Adaptive mesh modelling of the thermally driven annulus

James Roland Maddison

Wolfson College
Oxford University



A thesis submitted for the degree of Doctor of Philosophy
May 2011

Abstract

Numerical simulations of atmospheric and oceanic flows are fundamentally limited by a lack of model resolution. This thesis describes the application of unstructured mesh finite element methods to geophysical fluid dynamics simulations. These methods permit the mesh resolution to be concentrated in regions of relatively increased dynamical importance. Dynamic mesh adaptivity can further be used to maintain an optimised mesh even as the flow develops. Hence unstructured dynamic mesh adaptive methods have the potential to enable efficient simulations of high Reynolds number flows in complex geometries.

In this thesis, the thermally driven rotating annulus is used to test these numerical methods. This system is a classic laboratory scale analogue for large scale geophysical flows. The thermally driven rotating annulus has a long history of experimental and numerical research, and hence it is ideally suited for the validation of new numerical methods.

For geophysical systems there is a leading order balance between the Coriolis and buoyancy accelerations and the pressure gradient acceleration: geostrophic and hydrostatic balance. It is essential that any numerical model for these systems is able to represent these balances accurately. In this thesis a balanced pressure decomposition method is described, whereby the pressure is decomposed into a "balanced" component associated with the Coriolis and buoyancy accelerations, and a "residual" component associated with other forcings and that enforces incompressibility. It is demonstrated that this method can be used to enable a more accurate representation of geostrophic and hydrostatic balance in finite element modelling. Furthermore, when applying dynamic mesh adaptivity, there is a further potential for imbalance injection by the mesh optimisation procedure. This issue is tested in the context of shallow-water ocean modelling. For the linearised system on an f-plane, and with a steady balance permitting numerical discretisation, an interpolant is formulated that guarantees that a steady and balanced state remains steady and in balance after interpolation onto an arbitrary target mesh.

The application of unstructured dynamic mesh adaptive methods to the thermally driven rotating annulus is presented. Fixed structured mesh finite element simulations are conducted, and compared against a finite difference model and against experiment. Further dynamic mesh adaptive simulations are then conducted, and compared against the structured mesh simulations. These tests are used to identify weaknesses in the application of dynamic mesh adaptivity to geophysical systems. The simulations are extended to a more challenging system: the thermally driven rotating annulus at high Taylor number and with sloping base and lid topography. Analysis of the high Taylor number simulations reveals a direct energy transfer from the eddies to the mean flow, confirming the results of previous experimental work.

Contents

1. Introduction **2**

1.1. The Imperial College Ocean Model . . . 4
1.1.1. Unstructured mesh ocean modelling . . . 4
1.1.2. The finite element method . . . 6
1.1.3. Dynamic mesh adaptive ocean modelling . . . 8
1.1.4. Fluidity-ICOM and geophysical fluid dynamics modelling . . . 11
1.1.5. Fluidity-ICOM development status . . . 12
1.2. The thermally driven rotating annulus . . . 12
1.2.1. The rotating annulus as a laboratory scale analogue . . . 13
1.2.2. Baroclinic instability . . . 14
1.2.3. Rotating annulus flow regimes . . . 16
1.2.4. Numerical simulation of the rotating annulus . . . 19
1.3. Thesis outline . . . 20

2. Balance errors in finite element ocean modelling **22**

2.1. Origin of physical balance errors . . . 24
2.1.1. Continuous formulation . . . 24
2.1.2. Discrete formulation . . . 27
2.1.3. Balanced pressure decomposition methods . . . 34
2.2. Balanced pressure decomposition accuracy . . . 40
2.2.1. Helmholtz decomposition accuracy . . . 40
2.2.2. Geostrophic balance accuracy . . . 42
2.3. Application to the thermally driven rotating annulus . . . 48
2.3.1. Stabilised P_1P_1 finite element pair . . . 48
2.3.2. $P_{1\mathrm{DG}}P_2$ finite element pair . . . 58
2.3.3. Computational cost . . . 62
2.3.4. Model viscosity . . . 63
2.4. Summary . . . 64

3. Balance errors in mesh adaptive ocean modelling **68**

3.1. Galerkin projection . . . 69
3.2. Geostrophic balance preserving interpolants . . . 75
3.2.1. Discretised shallow-water equations . . . 75
3.2.2. Geostrophic balance preserving interpolants . . . 77
3.2.3. Implementation and boundary conditions . . . 80
3.3. Geostrophic balance preserving interpolation examples . . . 82
3.3.1. Preservation of balance . . . 82
3.3.2. Nearly balanced states . . . 87
3.3.3. Kelvin wave . . . 89
3.3.4. Accuracy . . . 95
3.3.5. Computational cost . . . 96
3.4. Application to other finite element pairs . . . 99
3.5. Application to Navier-Stokes . . . 104

3.6. Summary . . . 105

4. Application of dynamic mesh adaptivity to simulations of the thermally driven rotating annulus 107
4.1. Fixed mesh Fluidity-ICOM simulations . . . 107
4.1.1. Comparison with MORALS . . . 109
4.1.2. Comparison with laboratory data . . . 117
4.2. Adaptive mesh Fluidity-ICOM configuration . . . 126
4.2.1. Frequency of mesh optimisation . . . 126
4.2.2. Metric formulation . . . 127
4.2.3. Interpolant . . . 132
4.2.4. Configuration . . . 136
4.3. Adaptive mesh Fluidity-ICOM simulations . . . 137
4.4. Computational cost . . . 148
4.5. Summary . . . 149

5. High Taylor number simulations 151
5.1. Configuration . . . 152
5.2. Overview . . . 154
5.2.1. Computational meshes . . . 154
5.2.2. Flow features . . . 156
5.3. Analysis . . . 159
5.3.1. Eddy fluxes . . . 159
5.3.2. Spectral analysis . . . 162
5.4. Summary . . . 168

6. Conclusions 172
6.1. Summary . . . 172
6.2. Future work . . . 175

A. Unstructured mesh diagnostics 178
A.1. Fluidity-ICOM parallelism . . . 178
A.2. Heat flow . . . 179
A.3. Detectors . . . 180
A.4. Vertical integration . . . 182
A.5. Depth averaged stream function . . . 184
A.6. Azimuthally averaged stream function . . . 186

B. Fluidity-ICOM development 188
B.1. Contributions to Fluidity-ICOM . . . 188

Bibliography 193

Acknowledgements

Firstly I would like to thank my supervisors, Peter Read, David Marshall, and Matthew Piggott, for their guidance throughout this DPhil project. I would also like to thank Raymond Hide for providing early insight and motivation.

The Natural Environment Research Council (NERC) for funding this DPhil project.

The geophysical and planetary fluid dynamics research group of Atmospheric, Oceanic and Planetary Physics, in the Department of Physics at Oxford. Roland Young, for many excellent discussions, and for providing the MORALS data used in section 2.3 and chapter 4. Also Alfonso Castrejón-Pita for collecting laboratory data, which unfortunately could not be discussed in this thesis, and Tom Jacoby, for many useful conversations.

The physical oceanography group of Atmospheric, Oceanic and Planetary Physics, in the Department of Physics at Oxford. Particularly my fellow Fluidity-ICOM users: Amrita Shravat, Lucy Bricheno, Jemma Shipton, and Šárka Tuková. The Fluidity-ICOM developers for technical advice and for their dedication to the improvement of the numerical model. I would particularly like to thank Colin Cotter for his advice concerning chapter 3, and David Ham for driving continual improvements to the Fluidity-ICOM code base. Also Patrick Farrell, who has been an outstanding colleague, proving invaluable assistance and feedback.

Chris Pain for comments on chapter 2, and Hilary Weller for comments on chapter 3.

