 2011. Conservative interpolation between volume meshes by local Galerkin projection. Computer Methods in Applied Mechanics and Engineering 200 (1), 89–100.

Farrell, P., Piggott, M., Pain, C., Gorman, G., Wilson, C., 2009. Conservative interpolation between unstructured meshes via supermesh construction. Computer Methods in Applied Mechanics and Engineering 198 (33), 2632–2642.

Fitt, W. K., Brown, B. E., Warner, M. E., Dunne, R. P., 2001. Coral bleaching: interpretation of thermal tolerance limits and thermal thresholds in tropical corals. Coral Reefs 20 (1), 51–65.

Ford, R., Pain, C., Piggott, M., Goddard, A., De Oliveira, C., Umpleby, A., 2004. A nonhydrostatic finite-element model for three-dimensional stratified oceanic flows. part i: Model formulation. Monthly Weather Review 132 (12).

Frey, P.-J., Alauzet, F., 2005. Anisotropic mesh adaptation for CFD computations. Computer methods in applied mechanics and engineering 194 (48), 5068–5082.

Fringer, O., Gerritsen, M., Street, R., 2006. An unstructured-grid, finite-volume, nonhydrostatic, parallel coastal ocean simulator. Ocean Modelling 14 (3), 139–173.

Ganachaud, A., Wunsch, C., 2000. Improved estimates of global ocean circulation, heat transport and mixing from hydrographic data. Nature 408 (6811), 453–457.

Garrett, C., Kunze, E., 2007. Internal tide generation in the deep ocean. Annu. Rev. Fluid Mech. 39, 57–87.

Garrett, C., Munk, W., 1979. Internal waves in the ocean. Annual Review of Fluid Mechanics 11 (1), 339–369.

Gaspar, P., Gregoris, Y., Lefevre, J.-M., 1990. A simple eddy kinetic energy model for simulations of the oceanic vertical mixing: Tests at station papa and long-term upper ocean study site. Journal of Geophysical Research: Oceans (1978–2012) 95 (C9), 16179–16193.

Gayen, B., Sarkar, S., 2011a. Boundary mixing by density overturns in an internal tidal beam. Geophysical Research Letters 38 (14).

Gayen, B., Sarkar, S., 2011b. Boundary mixing by density overturns in an internal tidal beam. Geophysical Research Letters 38 (14).

Gent, P. R., McWilliams, J. C., 1990. Isopycnal mixing in ocean circulation models. Journal of Physical Oceanography 20 (1), 150–155.

George, P.-L., Borouchaki, H., 1998. Delaunay triangulation and meshing: application to finite elements. Hermes Paris.

Gerkema, T., Zimmerman, J., 2008. An introduction to internal waves. Lecture Notes, Royal NIOZ, Texel.

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.

Gill, A. E., 1982. Atmosphere-ocean dynamics. Vol. 30. Academic press.

Glynn, P. W., 1991. Coral reef bleaching in the 1980s and possible connections with global warming. Trends in Ecology & Evolution 6 (6), 175–179.

Goosse, H., Deleersnijder, E., Fichefet, T., England, M., 1999. Sensitivity of a global coupled ocean-sea ice model to the parameterization of vertical mixing. Journal of Geophysical Research: Oceans (1978–2012) 104 (C6), 13681–13695.

Gordon, A. L., 1986. Interocean exchange of thermocline water. Journal of Geophysical Research: Oceans (1978–2012) 91 (C4), 5037–5046.

Gorman, G., Pain, C., Piggott, M., Umpleby, A., Farrell, P., Maddison, J., 2009. Interleaved parallel tetrahedral mesh optimisation and dynamic load-balancing. Adaptive Modelling and Simulation, 101–104.

Green, J. M., Nycander, J., 2013. A comparison of tidal conversion parameterizations for tidal models. Journal of Physical Oceanography 43 (1), 104–119.

Gregg, M., 1989. Scaling turbulent dissipation in the thermocline. Journal of Geophysical Research: Oceans (1978–2012) 94 (C7), 9686–9698.

