

# An adaptive vertical coordinate for idealised and global ocean modelling

Angus Gibson

*Research School of Earth Sciences  
The Australian National University*

A thesis submitted for the degree of Doctor of Philosophy  
of The Australian National University

© Copyright by Angus Gibson 2019  
All Rights Reserved



---

# Preface

---

Thanks to the CFP (formerly GFD) group for being such a friendly hub within RSES. In particular, the student cohort is such a welcoming and diverse group of people. Of them, my officemates have been great company and great friends over the years.

Outside the CFP group, my friends, and especially my family have been endlessly supportive. They have kept me on-track, and more importantly, kept me sane by reminding me that a world exists outside of research.

Thanks to my advisory panel members, Callum Shakespeare, Andrew Kiss for their willingness to provide me with feedback despite our research specialisations not being completely aligned. My overseas collaborators, Alistair Adcroft, Bob Hallberg and Steve Griffies graciously hosted me on visits to their lab, and kindly devoted me attention and time from their busy schedules. Finally, my greatest thanks go to my supervisor, Andy Hogg, who has worked hard to support me on a project that is somewhat out of his comfort zone. His support, whether on scientific or administrative matters, has been invaluable, and his willingness to provide external resources was an immense asset and opportunity for development.

This research was supported by an Australian Government Research Training Program (RTP) Scholarship and the Australian Research Council (ARC) Centre of Excellence for Climate System Science (CE110001028). This research was undertaken with the assistance of resources and services from the National Computational Infrastructure (NCI), which is supported by the Australian Government.

Chapters 2 and 4 of this thesis comprise largely of work previously published by the author in Gibson et al. (2017). Except where otherwise indicated, this thesis is my own original work.



---

Angus Gibson  
March, 2019



---

# Abstract

---

Numerical models of ocean circulation are used in a wide variety of applications, for which speed, accuracy and memory requirements are a trade-off. The representations of physical flow features, such as deep ocean heat uptake and the meridional overturning circulation, are sensitive to the density structure of the ocean, which is determined by mixing and advection. While mixing itself is a physical phenomenon that is represented in the models, another source of mixing is numerical truncation errors, also known as spurious mixing. As a consequence, models with significant spurious mixing may poorly represent the flows being modelled.

We examine the separate contributions to spurious mixing from horizontal and vertical processes in an ocean model, MOM6, using reference potential energy (RPE). The RPE is a global diagnostic which can be used to estimate mixing between density classes. We extend this diagnostic to a sub-timestep timescale in order to individually separate contributions to spurious mixing through horizontal (tracer advection) and vertical (re-gridding/remapping) processes within the model. We both evaluate the overall spurious mixing in MOM6 against previously published output from other models (MOM5, MIT-gcm and MPAS-O), and investigate impacts on the components of spurious mixing in MOM6 across a suite of test cases: a lock exchange, internal wave propagation, and a baroclinically-unstable eddying channel.

Evaluating the independent contributions to spurious mixing in a model motivated development of a new vertical coordinate for ocean modelling. The goal of this coordinate is to provide local layers of approximately constant density to reduce numerical truncation errors in lateral advection, while maintaining sufficient vertical resolution to resolve surface interactions and vertical modes in weakly-stratified regions. This coordinate was then evaluated in two ways. First, spurious mixing was evaluated using the new split RPE diagnostic in idealised test cases. Using the new coordinate, spurious mixing was reduced in many of the idealised test cases compared to more commonly used coordinates, and in no cases was the spurious mixing increased. Second, the coordinate was evaluated in a more realistic global-scale simulation, comparing key physical metrics against pre-existing coordinates used in ocean modelling. This evaluation demonstrates the potential for the new coordinate to be used in ocean modelling, improving accuracy without a significant computational overhead.



---

# Contents

---

<b>Contents</b>	<b>7</b>
<b>List of Figures</b>	<b>9</b>
<b>List of Tables</b>	<b>11</b>
<b>1 Introduction</b>	<b>13</b>
1.1 Vertical coordinates . . . . .	14
1.2 Diagnosing spurious mixing . . . . .	17
1.3 Organisation of this thesis . . . . .	18
<b>2 The Split RPE method</b>	<b>21</b>
2.1 Reference potential energy . . . . .	21
2.2 The contributions to spurious mixing . . . . .	22
<b>3 AG: The adaptive grid coordinate</b>	<b>25</b>
3.1 Principles of AG . . . . .	27
3.2 Discretisation . . . . .	28
3.3 Density adaptivity . . . . .	29
3.4 Lateral smoothing . . . . .	33
3.5 Vertical restoring . . . . .	35
3.6 Convective grid adjustment . . . . .	37
3.7 Summary . . . . .	39
<b>4 Testing in idealised experiments</b>	<b>41</b>
4.1 Lock exchange . . . . .	42
4.2 Internal waves . . . . .	45
4.3 Baroclinic eddies . . . . .	48
4.4 Summary . . . . .	53
<b>5 Testing in a sector model</b>	<b>55</b>
5.1 Model configuration . . . . .	55
5.2 Stratification . . . . .	57
5.3 Meridional overturning circulation . . . . .	58
5.4 Circumpolar transport . . . . .	60
5.5 Computational performance . . . . .	63
5.6 Summary . . . . .	63
<b>6 Conclusions</b>	<b>67</b>
6.1 Summary . . . . .	67
6.2 Future work . . . . .	69
<b>7 References</b>	<b>71</b>





---

# List of Figures

---

1.1	Schematic of the quasi-Lagrangian ALE algorithm <i>Fig. credit: Alistair Adcroft</i> . . . . .	17
2.1	A schematic demonstrating the ability for regridding/remapping to cause a decrease in RPE when using a remapping scheme with a higher reconstruction order than piecewise constant. . . . .	23
3.1	Horizontal grid staggering in an Arakawa C-grid . . . . .	29
3.2	Schematic of density adaptivity in the adaptive grid coordinate . . . . .	30
3.3	Demonstration of interface tangling when limiting is not applied to regridding	32
3.4	Interface height flux due to lateral smoothing in the adaptive grid coordinate	34
3.5	Interface tangling in the adaptive grid without the barotropic smoothing term . . . . .	36
3.6	Demonstration of the effect of the adjustment term on a gravity current in the adaptive grid coordinate . . . . .	38
4.1	Snapshots of the idealised lock exchange test case in MOM6, using the $z$ -star coordinate . . . . .	43
4.2	(a) Normalised RPE evolution in the idealised lock exchange test case. (b) Instantaneous rate of RPE change at 17h in the lock exchange . . . . .	43
4.3	Horizontal and vertical contributions to the instantaneous rate of RPE change in MOM6 in the lock exchange at 17h. . . . .	44
4.4	Snapshots of the idealised internal wave test case initial conditions . . . . .	46
4.5	(a) Averaged rate of RPE change from 10 to 100 days in the internal waves test case. (b) Contributions to spurious mixing by horizontal and vertical processes in the internal waves test case in MOM6. . . . .	47
4.6	Snapshots of the idealised baroclinic eddies test case initial condition and surface state after 320 days . . . . .	49
4.7	Average rate of RPE change for all baroclinic eddy experiments over the entire 320 day run . . . . .	50
4.8	Average rates of RPE change for horizontal and vertical processes in MOM6 for the baroclinic eddy experiment . . . . .	52
5.1	Surface forcing and topography used for the sector simulation . . . . .	56
5.2	Zonal mean density structure (10-year average) in the sector simulation . . . . .	58
5.3	A schematic view of the meridional overturning circulation and the mechanisms that drive it <i>Fig. credit: Robert Marsh, National Oceanography Centre Southampton</i> . . . . .	59
5.4	Meridional overturning circulation (10-year average) in the sector simulation	61
5.5	Strength of circumpolar transport compared against coordinates in the sector simulation . . . . .	62



---

# List of Tables

---

4.1	Configuration for AG coordinate in the idealised experiments . . . . .	41
5.1	Configuration for AG coordinate in the sector model . . . . .	57
5.2	Walltime performance for running 24-months of the sector model simulation	63



---

# Introduction

---

One of the myriad uses of ocean models is in developing ocean heat uptake estimates and overturning circulation predictions (Armour et al., 2016). Additionally, the overturning circulation itself affects the wider climate, which manifests when ocean models are used as a component of coupled climate simulations. The strength of ocean heat uptake and the overturning circulation are both strongly controlled by the density structure of the ocean, which is determined by mixing and advection. For example, mixing at depth modifies the abyssal overturning cell that constitutes part of the meridional overturning circulation (Mashayek et al., 2015), while the time scale of adjustment of the overturning circulation toward equilibrium is sensitive to near-surface mixing (Vreugdenhil et al., 2015). A consequence of this sensitivity is that ocean models with significant mixing due to numerical truncation errors (spurious mixing) are unlikely to accurately constrain the abyssal overturning, heat transport, or heat uptake.

Numerical ocean models are governed by approximations to the Reynolds-averaged incompressible Navier-Stokes equations for momentum, also known as the primitive equations (Griffies, 2004). In these models, the vertical balance is hydrostatic, where the vertical pressure gradient force is matched by the gravitational force. The mixing of momentum by the unresolved eddy field from the mesoscale down to the Kolmogorov scale is parameterised by an explicit eddy viscosity term. Potential density of water parcels is a function of salinity and potential temperature through an equation of state. Alternative density representations may require different tracer quantities, such as conservative temperature, depending on the equation of state. These tracers are advected by the explicitly resolved eddy field, and mixed by the unresolved eddy field through a parameterised eddy diffusivity term. Due to the vast difference between lateral and vertical scales in models, the eddy viscosity and diffusivity terms are often separated into lateral and vertical components.

Although ocean models all fundamentally solve the same equations, there are still vast differences between the various models. Considering the primitive equations alone, numerical discretisations and algorithms often differ. For example, the partial differential equations are often solved in a finite volume framework to ensure volume conservation, the volume integral of tracers and occasionally energy. However, spectral or finite element methods can also be employed to solve the underlying equations. On top of the primitive equations may be various parameterisations that aim to represent the dynamical effects of unresolved (due to limited spatial or temporal resolution) processes. In this thesis, we focus on development on the MOM6 model, however previously published results from MOM5 (Griffies, 2012), MITgcm (Marshall et al., 1997) and MPAS-O (Ringler et al., 2013) are used in chapter 4 for the purpose of comparison.

MOM6 (Adcroft et al.) is the latest generation of the Modular Ocean Model, developed primarily by Alistair Adcroft (Program in Atmospheric and Oceanic Sciences, Princeton

University), Robert Hallberg and Stephen Griffies (National Oceanic and Atmospheric Administration/Geophysical Fluid Dynamics Laboratory). It differs significantly from the previous generation, MOM5, by using the Arbitrary Lagrangian-Eulerian algorithm, allowing the use of generalised vertical coordinates. Additionally, MOM6 uses the Arakawa C-grid, instead of the B-grid previously used in MOM5. MOM6 is intended for use in both idealised and realistic configurations, and may be run as a standalone ocean-only code, or as a component of a coupled model. Importantly, MOM6 is open source under the GNU LGPLv3 license, and publically developed (<https://github.com/NOAA-GFDL/MOM6>).

To solve the primitive equations, ocean models implement some form of discretisation, such as the finite volume method. This discretisation involves representing the computational domain as a series of grid cells in three-dimensional space, where each grid cell has associated velocities and tracer concentrations (e.g. temperature, salinity), and possibly higher moments (Prather, 1986), all of which are averaged over the spatial extent of the cell. Horizontal tracer advection schemes are discretisations of the advection equation that use information from neighbouring grid cells to create higher-order reconstructions of the tracer field than that which is stored directly in the cell. Flux of tracers between the cells occurs due to physical processes such as advection and mixing. In ocean models, mixing has two main causes, physical and numerical. The physical mixing comes from advection by numerically unresolved turbulence, which is typically parameterised as a diffusive process. On the other hand, numerical mixing arises from truncation errors in the discretisations and algorithms used by the ocean model to solve the governing equations. Numerical mixing is also known as spurious mixing and has no physical basis. For example, first-order upwind advection has numerical diffusion as the leading error term (Gentry et al., 1966).

Spurious mixing is undesirable in ocean models as it is unphysical and may add to the imposed and parameterised mixing to an unknown extent. Spurious mixing particularly affects numerical experiments which are contingent on the density structure of the ocean. Ocean heat uptake or overturning circulation strength in such experiments may be biased (Griffies et al., 2014). One of the considerations in model development and configuration is thus to ensure spurious mixing is minimised.

The magnitude of spurious mixing is strongly influenced by the choice of horizontal tracer advection scheme. Much of the focus in reducing spurious mixing has therefore been on tracer advection, through improving numerical accuracy or the model's tracer sub-grid scale representations. Some argue that a high-order advection scheme is sufficient to reduce the spurious mixing to acceptable levels (Daru and Tenaud, 2004). This is simply a matter of using a sufficiently high-order polynomial reconstruction to capture the structure of tracer distributions within grid cells. Other advection schemes attempt to preserve the sub-grid scale representation of a given field. For example, by carrying information about both first and second-order moments, the Prather (1986) method is able to reconstruct a field to second order. This second-order moment scheme must often be used in conjunction with a flux limiter to avoid the creation of spurious minima and maxima; these limiters in effect lead to a sub-cell diffusion (Morales Maqueda and Holloway, 2006). An alternative view is that the tracer advection scheme only needs sufficient accuracy before grid-scale noise in velocity becomes the dominant source of spurious mixing (Ilicak et al., 2012).

## 1.1 Vertical coordinates

A significant consideration in model configuration in order to minimise spurious mixing is the vertical coordinate. We first describe some of the main choices for the vertical coordinate in ocean models. The basis for the  $z$ -family of coordinates is the pure  $z$ -

level coordinate, where coordinate surfaces are simply fixed geopotential heights. The first extension to the  $z$ -level coordinate is  $z^*$  or  $z$ -star, which individually and uniformly expands or contracts water columns to accommodate changes in the free surface height (Adcroft and Campin, 2004).  $z$ -family coordinates allow for ahead-of-time specification of vertical resolution, which must be applied to the entire modelled domain and thus must be sufficiently general. Some disadvantages of these coordinates are poor representation of overflows (Legg et al., 2009), and the spurious diapycnal mixing associated with purely horizontal coordinate surfaces, e.g. isoneutral diffusion (Griffies et al., 2000).

An alternate choice of vertical coordinate is the isopycnal coordinate. Instead of being referenced to physical positions, isopycnal coordinate models use potential density as the vertical coordinate. This formulation provides enhanced vertical resolution where vertical density gradients are large. It also completely eliminates spurious diapycnal mixing, because there are no across-coordinate velocities (by construction), so no vertical inter-layer advection takes place. The diapycnal mixing is achieved through explicit mixing processes, and is thus not attributed to the isopycnal coordinate itself. Additionally, since density is uniform within layers there is no lateral density advection either. However, there are difficulties in representing the nonlinear equation of state, as there is no conservative density coordinate that is monotonic with depth (Griffies, 2004). Additionally, the surface mixed layer is essentially unstratified, and hence is a region of very low vertical resolution.

Similar to  $z$ -star is the terrain-following sigma coordinate, which is often used for coastal modelling. Instead of being defined relative to a fixed geopotential reference height, the sigma coordinate is scaled by the depth of local topography. By scaling in this way, all the available vertical resolution is used, even in shallow coastal shelves. However, as the coordinate surfaces slope along with topography, there have been problems with pressure gradient errors (Haney, 1991) and the implementation of neutral physics, which have hampered its adoption in large scale ocean modelling (Griffies, 2004).

Hybrid vertical coordinates combine or modify other vertical coordinates to optimise their performance, but incur additional complexity and computational cost. The coordinates may additionally be *adaptive*, where the structure is modified in response to the model state, taking into account factors such as stratification or velocity. One hybrid vertical coordinate is  $z$ -tilde (Leclair and Madec, 2011), which has Lagrangian behaviour (i.e., the grid is advected by the vertical velocity) for motions on short timescales, but relaxes to a target  $z$ -star grid over long timescales to prevent the grid from drifting. This scheme was demonstrated to reduce spurious mixing when modelling the propagation of internal gravity waves. Another example of hybrid coordinates is the continuous isopycnal coordinate (White et al., 2009), where instead of layers having a predefined density as in the pure isopycnal coordinate, interfaces have a target density. In this case, there must be dynamic adjustment of the coordinate surfaces in order to maintain the target density. The release of the constraint to layered isopycnals means that further physical processes can be more easily added to the model, such as geothermal heating or double diffusion (White et al., 2009). In isolation, each coordinate has strengths and weaknesses for ocean modelling, but the combination attempts to preserve the strengths of each.

A hybrid and adaptive vertical coordinate is used in production in MOM6. This scheme is called *HYCOM1*, as it aims to emulate the coordinate used in the HYCOM model (Bleck, 2002). This model solves the primitive equations for a generalised vertical coordinate, using a *grid generator* to determine the vertical coordinate for each column. The grid generator aims to form an isopycnal coordinate by moving interfaces and “unmixing” between adjacent layers. However, to make better use of vertical resolution, layers also have an enforced minimum thickness (Bleck, 2002). In particular, this minimum thickness prevents massless cells of zero- or near-zero thickness near the surface by sacrificing an

exactly isopycnal-following coordinate.


The HYCOM1 coordinate in MOM6 follows a similar principle, with two key differences. Firstly, instead of maintaining *per-layer* isopycnals, interfaces are moved such that they match a target density (by the same mechanism as the continuous isopycnal coordinate, above). Secondly, interfaces are enforced to have a minimum depth, in addition to a minimum layer thickness. However, the minimum thickness is implemented for numerical reasons rather than being a principle of the coordinate. For example, Bleck (2002) suggests 10 m as a minimum layer thickness, but a value on the order of 1 mm is more common in MOM6. Otherwise, minimum interface depths are equivalent to minimum layer thicknesses in the case where successive surface layers are massless. One of the main reasons for the differences in this coordinate between MOM6 and HYCOM is that the use of the ALE scheme (described below) prevents the need for an “unmixing” algorithm for adjusting the properties of adjacent layers to alter the vertical grid.

To allow generalised vertical coordinates, models can make use of an Arbitrary Lagrangian-Eulerian (ALE) scheme. There are two general implementations of ALE in ocean models, depending on the reference frame of the model (Margolin and Shashkov, 2003; Leclair and Madec, 2011). In quasi-Eulerian models, any changes in the vertical grid due to the choice of coordinate are incorporated into the solution of the primitive equations (Kasahara, 1974). Incorporating changes in the vertical grid is often done by calculating the motion of the new vertical grid relative to the old grid as a vertical velocity. As such, there could be an associated spurious mixing with advection in both the horizontal and vertical directions.