The staff of the Oxford Supercomputing Centre, for providing excellent and reliable service. The parallel rotating annulus simulations in section 2.3, chapter 4, and chapter 5 were conducted using Oxford Supercomputing Centre resources.

The many and varied developers of free and open source software, for providing high quality tools used throughout this project. In particular, the developers of the GNU Compiler Collection, Python, matplotlib, SciPy, VTK, and MayaVi.

Most of all I would like to thank my wife Ali for her tireless support and encouragement.

Chapter 1.

Introduction

Geophysical fluid dynamics modelling is fundamentally limited by a lack of model resolution. The physics of the small scale has a significant influence on the large scale mean dynamics, but it is generally challenging to formulate simple and accurate parameterisations to include their aggregated effects in coarse resolution simulations. Increasingly advanced parameterisation schemes are continually being developed in order to produce more physically realistic coarse resolution simulations, but high resolution simulations are, in most cases, of greater quality than their coarse resolution counterparts.

The current state-of-the-art global ocean models are based upon methods using computational meshes (or grids) that are inherently structured – that is, there are well-defined and usually simple adjacency relations between all points in the mesh. This inherent structure can then be exploited, due to data locality in low level computations, to ensure increased model efficiency per mesh degree of freedom. This thesis is concerned with the application of unstructured mesh numerical methods to the modelling of geophysical fluids. In a general implementation, these methods impose no such constraints upon the connectivity of the computational mesh, allowing arbitrary adjacency relations between all mesh points. Due to the loss of data locality this approach is inherently more computationally expensive per degree of freedom, but enables mesh resolution to be concentrated in regions of increased dynamical importance. This can be extended, via dynamic mesh adaptivity, to maintain mesh optimality even as a flow develops. The combination of unstructured mesh methods and dynamic mesh adaptivity therefore has the potential for increased accuracy for a given computational cost relative to structured mesh methods.

The Imperial College Ocean Model (Fluidity-ICOM) is an in-development finite element model, utilising fully unstructured meshes in all three dimensions, with the periodic application of dynamic

mesh adaptivity to ensure an optimised representation of solution fields. While unstructured mesh methods are relatively common in the field of computational fluid dynamics, they are uncommon in the field of geophysical fluid dynamics. Reasons for their rarity in this field include issues with representation of physical balance on unstructured meshes [Griffies et al., 2000] and their relative computational cost (see, for example, Danilov et al. [2008]). Hence, while there is undoubtedly potential for the methods implemented in Fluidity-ICOM to enable simulations of previously inaccessible phenomena, there is currently a lack of experience of their application to atmospheric or oceanic modelling.

As a result of this relative immaturity, there are number of fundamental issues to be tackled in order for geophysical simulations using Fluidity-ICOM to be viable. These include the accuracy of geostrophic balance representation on unstructured meshes, the configuration of dynamic mesh adaptivity and the interpolation of field data (including maintenance of physical balance), and the accuracy of dynamic mesh adaptive simulations relative to fixed mesh simulations. New issues encountered when applying unstructured and adaptive mesh techniques to geophysical systems require investigation, and if possible resolution, in order to enable any meaningful simulations of physical phenomena. Hence in this thesis model validation is undertaken, and possible issues with the numerical techniques identified, with the objective being to enable simulations of previously inaccessible geophysical phenomena using unstructured dynamic mesh adaptive methods.

In this thesis, a classic laboratory scale analogue of geophysical flows, the thermally driven rotating annulus, is used to test the numerical methods implemented in Fluidity-ICOM. Detailed laboratory data are available for this system, enabling a quantitative comparison of the performance of Fluidity-ICOM against carefully obtained laboratory measurements. Hence the rotating annulus can be used to identify difficulties or shortcomings associated with dynamic mesh adaptive finite element ocean modelling in a geophysical system. Furthermore, previous work on simulating the thermally driven rotating annulus has typically been limited to the simplest geometries and flow regimes. Hence the numerical methods implemented in Fluidity-ICOM have particular potential to enable simulations of previously challenging or inaccessible rotating annulus configurations. Fluidity-ICOM will therefore be used to simulate the rotating annulus with sloping lid and bottom topography and at relatively high Reynolds number.

In this chapter unstructured dynamic mesh adaptive numerical modelling, specifically using the Imperial College Ocean Model (Fluidity-ICOM), will be introduced. The thermally driven rotating annulus will be described, including a summary of previous numerical research using this system. At the end of this chapter an outline for the remainder of this thesis will be supplied.

1.1. The Imperial College Ocean Model

This section will introduce unstructured mesh numerical modelling using Fluidity-ICOM. Existing unstructured mesh ocean models will be described in section 1.1.1. A brief example of the finite element method will be supplied in section 1.1.2, and used to demonstrate the geometric flexibility inherent in the approach. In section 1.1.3 dynamic mesh adaptivity will be introduced, and an overview of previous dynamic mesh adaptive geophysical fluid dynamics modelling will be supplied. Finally, the application of unstructured dynamic mesh adaptive numerical methods to geophysical systems will be considered in section 1.1.4.

1.1.1. Unstructured mesh ocean modelling

The current state-of-the-art global ocean modelling is primarily based upon finite-difference numerical methods [Griffies et al., 2000]. The lack of freedom inherent in a structured mesh formulation leads to problems at domain boundaries – it is difficult to maintain structure while simultaneously conforming to a complicated boundary, such as an ocean coastline and bottom topography. In the horizontal dimensions, one approach is simply to use an axis aligned grid throughout the bulk of the domain, and then clip to form a "stair-cased" boundary, as in Bryan [1969]. In Adcroft and Marshall [1998] the use of stair casing is found to (quote) *"exert a spurious form stress"* at the domain boundary, dependent upon the details of the boundary condition and viscous stress implementations. More generally, conforming structured grids can be formed via the use of generalised curvilinear coordinates [Thompson et al., 1974] – for a simply connected domain in two dimensions this may be thought of as the application of a homeomorphism from a uniform structured axis-aligned grid of a square onto the domain of interest. This approach is taken in Russell and Eiseman [1998] for the global ocean. In the vertical dimension a general terrain following "σ-coordinate" system may be chosen [Arakawa and Suarez, 1983], albeit with a notorious issue associated with hydrostatic pressure gradient errors in the presence of steep topography [Mellor et al., 1994, Berntsen and Thiem, 2007] – such errors will be addressed in further detail in chapter 2. Alternatively, one may choose a z-coordinate mesh in the vertical, and apply the use of "shaved cells" (or "cut cells") in the vicinity of the boundary, as described in Adcroft et al. [1997] and Marshall et al. [1997a]. This enables a simple representation in the bulk, avoids the hydrostatic consistency issue, and enables an accurate representation of the influence of topography. Typical issues with this approach include additional memory requirements, and numerical stability issues associated with the creation of very small shaved cells.

This thesis is concerned with the application of unstructured mesh finite element methods to simulations of geophysical fluids. General unstructured meshes give almost total freedom over the mesh topology. Hence a mesh can conform to arbitrary bounding topography without any special considerations, as demonstrated in figure 1.1, and without the complications associated with generating the coordinate maps associated with curvilinear coordinates. Furthermore, the mesh resolution can be varied in space, to refine local regions of interest, and in time, to track features of the flow as they develop.