Griffies, S. M., Böning, C., Bryan, F. O., Chassignet, E. P., Gerdes, R., Hasumi, H., Hirst, A., Treguier, A.-M., Webb, D., 2000a. Developments in ocean climate modelling. Ocean Modelling 2 (3), 123–192.

Griffies, S. M., Pacanowski, R. C., Hallberg, R. W., 2000b. Spurious diapycnal mixing associated with advection in a z-coordinate ocean model. Monthly Weather Review 128 (3).

Guo, X., Gorman, G., Ashworth, M., Kramer, S., Piggott, M., Sunderland, A., 2010. High performance computing driven software development for next-generation modelling of the worlds oceans. Cray user group.

Gustafsson, M., Nordberg, K., 2002. The impact of climate and shore-level displacement on the late-holocene environmental development of Havstens Fjord and Koljö Fjord, Swedish west coast. The Holocene 12 (3), 325–338.

Haidvogel, D. B., Beckmann, A., 1999. Numerical ocean circulation modeling. Vol. 2. Imperial College Press London.

Haidvogel, D. B., Blanton, J., Kindle, J., Lynch, D., 2000. Coastal ocean modeling: processes and real-time systems. Oceanography Washington DC Oceanography Society 13 (1), 35–46.

Hassan, O., Probert, E., Morgan, K., 1998. Unstructured mesh procedures for the simulation of three-dimensional transient compressible inviscid flows with moving boundary components. International journal for numerical methods in fluids 27 (1-4), 41–55.

Hautmann, M., 2012. Extinction: End-Triassic Mass Extinction. eLS.

Hays, J. D., Imbrie, J., Shackleton, N. J., et al., 1976. Variations in the earth's orbit: pacemaker of the ice ages. American Association for the Advancement of Science.

Herndl, G. J., Reinthaler, T., 2013. Microbial control of the dark end of the biological pump. Nature Geoscience 6 (9), 718–724.

Hiester, H., Piggott, M., Allison, P., 2011. The impact of mesh adaptivity on the gravity current front speed in a two-dimensional lock-exchange. Ocean Modelling 38 (1), 1–21.

Hiester, H., Piggott, M., Farrell, P., Allison, P., 2014. Assessment of spurious mixing in adaptive mesh simulations of the two-dimensional lock-exchange. Ocean Modelling 73, 30–44.

Hiester, H. R., 2012. The application of adaptive mesh techniques to numerical simulations of gravity current flows. Ph.D. thesis, Imperial College London.

Hill, J., Piggott, M., Ham, D. A., Popova, E., Srokosz, M., 2012. On the performance of a generic length scale turbulence model within an adaptive finite element ocean model. Ocean Modelling 56, 1–15.

Hoegh-Guldberg, O., 1999. Climate change, coral bleaching and the future of the world's coral reefs. Marine and freshwater research 50 (8), 839–866.

Horn, D., Imberger, J., Ivey, G., 2001. The degeneration of large-scale interfacial gravity waves in lakes. Journal of Fluid Mechanics 434, 181–207.

Ilıcak, M., Adcroft, A. J., Griffies, S. M., Hallberg, R. W., 2012. Spurious dianeutral mixing and the role of momentum closure. Ocean Modelling 45, 37–58.

Ilıcak, M., Özgökmen, T. M., Özsoy, E., Fischer, P. F., 2009. Non-hydrostatic modeling of exchange flows across complex geometries. Ocean Modelling 29 (3), 159–175.

Imbrie, J., Hays, J. D., Martinson, D. G., McIntyre, A., Mix, A. C., Morley, J. J., Pisias, N. G., Prell, W. L., Shackleton, N. J., 1984. The orbital theory of pleistocene climate: Support from a revised chronology of the marine δ18 O record. In: Milankovitch and climate: Understanding the response to astronomical forcing. Vol. 1. p. 269.

Inall, M., Cottier, F., Griffiths, C., Rippeth, T., 2004. Sill dynamics and energy transformation in a jet fjord. Ocean Dynamics 54 (3), 307–314.
URL http://dx.doi.org/10.1007/s10236-003-0059-2

Inall, M., Rippeth, T., Griffiths, C., Wiles, P., 2005. Evolution and distribution of

tke production and dissipation within stratified flow over topography. Geophysical research letters 32 (8).