The quasi-Lagrangian algorithm (Hirt et al., 1974; Bleck, 2002) is for models which are primarily implemented in a Lagrangian frame of reference in the vertical (i.e. the interfaces primarily move with the vertical velocity). The algorithm is illustrated in fig. 1.1. In the Lagrangian frame of reference, the vertical grid may adjust during the dynamic solution of the primitive equations or as a consequence of parameterisations such as Gent-McWilliams thickness diffusion (Gent and McWilliams, 1990). This dynamical timestep incorporates most of the modelled processes, including the calculation and application of advective tracer fluxes. Typically, these fluxes are accumulated during the solution of the primitive equations and applied after the primitive equations are solved, on the updated grid. The model dynamics are then followed by the ALE timestep, which consists of two phases. In the first phase, referred to here as *regridding*, a new vertical grid is calculated using the current model state. This new grid may be as simple as a prescribed  $z$ -star coordinate, or could be a function of local density or depth. Secondly, the new grid is applied in the *remapping* phase, during which the model state is mapped onto the new grid. The remapping algorithm is often an adaptation of an advection scheme (Margolin and Shashkov, 2003), although other conservative algorithms may be used. Remapping differs from vertical advection in that the effective vertical velocities have a non-physical component to recover the new grid, as well as a physical component. Spurious mixing that occurs during the remapping phase therefore depends on the vertical dynamics, the new grid and the sub-gridscale reconstruction of tracers on the old grid.

The accuracy of the reconstruction scheme used in the remapping stage of ALE was investigated by White and Adcroft (2008) with their piecewise quartic method (PQM). PQM is the most accurate reconstruction method available in MOM6, and was found to significantly increase reconstruction accuracy for a small increase in computational cost compared to limited PPM (piecewise parabolic method). The impacts of different reconstruction schemes in regridding and remapping were considered by White et al. (2009), comparing their spurious mixing in terms of the change of volume distributions across density classes. Neither of these studies quantified the magnitude of spurious mixing in





**Figure 1.1:** The quasi-Lagrangian ALE algorithm. During *dynamics*, advection modifies layer thicknesses  $h^n$  (and therefore interface heights  $z^n$ ) and tracers  $\theta^n$ . The new grid  $z^{n+1}$  (red) is generated from the post-dynamics state,  $h^\dagger$  and  $\theta^\dagger$ . Remapping consists of *reconstruction* and *average* steps onto the new grid by generating a higher-order reconstruction (blue) and moving subcells (grey). *Fig. credit: Alistair Adcroft, Program in Atmospheric and Oceanic Sciences, Princeton University*

total, or as a comparison to the spurious mixing by horizontal advection. Formulating this comparison is one of the aims of this thesis.

## 1.2 Diagnosing spurious mixing

There is no consensus on the appropriate diagnostic technique to use to evaluate the performance of numerical schemes with regard to spurious mixing. The goal of diagnosing spurious mixing is to provide a quantification of the contribution of numerical errors to the overall solution. Griffies et al. (2000) used an effective diapycnal diffusivity, which allows for direct comparison between the spurious mixing and expected oceanic values. However, because it uses a reference density profile compiled from the entire domain, the effective diffusivity is only a single idealised vertical profile, and cannot be mapped back to real space.

An alternative to diagnosing spurious mixing from the model state is to calculate an analytical solution from the advection operator itself. Morales Maqueda and Holloway (2006) used this approach for upstream based schemes, such as the second-order moment method of Prather (1986), calculating a closed form expression for the implicit numerical diffusivity.

Substituting the second-order moment scheme for an arbitrary choice of horizontal advection scheme, Burchard and Rennau (2008) inferred the effect of sub-gridscale structure on the destruction of the spatial variance of a tracer, which was adapted to a finite volume framework by Klingbeil et al. (2014). A tracer variance diagnostic thus compares physical and numerical mixing through sub-gridscale changes. Tracer variance can be calculated for every model gridpoint, and its destruction gives information about the relative impact of physical and numerical mixing through full space.

An alternative diagnostic of spurious mixing is to measure its effect on the reference potential energy RPE (Winters et al., 1995) when all explicit diapycnal terms are set to

zero. This metric gives only timeseries data (no localised information), but allows for ready comparison across models with the same physical configuration. Additionally, it is desirable to use a density-based metric in the context of ocean modelling, especially because of its impact on key metrics such as heat uptake. The comparison of RPE across models allows for a quantitative evaluation of algorithms or parameters that differ across models, or within the same model. The global RPE diagnostic can be extended to provide a localised density of reference potential energy (Ilicak, 2016), alleviating one of the shortcomings of the global quantity, and showing, for example, that regions of high shear are the most susceptible to spurious mixing.

Ilicak et al. (2012) used the rate of change of RPE in analysing the role of momentum closure in different models (GOLD, MITgcm, MOM5 and ROMS). Comparisons were performed across a suite of test cases intended to stress different physical phenomena: a lock exchange, downslope flow, internal gravity waves, baroclinic eddies, and a global spindown. Ilicak et al. (2012) studied the dependence of spurious mixing on the lateral grid Reynolds number

$$\text{Re}_\Delta = \frac{U\Delta x}{\nu_h}, \quad (1.1)$$

where  $U$  is the characteristic horizontal velocity scale,  $\Delta x$  is the horizontal grid spacing and  $\nu_h$  is the horizontal viscosity. For the dissipation of spurious grid-scale noise in the velocity field, it can be shown the lateral grid Reynolds number should be less than 2 when a centred difference advection scheme is employed (Griffies, 2004, p. 410). By varying the horizontal viscosity, Ilicak et al. (2012) showed that spurious mixing increases with the lateral grid Reynolds number until it is dominated by grid-scale noise and becomes saturated at a high level of spurious mixing. It was recommended that  $\text{Re}_\Delta$  be constrained to be less than 10 to avoid the saturation threshold, although this was not a sufficient condition. Flow-dependent schemes such as Smagorinsky viscosity can be used to locally dissipate grid-scale noise, while maintaining a lower overall viscosity.

To look at the performance of a model with an ALE scheme, Petersen et al. (2015) extended the study of Ilicak et al. (2012). As well as the  $z$ -star and isopycnal coordinates, three additional vertical coordinates were used to demonstrate the ALE scheme in the MPAS-Ocean (Ringler et al., 2013, hereafter referred to as MPAS-O) model: terrain-following sigma coordinate,  $z$ -level, and  $z$ -tilde (Leclair and Madec, 2011). To compare with another model with a  $z$ -level vertical coordinate, POP was also added to the suite of models. As MPAS-O uses the quasi-Eulerian implementation of ALE, there is a resolved transport across vertical layer interfaces during tracer advection. Use of  $z$ -tilde as the vertical coordinate leads to a reduction in this transport, and thereby reduces spurious diapycnal mixing. However, the model timestep had to be halved in global simulations with  $z$ -tilde, which has a significant impact on computational cost. Additionally, the  $z$ -tilde coordinate was shown to be unsuitable for simulations with large, transient flows. Thus, while spurious mixing in ocean models can be reduced by the choice of vertical coordinate, further development and evaluation of the available vertical coordinates is required.

### 1.3 Organisation of this thesis

This thesis presents a new vertical coordinate for use in numerical ocean modelling. This coordinate aims to leverage the ability to impose generalised vertical coordinates in the ALE framework to provide a more accurate representation of flow features. We first derive a method for analysing spurious mixing in an ALE model before introducing the

---

coordinate and investigating it through a suite of idealised test cases. The coordinate is then validated in a realistic simulation against other vertical coordinates.

Chapter 2 presents a new method for diagnosing spurious mixing in an ALE model. This method is based on previously established diagnostics that attribute spurious mixing to increases in the reference potential energy (Ilıcak et al., 2012). The crucial underpinning of this thesis, the adaptive grid coordinate, is introduced in chapter 3. This coordinate is defined as an equation governing the evolution of model coordinate interfaces. The split RPE technique is demonstrated in chapter 4 in evaluating the performance of MOM6 as an ALE model, and the adaptive grid coordinate within the model. A final evaluation of the adaptive grid coordinate is then presented in chapter 5, which demonstrates the improvements in a realistic circulation using the adaptive grid coordinate compared to other coordinates used in ocean modelling. Conclusions follow in chapter 6.



---

# The Split RPE method

---

In this chapter we review the computation of reference potential energy and its application to diagnosing spurious mixing. We then expand on this theory, outlining a new method for separately calculating the individual contributions to spurious mixing within a model with the split RPE method.

## 2.1 Reference potential energy

A measurement of spurious mixing with a physical basis comes from the reference potential energy (RPE, known also as background potential energy; Winters et al., 1995). The RPE is the lowest potential energy attainable by rearrangement of a stratified fluid so that there is no energy available for motion. To calculate this state, the fluid must be adiabatically re-sorted to a stable one-dimensional stratification. The sum of the reference potential energy and the available potential energy (APE) gives the total potential energy in the system. Mathematically, the RPE is expressed as

$$\text{RPE} = g \int_{\Omega} z \rho^*(z) \, dV, \quad (2.1)$$

where  $g$  is the gravitational constant;  $z$  is the height, positive upward;  $\Omega$  is the domain;  $V$  is the volume; and  $\rho^*$  is the adiabatically re-sorted density profile, known as the reference density. There are some subtleties when using a nonlinear or compressible equation of state (Saenz et al., 2015), or domains with complex topography (Stewart et al., 2014). Here, we will only consider a linear and incompressible equation of state.

In an ideal numerical model of the hydrostatic primitive equations that exhibits identically zero mixing, with no buoyancy forcing, every fluid parcel should maintain its temperature and salinity. Under an incompressible equation of state, the density of any fluid parcel should thus be constant. In this case, the adiabatic resorting of the entire fluid yields an unchanging reference density profile and therefore constant RPE, regardless of the actual depth or location of the parcels. Thus, when all parameterised mixing is disabled, and there is no buoyancy forcing, a model with zero spurious mixing will have constant RPE.

Most ocean models experience spurious mixing under nontrivial conditions. By designing experiments with all explicit mixing disabled and without buoyancy forcing, any increase in RPE over time can be attributed to spurious mixing as a result of the model's numerics. For example, a limited advection scheme (which is unable to create water densities outside the existing range) will create some intermediate density class between two pre-existing densities through spurious mixing. The resulting reference density profile will have a higher centre of mass than prior to mixing, which manifests as an increase in RPE.

Conversely, RPE can only be decreased by lowering the centre of mass of the adiabatically re-sorted fluid. A disadvantage of the RPE technique is that in order to obtain a meaningful result, explicit mixing and buoyancy forcing must be disabled. Additionally, it gives only a global metric with no localisation information.

## 2.2 The contributions to spurious mixing

In many models, there is an explicit distinction between the horizontal and vertical dynamics, particularly in models with a generalised vertical coordinate, such as those employing the ALE algorithm. MOM6 performs its regridding/remapping implementation of ALE in a distinct part of a timestep, after the horizontal dynamics have been resolved. Taking advantage of this distinction, we can diagnose the instantaneous RPE at multiple points during a single timestep: at the beginning of a timestep, after horizontal dynamics, and after regridding/remapping has been performed. We denote these three calculations of instantaneous RPE as  $\text{RPE}_i$ ,  $\text{RPE}_h$  and  $\text{RPE}_v$ , respectively. The change in RPE due to horizontal dynamics is simply  $\text{RPE}_h - \text{RPE}_i$ , and similarly the change due to regridding/remapping is  $\text{RPE}_v - \text{RPE}_h$ .

Previous analyses of spurious mixing through changes in RPE have only used a time-series of RPE taken from a single stage in a timestep. This means that spurious mixing is only diagnosed for the model as a whole, and cannot be attributed to specific algorithms in the model. Our technique of separating RPE contributions allows us to address the interplay of different components, such as the order of horizontal advection scheme chosen, the order of accuracy of vertical reconstruction used for remapping, and the specific impact of the choice of vertical coordinate.

### 2.2.1 A caveat on the vertical component of spurious mixing


From a physical viewpoint, we expect mixing processes to produce a monotonic RPE increase with time. However, the vertical process of regridding/remapping can cause RPE to decrease in cases with sub-grid scale representation of tracer that is higher order than piecewise constant (PCM). We use a simple example to demonstrate how the combination of regridding/remapping may create a decrease in total potential energy which, for a statically stable single column case, is equivalent to the RPE.

Figure 2.1 shows a simple two-cell domain under regridding/remapping. For each cell  $i$ , the model carries both the mean density  $\bar{\rho}_i$  and thickness  $h_i$ . From this information, sub-grid reconstructions of density  $\rho_i(z)$  are computed. Regridding moves the interface between the cells from its initial position at  $z = h_1$  to the dashed line at  $z' = h_1 - \Delta h$ , and remapping mixes the integrated density  $M$  between the old and new interfaces, in this case from the lower cell to the upper cell. This integrated density is

$$M = \int_{h_1 - \Delta h}^{h_1} \rho_1(z) dz. \quad (2.2)$$

The potential energy of the domain is calculated using the information carried by the model, namely the cell mean density and cell thickness, and ignores the sub-grid reconstruction. Initially, the potential energy per unit area in the column in fig. 2.1 is

$$PE_i = \bar{\rho}_1 h_1 \frac{h_1}{2} + \bar{\rho}_2 h_2 \left( h_1 + \frac{h_2}{2} \right). \quad (2.3)$$



**Figure 2.1:** A schematic demonstrating the ability for regridding/remapping to cause a decrease in RPE when using a remapping scheme with a higher reconstruction order than piecewise constant.

After remapping, the potential energy becomes

$$PE_f = (\bar{\rho}_1 h_1 - M) \frac{h_1 - \Delta h}{2} + (\bar{\rho}_2 h_2 + M) \left( h_1 - \Delta h + \frac{h_2 + \Delta h}{2} \right). \quad (2.4)$$

Taking the difference between the final and initial potential energy gives the change due to regridding/remapping

$$\begin{aligned} PE_f - PE_i &= -\bar{\rho}_1 h_1 \frac{\Delta h}{2} - M \left( \frac{h_1 - \Delta h}{2} \right) \\ &\quad - \bar{\rho}_2 h_2 \frac{\Delta h}{2} + M \left( h_1 - \Delta h + \frac{h_2 + \Delta h}{2} \right) \\ &= M \left( \frac{h_1 + h_2}{2} \right) - (\bar{\rho}_1 h_1 + \bar{\rho}_2 h_2) \frac{\Delta h}{2}. \end{aligned} \quad (2.5)$$

PCM is the lowest order reconstruction, and gives  $M = \bar{\rho}_1 \Delta h$ , thus  $PE_f - PE_i \geq 0$  with the condition that  $\bar{\rho}_1 > \bar{\rho}_2$ . However, it is also possible for  $PE_f - PE_i < 0$  when the remapping is higher order (e.g. the piecewise quartic method; White and Adcroft, 2008) if the higher-order reconstruction is such that  $\frac{M}{\Delta h} < \frac{\bar{\rho}_1 h_1 + \bar{\rho}_2 h_2}{h_1 + h_2}$ , i.e. if the average density of the exchanged fluid is less than the density the two cells would have if they mixed completely. The right hand side of this expression lies in the range  $[\bar{\rho}_2, \bar{\rho}_1]$ , so it is clear that the PCM construction cannot give a negative PE change. Note that there may be cancellation between positive and negative PE changes in different cells, and so a net zero change in RPE for schemes that are second order or higher does not necessarily indicate zero spurious mixing.

### 2.2.2 Application to this thesis

The split RPE method described in this section was implemented in MOM6 by an online calculation of changes to RPE throughout the model timestep. By doing so, spurious mixing contributions can be separately attributed to horizontal and vertical processes

within the model. This is the first diagnostic providing this level of granularity, allowing in particular a more comprehensive analysis of the performance of ALE models with a strong distinction between these horizontal and vertical processes. Chapter 4 details the evaluation of both the split RPE method, and the overall spurious mixing in MOM6 compared with other models. Bearing in mind the caveats of RPE-based methods for diagnosing spurious mixing, this evaluation is performed in a suite of idealised test cases. The following chapter first details the implementation of the new *adaptive grid* vertical coordinate which is used throughout the rest of this thesis.



---

# AG: The adaptive grid coordinate

---

In this chapter, we introduce the formulation of a new vertical coordinate for an ocean model. This coordinate is spatially and temporally dynamic, and adjusts according to the local state. Following Burchard and Beckers (2004), we term this coordinate “adaptive”. As this coordinate is also intended for implementation in the regriding stage of an ALE model, we call it the *adaptive grid* (AG) coordinate.

An important consideration during the generation of a coordinate is that there are limits on the slope of a coordinate interface (R. Hallberg, personal communication, April 2018). These limits prevent errors in the pressure gradient acceleration term which lead to spurious or otherwise incorrect velocities, a well-known problem for terrain-following coordinates in the presence of steep topography (Haney, 1991). We construct an argument to show the role of coordinate slope in the pressure gradient acceleration,

$$1/\rho \nabla_s p + \nabla_s(gz) = \nabla_s(p/\rho + gz) + p \nabla_s(1/\rho), \quad (3.1)$$

where  $\nabla_s$  is the gradient along a surface,  $\rho$  is density,  $p$  is pressure,  $g$  is gravity and  $z$  is geopotential height. In pure density or pressure coordinates, there is only a single scalar potential whose gradient is the pressure gradient. Indeed, this also applies for a coordinate that is a function of *in-situ* density and pressure. However, for an arbitrary coordinate there are two terms that can contribute to the pressure gradient, with opposite signs. If both these terms are similar in magnitude, their compensation causes significant floating-point truncation errors, reducing the precision of the result. This case occurs when the coordinate slope is outside the bounds of isopycnal and isobaric, which occurs, for example, in terrain-following coordinates.

Adaptive coordinates have been a subject of research throughout the field of climate modelling. They see active use in atmospheric modelling as *hybrid* coordinates, switching from terrain-following sigma coordinates at low levels up to an isentropic coordinate at higher levels (Bleck, 1978). This scheme attempts to enhance resolution in regions of high static stability giving, for example, better representation of fronts. At higher levels, spurious mixing is avoided across isentropes due to the coordinate representation.

In addition to atmospheric modelling, adaptive coordinates have been used in ocean modelling. The HYCOM model (Bleck, 2002) uses a “grid generator” to blend between terrain-following sigma coordinates, geopotentials, and isopycnals, depending on the local state. There are also concessions specific to ocean modelling, such as the ability to reserve layers to provide resolution to the mixed layer as it deepens in winter.

A one-dimensional adaptive coordinate was proposed by Burchard and Beckers (2004). Instead of blending between different types of coordinates, resolution is shifted up or down the water column according to local state. For example, layers may accumulate in regions of high vertical shear in order to improve the model’s representation of the velocity field.

This scheme is based on the diffusion of interface positions within the column,

$$\partial_t z - \partial_\sigma(k^{\text{grid}}\partial_\sigma z) = 0, \quad (3.2)$$

where the normalised coordinate  $\sigma \in [-1, 0]$  and  $z(-1)$  and  $z(0)$  are the bottom and free surface of the column, respectively. The grid diffusivity coefficient,  $k^{\text{grid}}$ , is a function of  $\sigma$  and sets the resolution: stronger diffusivity leads to a more closely-spaced grid, and thus higher resolution in the steady state. The diffusivity coefficient is set by a linear combination of multiple parameters,

$$k^{\text{grid}} = \frac{cD}{T^{\text{grid}}} \left( c_\rho K_\rho^{\text{grid}} + c_u K_u^{\text{grid}} + c_d K_d^{\text{grid}} + c_b K_b^{\text{grid}} \right), \quad (3.3)$$

where the four constituent diffusivity components are stratification, shear, near-surface and background, respectively. The timescale  $T^{\text{grid}}$  applies over the full column depth  $D$  with a non-dimensional weight  $c$  to modify the overall strength and weights  $c_*$  for the individual terms. The first three diffusivity contributions are stratification ( $K_\rho^{\text{grid}}$ ) and shear ( $K_u^{\text{grid}}$ ) resolution enhancements, and a near-surface zooming  $K_d^{\text{grid}}$ . The background term in this scheme involves uniform diffusivity throughout the column,  $K_b^{\text{grid}} = 1/D$ .