Due to their additional geometric flexibility, unstructured meshes have seen some use in coastal ocean modelling, such as the finite element Advanced Circulation (ADCIRC) model [Luettich Jr et al., 1992] and the Finite Volume Coastal Ocean Model (FVCOM) [Chen et al., 2003]. In Chen et al. [2007] FVCOM is compared against finite difference coastal models, and the finite volume approach is found to yield more accurate results, particularly as compared against finite difference solutions on stair-cased meshes. While uncommon, unstructured meshes have seen some limited use for basin and global scale simulations. An early barotropic model of the North-Atlantic and Caribbean using a pseudo-isotropic triangulation of nodes lying on equally spaced h/f contours (where h is the depth and f is the Coriolis parameter) is described in Ford [1995]. The Finite Element Ocean Model (FEOM) solves the hydrostatic primitive equations using tetrahedral meshes that are unstructured in the horizontal and aligned in the vertical [Danilov et al., 2004, 2005].

Recently, unstructured mesh finite element ocean modelling has been proposed by Ford et al. [2004a], Ford et al. [2004b], Pain et al. [2005] and Piggott et al. [2008a], using the Imperial College Ocean Model (Fluidity-ICOM). Fluidity-ICOM is a non-hydrostatic model using anisotropic tetrahedral meshes and, in contrast to other unstructured mesh models, uses meshes that are unstructured in all three dimensions – i.e. the meshes are unstructured in both the horizontal and the vertical dimensions. Furthermore, Fluidity-ICOM implements three-dimensional anisotropic dynamic mesh adaptivity. Hence Fluidity-ICOM offers considerable additional geometric flexibility over many existing unstructured mesh ocean models, and is a significant change in approach as compared with most current state-of-the-art structured mesh ocean models.

Fluidity-ICOM is the primary model used for simulations conducted throughout this thesis. Hence the numerical methods implemented in Fluidity-ICOM, specifically the finite element method and mesh adaptivity methods, will be addressed in further detail in the following sections.

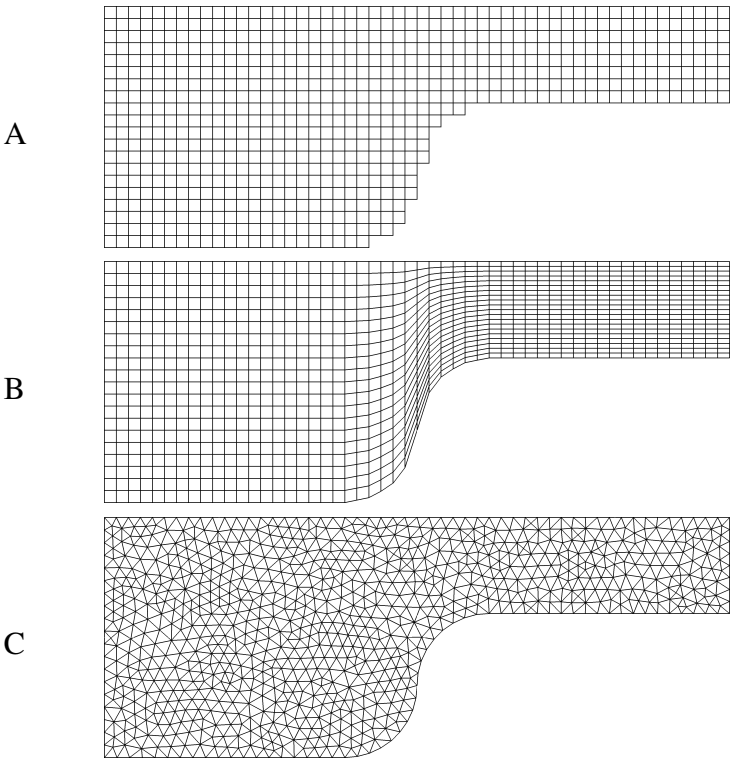

Figure 1.1.: Structured and unstructured meshes of a domain with a curved boundary. A: Structured uniform mesh with stair-casing. B: Structured non-uniform conforming mesh (as for a σ-coordinate mesh). C: Unstructured quasi-uniform mesh generated using Gmsh [Geuzaine and Remacle, 2009].

1.1.2. The finite element method

A finite element discretisation of a system consists of two stages: a covering partition of the domain into a finite number of discrete sub-volumes (the mesh), and the choice of finite-dimensional basis functions for the solution within each sub-volume. Throughout this thesis, we will limit ourselves to simplex meshes with exactly one element neighbour per simplex face (sometimes referred to as a *conforming* mesh). We further limit ourselves to Lagrange interpolating polynomial basis functions: polynomials of a given degree defined in terms of a uniform structured grid within each element, with each basis function equal to one at a single element node and zero at all other element nodes, as demonstrated in figures 1.2 and 1.3. The number of nodes per element are given by the triangular numbers in two dimensions, and the tetrahedral numbers in three dimensions. Hence the number of nodes per element scales (for sufficiently high degree) quadratically with degree for triangular elements, and as the third power with degree for tetrahedral elements. Following the usual finite element convention, a simplex mesh with degree N Lagrange interpolating polynomials is referred to as a P_N discretisation. We further follow the convention that a mesh for which element boundary nodes are distinct between element neighbours (i.e. a mesh with discontinuities at element boundaries) is denoted a $P_{N\mathrm{DG}}$ discretisation [1]. A finite element discretisation then

[1] A special case exists for piecewise constant basis functions, where the discontinuous notation can be neglected – a P_0 discretisation is always discontinuous at element boundaries

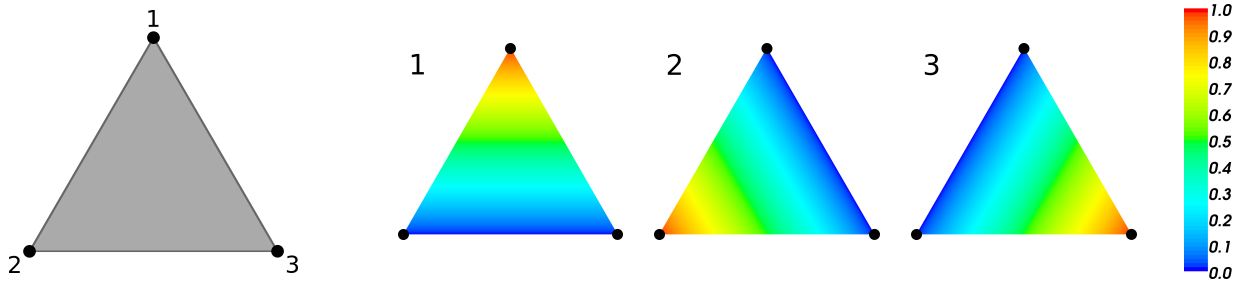

Figure 1.2.: Left: Nodal locations for a triangular element with piecewise linear Lagrangian interpolating polynomial (P_1) basis functions. Right: The three P_1 basis functions. Each basis function is equal to one at one element node and zero at all other nodes. A P_1 field within the element is described as a linear combination of these basis functions.