Inall, M. E., 2009. Internal wave induced dispersion and mixing on a sloping boundary. Geophysical Research Letters 36 (5).

Jacobs, D. K., Sahagian, D. L., 1995. Milankovitch fluctuations in sea level and recent trends in sea-level change: ice may not always be the answer. In: Sequence stratigraphy and depositional response to eustatic, tectonic and climatic forcing. Springer, pp. 329–366.

Johnsson, M., Green, J., Stigebrandt, A., 2007. Baroclinic wave drag from two closely spaced sills in a narrow fjord as inferred from basin water mixing. Journal of Geophysical Research: Oceans 112 (C11).

Joyce, S., 2000. The dead zones: oxygen-starved coastal waters. Environmental Health Perspectives 108 (3), A120.

Katz, B. J., 2005. Controlling factors on source rock development – a review of productivity, preservation, and sedimentation rate.

Khatiwala, S., 2003. Generation of internal tides in an ocean of finite depth: analytical and numerical calculations. Deep Sea Research Part I: Oceanographic Research Papers 50 (1), 3–21.

Kiessling, W., Roniewicz, E., Villier, L., Léonide, P., Struck, U., 2009. An early Hettangian coral reef in southern France: Implications for the end-Triassic reef crisis. Palaios 24 (10), 657–671.

Klemme, H., Ulmishek, G. F., 1991. Effective petroleum source rocks of the world: stratigraphic distribution and controlling depositional factors (1). AAPG Bulletin 75 (12), 1809–1851.

Kliem, N., Pietrzak, J. D., 1999. On the pressure gradient error in sigma coordinate ocean models: A comparison with a laboratory experiment. Journal of Geophysical Research: Oceans (1978–2012) 104 (C12), 29781–29799.

Klymak, J. M., Gregg, M. C., 2003. The role of upstream waves and a downstream density pool in the growth of lee waves: Stratified flow over the Knight Inlet sill. Journal of physical oceanography 33 (7).

Kramer, S. C., Cotter, C. J., Pain, C., 2010. Solving the poisson equation on small aspect ratio domains using unstructured meshes. Ocean Modelling 35 (3), 253–263.

Kuzmin, D., 2010. A vertex-based hierarchical slope limiter for p-adaptive discontinuous Galerkin methods. Journal of computational and applied mathematics 233 (12), 3077–3085.

Large, W. G., McWilliams, J. C., Doney, S. C., 1994. Oceanic vertical mixing: A review and a model with a nonlocal boundary layer parameterization. Reviews of Geophysics 32 (4), 363–403.

Lathuiliere, B., Marchal, D., 2009. Extinction, survival and recovery of corals from the Triassic to Middle Jurassic time. Terra Nova 21 (1), 57–66.

Laugier, M., Angot, P., Mortier, L., 1996. Nested grid methods for an ocean model: a comparative study. International journal for numerical methods in fluids 23 (11), 1163–1195.

Laurent, L. S., Stringer, S., Garrett, C., Perrault-Joncas, D., 2003. The generation of internal tides at abrupt topography. Deep Sea Research Part I: Oceanographic Research Papers 50 (8), 987–1003.

Ledwell, J. R., Watson, A. J., Law, C. S., 1998. Mixing of a tracer in the pycnocline. Journal of Geophysical Research: Oceans (1978–2012) 103 (C10), 21499–21529.

Legg, S., 2004a. Internal tides generated on a corrugated continental slope. part i: cross-slope barotropic forcing. Journal of Physical Oceanography 34 (1), 156–173.

Legg, S., 2004b. Internal tides generated on a corrugated continental slope. part ii: Along-slope barotropic forcing. Journal of Physical Oceanography 34 (8), 1824–1838.

Legg, S., Adcroft, A., 2003. Internal wave breaking at concave and convex continental slopes. Journal of Physical Oceanography 33 (11), 2224–2246.