The adaptive scheme of Burchard and Beckers (2004) was extended to three dimensions, adding three new properties (Hofmeister et al., 2010). These three properties are: to achieve Lagrangian vertical movements of the grid; for horizontal filtering of layer thicknesses; and for interfaces to have an isopycnal tendency. To introduce horizontal filtering, there is a lateral diffusion of layer thickness  $h_k = z_k - z_{k+1}$ ,

$$\partial_t h_k = \mathcal{A}_h (\partial_{xx} + \partial_{yy}) h_k, \quad (3.4)$$

where the diffusivity  $\mathcal{A}_h$  scales with  $\alpha_{\text{dif}}\Delta x^2/(4\Delta t_f)$ , for a strength  $\alpha_{\text{dif}}$  and diffusivity timescale  $\Delta t_f$  over a horizontal lengthscale  $\Delta x$ . Additionally, the isopycnal tendency term acts on interfaces themselves,

$$z_k^{n+1} = z_k^n + \alpha_{\text{iso}} \frac{\rho_k^* - \rho(z^n)}{\partial_z \rho}, \quad (3.5)$$

(superscripts represent discrete timesteps) where the target isopycnal values  $\rho_k^*$  may be prescribed statically, or dynamically determined from the local state.

The Hofmeister et al. (2010) adaptive coordinate has allowed for improved accuracy in coastal and basin-scale modelling. In a model of the inflow dynamics in the Baltic Sea, the adaptive coordinates produced a reduction in numerical mixing as effective as a doubling of the horizontal resolution (Hofmeister et al., 2011). Additionally, in a regional model of the North Sea and Baltic Sea, adaptive coordinates were estimated to provide a 20% reduction in numerical mixing compared to terrain-following coordinates, with only 5–8% increase in run time (Gräwe et al., 2015). However, despite the presence of an isopycnal tendency term, the Hofmeister et al. (2010) adaptive coordinate is intended for coastal ocean modelling, by producing a grid which is comparable to terrain-following sigma coordinates. These are not suitable for large-scale ocean modelling due to spurious velocities generated by the representation of the pressure gradient term (see Mellor et al., 1998, and discussion at beginning of this chapter). It may indeed be that the aforementioned adaptive coordinate is able to avoid this problem through the isopycnal tendency term, but this has not yet been demonstrated in the literature. We will now discuss a specific example of an adaptive coordinate in large-scale ocean modelling.

MOM6 is an ocean model employing a generalised vertical coordinate through the ALE algorithm that is easily able to accommodate adaptive vertical coordinates (see section 1.1

for details). MOM6 has an implementation of a modification of the HYCOM coordinate known as HYCOM1. In this implementation, the regridding routine yields isopycnals through most of the interior, and a pressure coordinate near the surface. To achieve this, every model interface has two quantities associated with it: a target density, and a minimum depth. The depth of the isopycnal at that density is found, and used for the interface unless this would make the interface higher in the water column than its associated minimum depth. In this way, there is a region near the surface in which interfaces follow a pressure coordinate, allowing sufficient resolution to represent the processes of the surface boundary layer. The resolution of this region is set by the spacing of minimum depths associated with successive interfaces.

An integral component of the HYCOM1 coordinate in MOM6 is the association of a single density value with every unique interface. This gives a strong link to the pure isopycnal coordinate and allows a maximum height to be chosen for each interface. However, because there are large changes in the vertical density distribution across the globe, this choice leads to non-optimal usage of the available resolution, particularly in weakly-stratified regions. This issue is addressed by the addition of an artificial compressibility into the equation of state during regridding only, which separates similar density classes as depth increases through the abyssal ocean. Despite the artificial compressibility, the chosen densities for interfaces must cover the full range expected during the simulation period. Additionally, maximum interface heights must be conservatively selected so that the surface mixing scheme (e.g., KPP or ePBL) does not attempt to mix into an isopycnal layer, which may be thicker than surface layers. If this mixing does occur, it manifests as a spurious mixing of surface properties into the interior. As a consequence of this compromise, some physical features of the ocean are poorly represented. In particular, some key overflows driving dense water formation in the North Atlantic mix too much and affect the strength of the Atlantic Meridional Overturning Circulation.

By utilising an adaptive coordinate, we aim to avoid these shortcomings. On small scales, the coordinate is isopycnal in order to reduce spurious diapycnal mixing from lateral advection. However, there is no requirement for an interface to maintain a single density value over its span, allowing for more effective utilisation of resolution over large scales. Weakly-stratified regions are afforded more resolution dynamically, thereby avoiding excessive mixing in overflows and exchanges. Finally, because of its adaptive nature, there is no need to manually specify the density and height of each interface, a process that must be tailored to each simulation and requiring intimate familiarity with expected properties over the entire domain.

### 3.1 Principles of AG

In the following sections, we propose the adaptive grid (AG) coordinate to improve upon the modified HYCOM coordinate in MOM6, by incorporating concepts from Hofmeister et al. (2010). Similar to the HYCOM1 coordinate, there are two key aims to the adaptive grid coordinate. Principally, the intention is to reduce spurious diapycnal mixing, particularly in the ocean interior and against topography. To this end, we target an isopycnal coordinate wherever stratification is high enough that isopycnals are well-defined. However, where stratification is weak, such as within the mixed layer, an isopycnal coordinate degrades due to poor spatial resolution, which may, for example under-resolve baroclinic modes (Stewart et al., 2017). In these regions, we instead target smooth interfaces and limit the vertical spacing between layers. This is a more relaxed constraint than HYCOM1, which blends between isopycnal and  $z$ -star, instead only requiring a certain minimum resolution to be achieved.

The adaptive grid coordinate consists of a single equation governing the time-evolution of interface heights applied during the regridding stage of the ALE algorithm. This equation was engineered to achieve the two key aims of the coordinate in as direct a manner as possible, while maintaining overall consistency of the grid. An algorithmic adjustment is applied as a source term in the equation, providing the ability to efficiently shift interfaces to ensure a minimum resolution if needed. Representing the regridding by a single equation in terms of physical quantities is intended to provide a solution independent of the horizontal resolution, the number of vertical layers or the timestep governing baroclinic dynamics. The governing equation describes the time-evolution of internal interfaces between the layers as


$$\begin{aligned} \frac{z_{kr} - z_{kd}}{\Delta t} = & \underbrace{-\nabla_H \cdot \left( \omega_\sigma \frac{\kappa \nabla_H \sigma}{\sqrt{(\partial_z \sigma)^2 + |\nabla_H \sigma|^2}} \right)}_{\text{density adaptivity}} \\ & + \underbrace{\nabla_H \cdot \left( \omega_z \kappa \nabla_H z_k \right)}_{\text{lateral smoothing}} + \underbrace{\tau_r^{-1} (z_k^* - z_k)}_{\text{vertical restoring}} + \underbrace{F_{\text{con}}}_{\text{convective adjustment}}, \end{aligned} \quad (3.6)$$

where  $z_{kr}$  and  $z_{kd}$  are the height of the interface  $k$  (between layers  $k$  and  $k + 1$ ) after the regridding and dynamics sub-steps, respectively;  $\nabla_H$  is the gradient operator acting along coordinate surfaces;  $\sigma$  is the locally-referenced potential density;  $\omega_\sigma$  is a weighting towards density adaptivity;  $\kappa$  is a lateral diffusivity that is constant within a given vertical column;  $\omega_z$  is a weighting towards lateral smoothing (typically  $1 - \omega_\sigma$ );  $\tau_r$  is a timescale of restoring toward a nominal coordinate  $z_k^*$ ; and  $F_{\text{con}}$  is the source term corresponding to the grid adjustment routine.

Equation 3.6 is structured as a time-evolution of internal interface heights, with contributions from three terms. The first of these provides the density adaptivity, adjusting the interface heights in order to minimise gradients of locally-referenced potential density on coordinate surfaces. This term is only non-zero when the local isopycnal slope is less than a specified threshold. The second and third terms provide lateral interface smoothing and vertical interface restoring, respectively. The combination of these terms is intended to provide sufficient vertical resolution with smoothed interfaces where stratification is too weak to perform density adaptivity. By evolving only internal interface heights, and designing limiters to ensure interfaces do not cross, the total column height is also preserved. The next four sections detail these terms and the additional adjustment step that cannot be represented in the physical equation.

## 3.2 Discretisation

In the implementation of the adaptive grid coordinate, it is important to consider how the different constituent quantities are spatially arranged. MOM6 uses the Arakawa C-grid (Arakawa and Lamb, 1977), which has the best representation of inertia-gravity waves of all grid staggering configurations. This comes with the requirement that the Rossby deformation radius is well-resolved, preventing spurious noise in the calculation of the Coriolis force (Adcroft et al., 1999). In a C-grid model, cell-mean tracer concentrations and layer thickness are stored at cell centres; zonal and meridional velocities,  $u$  and  $v$  are at zonal and meridional cell boundaries, respectively; and vorticity is at cell corners (fig. 3.1). These quantities are averaged vertically through the thickness of their containing



**Figure 3.1:** Horizontal grid staggering in an Arakawa C-grid. The cell centre  $h$  carries tracers and layer thickness. Velocities and transports are located at cell edges  $u$  and  $v$ . Cross-terms such as vorticity are on the cell corner  $q$ .

cell. Laterally, we adopt the convention of indexing with variables  $i$  and  $j$  for the zonal and meridional directions. Whole multiples of the index variables represent cell centres in the respective direction, and half-steps are the cell boundaries. For brevity, we adopt the convention that capitalised index variables refer to half-steps, e.g.,  $J = j + 1/2$ .

Vertically, we make a distinction between model *layers* and model *interfaces*. As a finite-volume model, there are no pointwise prognostic quantities in MOM6. Thus, layer thickness, defined as  $h_k = z_{K-1} - z_K$ , is a crucial quantity. However, the density adaptivity and horizontal smoothing terms of AG apply to interfaces directly. To unite the indexing of both layers and interfaces, we use  $k$  as the single vertical indexing variable. Whole multiples of  $k$  represent layers, and half-steps the interfaces between adjacent layers. The free surface is at  $z_{1/2}$ , and for a model with  $N$  layers, the bottom of a column is at  $z_{N+1/2}$ .

Due to the different spatial positions of the quantities involved in the adaptive coordinate calculation, averaging or interpolation is required to bring them to a common location. Laterally, the quantities are defined as a vertical cell average at a point within the cell. In one dimension, we can simply use an average to relocate a variable, i.e. for uniform grid spacing,  $\bar{\phi}_I = (\phi_i + \phi_{i+1})/2$  (here  $i$  may refer to either a cell centre or a cell boundary). Vertically, the calculation is more complex, as thicknesses may differ significantly between vertically adjacent cells in the same column. One technique to obtain values on interfaces is a simple linear interpolation,  $\phi_K = (h_{k+1}\phi_k + h_k\phi_{k+1})/(h_k + h_{k+1} + \epsilon_h)$ . The value  $\epsilon_h$  is chosen to prevent division by 0 if both layers are vanished, but affects the overall calculation at sub-roundoff precision. This guarantees the same result as though  $\epsilon_h$  was omitted, at a bitwise level. We will introduce specific discretisations for the adaptive grid coordinate in the following sections.


### 3.3 Density adaptivity

The most important component of the adaptive grid coordinate is the density adaptivity, provided by the first term on the right-hand side of eq. 3.6. Mathematically, this term takes the form of the divergence of the flux of vertical interface displacement. By construction, this term conserves the mean height of an interface, while locally minimising the gradient of locally-referenced potential density along the interface. The action of the density adaptivity term is detailed in the following paragraphs.

Because density adaptivity fluxes are calculated only between adjacent cells, the gradient minimisation only occurs locally. However, advective fluxes of tracers are also calculated locally (albeit with a potentially larger stencil), therefore any spurious mixing introduced by advection is affected by local gradients of properties within a layer. In

other words, a coordinate should need only be locally isopycnal, rather than globally isopycnal, in order to reduce spurious mixing. Additionally, by ensuring that the ALE algorithm is run on a timescale comparable to that of baroclinic dynamics, the coordinate should remain locally isopycnal.

A schematic to aid in understanding the action of the density adaptivity term is presented in fig. 3.2. At each of the four cell boundaries surrounding a tracer cell, the direction of the interfacial flux is determined from the sign of the difference of the locally-referenced potential density  $\sigma$ , on either side of the boundary. To calculate this difference, the local density differences for temperature and salinity are evaluated at the cell boundary. In the zonal direction,  $\delta_I \sigma = \alpha(T_{i+1} - T_i) + \beta(S_{i+1} - S_i)$ , where  $\alpha$  and  $\beta$  are the thermal expansion and haline contraction coefficients at the local pressure, and  $T$  and  $S$  are the relevant temperature and salinity tracers for the chosen equation of state. Instead of interpolating temperature and salinity onto the coordinate interfaces,  $\delta_I \sigma$  is evaluated within the two layers directly adjacent to each interface. If  $\delta_I \sigma$  differs in sign above and below the interface, either direction of interfacial flux would increase the density difference of an adjacent layer, making it further from isopycnal. In this case, the interfacial flux is set to zero. Otherwise, the smaller in magnitude of  $\delta_I \sigma$  above and below the interface is used for the subsequent calculations, ensuring that there are no overshoots. We call this the ‘‘twin gradient method’’, due to the use of gradient information in the two layers adjacent to an interface.



**Figure 3.2:** Schematic of density adaptivity in a 2D slice of two vertical layers in two adjacent columns. The interface height flux (horizontal arrow) due to the angle between the isopycnal surface (dashed) and coordinate interface (solid horizontal line), leads to vertical interface displacement (small vertical arrows) through its divergence (not pictured).

Now we consider the discretisation of the density adaptivity term. As shown above, the flux is decomposed into four contributions to each interface. Without loss of generality, we will examine only one of these fluxes,

$$F_{\sigma_{Ij}} = \omega_{\sigma} \kappa \frac{\partial_x \sigma}{\sqrt{(\partial_z \sigma)^2 + |\nabla_H \sigma|^2}}. \quad (3.7)$$

Here,  $F_{\sigma_{Ij}}$  is the zonal flux east of the tracer point  $(i, j)$  and  $\nabla_H$  is the along-coordinate lateral gradient operator, with  $\partial_x$  also acting in the along-coordinate direction. We now

introduce the fully-discretised version of this flux,

$$F_{\sigma_{I,j,K}} = \omega_{\sigma}(\Delta x_I^2/\Delta t)(\overline{h^?}^I/\Delta x_I) \frac{\overline{\delta_i\sigma}^K}{\sqrt{(\overline{h^?}^I\overline{\delta_i\sigma}^K/\Delta x)^2 + (\overline{h^?}^I\overline{\delta_j\sigma}^{IK}/\Delta y)^2 + (\overline{\delta_k\sigma}^I)^2}}. \quad (3.8)$$

This discretised form involves many new terms, which we will address in order. The diffusivity coefficient  $\kappa$  has been written as  $\Delta x_I^2/\Delta t$ , the zonal cell extent at the zonal boundary divided by the baroclinic timestep, in order to remove dependence on both lateral resolution and the baroclinic timestep. (In the meridional direction, the diffusivity coefficient uses  $\Delta y_j^2$ , the meridional cell extent at the meridional boundary, as the lengthscale.)  $\delta_i\sigma$  and  $\delta_j\sigma$  are lateral differences of locally-referenced potential density, and  $\delta_k\sigma$  is a vertical difference;  $h^?$  is an interpolated thickness quantity (with units of metres); and overbars denote interpolation to co-locate staggered quantities, with superscripts giving the location of the interpolated result. We will now discuss specific details of the discretised density adaptivity flux.

Due to the sensitivity of the density adaptivity term to layer thickness  $h$ , we take care in our discretisation to use the most dynamically-relevant quantity we can. Naïvely, we could average the four adjacent thicknesses onto the flux point. However, referring to fig. 3.2, the flux acts to “pivot” the interface about the cell boundary. In this sense, there are only two cells whose thickness is dynamically relevant to the density adaptivity. We define the “upstream thickness,”  $h^?$ , which depends on the sign of the product of the lateral flux difference and the vertical stratification. In the zonal direction, this becomes

$$\overline{h^?}^I = \begin{cases} 1/2(h_{i,j,k} + h_{i+1,j,k+1}), & \text{if } \overline{\delta_I\sigma}^K \overline{\delta_k\sigma}^I < 0 \\ 1/2(h_{i,j,k+1} + h_{i+1,j,k}), & \text{if } \overline{\delta_I\sigma}^K \overline{\delta_k\sigma}^I \geq 0. \end{cases} \quad (3.9)$$


A similar expression can be obtained for the upstream thickness on a meridional cell boundary.

A vertical gradient of locally-referenced potential density is required to compute a vertical interface displacement, given a lateral density gradient to minimise. The flux computed by the expression in eq. 3.8 is located at cell boundaries, however we calculate the vertical gradient  $\delta_k\sigma = \alpha(T_k - T_{k-1}) + \beta(S_k - S_{k=1})$  at tracer points, on coordinate interfaces (as previously,  $\alpha$  and  $\beta$  are the derivatives of locally-referenced potential density with respect to temperature and salinity). Therefore, this gradient must be interpolated onto the lateral cell boundary, which can be done with a simple lateral average,  $\overline{\delta_k\sigma}^I = (\delta_k\sigma_i + \delta_k\sigma_{i+1})/2$ . This quantity is, however, only the vertical gradient in *index space*. We define the full discretised vertical gradient in physical space at cell boundaries as  $\partial_z\sigma = \overline{\delta_k\sigma}^I/\overline{h^?}^I$ , where the exact form of the thickness is discussed above. However, for reasons of numerical stability where layers are thin, we multiply both the numerator and denominator in eq. 3.8 by the thickness so as to avoid division by potentially small quantities.

Finally, we discuss the discretisation involved in the squared terms in the denominator of eq. 3.8. This denominator normalises the flux, particularly when the vertical density gradient is small due to weak stratification. As above, all the terms comprising the denominator must be located at the cell boundary. The lateral gradient in the direction of the flux,  $\delta_I\sigma^2$  requires no interpolation and can be used directly. When the vertical gradient is interpolated to the correct location,  $\overline{\delta_k\sigma}^I$ , and the two contributions  $\delta_k\sigma_i$  and  $\delta_k\sigma_{i+1}$  are of opposite signs, the squared, interpolated value could be close to zero, leading the resultant flux to be erroneously large. Instead, we define  $(\overline{\delta_k\sigma}^I)^2 := 1/2((\delta_k\sigma_i)^2 + (\delta_k\sigma_{i+1})^2)$ . A similar case can be constructed for the other lateral gradient term in the denominator.