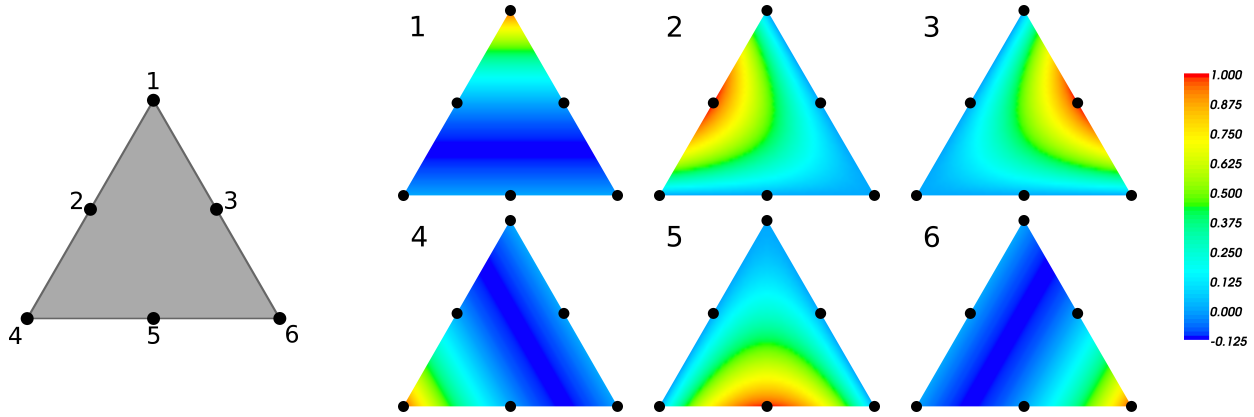

Figure 1.3.: Left: Nodal locations for a triangular element with piecewise quadratic Lagrangian interpolating polynomial (P_2) basis functions. Right: The six P_2 basis functions. Each basis function is equal to one at one element node and zero at all other nodes. A P_2 field within the element is described as a linear combination of these basis functions.

proceeds by projecting the continuous solution of the system onto the discrete space formed by the choice of mesh and basis functions, with a common approach being an L_2 optimal (or Galerkin) projection.

As an example, consider a finite element discretisation of the Poisson equation with Neumann boundary conditions:

$$\nabla^2 \phi = g, \tag{1.1}$$

$$\nabla \phi \cdot \hat{\boldsymbol{n}} = g_n \text{ on } \partial\Omega, \tag{1.2}$$

where $\partial\Omega$ bounds the domain Ω, and $\hat{\boldsymbol{n}}$ is a unit normal outward on $\partial\Omega$. The discretisation proceeds by multiplying by some test function ψ, integrating over the domain Ω and applying Green's theorem to yield the weak form equivalent equation:

$$\int_{\partial\Omega} \psi g_n - \int_\Omega \nabla\psi \cdot \nabla\phi = \int_\Omega \psi g \quad \forall\psi, \tag{1.3}$$

where the Neumann boundary condition has been directly substituted into the boundary term. Restricting ψ and ϕ to finite-dimensional spaces, to yield the finite-dimensional ψ^δ and ϕ^δ, completes the finite element discretisation:

$$\int_{\partial\Omega} \psi^\delta g_n - \int_\Omega \nabla\psi^\delta \cdot \nabla\phi^\delta = \int_\Omega \psi^\delta g \quad \forall\psi^\delta. \tag{1.4}$$

Choosing ψ^δ and ϕ^δ to be from the same function space defines a Galerkin discretisation. In this case, the final discretised equation is a Galerkin projection of the continuous system onto the discrete space.

In contrast to a finite difference discretisation, which proceeds by converting a continuous differential equation into differences between discrete nodal locations, a finite element discretisation proceeds by projecting the continuous system onto a discrete space. The former is defined in terms of differences between grid nodes, with an associated implicit grid structure, while the latter is defined in terms of integrals of bilinear forms over an arbitrary covering set of elemental sub-volumes. Hence there is no implied mesh structure in a finite element discretisation, and the method can be applied on arbitrary unstructured meshes.

1.1.3. Dynamic mesh adaptive ocean modelling

As well as possessing the ability to conform to arbitrary boundaries, the additional topological freedom offered by the use of arbitrarily unstructured meshes enables mesh resolution to be varied to concentrate on regions of relatively increased dynamical importance. The variation of resolution can be updated as a simulation develops, via dynamic mesh adaptivity, to maintain an optimal representation of the system state.

Dynamic mesh adaptivity can be broadly classified in terms of the optimisation procedure. *h*-adaptivity consists of element-wise topological operations with, in each optimisation, the nodal locations remaining fixed. Nodes can be added or removed in each element-wise operation. *r*-adaptivity consists of movement of the mesh nodes, such as via Lagrangian mesh movement, but

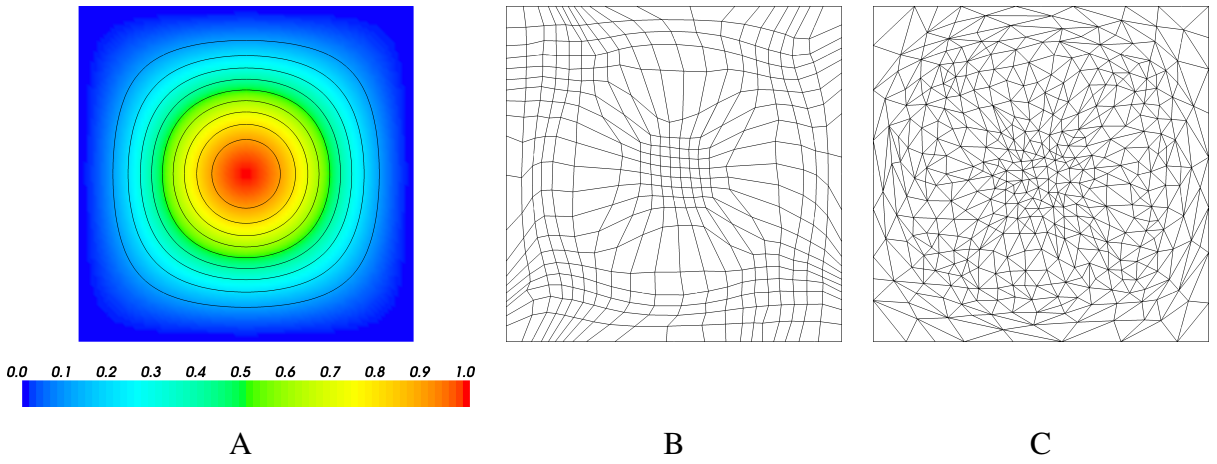

Figure 1.4.: Optimisation of a mesh to represent a field $\phi = \sin^2(\pi x)\sin^2(\pi y)$ in a domain $0 \leq x \leq 1, 0 \leq y \leq 1$. The meshes are generated using the Fluidity-ICOM mesh adaptivity library [Pain et al., 2001]. A: The field for which the mesh is optimised. B: A quadrilateral mesh with 441 nodes optimised via mesh movement (r-adaptivity). C: A triangular mesh with 417 nodes optimised via element-wise topological operations (h-adaptivity).

with the mesh topology remaining unchanged. p-adaptivity is performed via alteration of the element basis functions to be of locally higher or lower order, but with no change in the underlying domain partitioning.

Adaptive mesh refinement (AMR) [Berger and Oliger, 1984] uses patch-wise coarsening and refinement of a structured mesh to form a series of structured sub-meshes, and is a form of h-adaptivity. Since each sub-mesh is structured, AMR is suitable for use with structured mesh methods. In St-Cyr et al. [2008] two AMR shallow-water models are tested and, as compared against uniform mesh simulations, are found always to have a reduced computational cost with comparable model accuracy. In Blayo and Debreu [1999] AMR is applied to quasi-geostrophic ocean modelling of a baroclinic double gyre, and is found to yield superior results at comparable computational cost as compared to the use of nested grids. More general dynamic mesh adaptive unstructured meshes have seen some use in atmospheric modelling. In Bacon et al. [2003] a non-hydrostatic finite volume model is described with applications in local hazardous release [Bacon et al., 2000] and hurricane tracking [Bacon et al., 2007], implementing a mixture of h- and r-adaptivity via element-wise topological operations and node movement. In ocean modelling, unstructured dynamic mesh adaptive models are uncommon. In Remacle et al. [2006] an unstructured dynamic mesh adaptive shallow-water model is described, optimised via local element operations. The Second Generation Louvain-la-Neuve Ice-ocean Model (SLIM), a hydrostatic primitive equation solver using vertically aligned prismatic elements, can also apply dynamic mesh adaptivity [Bernard et al., 2007].