Legg, S., Huijts, K. M., 2006. Preliminary simulations of internal waves and mixing generated by finite amplitude tidal flow over isolated topography. Deep Sea Research Part II: Topical Studies in Oceanography 53 (1), 140–156.

Legg, S., Klymak, J., 2008a. Internal hydraulic jumps and overturning generated by tidal flow over a tall steep ridge. Journal of Physical Oceanography 38 (9), 1949–1964.

Legg, S., Klymak, J., 2008b. Internal hydraulic jumps and overturning generated by tidal flow over a tall steep ridge. Journal of Physical Oceanography 38 (9), 1949–1964.

Llewellyn Smith, S. G., Young, W., 2002. Conversion of the barotropic tide. Journal of Physical Oceanography 32 (5), 1554–1566.

Longhurst, A., Sathyendranath, S., Platt, T., Caverhill, C., 1995. An estimate of global primary production in the ocean from satellite radiometer data. Journal of Plankton Research 17 (6), 1245–1271.

Loseille, A., Alauzet, F., 2011. Continuous mesh framework part i: well-posed continuous interpolation error. SIAM Journal on Numerical Analysis 49 (1), 38–60.

Maddison, J. R., Cotter, C. J., Farrell, P. E., 2011. Geostrophic balance preserving interpolation in mesh adaptive linearised shallow-water ocean modelling. Ocean Modelling 37 (1), 35–48.

Maderich, V., Brovchenko, I., Terletska, K., Hutter, K., 2012. Numerical simulations of the nonhydrostatic transformation of basin-scale internal gravity waves and wave-enhanced meromixis in lakes. In: Nonlinear Internal Waves in Lakes. Springer, pp. 193–276.

Manolopoulos, Y., Nanopoulos, A., Papadopoulos, A. N., Theodoridis, Y., 2010. R-trees: Theory and Applications. Springer Science & Business Media.

Marshall, J., Adcroft, A., Hill, C., Perelman, L., Heisey, C., 1997a. A finite-volume, incompressible navier stokes model for studies of the ocean on parallel computers. Journal of Geophysical Research: Oceans (1978–2012) 102 (C3), 5753–5766.

Marshall, J., Hill, C., Perelman, L., Adcroft, A., 1997b. Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling. Journal of Geophysical Research 102, 5733–5752.

Marshall, J., Hill, C., Perelman, L., Adcroft, A., 1997c. Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling. Journal of Geophysical Research: Oceans (1978–2012) 102 (C3), 5733–5752.

Martin, B., 2014. A numerical investigation of internal gravity wave phenomena with adaptive mesh techniques. Ph.D. thesis, Imperial College London.

McDougall, T. J., 1988. Some implications of ocean mixing for ocean modelling. Elsevier oceanography series 46, 21–35.

Miller, G. R., 1966. The flux of tidal energy out of the deep oceans. Journal of Geophysical Research 71 (10), 2485–2489.

Moore, G. T., Hayashida, D. N., Ross, C. A., 1993. Late Early Silurian (Wenlockian) general circulation model-generated upwelling, graptolitic black shales, and organic-rich source rocks – an accident of plate tectonics? Geology 21 (1), 17–20.

Moum, J., Smyth, W., 2001. Upper ocean mixing processes. Encyclopedia of Ocean Sciences 6, 3093–3100.

Munk, W., Wunsch, C., 1998. Abyssal recipes ii: Energetics of tidal and wind mixing. Deep-Sea Research Part I 45 (12), 1977–2010.

Nash, J. D., Kunze, E., Toole, J. M., Schmitt, R. W., 2004. Internal tide reflection and turbulent mixing on the continental slope. Journal of Physical Oceanography 34 (5), 1117–1134.

NOAA, 2016. World ocean database. `https://www.nodc.noaa.gov/OC5/SELECT/dbsearch/dbsearch.html`.

Nycander, J., 2005. Generation of internal waves in the deep ocean by tides. Journal of Geophysical Research: Oceans 110 (C10).

Otis, R. M., Schneidermann, N., 1997. A process for evaluating exploration prospects. AAPG bulletin 81 (7), 1087–1109.