This involves interpolation of the gradient at a meridional cell boundary onto a zonal cell boundary, and therefore is the average of four adjacent quantities.

Another consideration in the implementation of this formulation for density adaptivity is to ensure that the resulting coordinate is well-defined. This means that for a given column, the top and bottom interfaces remain stationary, and interface heights within the column are monotonic (i.e. no interface “tangles” with another). There is, however, no restriction on the minimum thickness of a layer; interfaces are permitted to touch exactly, giving a zero-thickness layer. In order to preserve a well-defined coordinate, interfacial height fluxes must be limited. Consider the case of a relatively large lateral difference in  $\sigma$  in a region of weak stratification. In this scenario, the interfacial height flux may be very large, since a large vertical displacement is required to significantly alter the density of the interface. Applying this flux directly may result in tangling of the interface with its vertically adjacent neighbour. This scenario is similar to the case of an interface prevented from moving by topography, illustrated in fig. 3.3.



**Figure 3.3:** Applying an interfacial flux (vertical arrows) without limiting may cause an interface to intersect topography (hatched region). The displaced interface positions (dashed) cross adjacent interfaces in the same column, leading to tangling of the coordinate.

Limiting places a hard constraint on the interfacial flux to enforce a well-defined coordinate: the absolute value of the displacement of an interface must not exceed the thickness of the layers above or below the interfaces on either side of the cell boundary at which the flux is defined. We can refine this limiting to allow the maximum possible interfacial flux that results in a well-defined coordinate. This hinges on taking into account the direction of the flux, so that we only have to consider two thickness values in the calculation of the limited value. Figure 3.2 shows that for a positive flux, the interface to the right of the cell boundary shall be displaced upwards, and the interface to the left displaced downwards (vice versa for a negative flux). Then the flux can be limited such that the resulting displacement does not exceed  $\min\{h_{\text{up,right}}, h_{\text{down,left}}\}$ , which we call the “minimum downwind thickness.” To express this as a limit on the flux itself, as opposed to the displacement, we incorporate dimensions of the lateral grid to give a limited flux. In the



zonal direction, and for a positive flux, this limited flux is

$$(F_A^{\text{lim}})_{I,j,K} = \min \left\{ (F_A)_{I,j,K}, \right. \\ \left. \frac{1}{8} \frac{A_{i,j}}{\Delta y_{I,j}} h_{i,j,k}, \right. \\ \left. \frac{1}{8} \frac{A_{i+1,j}}{\Delta y_{I,j}} h_{i+1,j,k+1} \right\}, \quad (3.10)$$

where  $F_A$  is the un-limited interfacial flux,  $A$  is the area of a cell,  $h$  is the thickness of a layer,  $\Delta y$  is the width of a cell (at the zonal cell boundary). The change in the thickness of a cell has eight separate flux contributions: there is an interfacial flux at each of the four cell boundaries, and these contributions come from the interfaces both above and below the cell. If all these flux contributions were acting to reduce the thickness of the cell, we need to prevent them from causing a negative final cell thickness. Introducing the factor  $1/8$  ensures that in the worst case, thickness may go to zero, but never negative. The limiting expressions for negative fluxes are similar, but consider different thickness values.

There are two final parts of the density adaptivity term that have not been covered. The first of these is the diffusivity coefficient  $\kappa$ , which takes the form of a squared length-scale divided by a timescale. Fluxes in the x-direction use  $\Delta x$  at the cell boundary as the lengthscale, and similarly  $\Delta y$  for the y-direction fluxes. Using local cell dimension ensures that the behaviour of the coordinate is minimally affected by lateral resolution. The timescale may be chosen to set the strength of the interfacial fluxes. At the timescale of baroclinic dynamics or faster, the coordinate is permitted to adjust as much as possible (restricted by the limiter) at each step of the ALE algorithm. Longer timescales reduce the amount of adjustment, tending towards more Lagrangian behaviour. This is similar to the spirit of the z-tilde coordinate (Leclair and Madec, 2011).

Finally, the weight of the density adaptivity term is set by the non-dimensional coefficient  $\omega_\sigma$ . Its complement,  $\omega_z$ , sets the weight of the lateral smoothing term, active when stratification is too weak to use density adaptivity, and the subject of the following section.

### 3.4 Lateral smoothing

In unstratified or weakly-stratified regions of the ocean, density adaptivity is insufficient to yield a sensible coordinate. Despite the extreme aspect ratio of the oceans, in these regions lateral differences on the grid exceed vertical differences, due to the significantly smaller vertical grid spacing than horizontal. As a result, minimising lateral density differences may require extreme vertical displacements. This also brings into question the ability of the coordinate to accurately represent the underlying dynamics when interfaces are displaced by large distances. The following two sections provide a mechanism for ensuring resolution in these regions.


In regions where density adaptivity is compromised by the lack of stratification, and particularly at the boundary of a locally-isopycnal region, the coordinate should smoothly transition toward being nearly flat. To do this, we introduce ‘‘lateral smoothing’’ as the divergence of a flux of interface height displacement, similar to density adaptivity. Instead of calculating the lateral smoothing flux from the density field, the interface heights themselves are used. Indeed, isolating the lateral smoothing term gives the heat equation  $\partial_t z_k - \alpha_z \nabla^2 z_k = 0$ , which, at steady state and with a zero-flux boundary condition results

in a flat interface  $z_k$ . As with the density adaptivity flux, we present the discretised form of the lateral smoothing flux in one direction only, without loss of generality. The zonal flux is then

$$F_{z_{I,j,K}} = \omega_z (\Delta x_I^2 / \Delta t) \delta_i z_k. \quad (3.11)$$

To avoid a potential shock when the coordinate switches between the density adaptivity and lateral smoothing regimes, the smoothing term is also framed as the divergence of a flux. Then the transition between the regimes, according to an isopycnal slope metric, defined below, is simply a weighting between the two fluxes. However, this formulation requires once again the use of a flux limiter to ensure the resultant coordinate is well-defined. The limiter described for the density adaptivity term in eq. 3.10 is sufficiently general that it can be applied to the lateral smoothing without modification.

There is, however, a further observation that permits the use of a less restrictive flux limiter. As the smoothing flux is computed on velocity points (i.e., lateral cell boundaries) by taking the difference in interface height between the adjacent cells, it gives an effective interfacial slope. Taking the divergence of this flux acts to reduce the magnitude of the interfacial slope (fig. 3.4). If the diffusivity applied to the lateral smoothing,  $\omega_z \kappa$  is constant for all interfaces  $k$  within a column, no limiting is required between adjacent interfaces in the column: a coordinate that begins untangled is unable to tangle under the lateral smoothing operator since all slopes are reduced in magnitude. Under these conditions, the only required limiting is to prevent interior interfaces from crossing the top or bottom of the column. We call this the “unrestrictive” limiter.



**Figure 3.4:** Interface height flux (horizontal arrow) due to lateral smoothing, tends to flatten interfaces through interface displacement (small vertical arrows).

To use the unrestrictive limiter, the smoothing flux must be applied with a constant diffusivity within the column. We introduce a cutoff condition that applies a minimum diffusivity to a user-specified fraction of the smoothing flux. The rest of the flux contribution depends on the local stratification; smoothing flux is applied where stratification is weak, and density adaptivity when stratification is sufficiently strong. This gives a cutoff condition in terms of the isopycnal slope  $S$ ,

$$\omega_\sigma = \begin{cases} 1 - \omega_z^{\min}, & \text{if } S \leq S_{\max} \\ 0, & \text{otherwise} \end{cases} \quad \text{and } \omega_z = 1 - \omega_\sigma. \quad (3.12)$$

This splits the lateral smoothing term into two contributions. The first component is always active, giving  $-\nabla_H \cdot \omega_z^{\min} \kappa \nabla_H z_k$ , governed by the unrestrictive limiter which prevents any interface from crossing the top or bottom of the column. The second component is

the remainder  $-\nabla_H \cdot (\omega_z - \omega_z^{\min}) \kappa \nabla_H z_k$ , which is governed by the same generalised limiter as the density adaptivity term.


An alternative to this scheme would be to blend smoothly between density adaptivity and lateral smoothing, depending on the local isopycnal slope. However, in an idealised sector simulation (chapter 4), the isopycnal slope distribution was strongly bimodal, clearly separating the two regimes.

To demonstrate the impact of the unrestrictive limiter, an experiment was performed where the regridding routine was run iteratively without dynamics enabled, from a uniform initial grid of 75 layers at  $0.25^\circ$  horizontal resolution. Figure 3.5a shows a section of the initial density space, which was taken from a zonal average of density from the MOM6 example “Global ALE Experiment”, where a lateral density gradient transitions to a vertical density gradient. To avoid cumulative errors from remapping, after each iteration of the regridding routine, this initial state was remapped onto the newly-generated grid. The regridding routine was iterated upon 1000 times to push the grid towards a converged state. First, this iterative procedure was run with only the density adaptivity and lateral smoothing terms enabled, with  $\omega_z^{\min} = 0$ , i.e. no contribution from lateral smoothing with the unrestrictive limiter. This case is shown in fig. 3.5b, which overlays the interface positions on the square of the along-coordinate density slope. Through the transition from a lateral to vertical density gradient, the coordinate “tangles” into a cluster of massless layers. The formulation of the generalised limiter governing the density adaptivity and lateral smoothing terms precludes the breaking of these tangles. By introducing lateral smoothing with the unrestrictive limiter (fig. 3.5c) where  $\omega_z^{\min} = 0.1$ , these tangles are broken and unable to persist through the iterated regridding.

### 3.5 Vertical restoring

The previous two terms, density adaptivity and lateral smoothing conserve average layer height, i.e.,  $\partial_t \langle z_k \rangle_k = 0$ . This constraint ties the resulting grid to the initial condition, preventing free evolution according to the state itself. At the same time, evolution of the state can push the grid to an uncorrectable and undesirable state. For example, buoyancy forcing may inflate the surface layer, lowering the average layer height and reducing resolution. In order to break this constraint, and allow for the redistribution of resolution within a column, we introduce a source term. By doing so, the displacement of an interface within a column does not require a commensurate displacement elsewhere along the coordinate interface. The source term takes the form of a vertical restoring  $\tau_r(z)^{-1}(z_k^* - z_k)$ , with a timescale  $\tau_r(z)$ , where the target grid for restoring is  $z_k^*$ , and is a free choice in both time and space. The target grid could, for example, be used to enhance vertical resolution in the upper portion of the water column, improving the representation of mixed-layer physics.

Maintaining surface resolution is important for the representation of the surface boundary layer which affects the representation of surface fluxes and mixed layer physics. This is associated with maintaining a minimum resolution throughout the entire water column, ensuring that the laterally-permitted baroclinic modes may be vertically resolved (Stewart et al., 2017). This argument suggests that a coordinate used for ocean modelling should have high vertical resolution near the surface. The density adaptivity mechanism proposed in this chapter is insufficient to achieve this vertical resolution, as the density structure in the near-surface region is often well-mixed, such that following surfaces of constant density is likely to yield a boundary layer of similar thickness to the mixed layer. Therefore, the vertical restoring timescale  $\tau_r(z)$  is chosen to be a function of depth: a very short timescale can be used near the surface, so that coordinates closely follow a pre-



**Figure 3.5:** Demonstration of interface tangling (lines in panel (b)) without the barotropic smoothing term. (a) Density initial condition; (b) interfaces (lines) and squared along-coordinate density slope (colours) with only density adaptivity and limited smoothing ( $\omega_z^{\min} = 0$ ); (c) as above, with barotropic smoothing ( $\omega_z^{\min} = 0.1$ ) which eliminates tangling.

scribed vertical structure; at depth, a weak timescale ensures that density adaptivity is the predominant affector of layer structure, but a minimum resolution can still be maintained.


Vertical restoring is applied as an adjustment to the *intermediate* grid after calculation and application of density adaptivity and lateral smoothing with flux limiting, taking the form of a function  $z_i(k) \mapsto z_r(k)$ . This operation essentially applies the restoring to each separate interface. Following Adcroft (2015), we define two constants to specify the operation of the vertical restoring: a “shallow” depth  $z_s$ ; and a “deep” depth  $z_d$ . The shallow depth defines the point above which the restoring timescale is shorter than the baroclinic timestep, i.e. effectively instant. By restoring instantly toward the target coordinate, the surface resolution can be well-specified. On the other hand, interfaces deeper than the  $z_d$  threshold use a user-provided parameter  $\tau_0$  for a long timescale. This parameter should be chosen such that restoring does not dominate the processes of density adaptivity and lateral smoothing in the deep ocean. In the intermediate region between  $z_s$  and  $z_d$ , the restoring timescale  $\tau_r(z)$  takes a cubic profile.

### 3.6 Convective grid adjustment

As described thus far, the adaptive grid coordinate is applicable in many dynamic regimes. However, there are two similar processes that require an explicit adjustment: surface buoyancy loss and lateral exchanges (or gravity currents). By way of a motivating illustration, we first demonstrate a scenario where density adaptivity is unable to sufficiently adjust the coordinate to a consistent state. Figures 3.6a and 3.6c show a simple two-dimensional overflow simulation without the adjustment enabled. A pulse of dense water at 35 psu is released from the top of a slope into stably-stratified ambient water, with a linear gradient of  $0.25 \text{ psu km}^{-1}$ . The domain is 4 km deep and 800 km long, with a shelf of 800 m depth and 80 km width. A linear slope of 480 km with connects the shelf to the deep ocean basin. On the shelf, a single layer contains the entirety of the dense pulse, and all other layers have zero thickness. The initial condition for the coordinate is a uniform z-coordinate. We see that by simply varying which layer initially contains the dense water, but not changing any physical parameters, drastically different physical outcomes can be obtained (compare fig. 3.6b to fig. 3.6a).

The two processes of surface buoyancy loss and lateral exchange require grid adjustments to move water masses to a different layer, accommodating the density modification. This is because the density adaptivity term is limited in how quickly interfacial fluxes are able to adjust the coordinate in response to strong lateral density gradients: the strong limiting constraint requires many iterations to propagate the signal vertically. However, we are able to take advantage of the fact that we are free to alter the coordinate beyond the physical evolution in eq. 3.6. As such, we propose an explicit adjustment that aims to make the coordinate consistent with model isopycnals when density adaptivity alone is insufficient (fig. 3.6d).

To begin describing the adjustment term, we relate the processes of coordinate generation (regridding) and advection. Algorithmically, full three-dimensional advection in MOM6 is decomposed into a vertical component that manifests during vertical remapping, and a purely along-layer lateral component. The shortcomings illustrated in the example above are associated with this along-layer advection. As implemented in numerical models, advection modifies water properties downstream, but does not change tracer concentrations in the upstream cell. As such, when adjusting the grid we modify only the upstream (source) column so that it may impact subsequent advection operations. For every cell that is a source for advection, we define the “k-space isopycnal slope” in the direction of advection. This slope is defined as the number of layers spanned between an isopycnal



**Figure 3.6:** Demonstration of the effect of the adjustment term on a gravity current. Panels show salinity (colours) and interfaces (solid lines) mid-way down a slope, at the head of a gravity current. The dense watermass starts at the far left of the domain, occupying the full depth in the bottom layer in panels a) and c); and in the middle layer in panels b) and d). All other layers in the columns containing the dense watermass are massless. Panels a) and b) were run without the adjustment term, which was then enabled for c) and d). Density inversions are significantly reduced after adjustment is enabled.

in the two adjacent columns,  $S_k = k_t - k$ , where  $k$  is the index of the source layer and  $k_t$  is the index of the lowest or highest layer spanned by the isopycnal, depending on the sign of the isopycnal slope. The magnitude of the grid adjustment is proportional to the k-space slope and the lateral advective CFL number,  $u\Delta t/\Delta x$ . Intuitively, this is because an isopycnal that is close in index space does not require significant adjustment, and because smaller velocities affect water properties downstream to a lesser degree. During the actual calculation of the k-space slope, we make a deliberate underestimate so as not to overshoot the adjustment.

Finally, to apply the grid adjustment as an actual interface displacement, we must introduce a lengthscale. For this, we use the minimum of the source cell thickness and the summed thicknesses of the cells spanned by the isopycnal in the destination column,  $L = \min \left\{ h_{i,j,k}, \sum_{k'=k+1}^{k_t} h_{i,j,k'} \right\}$ . This expression for  $L$  assumes that the isopycnal slopes downwards (towards higher  $k$  indices) from the source cell; a similar construction is possible for an upward-sloping isopycnal. Using this definition for the lengthscale avoids two negative consequences associated with thin layers. Firstly, if the destination layers are already relatively thin, they would be significantly modified by advection, bringing them to a state consistent with the isopycnal. On the other hand, we avoid a disproportionately strong adjustment from a thin source layer that would not appreciably modify the destination. This gives the final expression for the grid adjustment term due to a single advective contribution with a downward-sloping isopycnal,

$$\Delta z_{i,j,K} = \alpha_z \frac{u_{I,j,k}\Delta t}{\Delta x} (k_t - k)L, \quad (3.13)$$

where  $\alpha_z$  is a non-dimensional weight. A full grid adjustment routine is achieved by applying the above expression for all advective contributions in both the zonal and meridional directions.

The result of applying the adjustment routine in the downslope flow example is shown in figs. 3.6c and 3.6d. In the first case, the dense watermass starts in the lowest layer and should not be adjusted. We can see that this is indeed the case, although some massless layers have been moved around, which does not affect the overall result. The second case (where the dense initial condition starts in a middle layer) shows a significant improvement. A density inversion and strong lateral mixing can be seen in the non-adjusted case in fig. 3.6b. Enabling adjustment (fig. 3.6d) has the most important effect of eliminating the density inversion. Unfortunately, there is still some lateral mixing present, as the adjustment is limited in being able to shift massless layers only. Regardless, the adjustment term improves on the inconsistent grid condition without negatively impacting the grid where it is consistent, as well as bringing the solution closer to the more physically correct structure (figs. 3.6a and 3.6c).

### 3.7 Summary

In this chapter, we have introduced the ‘‘adaptive grid’’ coordinate. Using online density information, this coordinate aims to generate a grid that can optimally represent the current state. This optimisation is guided by three principles: density adaptivity, lateral smoothing, and vertical restoring. The coordinate is presented in eq. 3.6, however this chapter also provides details of the discretisation such that the coordinate is robust. In addition to the main formulation (eq. 3.6), an empirical adjustment routine has been proposed to handle grid inconsistencies that cannot otherwise be removed by the other terms.

Density adaptivity is the core principle of AG, guided by the strengths of isopycnal coordinates. However, the requirement for an entire model interface to follow an isopycnal is relaxed, optimising only for local uniform density. This permits AG to be more flexible, allowing, for example, maximum use of available resolution across basins of vastly different densities. When local stratification is weak, density adaptivity is unable to optimise for local isopycnal structure and the lateral smoothing term is used instead. The weighting between these terms is governed by the cutoff condition eq. 3.12. Finally, as the density adaptivity and lateral smoothing terms take the form of fluxes of interface height, they preserve the mean height of interfaces across the model domain. To break the constraint of unvarying mean interface height, a source term in the form of a vertical restoring towards a target depth profile is added. The restoring timescale varies through depth: a short timescale at the surface ensures minimum resolution is maintained; and a longer timescale at depth allows the other two terms to predominantly dictate the coordinate structure.