Fluidity-ICOM implements dynamic mesh adaptivity using a mixture of *h*- and *r*- adaptivity, using element-wise topological operations and nodal perturbations as demonstrated in figure 1.4. The mesh optimisation procedure utilised by Fluidity-ICOM is composed of three steps: the computation of a mesh quality metric, the mesh optimisation procedure, performed using local element operations and nodal perturbations to improve the quality of the mesh as defined by this metric, and the interpolation or projection of the model fields from the old donor mesh to the new optimised target mesh.

Fluidity-ICOM mesh adaptivity is controlled with a mesh adaptivity metric, specified using a tensor field M defining a Riemannian metric for the size of the element in an optimised mesh space. In this metric space a displacement $\boldsymbol{v}$ has length $\sqrt{(\boldsymbol{v}^T M \boldsymbol{v})}$, and the ideal simplex has edge lengths all equal to unity. From this metric one may define an objective functional defining the quality of an element. This functional can then be used to test for the improvement in mesh quality of element-wise topological operations and nodal perturbations. A key property of this approach is that the use of a general Riemannian metric allows the optimised mesh to have an arbitrary anisotropy, with elements stretched to conform to the curvature specified by the metric.

Hence, if one can derive a metric tensor M based upon estimates of the model error, then this can be used to generate a mesh with the model error equally distributed throughout the computational domain, thereby increasing resolution in regions of increased dynamical importance. In Pain et al. [2001] and Piggott et al. [2005] the formulation of a mesh quality metric derived from estimation of the current interpolation error of simulation fields is described. A more advanced error metric formulation, controlling the error in a specified quantity (or "goal") and derived via the use of an adjoint model, is described in Power et al. [2006]. The issue of error metric formulation will be returned to in section 4.2.2.

The donor and target meshes generated by the Fluidity-ICOM mesh adaptivity procedure have, in general, no relationship to each other – they are each some unrelated covering partition of the same domain. Hence the final interpolation step is challenging, particularly if one wishes to preserve key system properties, such as integral quantities, or if one wishes to formulate an interpolant that is optimal in some sense, such as an L_2 optimal (Galerkin) projection. This problem will be further discussed in chapter 3 and section 4.2.3.

Several examples of the application of Fluidity-ICOM to geophysical systems are given in Piggott et al. [2008a]. These examples include simulations of a barotropic gyre, an early simulation of the thermally driven rotating annulus, and 2D and 3D density currents. No quantitative comparisons against existing models or available data is supplied. A much more thorough application of

Fluidity-ICOM using dynamic mesh adaptivity is described in Munday et al. [2010]. Here Fluidity-ICOM is used to simulate the flow past a cylinder as an idealised analogue for flow past an island. Dynamic mesh adaptivity is used to resolve the development of small scale detached shear layers within a much larger, more coarsely resolved, background flow.

1.1.4. Fluidity-ICOM and geophysical fluid dynamics modelling

Geophysical systems are, to leading order, in thermal wind balance: there is a leading order balance between the pressure gradient acceleration and the buoyancy and Coriolis accelerations. A numerical method that fails to represent this balance relation accurately can be susceptible to pressure gradient errors, similar in nature to the hydrostatic pressure gradient errors associated with σ-coordinate modelling, that pollute the resulting solution. A key question in applying Fluidity-ICOM to geophysical systems is therefore whether accurate physical balance representation can be maintained on arbitrary unstructured meshes.

Furthermore, when applying dynamic mesh adaptivity, there is the additional potential for imbalance injection by the mesh optimisation procedure itself. A system that is in balance on an old donor mesh need not remain balanced after interpolation onto a new optimised target mesh. This issue may also arise in when interpolating data between fully structured meshes, and hence a study of the problem (and techniques for tackling it) may be relevant for more conventional structured mesh formulations. Since Fluidity-ICOM utilises general dynamic mesh adaptivity, there is in general no relationship between the donor and target meshes, other than that they are covering partitions of the same domain. This presents a particularly challenging interpolation problem. Hence a further question in applying Fluidity-ICOM to geophysical systems is whether physical balance can be preserved in the interpolation of solution fields when applying dynamic mesh adaptivity. In this thesis this issue is discussed in the context of general unstructured meshes.

Finally, there is a significant lack of detailed validation of the numerical methods implemented in Fluidity-ICOM. Tidal modelling using Fluidity-ICOM has previously been validated against tide gauge data and other models (see, for example, Wells et al. [2005a] and Wells et al. [2005b]). However, previous applications of the model to large-scale geophysical flows have been used primarily as illustrative examples, with few practical applications to physical oceanography and little comparison with existing models or available data. The degree to which dynamic mesh adaptivity leads to improved solutions in geophysical systems requires testing and evaluation.

1.1.5. Fluidity-ICOM development status

Fluidity-ICOM is an in-development code. This means that the features and technologies implemented in Fluidity-ICOM are in varying states of maturity and have varying stability. In addition, significant sections of the Fluidity-ICOM code base were re-engineered or extended over the course of this DPhil project. These developments have included: a new model interface, re-implementation of the dynamical core (including the Navier-Stokes solver), re-implementation of the low level Fluidity-ICOM parallelisation scheme, and new interfaces to model libraries (including the mesh adaptivity and parallel load balancing interfaces).

Due to the lack of model maturity, significant development work was required, as part of this DPhil project, in order to enable the research described in this thesis to be conducted. 17% of all revisions made to the main Fluidity-ICOM source code repository have been made directly as part of this DPhil project [2].

Details of the development of the new model interface and the related options system back-end (the "Spud-Diamond" system) are described in Ham et al. [2009]. More significant Fluidity-ICOM development contributions made as part of this project are listed in appendix B.

A related issue is that there is a relative lack of experience in applying unstructured mesh numerical methods to geophysical problems. In particular, it is challenging to compute many standard diagnostics on arbitrary (and possibly dynamically adapting) unstructured meshes. For example, it is challenging to compute vertical integrals of quantities for meshes that have no inherent vertical structure. A number of unstructured mesh diagnostics were utilised or developed as part of this project, the more significant of which are described in appendix A.

1.2. The thermally driven rotating annulus

The rotating annulus is a classic laboratory scale analogue for large scale geophysical flows, with over a half century of research already conducted. In this section the justifications for the use of annulus experiments as an analogue for geophysical flows will be given. The key flow regimes will be described, and an overview of previous numerical work using this system will be supplied. To conclude, a summary of key avenues that are open for numerical research will be outlined.

[2] Percentage of all commits to the main Fluidity-ICOM subversion repository, the "trunk", including all revisions from revision r1 to revision r14166 on 29/08/10.

1.2.1. The rotating annulus as a laboratory scale analogue

The goal of an experimental analogue for geophysical flows is to capture essential features of the large scale while enabling less essential features to be suppressed or neglected. In developing an analogue for the dynamics of the atmosphere, as a first approximation, the key properties to capture are the thermal forcing and the constraint imposed by the strong influence of rotation [Lorenz, 1967].

More precisely, the equations governing the dynamics of a system can be re-expressed in terms of non-dimensional variables (normalisation) generating a series of non-dimensional parameters expressing the scale of the system relative to the non-dimensional system coordinates. The dynamics of two systems with matched dimensionless parameters in the non-dimensional coordinates (more precisely, their trajectory in phase space), which have an identical initial condition (starting point in phase space) and boundary conditions, are necessarily identical, and are said to be "dynamically similar" [Batchelor, 1967]. Hence, if the physical parameters in an experiment are matched to those found in a large scale system, the experiment can be used as a small scale analogue for the large scale processes.