Pacanowski, R. C., Philander, S. G. H., 1981. Parameterization of vertical mixing in numerical models of tropical oceans. Journal of Physical Oceanography 11 (11), 1443–1451.

Page, A., Williams, J. Z. M., Popov, L., 2007. Were transgressive black shales a negative feedback modulating glacioeustasy in the Early Palaeozoic Icehouse? Deep-time perspectives on climate change: Marrying the signal from computer models and biological proxies, 123.

Pain, C., Piggott, M., Goddard, A., Fang, F., Gorman, G., Marshall, D., Eaton, M., Power, P., De Oliveira, C., 2005. Three-dimensional unstructured mesh ocean modelling. Ocean Modelling 10 (1), 5–33.

Parkinson, S., Hill, J., Piggott, M., Allison, P., 2014. Direct numerical simulations of particle-laden density currents with adaptive, discontinuous finite elements. Geoscientific Model Development 7 (5), 1945–1960.

Parrish, J. T., 1982. Upwelling and petroleum source beds, with reference to Paleozoic. AAPG Bulletin 66 (6), 750–774.

Peraire, J., Persson, P.-O., 2008. The compact discontinuous Galerkin (CDG) method for elliptic problems. SIAM Journal on Scientific Computing 30 (4), 1806–1824.

Peraire, J., Vahdati, M., Morgan, K., Zienkiewicz, O. C., 1987. Adaptive remeshing for compressible flow computations. Journal of computational physics 72 (2), 449–466.

Piggott, M., Farrell, P., Wilson, C., Gorman, G., Pain, C., 2009. Anisotropic mesh adaptivity for multi-scale ocean modelling. Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 367 (1907), 4591–4611.

Piggott, M., Gorman, G., Pain, C., Allison, P., Candy, A., Martin, B., Wells, M., 2008a. 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.

Piggott, M. D., Pain, C. C., Gorman, G. J., Marshall, D. P., Killworth, P. D., 2008b. Unstructured adaptive meshes for ocean modeling. Ocean modeling in an eddying regime, 383–408.

Pingree, R., Griffiths, D., 1981. S 2 tidal simulations on the north-west european shelf. Journal of the Marine Biological Association of the United Kingdom 61 (03), 609–616.

Pittet, B., Strasser, A., 1998. Long-distance correlations by sequence stratigraphy and cyclostratigraphy: examples and implications (oxfordian from the swiss jura, spain, and normandy). Geologische Rundschau 86 (4), 852–874.

Polzin, K., Toole, J., Ledwell, J., Schmitt, R., 1997. Spatial variability of turbulent mixing in the abyssal ocean. Science 276 (5309), 93–96.

Qian, H., Shaw, P.-T., Ko, D. S., 2010. Generation of internal waves by barotropic tidal flow over a steep ridge. Deep Sea Research Part I: Oceanographic Research Papers 57 (12), 1521–1531.

Rahmstorf, S., 2002. Ocean circulation and climate during the past 120,000 years. Nature 419 (6903), 207–214.

Rahmstorf, S., 2003. Thermohaline circulation: The current climate. Nature 421 (6924), 699–699.

Redfield, A. C., 1934. On the proportions of organic derivatives in sea water and their relation to the composition of plankton. University Press of Liverpool James Johnstone memorial volume.

Redfield, A. C., 1963. The influence of organisms on the composition of sea-water. The sea, 26–77.

Satterley, A., 1996. The interpretation of cyclic successions of the Middle and Upper Triassic of the Northern and Southern Alps. Earth-Science Reviews 40 (3), 181–207.

Scourse, J., Austin, W., 2002. Quaternary shelf sea palaeoceanography: recent developments in europe. Marine Geology 191 (3), 87–94.

Shackleton, N., 1987. Oxygen isotopes, ice volume and sea level. Quaternary Science Reviews 6 (3), 183–190.

Sharples, J., Tweddle, J. F., Mattias Green, J., Palmer, M. R., Kim, Y.-N., Hickman, A. E., Holligan, P. M., Moore, C., Rippeth, T. P., Simpson, J. H., et al., 2007. Spring-neap modulation of internal tide mixing and vertical nitrate fluxes at a shelf edge in summer. Limnology and Oceanography 52 (5), 1735–1747.