Having fully implemented AG in the MOM6 model, we must evaluate its performance and ensure robustness. In chapter 4, the coordinate is evaluated in idealised test cases against other models and coordinates. To do so, the split RPE technique (introduced in chapter 2) is used to diagnose the spurious mixing due to horizontal and vertical processes separately. After evaluation in the idealised context, the coordinate is tested for its robustness and effects on the circulation itself in a sector model, in chapter 5.



---

# Testing in idealised experiments

---

Having defined the formulation of the adaptive grid coordinate in the previous chapter, we now evaluate it in the context of other coordinates and other ocean models. The objectives of this chapter are twofold: to evaluate the adaptive grid coordinate for its spurious mixing characteristics in the RPE framework, and to demonstrate the ability of the split RPE method (introduced in chapter 2) to attribute contributions to spurious mixing to vertical and horizontal processes across a range of vertical coordinates. Using the split RPE method, we can for the first time interrogate spurious mixing in a model in terms of horizontal and vertical contributions. In particular, the effects of horizontal and vertical model resolution, as well as the vertical coordinate are investigated.

The split RPE method is designed to attribute the spurious mixing in an ocean model to specific processes, by breaking down to horizontal and vertical contributions. A caveat of this method is that explicit mixing and buoyancy forcing must be disabled, as they contaminate the diagnostic. This restricts the scope of simulations that can be analysed using the technique. For this reason, we use the idealised test cases presented by Ilıcak et al. (2012) that were designated for analysis using a global RPE diagnostic, and as such are transient simulations without explicit mixing or buoyancy forcing. This choice allows comparison of MOM6 performance (including the adaptive grid coordinate outlined in chapter 3) to existing MITgcm and MOM5 simulations (Ilıcak et al., 2012) and MPAS-O (Petersen et al., 2015).

The idealised experiments are run in MOM6 across three test cases: a lock exchange, internal waves, and baroclinic eddies, chosen to match Ilıcak et al. (2012). The model is run using the piecewise parabolic (PPM) horizontal tracer advection scheme, and the PPM  $h_4$  vertical remapping scheme. Four vertical coordinates are used:  $z$ -star and  $z$ -tilde with uniform grid spacing; isopycnal ( $\rho$ ) with linear layer densities; and the adaptive grid coordinate (chapter 3; configuration shown in table 4.1). When not specified,  $z$ -star is used. The other coordinates are referred to by their name. A linear equation of state is used, with constant salinity such that density is a function only of temperature,  $\rho = \rho_0 - \alpha\theta$ . Here,  $\rho_0 = 1001 \text{ kg m}^{-3}$  is the reference density,  $\alpha = 0.2 \text{ kg m}^{-3} \text{ K}^{-1}$  is the thermal expansion coefficient, and  $\theta$  is the temperature in degrees Celsius.

**Table 4.1:** Configuration for AG coordinate in the idealised experiments

Parameter	Value	Notes
$S_{\max}$	0.005	Isopycnal slope cutoff
$\omega_z^{\min}$	0.1	Barotropic smoothing component
$\alpha_z$	0.1	Adjustment strength (eq. 3.13)
$\tau_0$	$10^3 \Delta t$	Abyssal restoring timescale
$z_s$	0	Shallow restoring threshold
$z_d$	$D$	Deep restoring threshold

Each of the test cases is run multiple times, varying the horizontal viscosity  $\nu_h$ . This variation in  $\nu_h$  gives control over the lateral grid Reynolds number  $\text{Re}_\Delta = U\Delta x/\nu_h$ , for characteristic velocity  $U$  and horizontal grid spacing  $\Delta x$ . For high enough values of  $\text{Re}_\Delta$ , grid-scale noise saturates at a high level of spurious mixing (see chapter 1). Therefore, by controlling this parameter, we are able to investigate spurious mixing across the spectrum of being dominated by grid-scale noise through to higher viscosities where this is not the case.

All the chosen test cases use flat-bottomed domains, so we present no evaluation of the split RPE method in the presence of topography. The test case from Ilıcak et al. (2012) that does include topography is a downslope flow that involves convective instabilities, which are often parameterised with some form of explicit mixing. However, this is not a limitation of the split RPE method, which is extensible to arbitrary domains using techniques described by Stewart et al. (2014). The downslope flow test case was not selected for this thesis due to concerns by the author regarding the time period over which RPE was calculated: within this time period, the plume had already flowed off the slope and was interacting with the solid end-wall of the domain. An earlier time period would capture the spurious mixing due to the actual overflow process, but the inability for the author to re-run the experiment across all the models led to its omission.

## 4.1 Lock exchange

The lock exchange test case (fig. 4.1) is a simple configuration that shows the creation of intermediate densities by spurious mixing. The test case takes place in a two-dimensional domain 64km long and 20m deep. Only the highest resolution experiments investigated by Ilıcak et al. (2012) are chosen, with horizontal and vertical grid spacings of  $\Delta x = 500$  m and  $\Delta z = 1$  m, respectively. The lock exchange is defined by an initial temperature distribution comprised of one density class on each side of the domain,

$$\Theta(x) = \begin{cases} 5^\circ\text{C} & x < 32 \text{ km} \\ 30^\circ\text{C} & x \geq 32 \text{ km}. \end{cases} \quad (4.1)$$


This case is equivalent to two adjacent basins, each at constant temperature, with a dam between them that is removed at  $t = 0$ . The warm water from the right basin flows from right-to-left above cold water, while conversely cold water from the left basin flows underneath the warm water from left-to-right (fig. 4.1). The flow is a gravity current, for which we have a theoretical prediction for the front velocity in a rectangular channel, given by

$$u_f = \frac{1}{2} \sqrt{gH\Delta\rho/\rho_0}, \quad (4.2)$$


where  $H$  is the domain depth and  $\Delta\rho$  is the density difference across the front, normalised by a reference density,  $\rho_0$  (Benjamin, 1968). However, the expected solution of this test case differs from a physical gravity current in that we expect there to be no region of turbulent mixing, and therefore no intermediate densities, behind the head of the current, as we have disabled all explicit and parameterised mixing.

When calculating the grid Reynolds number, the theoretical front velocity is used as the characteristic velocity scale, to match Ilıcak et al. (2012). All runs were carried out for 17 hours using a baroclinic timestep that satisfied CFL conditions across the range of horizontal viscosities  $\nu_h$  (0.01, 0.1, 1, 10, 100, and  $200 \text{ m}^2 \text{ s}^{-1}$ ).

The time series of normalised RPE (normalised by initial condition RPE) in fig. 4.2a shows MOM6 with the  $z$ -star coordinate having a similar shape to MITgcm and MOM5.




**Figure 4.1:** Snapshots of lock exchange in MOM6 at 6 hours (a) and 17 hours (b) with the  $z$ -star vertical coordinate, and at 17 hours (c) with the AG coordinate, at  $\nu_h = 0.01 \text{ m}^2 \text{ s}^{-1}$  ( $\text{Re}_\Delta = 2.5 \times 10^4$ ). Temperature ( $^\circ\text{C}$ ) is shown in colours.



**Figure 4.2:** (a) Normalised RPE evolution for  $\nu_h = 0.01 \text{ m}^2 \text{ s}^{-1}$  ( $\text{Re}_\Delta = 2.5 \times 10^4$ ) in the lock exchange. MITgcm and MOM5 results come from Ilıcak et al. (2012), while MPAS-O results come from Petersen et al. (2015). (b) Instantaneous rate of RPE change at 17h in the lock exchange for MPAS-O, MITgcm and MOM6 using the  $z$ -star and AG vertical coordinates. The dashed line is at  $\text{Re}_\Delta = 2$  showing the necessary condition for viscosity to damp grid-scale noise.

The actual values of the normalised RPE are much lower in MOM6, showing that less spurious mixing occurs in MOM6. Using the adaptive grid coordinate, the spurious mixing is negligible. Comparing the  $z$ -star and AG snapshots in fig. 4.1, this can be seen through the much narrower front between the density classes, and the significantly reduced volume of mixed fluid at the front. Figure 4.2b shows the dependence on grid Reynolds number  $Re_\Delta$  of the rate of change of RPE at 17h. The instantaneous rate of change of RPE using the adaptive grid coordinate is 1 to 2 orders of magnitude lower than the rest of the runs. This is due to the ability of the adaptive grid coordinate to represent the layered flow in terms of its density classes, which remain separate during tracer advection. By doing so, there is only a limited chance for the separate density classes to mix through vertical remapping.

Using the  $z$ -star vertical coordinate, MOM6 performs better than the other models over the range  $Re_\Delta > 2$  where we expect a saturation of spurious mixing (Ilıcak et al., 2012). However, spurious mixing is further reduced with the use of the adaptive grid coordinate. To better understand why MOM6 performs so well in this case, we next look at the sensitivity to advection order and the horizontal/vertical contributions to spurious mixing.



**Figure 4.3:** Horizontal and vertical contributions to the instantaneous rate of RPE change in MOM6 in the lock exchange at 17h.

The contributions to spurious mixing in MOM6 are broken down in fig. 4.3. This figure shows the instantaneous rate of change of RPE due to separate horizontal and vertical processes, for the two tested vertical coordinates across the range of grid Reynolds numbers. Figure 4.3 shows that for the  $z$ -star coordinate, spurious mixing is predominantly due to horizontal processes. This appears to be at odds with the conclusions of Ilıcak et al. (2012), who found that “high spatial variability in the vertical velocities between the fronts is therefore the dominant cause of spurious diapycnal mixing in this test case.” However, MOM6 uses an upstream momentum advection scheme that reduces spatial variability in the velocity field between the fronts. Vertical velocities in a layered model manifest as changes in local layer thickness, which are redistributed in the regridding/remapping step. However, the spurious mixing across the fronts due to horizontal tracer advection is the leading-order effect. Indeed, for all of the experiments, the average RPE change due to regridding/remapping is negative. Physically, this means that regridding/remapping tends to slightly lower the centre of mass of the domain, counteracting some of the centre of mass increase due to mixing by the advection scheme. The magnitude of this compensation by regridding/remapping is negligible, so the majority of the spurious mixing still occurs

during horizontal tracer advection.

The horizontal tracer advection also dominates the spurious mixing for the adaptive grid coordinate, albeit at a significantly lower order of magnitude. Although the coordinate aims to fit density surfaces, there is some smoothing due to the finite adjustment of layers. There is a small vertical contribution to spurious mixing as interfaces slightly lag the movement of isopycnals with the propagation of the front. This lag causes finite, but small along-layer temperature gradients across the front, which are then mixed to intermediate densities during the horizontal tracer advection step causing the largest increase in RPE. Nonetheless, the horizontal spurious mixing in the adaptive grid coordinate is more than ten times smaller than the  $z$ -star coordinate.

## 4.2 Internal waves


The breaking of nonlinear internal waves in the ocean is a significant source of abyssal mixing, and thus is an important process contributing to the abyssal ocean circulation (Waterhouse et al., 2014). While ocean models do not, in general, represent the breaking of nonlinear internal waves, the propagation of linear internal waves is an important process, for example contributing to the variability of the Atlantic meridional overturning circulation (Blaker et al., 2012) through near-inertial gravity waves. The propagation of internal waves can produce spurious vertical mixing in ocean models, particularly with vertical grids that do not directly represent their propagation, such as  $z$ -star (Gouillon, 2010), which is fixed with respect to non-barotropic motion. However, other vertical coordinates, such as  $z$ -tilde, permit layers to move with the waves, thereby restricting transport between layers and ultimately reducing spurious mixing.

This test case is configured as presented by Ilıcak et al. (2012). It consists of a linearly stratified background temperature distribution in a domain 500 m deep and 250 km wide (fig. 4.4a). The horizontal grid spacing is 5 km, and the vertical grid spacing  $\Delta z$  is 25 m. A wave perturbation is superimposed, lifting the isopycnals in the centre of the domain to set up counter-propagating internal waves towards the left and right horizontal boundaries. This initial condition can be achieved in two separate ways: by altering the thicknesses of constant temperature layers (layered IC), and by altering the temperature in constant thickness layers (level IC).

Experiments were run with horizontal viscosities  $\nu_h$  of 0.01, 1, 15, and 150 m<sup>2</sup> s<sup>-1</sup>, for  $z$ -star,  $z$ -tilde (using level and layered initial conditions), adaptive (using level and layered initial conditions), and continuous isopycnal coordinates. The test case was run for 100 days to allow the waves to propagate many times across the full extent of the domain. The first 10 days are treated as spin-up, however, so the characteristic velocity scale was calculated from the mean domain-averaged kinetic energy from 10 to 100 d. Similarly, the rate of change of RPE was averaged over this same period.

Comparing the model state after the full 100 d evolution (figs. 4.4b to 4.4d) reveals the impacts of both the choice of vertical coordinate and initial condition. Coordinate surfaces are misaligned with temperature surfaces in  $z$ -star, thus lateral advection leads to diffusion associated with the gradient within layers. The  $z$ -tilde and adaptive coordinates successively reduce along-layer temperature gradients and therefore lead to reduced spurious mixing and better preserve the vertical stratification.

It is important to note here that the solution varies substantially depending on the choice of coordinate and initial condition (figs. 4.4b to 4.4d). Due to the dispersive nature of the internal wave problem, we can't expect the solution to remain fully monochromatic when integrated by numerical methods. Using an isopycnal coordinate would give a solution accurate to second order. Additionally, when the  $z$ -tilde or adaptive (fig. 4.4d)




**Figure 4.4:** Snapshot of the internal wave level initial condition (a), and state in MOM6 after 100 days for different vertical coordinates:  $z$ -star (b),  $z$ -tilde (layered initial condition) (c) and adaptive (level initial condition) (d). Temperature ( $^{\circ}\text{C}$ ) is shown in colours, and interfaces are shown with black lines.

simulations are initialised from a level initial condition, temperature within layers may be non-uniform in the final solution, highlighting the strong influence of initial condition. We present the results from the  $z$ -tilde, adaptive and continuous isopycnal coordinates with the layered initial condition only, as we expect that internal wave generation would be associated with isopycnal motion rather than a representative gradient of temperature within layers in these coordinates.

With regard to the total rate of spurious mixing, MOM6 performs better than all the other tested models across the full range of horizontal viscosities with all of the tested vertical coordinates (fig. 4.5a). As a quasi-Lagrangian model, vertical layers in MOM6 are able to move within their column as waves pass through during the dynamics phase of a single timestep, regardless of the chosen coordinate. As this is an adiabatic process, there is no spurious mixing from vertical motions before horizontal tracer advection is performed. The temperature structure that existed at the beginning of the dynamics phase is retained, thus the local lateral gradients along layers are weaker than they would be, had vertical advection been performed. As a consequence, the spurious mixing due to horizontal tracer advection itself is reduced. The more dynamically-relevant temperature gradient in this test case exists in the vertical direction, and we can conclude from this result that the Lagrangian tendency of ALE permits lower spurious mixing in the presence of vertical motions. Indeed, we will see that the contribution from horizontal tracer advection becomes significantly smaller than that of remapping. This inference is further supported by the reductions in spurious mixing with the  $z$ -tilde, adaptive and continuous isopycnal ( $\rho$ ) vertical coordinates, which are all more Lagrangian than  $z$ -star.

Along with the  $z$ -tilde and continuous isopycnal coordinates, which are partially Lagrangian by their definition, it is perhaps less clear that this behaviour also governs the adaptive grid coordinate. Referring to the previous chapter, we can see that the regriding routine for the adaptive grid coordinate is defined only as adjustments applied to the initial grid; there is no absolute motion of coordinate interfaces. As such, the adaptive grid coordinate accommodates the propagation of internal waves primarily through the movement of these coordinate interfaces, leading to weak spurious mixing, similar to the



**Figure 4.5:** (a) Averaged rate of RPE change from 10 to 100 days in the internal waves test case. MITgcm and MOM5 results come from Ilıcak et al. (2012), while MPAS-O results come from Petersen et al. (2015). MPAS-O  $z$ -tilde uses  $\tau_{Dlf} = 100$  d. (b) Contributions to spurious mixing by horizontal and vertical processes in the internal waves test case in MOM6. Each contribution is the time-averaged total rate of RPE change from 10 to 100 days.

$z$ -tilde coordinate.

The full  $z$ -tilde coordinate, defined by Leclair and Madec (2011) and implemented in MPAS-O (Petersen et al., 2015) is governed by a simple philosophy. This is to treat the vertical motions associated with barotropic and high-frequency baroclinic motions as Lagrangian. This is an extension of  $z$ -star behaviour, allowing, for example, the propagation of internal waves. As a corollary of this, the remaining low-frequency baroclinic motion is treated as Eulerian, as usual. There are two key timescales in the  $z$ -tilde coordinate:  $\tau_{Dlf}$  is the cutoff period of a first-order low-pass filter for defining the low-frequency baroclinic motion. As this period increases, the coordinate converges to being purely Lagrangian. The second timescale,  $\tau_{hhf}$  governs restoring back to a  $z$ -star coordinate, preventing long-term drift. This restoring acts as a first-order high-frequency filter. In combination, the two timescales act to provide tuning for the range of baroclinic frequencies treated as Lagrangian.

As MOM6 uses a quasi-Lagrangian ALE implementation, vertical velocities are implicit in the regridding/remapping stage. Additionally, motions are Lagrangian within a single timestep. The simplified  $z$ -tilde coordinate in the model can then be described as only the restoring term

$$\tilde{z}_k^{n+1} = \tilde{z}_k^n + \frac{\Delta t}{\tau_{hhf}} (z_k^* - \tilde{z}_k^n), \quad (4.3)$$

where  $\tilde{z}_k^{n+1}$  are interface positions after regridding,  $\Delta t$  is the regridding/remapping timestep (taken here as the baroclinic timestep), and  $\tau_{hhf}$  is the high-frequency thickness restoring timescale to the usual  $z$ -star coordinate  $z_k^*$ . To match the MPAS-O configuration, we have taken  $\tau_{hhf} = 30$  d. The full  $z$ -tilde coordinate also allows for low-frequency motions to be treated in an Eulerian manner. However, as the low-pass filter timescale

$\tau_{Dif} = 100$  d used in MPAS-O spans the entire length of the simulation, we haven't included it in the formulation here.

For all of the coordinates tested in MOM6 (shown in red in fig. 4.5a), we compute the horizontal and vertical contributions to the spurious mixing (fig. 4.5b). Consistent across all the MOM6 configurations is the dominance of spurious mixing by vertical processes in this test case, shown by orange points in fig. 4.5b exceeding their corresponding blue points. The continuous isopycnal coordinate is the most Lagrangian, so that the movement of the grid due to regridding under internal wave forcing is the smallest. In turn, remapping is less destructive and causes less mixing than in the other coordinates.