Early attempts to produce such analogues were conducted by Vettin [1857]. These experiments used point-wise heating of a rotating disk of air, with smoke used as a passive tracer. Later experiments used ice at the disc centre, giving a more atmosphere-like thermal forcing and constituting an early "dish-pan" type experiment. Dish-pan type experiments are discussed in detail in Fultz [1951] and Fultz et al. [1959]. The dish-pan experiment is a logical first step in the construction of a laboratory analogue for the atmosphere. However the early experiments exhibited significant time variability, requiring long duration experiments, and experimental measurements that could be taken were limited [Lorenz, 1967].

The rotating annulus originates in the PhD work of R. Hide [Hide, 1953, 2006] as an analogue for the dynamics of the Earth's outer core, and was only later recognised for its relevance to atmospheric dynamics. In the very first experiments conducted using the system, as described in Hide [2006], a distinct wavenumber four type disturbance was observed within the chamber. This has since been recognised as a manifestation of a instability known as baroclinic instability.

The standard annulus system consists of an annular chamber with the outer and inner wall differentially heated and cooled and the system rotated about its axis of symmetry, as shown in figure 1.5. Higher order effects can be introduced via modifications to the standard system geometry or forcing. The effect of variation of rotation with latitude (the β-effect) can be introduced via sloping

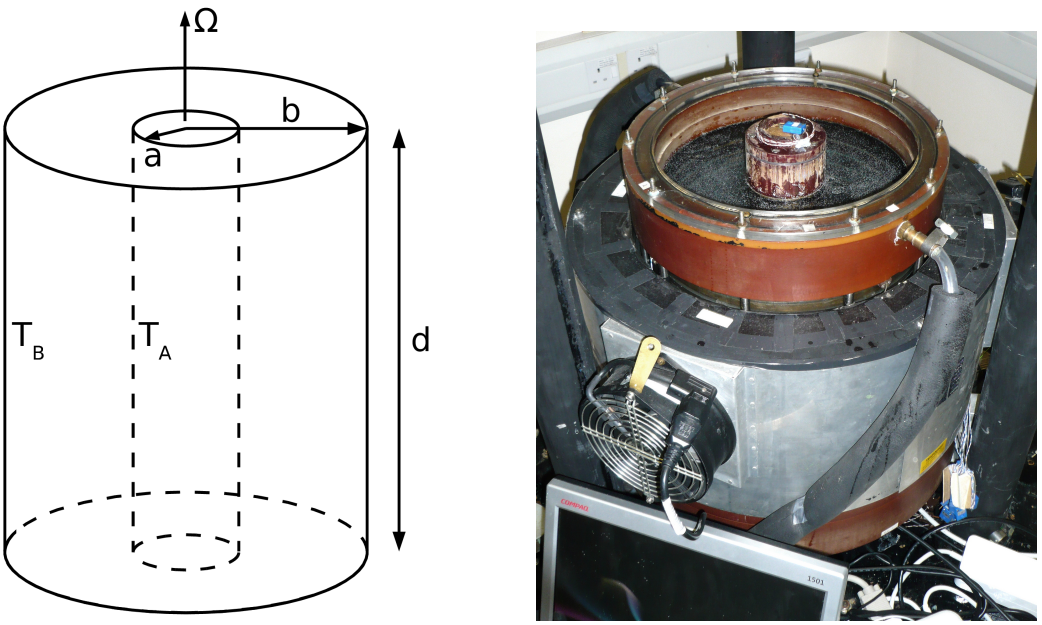

Figure 1.5.: Left: A standard annulus geometry. The chamber has major radius b, minor radius a and depth d. The outer and inner side-walls are differentially heated and cooled at temperatures T_b and T_a respectively, with typically $T_b > T_a$. The system is rotated about the axis of symmetry at a rotation rate ω. The upper surface may be free or in-contact with a rigid lid. Right: Photograph of a thermally driven rotating annulus experiment in the Geophysical Fluid Dynamics Laboratory (Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford). The upper lid has been removed, and in-situ tracer particles are visible.

top and bottom boundaries [Mason, 1975], although a direct analogy with the topographic β-plane for a barotropic fluid is complicated by the inherent baroclinic nature of the annulus flows. The effect of topography can be introduced with still more complicated upper and lower boundaries [Risch, 1999], and the effect of continental boundaries can be simulated via the insertion of radial barriers [Hide, 1998, Rayer et al., 1998, Wordsworth, 2008].

1.2.2. Baroclinic instability

One of the most striking features of a typical annulus experiment is a wave, or series of waves, dominating the interior of the flow, an example of which is shown in figure 1.6. These waves are a manifestation of baroclinic instability.

Consider a simplified 2D case: fluid in a rectangular cavity subject to gravity and thermally forced via differentially heated side-walls. A radial overturning is established, and at large time, and for sufficiently high Rayleigh number, the flow develops two distinct regions: an interior in which the fluid relaxes towards a stably stratified temperature profile with quasi-horizontal temperature contours [Read, 1992], and a diffusion dominated boundary layer region with near vertical temperature contours [Gill, 1966, Read, 1992].

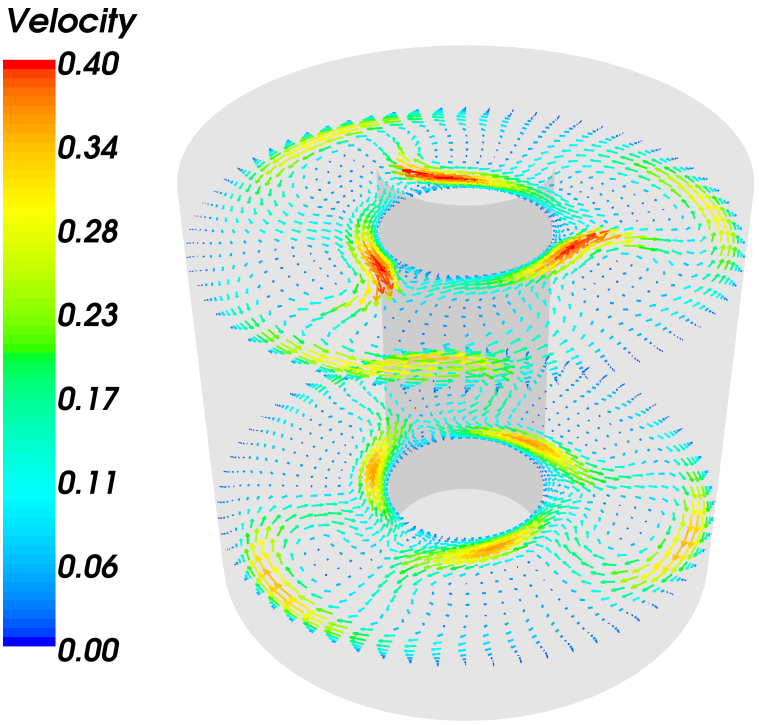

Figure 1.6.: Horizontal velocity field at two levels in a simulation of the thermally driven rotating annulus, with a fully developed baroclinic wave (from the $\omega = 2.764$ rad s^{-1} fixed structured mesh Fluidity-ICOM simulation described in section 4.1 and with parameters as given in table 2.1). Velocity units are cm s^{-1}.

Now consider an extension of this 2D domain in the third dimension to form an annular chamber. In the absence of rotation the dynamics are as in the 2D case, albeit with a geometric modification. If the annulus is now rotated about its axis of symmetry, and for sufficiently small Rossby number, the radial overturning is inhibited by the constraint of rotation – an axisymmetric geostrophically balanced flow in the radial direction is not permitted in the interior as no single valued azimuthal pressure gradient can exist to support it. Hence the thermal transport from the heated wall to the cooled wall is suppressed, leading to a tilting of temperature contours. Isopycnals become tilted away from the horizontal, and hence become tilted away from the isobars, leading to a baroclinic flow. In this case there is a region of instability, shown in figure 1.7, within which there is an opportunity for fluid parcel exchange to lead to potential energy release. The system is conditionally unstable against "sloping convection".