Siedler, G., Church, J., Gould, J., 2001. Ocean Circulation and Climate: Observing and Modelling the Global Ocean. Academic Press.

Smethier Jr, W. M., 1981. Vertical mixing rates in fjords determined using radon and salinity as tracers. Estuarine, Coastal and Shelf Science 12 (2), 131–153.

Smith, R. C., Hill, J., Collins, G. S., Piggott, M. D., Kramer, S. C., Parkinson, S. D., Wilson, C., 2016. Comparing approaches for numerical modelling of tsunami generation by deformable submarine slides. Ocean Modelling 100, 125–140.

Spiegel, E., Veronis, G., 1960. On the boussinesq approximation for a compressible fluid. The Astrophysical Journal 131, 442.

Stacey, M. W., Zedel, L. J., 1986. The time-dependent hydraulic flow and dissipation over the still of observatory inlet. Journal of physical oceanography 16 (6), 1062–1076.

Stashchuk, N., Inall, M., Vlasenko, V., 2007. Analysis of supercritical stratified tidal flow in a scottish fjord. Journal of physical oceanography 37 (7), 1793–1810.

Stewart, R. H., September 2008a. Chapter 8 - equations of motion with viscosity. `http://oceanworld.tamu.edu/resources/ocng_textbook/chapter08/chapter08_05.htm`.

Stewart, R. H., 2008b. Introduction to physical oceanography. Citeseer.

Stigebrandt, A., 1999. Resistance to barotropic tidal flow in straits by baroclinic wave drag. Journal of Physical Oceanography 29 (2), 191–197.

Stigebrandt, A., Aure, J., 1989. Vertical mixing in basin waters of fjords. *Journal* of Physical Oceanography 19, 917–926.

Stüben, K., 2001. A review of algebraic multigrid. Journal of Computational and Applied Mathematics 128 (1), 281–309.

Thiem, Ø., Berntsen, J., 2009. Numerical studies of large-amplitude internal waves shoaling and breaking at shelf slopes. Ocean dynamics 59 (6), 937–952.

Thorpe, S. A., 2005. The turbulent ocean. Cambridge University Press.

Tissot, B., Deroo, G., Herbin, J., 1979. Organic matter in cretaceous sediments of the north atlantic: contribution to sedimentology and paleogeography. Deep Drilling Results in the Atlantic Ocean: Continental Margins and Paleoenvironment, 362–374.

Toggweiler, J., Key, R., 2003. Thermohaline circulation. Encyclopedia of Ocean Sciences.

Trottenberg, U., Oosterlee, C. W., Schuller, A., 2000. Multigrid. Academic press.

Tseng, Y.-h., Ferziger, J. H., 2001. Mixing and available potential energy in stratified flows. Physics of Fluids (1994-present) 13 (5), 1281–1293.

Tyrrell, T., 1999. The relative influences of nitrogen and phosphorus on oceanic primary production. Nature 400 (6744), 525–531.

Tyson, R., 2005. The" productivity versus preservation" controversy: cause, flaws, and resolution. Special Publication - SEPM 82, 17.

Van der Zwan, C., 2002. The impact of Milankovitch-scale climatic forcing on sediment supply. Sedimentary Geology 147 (3), 271–294.

Vasilevskii, Y. V., Lipnikov, K., 1999. An adaptive algorithm for quasioptimal mesh generation. Computational mathematics and mathematical physics 39 (9), 1468–1486.

Vellinga, M., Wood, R. A., 2002. Global climatic impacts of a collapse of the atlantic thermohaline circulation. Climatic change 54 (3), 251–267.

Venayagamoorthy, S., Fringer, O., 2006. Numerical simulations of the interaction of internal waves with a shelf break. Physics of Fluids (1994-present) 18 (7), 076603.

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

Vitousek, S., 2013. Towards internal wave resolving simulations of the ocean. Ph.D. thesis, Stanford University.