Following the continuous isopycnal coordinate, the adaptive and  $z$ -tilde coordinates are also Lagrangian. However, in both cases, when initialised from a level initial condition where layers have varying density, total spurious mixing is higher than with a layered initial condition (not pictured). This is because there is no provision in the coordinate itself to enforce layers of constant density. In effect, its Lagrangian nature leads to a significant improvement in the horizontal spurious mixing, accompanied by only a modest improvement in the vertical contribution. As this test case is dominated by the vertical spurious mixing, it is only this modest improvement that shows up in the total spurious mixing. Initialising this test case from a layered initial condition, such as that used for the continuous isopycnal coordinate, has a more significant impact on spurious mixing.

Although it is also a Lagrangian coordinate, the breakdown between horizontal and vertical contributions to spurious mixing in the adaptive grid coordinate leads to an interesting story. As we have seen above, the expected behaviour for a Lagrangian coordinate is a large component of vertical spurious mixing, as regridding/remapping must accommodate interface motions due to the propagation of internal waves. However, the competing adjustments of density adaptivity and vertical restoring lead to a slow erosion of the density structure maintained by layers: density adaptivity is performed first on a nearly-isopycnal state, and so gives a minimal adjustment; vertical restoring occurs regardless and moves the interfaces away from this near-isopycnal state. This is evident in the split RPE contributions, as there is a larger horizontal contribution associated with the along-layer temperature gradients introduced by the vertical restoring term.


As the internal waves test case involves primarily vertical motions, the along-layer gradients remain small regardless of the chosen vertical coordinate. The spurious mixing due to horizontal tracer advection is approximately an order of magnitude lower than that due to regridding/remapping. The  $z$ -tilde and continuous isopycnal coordinates are strongly Lagrangian in this test case, which, without any explicit mixing leads to very little along-layer advection of temperature (equivalent to density in this test case). As expected, the spurious mixing due to horizontal tracer advection with the  $z$ -star coordinate is much higher, up to 5 times that of the more Lagrangian coordinates.

### 4.3 Baroclinic eddies

The previous two test cases were two-dimensional and also did not incorporate the influence of the Coriolis force. We now introduce a third test case from Ilıcak et al. (2012) that involves a baroclinically unstable temperature front in a periodic channel with rotation, described by Petersen et al. (2015). This front generates vigorous eddying without either mechanical or buoyancy forcing, thus it is a closed system suitable for analysis by changes in RPE. The domain is a periodic channel in the  $x$  direction, 160 km wide by 500 km long, with a depth of 1000 m at vertical resolution  $\Delta z = 50$  m (fig. 4.6a).

In order to enhance baroclinicity, a strong quadratic bottom drag with drag coefficient  $C_D = 0.01$  is used. Experiments were performed at horizontal resolutions of 1, 4, and






**Figure 4.6:** (a) Snapshot of initial condition of surface temperature for the baroclinic eddies test case in MOM6 at 1 km horizontal resolution. The temperature perturbation can be seen at the third trough in the sinusoidal front. Snapshots of surface temperature ( $^{\circ}\text{C}$ ) after 320 days of simulation at 1 km horizontal resolution with the adaptive grid coordinate, for (b)  $\nu_h = 1 \text{ m}^2 \text{ s}^{-1}$  ( $\text{Re}_{\Delta} = 88$ ) and (c)  $\nu_h = 200 \text{ m}^2 \text{ s}^{-1}$  ( $\text{Re}_{\Delta} = 0.25$ ).

10 km. For each choice of horizontal resolution, the horizontal viscosities  $\nu_h$  were 1, 5, 10, 20, and  $200 \text{ m}^2 \text{ s}^{-1}$ , giving a range of lateral grid Reynolds numbers from  $O(1)$  at the highest viscosity, through to  $O(1000)$  for the lowest viscosity. To calculate the grid Reynolds number, the characteristic velocity scale was calculated from the mean of the domain-averaged kinetic energy.

Figures 4.6b and 4.6c show the surface temperature after the full 320 days of simulation at 1 km horizontal resolution, at the lowest and highest viscosities,  $1 \text{ m}^2 \text{ s}^{-1}$  and  $200 \text{ m}^2 \text{ s}^{-1}$ , respectively, using the adaptive grid coordinate. In the low viscosity case, stronger spurious mixing has occurred, but finer-scale features are also evident. Conversely, the volume of mixed fluid is significantly less with a higher horizontal viscosity, but the eddies are much weaker due to the momentum damping by the viscosity. Due to the eddying nature of this simulation, the solution is strongly affected by the momentum closure, which is set by the viscosity. Identifying any changes in the solution due to spurious mixing itself is a secondary concern in this study.

At all horizontal resolutions, the spurious mixing in MOM6 with the  $z$ -star coordinate, MOM5 and MITgcm is very similar (fig. 4.7), with a modest improvement in MOM6 over MITgcm in many cases. However, the differences between these models are marginal compared to MPAS-O, as discussed by Petersen et al. (2015).

The adaptive grid coordinate has significantly different characteristic spurious mixing behaviour to all models using the  $z$ -star coordinate. At the lowest horizontal resolution, spurious mixing in the adaptive grid coordinate is on par with, or below that of MPAS-O. At 4 km horizontal resolution, mixing is still improved compared to the  $z$ -star models, and at the highest horizontal resolution, mixing is slightly higher than  $z$ -star in MOM6. There are two observations from these results that we will discuss in further detail: the relative performance of the adaptive and  $z$ -star coordinates as horizontal resolution changes; and the flat (or negative) trend of spurious mixing with grid Reynolds number.



**Figure 4.7:** Average rate of RPE change for all baroclinic eddy experiments over the entire 320 day run. MITgcm and MOM5 results come from Ilcak et al. (2012), while MPAS-O results are from Petersen et al. (2015). (a)  $\Delta x = 10$  km, (b)  $\Delta x = 4$  km, (c)  $\Delta x = 1$  km.

Firstly, the magnitude of the total spurious mixing in MOM6 with the adaptive coordinate is insensitive to horizontal resolution compared to the other models. As such, the reduction in the relative improvement in spurious mixing over the  $z$ -star coordinate can be attributed to an improvement of  $z$ -star with resolution, rather than a regression with the adaptive grid coordinate. This suggests that the spurious mixing in the adaptive grid coordinate isn't controlled by horizontal processes as we may intuitively expect for  $z$ -star.


The second observation is that spurious mixing with the adaptive grid coordinate is insensitive to grid Reynolds number. Indeed, the trend is slightly negative at 4 km horizontal resolution, in contrast to the other results. If truncation errors in advection were the dominant contribution to the spurious mixing, we would expect the increased grid-scale variance in velocity due to reduced viscosity to saturate at sufficiently high grid Reynolds numbers (Ilıcak et al., 2012). In combination with the insensitivity to horizontal resolution, we can conclude that spurious mixing with the adaptive grid coordinate is controlled differently than with the  $z$ -star coordinate.

The final coordinate in MOM6 is the continuous isopycnal coordinate, for which the total average rate of RPE change is negative for all configurations, so it is not shown in fig. 4.7. With this coordinate, the magnitude of the RPE change varies from half to an order of magnitude lower than the  $z$ -star configuration. As the difference in the models except MPAS-O is slight, we instead turn our attention to the contributions to the spurious mixing in MOM6 and the impact of resolution.

We consider the horizontal and vertical contributions to spurious mixing at each of the tested horizontal resolutions (1, 4, and 10 km) with the  $z$ -star coordinate in fig. 4.8a. As horizontal resolution increases (grid spacing decreases), the spurious mixing present in horizontal tracer advection decreases, accompanied by a smaller increase in the spurious mixing by regridding/remapping. Therefore the fraction of the total spurious mixing contributed by regridding/remapping increases with horizontal resolution. Indeed, at 1 km horizontal resolution, regridding/remapping contributes the majority of the total spurious mixing across the full range of viscosities. The fact that the spurious mixing contributed by regridding/remapping stays relatively constant suggests that the implicit vertical velocities are physical, occurring with larger horizontal scales than the grid-scale. However, we speculate that the increase of this contribution as horizontal resolution decreases can be attributed in part due to the larger number of water columns in which regridding/remapping must be applied. A lower bound on the per-column baseline spurious mixing due to remapping, would then show up as a nonlinear contribution which is affected by horizontal resolution.

As shown in fig. 4.8a, the vertical component of spurious mixing only approximately doubles for a tenfold increase in horizontal resolution when using the  $z$ -star coordinate. Figure 4.8b shows the effect of doubling ( $\Delta z = 25$  m) and halving ( $\Delta z = 100$  m) the vertical resolution, at a constant 1 km horizontal resolution. With insufficient vertical resolution to fully resolve the baroclinic structure of the flow ( $\Delta z = 100$  m), spurious mixing due to horizontal tracer advection is elevated, whereas the convergence of the horizontal contributions at  $\Delta z = 25$  and 50 m shows that excess vertical resolution has little impact. However, the vertical contribution continues to improve as resolution increases. This benefit is marginal, as the flow is already resolved in both the lateral and vertical grids.

Turning our attention to the adaptive grid coordinate, we can see markedly different behaviour in the horizontal/vertical split for varying horizontal resolution (fig. 4.8c). The most striking feature here is the negative trend in vertical spurious mixing with grid Reynolds number at the 4 and 10 km resolutions. While the split RPE diagnostic allows us to perform the decomposition to see this trend, it does not explain its origin. A



**Figure 4.8:** Average rates of RPE change for horizontal and vertical processes over the full simulation of the baroclinic eddies test case, showing spurious mixing contributions in MOM6 across the range of horizontal viscosities with the  $z$ -star (a,b) and AG (c,d) coordinates: (a,c) each horizontal resolution at 50 m vertical resolution, (b,d) varying vertical resolution at 1 km horizontal resolution.

possibility is that the adaptive grid coordinate performs poorly at these low horizontal resolutions: the adjustment of an interface within a single column affects a larger area at these low horizontal resolutions. The coordinate may barely adjust interface positions, so the regridding/remapping operation does not significantly contribute to the overall magnitude of spurious mixing. However, if the coordinate is unable to exactly maintain isopycnal coordinate surfaces, the resultant lateral gradients affect a larger water mass than they would for a higher horizontal resolution. This behaviour is essentially the same as we see in the  $z$ -star coordinate (fig. 4.8a), which by definition does not adjust according to the local density structure.

At a high horizontal resolution, the adaptive grid coordinate has a more significant impact (fig. 4.8d). Here, there is sufficient horizontal resolution for the coordinate to adjust towards variations in the local density structure, so we see a higher contribution from the vertical component. Commensurate with this adjustment to density structure is a reduction in the horizontal spurious mixing, as the along-layer temperature gradients are reduced by finding an arrangement of coordinate surfaces that more closely aligns with isopycnals. The increase in vertical resolution also has the expected effect of limiting any spurious mixing to a smaller volume, leading to an improvement at the expense of higher computational and storage demands.

This result demonstrates that increasing vertical resolution provides another way of reducing spurious mixing for a given vertical coordinate, which is particularly important at high horizontal resolutions, such as the transition to eddy resolving in the horizontal grid, shown here.

## 4.4 Summary

In this chapter, we have outlined the spurious mixing performance of MOM6 against other models in three idealised test cases at various resolutions, with different vertical coordinates. The split RPE method introduced in chapter 2 allowed for the separation of contributions to spurious mixing by horizontal and vertical processes. Across the suite of test cases, spurious mixing in MOM6 was, on the whole, reduced compared to the other tested models. This result can be attributed mainly to the strong performance of the regridding/remapping implementation of ALE, which significantly reduces the vertical contribution to spurious mixing.

The other major advantage of ALE is the flexibility to implement arbitrary vertical coordinates. The  $z$ -tilde and continuous isopycnal coordinates both showed improvements over  $z$ -star in sustaining the propagation of internal waves while reducing spurious mixing resulting from their motion. Additionally, the adaptive grid coordinate introduced in chapter 3 reduced spurious mixing in the three idealised test cases, although the improvement over the other models reduced with higher resolution in the baroclinic eddies test case. This result can be attributed to a significantly higher contribution by vertical processes at this resolution. This baroclinic eddies test case is the most realistic of the three idealised cases presented here, but motivates a comparison of the adaptive grid coordinate in an experiment more comparable to the real-world ocean. In particular, we now look to determine whether the improvement in spurious mixing when using the adaptive grid coordinate manifests as an improvement of the circulation itself.



---

# Testing in a sector model

---


A preliminary evaluation of the adaptive grid coordinate was performed in the previous chapter. Using idealised test cases, spurious mixing was diagnosed through increases in the reference potential energy. The split RPE technique (introduced in chapter 2) then allowed for the attribution of spurious mixing to horizontal and vertical processes separately. However, accurate diagnosis of spurious mixing through increases in reference potential energy requires that these idealised simulations are conducted in the absence of buoyancy forcing, and with all explicit mixing sources disabled.

The adaptive grid coordinate performed well in these idealised test cases (chapter 4). However, more comprehensive tests are required to document its performance in realistic basin- to ocean-scale simulations. The restrictions of the RPE-based methods preclude more realistic simulations, as key circulation features (such as a meridional overturning circulation) could not be obtained without buoyancy forcing or a background diapycnal mixing. To further evaluate the adaptive grid coordinate, we make use of a more realistic model that includes surface buoyancy forcing and explicit diapycnal mixing. We are unable to use the RPE method to directly evaluate spurious mixing in this case; instead, we evaluate physical properties of the circulation, such as its strength or structure, and confirm numerical stability in a more comprehensive configuration. Diagnostics include the meridional overturning circulation streamfunction, density structure, and circumpolar transport.

## 5.1 Model configuration

The model we use is an idealised representation of the Atlantic Ocean (Hogg et al., 2013). The domain is  $40^\circ$  wide, and spans the latitude range  $70^\circ\text{N}$  to  $70^\circ\text{S}$ . Horizontal grid spacing is  $1/4^\circ$ . While this resolution is only marginally eddy-resolving, we don't use an explicit eddy parameterisation. The domain is 4 km deep through the ocean basin.

As an idealised sector model, the bathymetry is simplified compared to the real-world Atlantic Ocean (fig. 5.1d). The side walls run directly north-south, rather than following continental boundaries, and there is no representation of the continental shelves along these walls. Instead, they slope toward the bottom of the basin in a Gaussian profile over 5 degrees of longitude. In contrast, the Antarctic continental shelf is represented, at 900 m deep. This shelf is designed to allow the formation of dense Antarctic Bottom Water that is essential for the abyssal cell of the meridional overturning circulation (Orsi et al., 1999). Between  $55^\circ\text{S}$  and  $65^\circ\text{S}$  there is a reentrant channel with a sill of 2000 m depth, representing the Drake Passage which permits a zonally-unbounded circumpolar current.



**Figure 5.1:** Surface forcing and topography used for the sector simulation: (a) Zonally-uniform zonal wind stress profile. (b) Surface temperature restoring profile. (c) Surface freshwater forcing through an evaporation profile. (d) Bottom topography field.

The model is run with a linear equation of state,

$$\rho = \rho_0 - \alpha\Delta T + \beta\Delta S, \quad (5.1)$$

with  $\alpha = 0.2$  and  $\beta = 0.8$ . Both mechanical and buoyancy forcing are present, applied over a 0.5 m deep surface boundary layer. A zonally-uniform, time-independent, zonal wind stress is applied with the profile shown in fig. 5.1a. Temperature is applied as a restoring toward the surface state shown in fig. 5.1b, using a piston velocity of  $0.5 \text{ m d}^{-1}$ . Finally, salinity is forced by a globally-balanced freshwater flux, shown in fig. 5.1c. Imposing this form of buoyancy forcing through a flux rather than a restoring allows for the driving of continuous dense water formation processes at the poles, while not altering the global mass balance (achieved by the global balance of the flux).

In contrast to the idealised experiments presented in the previous chapter, explicit diapycnal mixing is imposed through a vertical diffusivity. A background diffusivity of  $1 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$  is applied, however diffusivity near the surface is also elevated by the energetically-constrained ePBL scheme (Reichl and Hallberg, 2018).

Spinup for the model occurred in two stages. Initially, the  $z$ -star coordinate model was spun up. The configuration initially matched that specified above, but with a meridionally-symmetric freshwater flux. After approx. 200 years, the overturning circulation completely shut off as very light water formed in the south, unable to form the abyssal overturning cell. To rectify this, the simulation was rolled back to year 143, where the overturning circulation was still present. The freshwater flux was modified to the profile shown in fig. 5.1c by reducing the strength of the flux in the southern hemisphere, while retaining a global balance. This allowed for the formation of sufficiently dense water on the southern shelf, maintaining the global overturning. This stage of the spinup was then run for a further 70 years to settle into a new equilibrium, giving 213 total years of spinup. Each of the three experiment configurations, described below, was run for a further 110 years after this point, with the final 10 years of this period being averaged for the results.

This idealised sector model is intended to provide a framework in which multiple parameter configurations can be compared. In particular, we compare two vertical coordinates that are already used in ocean modelling,  $z$ -star and the MOM6 implementation



of the HYCOM coordinate, HYCOM1 (chapter 1), against the adaptive grid (AG) coordinate (chapter 3). We treat  $z$ -star as the control against which the other two coordinates are compared. The parameters governing AG are shown in table 5.1. To ensure as fair a comparison as possible, each coordinate configuration is based on a common vertical resolution profile. This profile is defined as  $\Delta z(k) = k^{4.5}$ , scaled such that it spans 4000 m (the full depth of the domain), with a minimum grid spacing of 2 m at the surface. This vertical resolution profile is used as the base profile for  $z$ -star directly, where each individual column is scaled according to the free-surface height.

**Table 5.1:** Configuration for AG coordinate in the sector model

Parameter	Value	Notes
$S_{\max}$	0.01	Isopycnal slope cutoff
$\omega_z^{\min}$	0.1	Barotropic smoothing component
$\alpha_z$	0.1	Adjustment strength (eq. 3.13)
$\tau_0$	$10^6 \Delta t$	Abyssal restoring timescale
$z_s$	50 m	Shallow restoring threshold
$z_d$	500 m	Deep restoring threshold


The HYCOM1 coordinate configuration uses the base vertical resolution profile to define the minimum spacing between interfaces. This spacing prevents massless layers near the surface, which is one of the difficulties faced by pure isopycnal coordinates. In addition to the minimum spacing, and in the same manner as isopycnal coordinates, HYCOM1 requires nominal densities for each interface. To generate the density profile, diagnostic output of potential density was taken from the end of the spinup period. From this density profile, a global histogram of the density distribution was generated, and then split into even bins according to the number of vertical levels required. This distribution was found to be insufficient, however. At higher densities, small changes in density span large volumes. Additionally, it is less important that exact densities are followed at the surface, as the minimum spacing mechanism tends to dominate. Because of this, the edges of the bins were tweaked to favour higher resolution in density space toward the bottom.

## 5.2 Stratification

We first consider the density structure of the model ocean with the different coordinate configurations. This is the field in which we expect different representations by the different coordinates. In order to yield a direct comparison between the coordinate configurations, density was remapped online onto a common  $z$ -space grid, regardless of the actual coordinate used for computation. Figure 5.2 shows a 10-year zonal mean of the potential density as a transect. The actual density values are shown for the control  $z$ -star configuration in fig. 5.2a, with the other subpanels showing anomalies from this control.