A tilt between the isobars and isopycnals is a necessary but not sufficient condition for the onset of baroclinic instability. It is a simple explanation of the energetics of the instability, but neglects physical constraints, in particular the conservation of potential vorticity. More rigorous theoretical treatments include these conservation principles, and are often conducted using a quasi-geostrophic framework. Two early models of this type are described in Charney [1947] and Eady [1949], with extensions for the thermally driven rotating annulus in Hide [1969] and Hide and Mason [1975]. For the rotating annulus, the Ekman layers have a significant influence upon the interior dynamics.

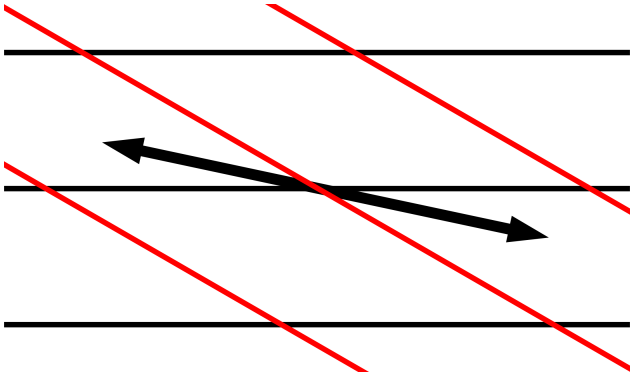

Figure 1.7.: Sloping convection – geopotential contours in black (parallel to isobars for a system in hydrostatic balance) and isopycnals in red (after [Hide and Mason, 1975] figure 1b). In fluid parcel exchange between the regions indicated by arrows there is a net potential energy release.

Hence the inclusion of Ekman layer boundary conditions is essential for Eady-type models to give a realistic description of the onset of baroclinic instability for this system [Barcilon, 1964, Hide, 1969].

It should be noted that at first glance Eady-type models appear to provide a rather unsatisfactory explanation of the annulus system, with most of the built-in assumptions being generally very poorly satisfied. Also, these models make an assumption of zero bulk potential vorticity, with the instability originating from edge-wave interaction rather than resulting from a sign change in the meridional potential vorticity gradient (a condition more commonly met in the atmosphere) [Vallis, 2006]. The problem can be re-formulated via more exotic boundary modified potential vorticity formulations such as described in Bretherton [1966] in order to avoid this issue. In Bell and White [1988], an alternative approach is outlined, considering the stability of an internal baroclinic jet. Despite the apparent issues, zero potential vorticity Eady models are known to provide a remarkably accurate description of wavenumber transitions in the annulus [Hide and Mason, 1975]

1.2.3. Rotating annulus flow regimes

The general annulus flow regimes are, in order of increasing rotation:

1. Axisymmetric: There are two distinct axisymmetric regimes, a super-stratified "upper-symmetrical" regime and a diffusion dominated "lower-symmetrical" regime.

2. Weakly non-linear: The axisymmetric regime is unstable and baroclinic instability is observed, with finite amplitude baroclinic waves.

3. Irregular: Transition to turbulent dynamics.

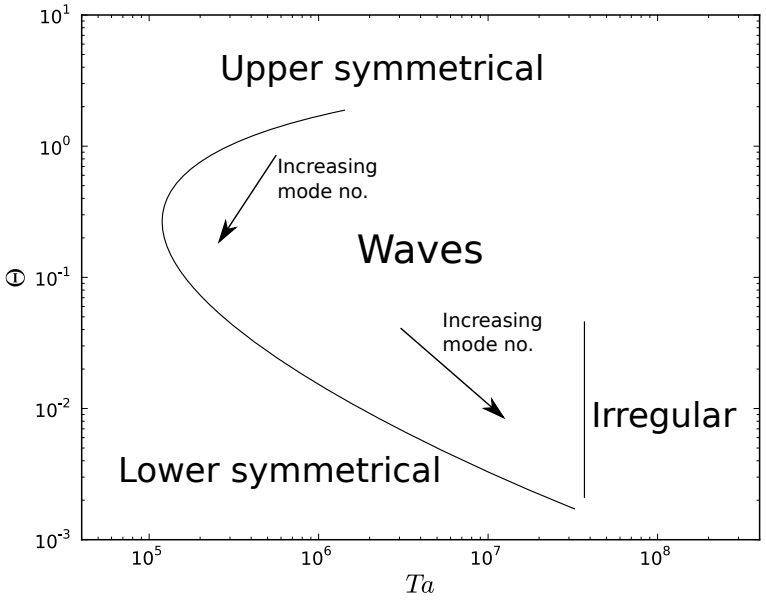

Figure 1.8.: Schematic of a typical regime diagram for the thermally driven rotating annulus, after figure 7 of Hide and Mason [1975] (for a Prandtl number $Pr = \nu/\kappa = 7.2$ fluid, where ν is the kinematic viscosity and κ the thermal diffusivity). The abscissa is the Taylor number as per equation (1.6), and the ordinate is the thermal Rossby number as per equation (1.5). The "anvil of instability" region, where baroclinic waves are observed, lies between two axisymmetric regimes: the super-stratified upper symmetric region, and the diffusion dominated lower symmetrical region. The anvil "knee" is the turning point in the anvil of instability, at the minimum Taylor number for which instability is observed.

Within the weakly non-linear regime there are a series of wavenumber mode transitions together with a number of observed sub-regimes corresponding to higher order instabilities. These include amplitude vacillation (vacillation in wave amplitude) and structural vacillation (vacillation in wave shape). In addition to this, the system is not entirely a boundary value problem – hysteresis is observed, with wavenumber transition points being a function of the direction from which they are approached in phase space [Read, 1992, Sitte and Egbers, 2000]. A more complete description of higher order annulus regimes can be found in Früh and Read [1997].

The dynamics of the rotating annulus are determined by a number of non-dimensional parameters. For a Boussinesq fluid with constant viscosity, diffusivity, and a linear equation of state, the five relevant non-dimensional parameters are the two aspect ratios $r_H = b/a$ and $r_V = (b-a)/d$, the Prandtl number $Pr = \nu/\kappa$, the thermal Rossby (or Hide) number:

$$\Theta = \frac{g\alpha\Delta T}{\omega^2}\frac{d}{(b-a)^2}, \tag{1.5}$$

and the Taylor number:

$$Ta = \frac{\omega^2}{\nu^2}\frac{4(b-a)^5}{d}. \tag{1.6}$$

Here a is the minor annulus radius, b the major annulus radius, d the tank depth, g the gravitational acceleration, α the thermal expansion coefficient, ν the kinematic viscosity, ω the magnitude of rotation, and $\Delta T = T_B - T_A$, where T_A and T_B are the temperatures of the inner and outer sidewalls respectively. A more complete description of the non-dimensional parameters for the rotating annulus is given in Fowlis and Hide [1965].

The dynamics of the rotating annulus are primarily determined by the thermal Rossby and Taylor numbers, with the Prandtl number playing a role in the more detailed structure of the resulting flow regimes [Castrejón-Pita and Read, 2007]. This is particularly the case when the Taylor number is sufficiently large for instability to be permitted [Hide, 1969]. A typical annulus regime diagram, showing the location of the general annulus regimes in terms of the thermal Rossby number and Taylor number, is shown schematically in figure 1.8.