Vlasenko, V., Hutter, K., 2002. Numerical experiments on the breaking of solitary internal wavesover a slope-shelf topography. Journal of Physical Oceanography 32 (6), 1779–1793.

Vlasenko, V., Stashchuk, N., 2007. Three-dimensional shoaling of large-amplitude internal waves. Journal of Geophysical Research: Oceans 112 (C11).

Vlasenko, V., Stashchuk, N., Hutter, K., 2002. Water exchange in fjords induced by tidally generated internal lee waves. Dynamics of Atmospheres and Oceans 35, 63–89.

Vlasenko, V., Stashchuk, N., Inall, M. E., Hopkins, J. E., 2014. Tidal energy conversion in a global hot spot: On the 3-D dynamics of baroclinic tides at the celtic sea shelf break. Journal of Geophysical Research: Oceans 119 (6), 3249–3265.

Wall, M., Putchim, L., Schmidt, G., Jantzen, C., Khokiattiwong, S., Richter, C., 2015. Large-amplitude internal waves benefit corals during thermal stress. Proceedings of the Royal Society of London B: Biological Sciences 282 (1799), 20140650.

Walsh, J. J., 1988. On the nature of continental shelves. Vol. 508. Academic Press New York.

Walsh, J. J., Biscaye, P., Csanady, G., 1991. Importance of continental margins in the marine biogeochemical cycling of carbon and nitrogen. Nature 350 (6313), 53–55.

Walsh, J. J., Biscaye, P. E., Csanady, G. T., 1988. The 1983–1984 shelf edge exchange processes (SEEP) – i experiment: hypotheses and highlights. Continental Shelf Research 8 (5-7), 435–456.

Webb, D., Suginohara, N., 2001. Oceanography: Vertical mixing in the ocean. Nature 409 (6816), 37–37.

Wells, M. R., Allison, P. A., Piggott, M. D., Gorman, G. J., Hampson, G. J., Pain, C. C., Fang, F., 2007. Numerical modeling of tides in the late pennsylvanian midcontinent seaway of north america with implications for hydrography and sedimentation. Journal of Sedimentary Research 77 (10), 843–865.

Wells, M. R., Allison, P. A., Piggott, M. D., Pain, C. C., Hampson, G. J., De Oliveira, C. R., 2005. Large sea, small tides: the late carboniferous seaway of nw europe. Journal of the Geological Society 162 (3), 417–420.

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, 179–193.

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.

Wunsch, C., 1996. The ocean circulation inverse problem. Cambridge University Press.

Wunsch, C., Ferrari, R., 2004. Vertical mixing, energy, and the general circulation of the oceans. Annu. Rev. Fluid Mech. 36, 281–314.

Xing, J., Davies, A., 2006. Processes influencing tidal mixing in the region of sills. Geophysical Research Letters 33 (4).

Xing, J., Davies, A., 2009a. Influences of multiple sills upon internal wave generation and the implications for mixing. Geophysical Research Letters 36 (13), L13602.

Xing, J., Davies, A., 2010. The effects of large and small-scale topography upon internal waves and implications for tidally induced mixing in sill regions. Ocean Dynamics 60 (1), 1–25.

Xing, J., Davies, A. M., 2007. On the importance of non-hydrostatic processes in determining tidally induced mixing in sill regions. Continental Shelf Research 27 (16), 2162–2185.

Xing, J., Davies, A. M., 2009b. Influence of bottom frictional effects in sill regions upon lee wave generation and implications for internal mixing. Ocean dynamics 59 (6), 837–861.

Xing, J., Davies, A. M., 2011. On the interaction of internal tides over two adjacent sills in a fjord. Journal of Geophysical Research: Oceans 116 (C4).