There are some striking similarities in the stratification anomaly between the HYCOM1 and adaptive grid coordinates. The very southern edge of the domain is more dense than the control, and in a narrower band. This density anomaly is stronger in HYCOM1 than adaptive grid. This mirrors the changes seen below in the dense water formation region of the lower cell of the overturning circulation. The characteristic narrowing structure that was seen in the upwelling region of the lower cell can also be seen clearly in the midlatitudes for both coordinates. A feature similar to this is mirrored in the northern hemisphere.

Around the equator there appears to be significantly denser water in HYCOM1 between 20°S and 20°N, at about  $0.3 \text{ kg m}^{-3}$  denser than both the control and AG. However, this



**Figure 5.2:** (a) Zonal mean density structure (10-year average) in  $z$ -star control configuration; (b) Anomaly from control with adaptive grid coordinate; (c) Anomaly with HYCOM1 coordinate.


coincides with the lightest water over the entire domain, and in this context is a relatively small change. Because the HYCOM1 coordinate requires specification of density levels for each coordinate interface, changes in surface density due to the surface buoyancy forcing may diverge from the original density levels. At higher latitudes, where density gradients are weaker, overall changes in density have a greater impact.

The mid-to-high latitude northern hemisphere sees the greatest differences between the HYCOM1 and adaptive grid coordinates. In HYCOM1 (fig. 5.2c), water sinking right along the northern boundary is more dense than in the control. Curiously, this feature isn't replicated at all when using the adaptive grid coordinate. However, looking further south, between approximately 30°N and 60°N, the lateral density gradient has been significantly eroded in HYCOM1, whereas the sharp front present in the control case can also be seen in the adaptive grid coordinate. As we will see in the section below, the specifics of the coordinate representations in the northern hemisphere are altering the dynamics slightly, but it is unclear if either is “more correct” than the other.

### 5.3 Meridional overturning circulation

We have seen in the previous section that each of the different coordinates represents the stratification of the model ocean differently. To show the impact that the stratification has on the circulation, we will now look at the meridional overturning circulation. This is an averaged north-south circulation that arises due to water mass transformations and latitudinal variation in wind stress curl, which causes divergent Ekman transport at the

surface. A schematic for the overturning circulation in the real ocean is shown in fig. 5.3, and one of the key features of the sector model is that it retains the mechanisms that drive this circulation. It is expected that smaller spurious density transformations in the interior result in a stronger overturning circulation. This particularly targets processes such as dense water overflows and upwelling. The meridional overturning circulation is presented as a latitude-density streamfunction,  $\Psi(y, \sigma)$ , showing the strength and direction of the circulation as it travels along isopycnals or is transformed to different densities. In order to perform these density-space calculations, diagnostic output is remapped from the coordinate used for calculation (e.g.  $z$ -star, adaptive or HYCOM1) to a common density-space grid online. This uses the same high-order interpolation and remapping routines that are used in the ALE core of MOM6.



**Figure 5.3:** A schematic view of the meridional overturning circulation and the mechanisms that drive it *Fig. credit: Robert Marsh, National Oceanography Centre Southampton*

Figure 5.4a shows the latitude-density streamfunction for the control configuration. The characteristic two-cell structure in this model is relatively weak, however there are still two clearly-defined cells from which we can infer sensitivity to coordinate choice. The southern/lower cell (in blue) has a counter-clockwise sense of circulation, bringing dense waters formed at the high-latitude surface downward and northward, with a maximum transport of 13.8 Sv. These waters upwell as they transit back southward, getting lighter until they reach the surface and close the cell. A similar process happens for the northern/upper cell (in red), which is also stronger at 15.8 Sv. Notably, this cell is less dense (as seen in fig. 5.2a), so it overlies the lower cell when they meet at the equatorial abyss.

Now consider the adaptive grid coordinate (fig. 5.4b), shown as an anomaly from the control overturning streamfunction. The most interesting and relevant differences occur in the lower cell. At the dense water formation region, the downwelling water mass is constrained to a narrower band of density and reducing from 20° to 10° latitude over the control, however it remains at a lower density than in the control. This suggests that the adaptive grid coordinate’s density adaptivity regime captures this downwelling dense water with less diffusivity than the control using the  $z$ -star coordinate, but that it is also modified to lighter densities at the surface.

Following the downwelling dense water northward along the basin bottom, we can see further differences due to the adaptive grid coordinate. This abyssal water flows in a model layer of more uniform density than with the  $z$ -star coordinate, so that it remains closer to the original density until mixing with the upper cell modifies the properties from 50°S to the equator. Broadly however, the lower cell is weaker in the adaptive grid coordinate as it travels northward, and has a maximum transport of 12.5 Sv.

The final major difference in the representation of the lower overturning cell with the adaptive grid coordinate can also be attributed to the density-following behaviour. In

particular, the upwelling arm of the lower cell (densities lighter than  $1036.5 \text{ kg m}^{-3}$ ) is also constrained to a tighter band of the overturning streamfunction. Alongside this, the circulation is also stronger in this region of the overturning streamfunction, at  $20.5 \text{ Sv}$ , a significant increase over the  $z$ -star control. The narrower, stronger circulation is expected as a consequence of more closely matching the circulation structure with the coordinate itself, reducing the diffusion that arises from spurious numerical mixing.

The reduced strength of the lower overturning cell in the adaptive grid coordinate is a surprising result. By more closely representing the density structure, the spurious mixing should decrease, leading to lower overall diapycnal mixing. Operating in a weak mixing regime, the overturning strength is proportional to the diapycnal mixing (Nikurashin and Vallis, 2011). However, this should be accompanied by an increase in the stratification, which we do not clearly see (fig. 5.2b). One possible explanation is that the combination of a weaker overturning and lighter surface density (due to different surface representation in the adaptive grid coordinate) propagates these light waters into the interior, eroding the stratification. Further exploration would be required to determine the full mechanism behind the altered overturning circulation with the adaptive grid coordinate, but that is beyond the scope of this study.


The goals of the HYCOM1 coordinate (fig. 5.4c) are similar to the adaptive grid coordinate: maintain sufficient resolution at the surface to resolve the physics occurring in the region, but follow density surfaces in the abyss where the circulation tends to be adiabatic. As such, the differences in overturning circulation when using this coordinate are similar to the adaptive grid coordinate, albeit not as pronounced. The downwelling dense water is again less modified as it travels from the surface to the ocean floor, and there is a slight sharpening of the circulation in the upwelling region. This causes a slightly stronger abyssal cell than the control and adaptive grid cases, with a maximum transport of  $14.2 \text{ Sv}$ . These effects can be attributed to the predominantly isopycnal behaviour of the HYCOM1 coordinate for these parts of the circulation.

Curiously, the HYCOM1 coordinate strongly affects the downwelling region of the upper cell. This manifests as a stronger circulation at a higher density at  $\sim 50^\circ \text{N}$  than the control case. One possibility is that the sharp slope at the northern boundary permits the water mass to flow downward and southward without being impeded by topography. The stronger downwelling here leads to a strengthened upper cell ( $22.9 \text{ Sv}$ ) as it flows southward until its interaction with the lower cell. A similar sort of modification to the upper cell can be seen with the adaptive grid coordinate, but to a reduced extent compared with HYCOM1. These coordinate comparisons show the impact of resolving abyssal flow with coordinates that more closely match the isopycnal structure of the ocean.

## 5.4 Circumpolar transport

In addition to meridional circulation, the zonal transport in the Antarctic Circumpolar Current (ACC) plays an important role in the global oceans. The ACC is unique in that it is zonally-unbounded, impacting the dynamics of the Southern Ocean by preventing mean meridional flow across its span. It is also linked to the stratification of the meridional overturning circulation (Hogg, 2010) and feeds back via the Southern Ocean eddy field (Marshall and Radko, 2003). To get a sense of the representation of the ACC in a model, the strength of its transport can be measured through the constriction of Drake Passage. We ensure that there is a zonally-unbounded channel in the sector model so that these mechanisms can be explored.

In the sector model, the zonally-unbounded region between  $55^\circ \text{S}$  and  $65^\circ \text{S}$  leads to the formation of an eastward jet, similar to the Antarctic circumpolar current (ACC)




**Figure 5.4:** (a) Meridional overturning circulation (10-year average) in control configuration; (b) Anomaly from control with adaptive grid coordinate (colours) overlaid with contours of total streamfunction from  $-16\text{ Sv}$  to  $16\text{ Sv}$ ; (c) Anomaly with HYCOM1 coordinate (colours) overlaid with contours of total streamfunction, as above. Note that the vertical axis is non-linear to give higher resolution toward higher densities.

of the ocean. The transport in Sverdrups through the strait for each of the coordinate configurations is shown in fig. 5.5. With a 6-month moving average, transport in the control configuration was steady at around 24 Sv. After a few years' equilibration, the transport increases with either of the more density-following coordinates, HYCOM1 or adaptive grid.

The exact mechanisms governing the ACC are not known, but for simulations of the real ocean, it can be linked to the overturning circulation (Gent et al., 2001). In addition, in the idealised sector model itself, the ACC transport scales linearly to first order with the sum of the Ekman transport and the abyssal overturning (Shakespeare and Hogg, 2012). Given that the horizontal resolution remained constant across the three experiments, and that vertical resolution near the surface should not be significantly altered in HYCOM1 or adaptive grid compared to the control, we expect that the Ekman transport also does not vary strongly across the experiments.

However, if the Ekman transport is approximately constant across experiments, then the small differences in abyssal cell transport alone are insufficient to explain the significant variations in ACC transport. We make the observation that Shakespeare and Hogg (2012) used the MITgcm numerical model, which developed a three-layer meridional overturning structure. Using MOM6, only two layers can be identified, independent of the vertical coordinate (fig. 5.4). Due to this, it is unclear whether the scaling of Shakespeare and Hogg (2012) is applicable in this case. Regardless, the HYCOM1 and adaptive grid coordinates lead to increases in ACC transport due to differences in representing the physical flow field.



**Figure 5.5:** Comparison of strength of circumpolar transport (“ACC”) across the three experiment configurations (6-monthly moving average). Both the adaptive and HYCOM1 coordinates increase the transport strength over the control.

## 5.5 Computational performance

The added complexity of the adaptive grid coordinate comes at the cost of requiring extra computation during the regridding phase. In particular, computing the new grid for a given column requires incorporating data from adjacent columns. This computation requires in turn that the “halo” points surrounding a given processing element’s computational domain are up to date, incurring a possible communication overhead. The sector model allows us to evaluate the computational impact of the adaptive grid coordinate in terms of additional walltime required.

The walltime performance for the regridding and remapping routines is shown in table 5.2. This table shows the average time taken in the routines across all 280 processing elements, over 20 runs of 24 months each. The data is separated by the coordinate used for the run configuration. The first observation is that the time for remapping is insensitive to the coordinate. This result is expected, as the remapping routine is called for every model column during ALE, and must always perform interpolation and redistribution from the old grid to the new grid generated during regridding. In that sense, the time taken in the remapping routine is insensitive to the input data.

The more interesting performance data comes from the regridding routine.  $z$ -star is the fastest coordinate for regridding calculations: it takes into account only the target grid, local bathymetry depth and local free surface height. As such, it is straightforward and parts of the algorithm may be easily vectorised by a compiler. The HYCOM1 regridding routine is almost exactly four times slower than  $z$ -star, requiring a sweep through the column to determine the depth of each isopycnal. Using a non-linear equation of state could increase the time taken in this routine, as the problem of finding isopycnal depths requires iteration in this case. Finally, the regridding routine using the adaptive grid coordinate is almost ten times slower than  $z$ -star. However, even in this worst case, regridding with the adaptive grid coordinate makes up only approximately 3-4% of the total runtime. This ends up as about an extra 3 minutes on top of the 1.75 hours to run the entire 24 months. Relative to the potential increases in numerical performance, this may be an acceptable tradeoff. Indeed, compared to other coordinates, such as  $z$ -tilde, which required a halving of the model timestep (Leclair and Madec, 2011), the additional cost of adaptive grid is minor.

**Table 5.2:** Walltime performance for running 24-months of the sector model simulation across the three coordinates. Averages are taken over all 280 CPUs over 20 runs. The increase in regridding time is calculated relative to the  $z$ -star coordinate.

Coordinate	Routine	Average time (s)	Standard deviation	Relative increase
$z$ -star	Regridding	22.72	2.795	1.00
	Remapping	267.2	21.56	
HYCOM1	Regridding	90.71	11.72	3.99
	Remapping	273.0	21.56	
AG	Regridding	214.3	57.33	9.43
	Remapping	289.6	25.75	

## 5.6 Summary

Having evaluated the adaptive grid coordinate for its performance in idealised test cases in chapter 4, we demonstrated the coordinate in a more realistic configuration. This config-

---

uration was modelled as an idealised representation of the Atlantic Ocean, capturing key features such as a meridional overturning circulation and circumpolar current. Although this configuration could not be evaluated analytically (e.g. quantifying spurious mixing), physical diagnostics were computed from the simulations in different coordinates. These diagnostics depend strongly on stratification and the representation of adiabatic flows.

In the control configuration, the  $z$ -star coordinate was used, which does not directly represent the density structure of the ocean at all, and as such suffers spurious mixing in the representation of adiabatic flows. Using either of the HYCOM1 or adaptive grid coordinates, which attempt to follow isopycnal structure in the abyssal ocean, increases were seen in the strength of circumpolar transport and mean stratification. As spurious mixing tends to weaken diagnostics such as these, we interpret their strengthening as an improvement due to the coordinate. Further, the adaptive grid coordinate's flexibility lies in following density structure only locally, where it has the most significant impact on spurious mixing during advection. This can be seen particularly in the stronger circumpolar transport. The lower meridional overturning cell likely weakens due to the reduced spurious mixing. However, more information is required to determine if this is indeed the case. An example experiment may be to investigate the impact of imposed mixing on the overturning circulation with different vertical coordinates. With stronger explicit mixing, reductions in spurious mixing have a smaller overall impact.

The changes seen in the circulation and transport can be attributed to the different representations of the ocean stratification with the different coordinates. The HYCOM1 and adaptive grid coordinates are able to represent density-dependent features such as regions of sinking dense water or upwelling lighter water. We suggest that they do so by maintaining higher resolution in density space, as well as limiting the crossing of coordinate interfaces in the adiabatic regime. This means that spurious mixing of these watermasses is reduced, so density gradients remain sharper, leading to stronger circulation. The localisation of spurious mixing was outside the scope of this study, but may be achieved through techniques such as diagnosing tracer variance decay (Burchard and Rennau, 2008) or quantifying the reference potential energy density changes (Ilıcak, 2016).

An important consequence of the reduced spurious mixing observed with AG is its impact on parameterisations. In particular, there is more control over mixing processes. When a coordinate exhibits spurious mixing, this manifests as an unavoidable background level of mixing. This background level of mixing is spatially varied, and unknown in magnitude. Mixing parameterisations may not reduce mixing below this threshold, and are likely to cause excessive mixing due to its presence. By reducing spurious mixing, we have more precise control over the overall mixing through both the explicit mixing terms and parameterisations.

The adaptive grid coordinate is not a silver bullet, however. While it performs well at representing the physical circulation, it comes with a higher computational cost. Compared with the  $z$ -star coordinate, regridding with the adaptive grid coordinate takes approximately ten times longer. However, this still makes up only 3-4% of the entire runtime of the model, so in absolute terms the additional time is a minimal cost.

Demonstrating the adaptive grid coordinate in a sector configuration shows that it has potential applicability for further use in the ocean modelling domain. The ability for the coordinate to follow local isopycnals in the density adaptivity regime leads to improvements in the representation of stratification, which manifests in circulation and transport diagnostics. The benefits of this coordinate over  $z$ -star are comparable to the HYCOM1 coordinate with its similar philosophy, however with a less prescriptive configuration. This means that the adaptive grid coordinate should lend itself to representing basins of different density compositions while maximising the available resolution, compared to



---

HYCOM1 for which density levels must be chosen as a global compromise. The increased complexity of the adaptive grid coordinate does come with an increased computational cost, but this is minor in the context of the total runtime of the model.



---

# Conclusions

---

In numerical ocean modelling, the primitive equations for fluid dynamics are solved, often in combination with a suite of parameterisations for unresolved physics. The models are used for example to gain understanding of oceanic dynamics, or to model real-world scenarios as the basis for forecast systems. Key dynamics in the ocean, such as the overturning circulation are strongly controlled by the density structure of the ocean. Across the many ocean models that exist today, the primitive equations are solved with different algorithms and different discretisations, which may be sources of numerical truncation errors. These errors often cause spurious mixing events, altering the density structure and ultimately the circulation itself.

Mixing in the ocean is an important physical process that must be explicitly parameterised in the majority of ocean modelling. However, due to the potential undesired effects of additional spurious mixing, methods have been devised to attempt to diagnose its magnitude and/or its physical location. For example, using the decay of tracer variance through model evolution (Burchard and Rennau, 2008) or the density of reference potential energy (Ilıcak, 2016), information can be gleaned about both the magnitude and location of spurious mixing. Focusing on the MOM6 ocean model, which implements the Arbitrary Lagrangian-Eulerian (ALE) algorithm, we instead focus on the relative magnitudes of spurious mixing due to horizontal and vertical processes within the model.

## 6.1 Summary

To address the problem of investigating the relative contributions to spurious mixing by various processes within MOM6, or indeed other ALE models, we have proposed a new diagnostic. The split RPE method breaks contributions to the global RPE (Winters et al., 1995) budget on a sub-timestep scale, in a manner similar to Ilıcak et al. (2012). This fine-grained budgeting allows for the attribution of increases to RPE in non-mixing simulations without buoyancy forcing to spurious mixing present in horizontal and vertical processes. Using the split RPE method is particularly useful in an ALE model as it sheds light on the role played by the vertical coordinate, implemented by the vertical regridding/remapping phase of ALE.

As an evaluation of MOM6 and the split RPE method, we performed a suite of experiments in three idealised test cases. The test cases were chosen from previous studies by Ilıcak et al. (2012) and Petersen et al. (2015), allowing MOM6 to be added to the pre-existing results. Overall, MOM6 performed well compared to the other models (MOM, MITgcm and MPAS-O). The split RPE method demonstrated that this is likely due to the accuracy of regridding/remapping, leaving the majority of spurious mixing to occur during horizontal tracer advection. The flexibility to implement arbitrary vertical coordinates in ALE models allows for influence of the along-layer properties. For example, in the

extreme case of an isopycnal vertical coordinate, density is uniform within layers, which reduces the spurious mixing due to horizontal tracer advection. The results of idealised testing, along with this observation regarding vertical coordinates led to the formulation and development of a new vertical coordinate for ocean modelling.

To improve accuracy and make the most efficient use of available resolution, we formulated a new vertical coordinate, the Adaptive Grid (AG) coordinate. This coordinate is based on the principle of following local density structure where possible, while retaining sufficient vertical resolution such that relevant physical processes may be accurately resolved. To do so, a decision is made to seek density adaptivity or simply smoothing of an interface, based on the local stratification. In order to ensure that a minimum resolution is maintained, and that the grid cannot get in an inconsistent state, two additional terms were added. The first of these is a depth-dependent restoring toward a target grid, giving minimum resolution over a chosen timescale. This allows us to take advantage of resolution requirements such as those proposed by Stewart et al. (2017) for vertically resolving the baroclinic modes which are resolved by the horizontal grid, or resolving processes within the mixed layer.