APPENDIX A

THE INFLUENCE OF THE INCLUSION OF THE CORIOLIS FORCE FOR LOCH ETIVE

Contents

Results presented (Fig. A.0.1) compare Fluidity simulations run with and without the inclusion of the Coriolis force with viscosity ν_a (which results in the most complex dynamics studied in chapter 3). These simulation show that the main structure of the dynamics is very similar, with the same number of overturning rolls which propagate the same distance from the sill. The depth of these rolls is also very similar and the areas where rolls begin to converge is also comparable. However, there are small differences in the finer structures of the dynamics as simulations with the Coriolis force included seem to show more detailed dynamics with multiple vertical layers of the overturning rolls, which leads to a smaller area of $\overline{A}_{\mathrm{Ri}}$ at the lower extent of the rolls. Again in general the features of $\overline{A}_{\mathrm{Ri}}$ are very similar, however a slightly smaller critical area is present when the Coriolis force is included due to finer details of the dynamics. Horizontal and vertical velocities

show the same features, however amplitudes vary. For results without the Coriolis force the vertical velocities decrease in strength at medium extent from the sill before reaching a similar signal strength at the furthest extent of the rolls.

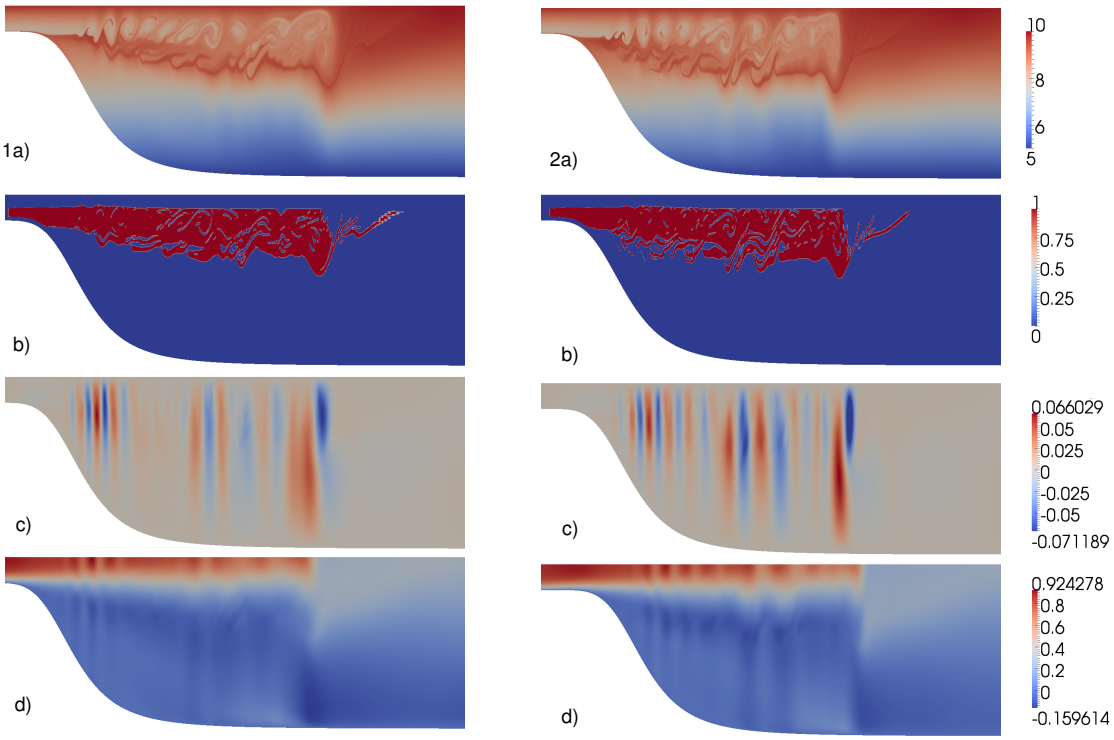


Figure A.0.1.: Comparison of Fludity results for Loch Etive without (1) and with (2) the application of the Coriolis force. Results are presented at the resolution required for convergence ($\Delta x = 1.25$ m, 250 layers and a timestep of $T/28800$ s) for temperature (a), $\overline{A}_{\mathrm{Ri}}$ (b), vertical velocity (c) and horizontal velocity (d) at $T/4$.

Since moderate and large scale features are comparable with and without the use of the Coriolis force, for further scenarios modelled it is possible to neglect the Coriolis force and still analyse changes in dynamics due to the variables explored in chapter 5.