The final term added to the adaptive grid coordinate is an empirical adjustment. This was required because density adaptivity occurs over short lengthscales, and may lead to larger-scale inconsistencies, for example a strong lateral gradient within a layer in a region of overflow. To account for inconsistencies of this kind, an adjustment is made by shifting massless, or very thin layers within the water column. Moving water into these thin layers can be done very accurately, and may result in significantly reduced spurious mixing, or even avoid unphysical layer configurations like strong, spurious density inversions entirely.

After development and implementation of the adaptive grid coordinate in MOM6, a first evaluation of the adaptive grid coordinate in MOM6 was performed using a suite of idealised tests. In the specific process-based tests of a lock exchange and propagation of internal waves, the adaptive grid coordinate performed strongly by reducing spurious mixing compared to other models and other coordinates. In the more realistic baroclinic eddies test case, the adaptive grid coordinate showed large improvements over the other models and coordinates at low horizontal resolution, but the improvement diminished as the horizontal resolution increased. This change was attributed to a greater contribution to spurious mixing by the vertical regridding/remapping processes, where the change in along-layer structure due to vertical coordinate choice becomes less important. Despite the diminishing returns with horizontal resolution, the adaptive grid coordinate performed well, and a more realistic configuration was devised to test its influence on actual circulations as opposed to solely considering spurious mixing.

Using a diagnostic of spurious mixing based on reference potential energy necessitated that the idealised test cases were run with zero explicit mixing, and no buoyancy forcing. These conditions limit the scope of investigation into the performance of the adaptive grid coordinate. To assess the impact of the adaptive grid coordinate on circulation, an idealised sector model was used. This is a rectangular domain with sloped sidewalls, intended as an analogue of the Atlantic Ocean. In particular, this model has both mechanical and buoyancy forcing at the surface, and sets up a meridional overturning circulation with a zonally-unbounded surface current in the south. The flow features present in the model domain are important in setting the circulation in the real ocean, and are particularly impacted by the representation of stratification through the water column. This experiment was spun-up with a  $z$ -star coordinate, then branched off with the  $z$ -star coordinate as a control, and both the MOM6 HYCOM1 coordinate and the adaptive grid coordinate as variations.

Both the HYCOM1 and adaptive grid coordinates are similar, in that they use density-

following coordinates in the deep ocean, but have a fixed resolution near the surface. However, the specifics of the implementations vary significantly: HYCOM1 achieves this by following specific isopycnal values and enforcing minimum layer thicknesses; the adaptive grid coordinate uses the local density adaptivity term in regions of sufficiently large isopycnal slope relative to the grid. In the sector model, both coordinates improve the representation of the stratification, capturing dense water formation with less mixing and maintaining a less diffusive upwelling region of the meridional overturning circulation. These improvements in stratification impact both the meridional overturning circulation and the analogue of the circumpolar current, which is about 50% stronger than the control for both the adaptive grid and HYCOM1 coordinates.

Using more complex vertical coordinates is a trade-off in terms of computational resources and time. Performing regridding to the HYCOM1 coordinate is approximately four times slower than  $z$ -star. Due to the significantly greater amount of computation in the adaptive grid coordinate, and its nonlocal nature, regridding time increases even further. Compared to  $z$ -star, the adaptive grid coordinate is almost ten times slower to regrid. However, this may be a worthwhile tradeoff considering the improvements in the representation of stratification and circulation. Additionally, regridding makes up only 4% of the entire runtime of a 24-month run of the model when the adaptive grid coordinate is used, and is therefore not a large cost in absolute terms.

## 6.2 Future work

This thesis has outlined the development of a new diagnostic for spurious mixing, the split RPE method, particularly applicable to the ALE models that see widespread use today. By using this diagnostic, a need was identified for further research into the area of vertical coordinates. To address this, a new adaptive grid coordinate was developed, which makes use of the local density structure to make informed adjustments to the grid according to the principles of density adaptivity, lateral smoothing, vertical restoring and finally adjustment. Testing and evaluation of this coordinate in highly idealised, and more realistic model configurations was positive, showing promising improvements over existing coordinates. The development of the coordinate should not be considered complete at this stage: further testing in ocean modelling should still be performed, and exploration into the choices of parameters could yield further improvements.

Although the adaptive grid coordinate was designed to be governed by only a few general parameters, investigation of sensitivity to these parameters is outside the scope of this thesis. In particular, a few sensitivity studies could be conducted in future. The balance of smoothing and restoring against density adaptivity, through the  $\omega_z^{\min}$  and  $\tau_0$  parameters would demonstrate how closely the coordinate may follow density structure without being regulated back to a fixed grid. Additionally, the impact of the vertical profile of restoring through the  $z_s$  and  $z_d$  parameters is unknown: is density adaptivity at all advantageous near the surface? These parameters could be examined in a similar manner to that presented in this thesis, by investigating the change of the modelled stratification and circulation in experiments branched from a common starting point.

Another outstanding question about the AG coordinate is its suitability for model spinup. Leclair and Madec (2011) found that the Lagrangian nature of the  $z$ -tilde coordinate rendered it unsuitable for spinning up a model from rest. The AG coordinate also exhibits Lagrangian behaviour, particularly in the abyss, however the surface is more closely regulated. The coordinate may require relaxation towards the final parameters of maximal density adaptivity, or indeed it may be the case that initial spinup must be undertaken with  $z$ -star, as done in this thesis. Further, the coordinate needs to be demon-

strated in realistic configurations with complex bathymetry if it is to be used beyond idealised simulations. Other, realistic test cases of interest would be the Beckmann and Haidvogel (1993) seamount, or a near-equatorial current.

MOM6 was chosen as the target for the AG coordinate in part due to its open source and open development model. It is envisaged that as the community uses and learns about the AG coordinate, they may feed back their experience to further its development. This may encompass testing and gaining more knowledge about parameter selection, or even addition of new terms and principles to the coordinate formulation itself.

---

# References

---

- Adcroft, A. (2015). MOM6 source code. <https://github.com/NOAA-GFDL/MOM6/commit/8fcc5c>.
- Adcroft, A., Anderson, W., Balaji, V., Blanton, C., Bushuk, M., Dufour, C. O., Dunne, J. P., Griffies, S. M., Hallberg, R. W., Harrison, M. J., Held, I., Jansen, M. F., John, J., Krasting, J. P., Langenhorst, A., Legg, S., Liang, Z., McHugh, C., Radhakrishnan, A., Reichl, B. G., Rosati, T., Samuels, B. L., Shao, A., Stouffer, R., Winton, M., Wittenberg, A. T., Xiang, B., Zadeh, N., and Zhang, R. The gfdl global ocean and sea ice model om4.0: Model description and simulation features. *Journal of Advances in Modeling Earth Systems*, 0(ja).
- Adcroft, A. and Campin, J.-M. (2004). Rescaled height coordinates for accurate representation of free-surface flows in ocean circulation models. *Ocean Modelling*, 7(3–4):269–284.
- Adcroft, A. J., Hill, C. N., and Marshall, J. C. (1999). A New Treatment of the Coriolis Terms in C-Grid Models at Both High and Low Resolutions. *Monthly Weather Review*, 127(8):1928–1936.
- Arakawa, A. and Lamb, V. R. (1977). Computational Design of the Basic Dynamical Processes of the UCLA General Circulation Model. In Chang, J., editor, *Methods in Computational Physics: Advances in Research and Applications*, volume 17 of *General Circulation Models of the Atmosphere*, pages 173–265. Elsevier.
- Armour, K. C., Marshall, J., Scott, J. R., Donohoe, A., and Newsom, E. R. (2016). Southern Ocean warming delayed by circumpolar upwelling and equatorward transport. *Nature Geoscience*, 9(7):549–554.
- Beckmann, A. and Haidvogel, D. B. (1993). Numerical Simulation of Flow around a Tall Isolated Seamount. Part I: Problem Formulation and Model Accuracy. *Journal of Physical Oceanography*, 23(8):1736–1753.
- Benjamin, T. B. (1968). Gravity currents and related phenomena. 31(2):209–248.
- Blaker, A. T., Hirschi, J. J.-M., Sinha, B., de Cuevas, B., Alderson, S., Coward, A., and Madec, G. (2012). Large near-inertial oscillations of the Atlantic meridional overturning circulation. *Ocean Modelling*, 42:50–56.
- Bleck, R. (1978). On the Use of Hybrid Vertical Coordinates in Numerical Weather Prediction Models. *Monthly Weather Review*, 106(9):1233–1244.
- Bleck, R. (2002). An oceanic general circulation model framed in hybrid isopycnic-Cartesian coordinates. *Ocean Modelling*, 4(1):55–88.

- 
- Burchard, H. and Beckers, J.-M. (2004). Non-uniform adaptive vertical grids in one-dimensional numerical ocean models. *Ocean Modelling*, 6(1):51–81.
- Burchard, H. and Rennau, H. (2008). Comparative quantification of physically and numerically induced mixing in ocean models. *Ocean Modelling*, 20(3):293–311.
- Daru, V. and Tenaud, C. (2004). High order one-step monotonicity-preserving schemes for unsteady compressible flow calculations. *Journal of Computational Physics*, 193(2):563–594.
- Gent, P. R., Large, W. G., and Bryan, F. O. (2001). What sets the mean transport through Drake Passage? *Journal of Geophysical Research: Oceans*, 106(C2):2693–2712.
- Gent, P. R. and McWilliams, J. C. (1990). Isopycnal Mixing in Ocean Circulation Models. *Journal of Physical Oceanography*, 20(1):150–155.
- Gentry, R. A., Martin, R. E., and Daly, B. J. (1966). An Eulerian differencing method for unsteady compressible flow problems. *Journal of Computational Physics*, 1(1):87–118.
- Gibson, A. H., Hogg, A. M., Kiss, A. E., Shakespeare, C. J., and Adcroft, A. (2017). Attribution of horizontal and vertical contributions to spurious mixing in an Arbitrary Lagrangian–Eulerian ocean model. *Ocean Modelling*, 119:45–56.
- Gouillon, F. (2010). *Internal Wave Propagation and Numerically Induced Diapycnal Mixing in Oceanic General Circulation Models*. PhD thesis, Florida State University.
- Gräwe, U., Holtermann, P., Klingbeil, K., and Burchard, H. (2015). Advantages of vertically adaptive coordinates in numerical models of stratified shelf seas. *Ocean Modelling*, 92:56–68.
- Griffies, S. M. (2004). *Fundamentals of Ocean Climate Models*. Princeton University Press, Princeton, N.J. OCLC: ocm54487392.
- Griffies, S. M. (2012). Elements of the modular ocean model (mom). *GFDL Ocean Group Tech. Rep*, 7:620.
- Griffies, S. M., Pacanowski, R. C., and Hallberg, R. W. (2000). Spurious Diapycnal Mixing Associated with Advection in a z-Coordinate Ocean Model. *Monthly Weather Review*, 128(3):538–564.
- Griffies, S. M., Winton, M., Anderson, W. G., Benson, R., Delworth, T. L., Dufour, C. O., Dunne, J. P., Goddard, P., Morrison, A. K., Rosati, A., Wittenberg, A. T., Yin, J., and Zhang, R. (2014). Impacts on Ocean Heat from Transient Mesoscale Eddies in a Hierarchy of Climate Models. *Journal of Climate*, 28(3):952–977.
- Haney, R. L. (1991). On the Pressure Gradient Force over Steep Topography in Sigma Coordinate Ocean Models. *Journal of Physical Oceanography*, 21(4):610–619.
- Hirt, C. W., Amsden, A. A., and Cook, J. L. (1974). An arbitrary Lagrangian-Eulerian computing method for all flow speeds. *Journal of Computational Physics*, 14(3):227–253.
- Hofmeister, R., Beckers, J.-M., and Burchard, H. (2011). Realistic modelling of the exceptional inflows into the central Baltic Sea in 2003 using terrain-following coordinates. *Ocean Modelling*, 39(3):233–247.



- 
- Hofmeister, R., Burchard, H., and Beckers, J.-M. (2010). Non-uniform adaptive vertical grids for 3D numerical ocean models. *Ocean Modelling*, 33(1–2):70–86.
- Hogg, A. M. (2010). An Antarctic Circumpolar Current driven by surface buoyancy forcing. *Geophysical Research Letters*, 37(23).
- Hogg, A. M., Dijkstra, H. A., and Saenz, J. A. (2013). The Energetics of a Collapsing Meridional Overturning Circulation. *Journal of Physical Oceanography*, 43(7):1512–1524.
- Ilıcak, M. (2016). Quantifying spatial distribution of spurious mixing in ocean models. *Ocean Modelling*, 108:30–38.
- Ilıcak, M., Adcroft, A. J., Griffies, S. M., and Hallberg, R. W. (2012). Spurious diapycnal mixing and the role of momentum closure. *Ocean Modelling*, 45–46:37–58.
- Kasahara, A. (1974). Various Vertical Coordinate Systems Used for Numerical Weather Prediction. *Monthly Weather Review*, 102(7):509–522.
- Klingbeil, K., Mohammadi-Aragh, M., Gräwe, U., and Burchard, H. (2014). Quantification of spurious dissipation and mixing – Discrete variance decay in a Finite-Volume framework. *Ocean Modelling*, 81:49–64.
- Leclair, M. and Madec, G. (2011). Z-Coordinate, an Arbitrary Lagrangian–Eulerian coordinate separating high and low frequency motions. *Ocean Modelling*, 37(3–4):139–152.
- Legg, S., Ezer, T., Jackson, L., Briegleb, B., Danabasoglu, G., Large, W., Wu, W., Chang, Y., Özgökmen, T. M., Peters, H., Xu, X., Chassignet, E. P., Gordon, A. L., Griffies, S., Hallberg, R., Price, J., Riemenschneider, U., and Yang, J. (2009). Improving Oceanic Overflow Representation in Climate Models: The Gravity Current Entrainment Climate Process Team. *Bulletin of the American Meteorological Society*, 90(5):657–670.
- Margolin, L. G. and Shashkov, M. (2003). Second-order sign-preserving conservative interpolation (remapping) on general grids. *Journal of Computational Physics*, 184(1):266–298.
- Marshall, J., Adcroft, A., Hill, C., Perelman, L., and Heisey, C. (1997). A finite-volume, incompressible navier stokes model for studies of the ocean on parallel computers. *Journal of Geophysical Research: Oceans*, 102(C3):5753–5766.
- Marshall, J. and Radko, T. (2003). Residual-Mean Solutions for the Antarctic Circumpolar Current and Its Associated Overturning Circulation. *Journal of Physical Oceanography*, 33(11):2341–2354.
- Mashayek, A., Ferrari, R., Nikurashin, M., and Peltier, W. R. (2015). Influence of Enhanced Abyssal Diapycnal Mixing on Stratification and the Ocean Overturning Circulation. *Journal of Physical Oceanography*, 45(10):2580–2597.
- Mellor, G. L., Oey, L.-Y., and Ezer, T. (1998). Sigma Coordinate Pressure Gradient Errors and the Seamount Problem. *Journal of Atmospheric and Oceanic Technology*, 15(5):1122–1131.
- Morales Maqueda, M. A. and Holloway, G. (2006). Second-order moment advection scheme applied to Arctic Ocean simulation. *Ocean Modelling*, 14(3–4):197–221.

- 
- Nikurashin, M. and Vallis, G. (2011). A Theory of Deep Stratification and Overturning Circulation in the Ocean. *Journal of Physical Oceanography*, 41(3):485–502.
- Orsi, A. H., Johnson, G. C., and Bullister, J. L. (1999). Circulation, mixing, and production of Antarctic Bottom Water. *Progress in Oceanography*, 43(1):55–109.
- Petersen, M. R., Jacobsen, D. W., Ringler, T. D., Hecht, M. W., and Maltrud, M. E. (2015). Evaluation of the arbitrary Lagrangian–Eulerian vertical coordinate method in the MPAS-Ocean model. *Ocean Modelling*, 86:93–113.
- Prather, M. J. (1986). Numerical advection by conservation of second-order moments. *Journal of Geophysical Research: Atmospheres*, 91(D6):6671–6681.
- Reichl, B. G. and Hallberg, R. (2018). A simplified energetics based planetary boundary layer (ePBL) approach for ocean climate simulations. *Ocean Modelling*, 132:112–129.
- Ringler, T., Petersen, M., Higdon, R. L., Jacobsen, D., Jones, P. W., and Maltrud, M. (2013). A multi-resolution approach to global ocean modeling. *Ocean Modelling*, 69:211–232.
- Saenz, J. A., Tailleux, R., Butler, E. D., Hughes, G. O., and Oliver, K. I. C. (2015). Estimating Lorenz’s Reference State in an Ocean with a Nonlinear Equation of State for Seawater. *Journal of Physical Oceanography*, 45(5):1242–1257.
- Shakespeare, C. J. and Hogg, A. M. (2012). An Analytical Model of the Response of the Meridional Overturning Circulation to Changes in Wind and Buoyancy Forcing. *Journal of Physical Oceanography*, 42(8):1270–1287.
- Stewart, K. D., Hogg, A. M., Griffies, S. M., Heerdegen, A. P., Ward, M. L., Spence, P., and England, M. H. (2017). Vertical resolution of baroclinic modes in global ocean models. *Ocean Modelling*, 113:50–65.
- Stewart, K. D., Saenz, J. A., Hogg, A. M., Hughes, G. O., and Griffiths, R. W. (2014). Effect of topographic barriers on the rates of available potential energy conversion of the oceans. *Ocean Modelling*, 76:31–42.
- Vreugdenhil, C. A., Hogg, A. M., Griffiths, R. W., and Hughes, G. O. (2015). Adjustment of the Meridional Overturning Circulation and Its Dependence on Depth of Mixing. *Journal of Physical Oceanography*, 46(3):731–747.
- Waterhouse, A. F., MacKinnon, J. A., Nash, J. D., Alford, M. H., Kunze, E., Simmons, H. L., Polzin, K. L., St. Laurent, L. C., Sun, O. M., Pinkel, R., Talley, L. D., Whalen, C. B., Huussen, T. N., Carter, G. S., Fer, I., Waterman, S., Naveira Garabato, A. C., Sanford, T. B., and Lee, C. M. (2014). Global Patterns of Diapycnal Mixing from Measurements of the Turbulent Dissipation Rate. *Journal of Physical Oceanography*, 44(7):1854–1872.
- White, L. and Adcroft, A. (2008). A high-order finite volume remapping scheme for nonuniform grids: The piecewise quartic method (PQM). *Journal of Computational Physics*, 227(15):7394–7422.
- White, L., Adcroft, A., and Hallberg, R. (2009). High-order regridding–remapping schemes for continuous isopycnal and generalized coordinates in ocean models. *Journal of Computational Physics*, 228(23):8665–8692.
- Winters, K. B., Lombard, P. N., Riley, J. J., and D’Asaro, E. A. (1995). Available potential energy and mixing in density-stratified fluids. *Journal of Fluid Mechanics*, 289:115–128.