The Taylor number expresses the strength of rotation to viscosity [Lorenz, 1967, Hide and Mason, 1975]. Its square root is sometimes referred to as a rotational Reynolds number [Lorenz, 1967]. The Taylor number is empirically found to be the abscissa that best collapses annulus regimes onto a single regime diagram in experiment [Fowlis and Hide, 1965] and is a one parameter defining the location of the regime diagram "knee" (the turning point in the anvil of instability in an annulus regime diagram, illustrated in figure 1.8) when performing linear stability analysis [Hide, 1969]. The thermal Rossby number defines the relative strength of buoyancy to rotation. Hence, combined, the Taylor number and thermal Rossby number give an indication of the two dominating features of geophysical flows.

Other important non-dimensional parameters can be expressed in terms of the Prandtl, thermal Rossby, and Taylor numbers. These include the Ekman number, characterising the strength of viscosity relative to rotation:

$$Ek = 2\sqrt{Ta^{-1}r_V}. \tag{1.7}$$

Assuming leading order thermal wind balance leads to a characteristic horizontal velocity scale $U \sim g\alpha\Delta T/\omega$. Hence the Rossby number $Ro = U/\left[(b-a)\,f\right]$ is related to the thermal Rossby number via:

$$Ro = \Theta r_v. \tag{1.8}$$

1.2.4. Numerical simulation of the rotating annulus

Early numerical simulations of the thermally driven rotating annulus were conducted by Williams [1969], in which the Navier-Stokes equations were integrated on a uniform finite difference C-grid. These simulations reproduced qualitative features of the flow, including a (quote) *"finite amplitude Eady wave"*. A somewhat more advanced C-grid model for the annulus of similar type, the Met Office / Oxford Rotating Annulus Simulation (MORALS) code [Farnell and Plumb, 1975, 1976], uses a hyperbolic tangent stretched mesh to concentrate computational resources into the side-wall (thermal) and lid and base (Ekman) boundary layers. MORALS was compared against experiment in James et al. [1981], where it was noted that the model gave a good representation of low wavenumber flows, while at higher wavenumber discrepancies were attributed to low grid resolution. Also, the wavenumber transition points were found to not correspond exactly with experiment. Further experimental comparison was conducted in Hignett et al. [1985], in which the accuracy of this same model was tested over a wider range of regimes, including a comparison of the system heat transport against experimental measurements. The heat transport of the system was found to be poorly represented by the model, attributed to limited model resolution. Still more recently, MORALS has been used to investigate heat transport in the context of eddy parameterisations in Read [2003], for more detailed regime diagram tracing in Young and Read [2008a] and for testing of data assimilation and forecasting techniques in Young and Read [2008b].

Simulations of non-standard annulus geometries have been attempted, but are of limited scope. More general mesh representations were attempted in the DPhil work of Elliott [1995] using a control volume finite element model. Technical issues prevented this model being applied to full 3D annulus simulations. In Risch [1999] topographic simulations were conducted using a modified version of MORALS with boundary stair-casing, and with a high-viscosity zero-diffusivity sponge region approach. The former led to spurious eddies on the sloping boundary, while the latter suffered from heat loss into the sponge region. Further simulations were also conducted using a modified version of an early MITgcm model with stair-cased boundaries. Simulations of the full

radial barrier annulus were conducted using another modified version of MORALS in Rayer [1992] and Rayer [1994], although the model gave a very poor representation of the horizontal flow field in the region of the tank inner wall.

Taking a slightly different approach, a two-layer quasi-geostrophic model for the annulus is described in Williams et al. [2009], and is used to investigate the emergence and impact of inertia-gravity waves in a two-layer system in Williams et al. [2003], Williams [2003] and Williams et al. [2004].

Hence, while the thermally driven rotating annulus has a long history of experimental research, numerical research has been primarily limited to the simplest geometries and to lower Taylor number flows. Simulations of non-standard geometries, such as the inclusion of a full radial barrier, or base and lid topography, have been limited in scope and success. Previous experimental research makes the rotating annulus ideally suited to the validation of new numerical methods, and the availability of new numerical methods makes possible the simulation of new annulus geometries and regimes.

1.3. Thesis outline

The aim of this DPhil project is to assess the applicability of unstructured dynamic mesh adaptive numerical modelling to geophysical fluid flows. Based upon this assessment, this project has the further aim of applying these methods to simulate previously challenging or inaccessible rotating annulus systems, particularly at higher Reynolds numbers and in non-standard domains.

In particular, the issue of accurate balance representation in unstructured mesh finite element numerical modelling will be addressed. Further, the problem of field transfer in applying dynamic mesh adaptivity to geophysical systems will be discussed, with emphasis on the implications for the maintenance of physical balance when performing mesh optimisation. Following these studies, unstructured dynamic mesh adaptive simulations of the thermally driven rotating annulus will be conducted, and used to assess the performance of the method. These simulations will be extended to annulus regimes and geometries that have been previously challenging to simulate numerically.

The thesis will proceed as follows. Chapter 2 will discuss the origin of physical balance errors in numerical modelling of geophysical systems. The finite element discretisation of the Navier-Stokes equations using a continuous Galerkin finite element method will be outlined. Sources of physical balance errors in the discrete system will be identified, and used to motivate a balanced

pressure decomposition method for accurate physical balance representation. The accuracy of this method will be quantified, and validated by comparison of numerical simulations of the thermally driven rotating annulus against experiment. Chapter 3 will further consider the origin of physical balance errors when applying dynamic mesh adaptivity. In the context of shallow-water ocean modelling, an interpolant will be developed which, for the linearised system on an f-plane with a steady balance permitting numerical discretisation, guarantees that a steady and balanced flow on a donor mesh remains steady and balanced after interpolation onto an arbitrary target mesh. Chapter 4 will consider the application of dynamic mesh adaptivity to simulations of the thermally driven rotating annulus. Fixed mesh and dynamic mesh adaptive simulations will be compared against each other, against MORALS, and against experiment, and used to identify advantages or shortcomings of the approach. Chapter 5 will extend these simulations to a high Taylor number thermally driven rotating annulus configuration with sloping end walls, and use these simulations to test theories concerning energy transfer locality (in Fourier space) in geostrophic turbulence. Finally, the thesis will conclude in chapter 6 with a summary of the DPhil project, and with possible topics for future research in this area.

Chapter 2.

Balance errors in finite element ocean modelling

This chapter is based and expands upon Maddison et al. [2011b].

The application of unstructured dynamic mesh adaptive finite element modelling to geophysical systems is, particularly on meshes that are unstructured in all three dimensions, relatively untested. This chapter is concerned with the application of unstructured mesh finite element numerical modelling to atmospheric and oceanic simulations. For these applications there is a leading order balance between Coriolis and buoyancy accelerations and the pressure gradient acceleration: geostrophic and hydrostatic balance. In order for a numerical model to be applicable to atmospheric or oceanic modelling, it is essential that the model be able to represent these physical balances accurately.

The issue of balance accuracy is acknowledged in terrain following σ-coordinate modelling. This class of models define a structured mesh with nodal spacings in a stretched vertical σ-coordinate, usually defined:

$$\sigma = \frac{z - \eta}{H + \eta}, \tag{2.1}$$

where η is the free surface height and H the fluid depth. In an x-y-σ coordinate system the horizontal pressure gradient becomes [Mellor et al., 1994]:

$$\left.\frac{\partial p}{\partial x}\right|_z = \left.\frac{\partial p}{\partial x}\right|_\sigma - \frac{\sigma}{H}\frac{\partial p}{\partial \sigma}\frac{\partial H}{\partial x}. \tag{2.2}$$

Due to the non-orthogonality of the σ-coordinate system a vertical (z-direction) pressure gradient contributes to each of the right-hand-side terms in equation (2.2). In a system close to hydrostatic balance, and for sufficiently large $\partial H/\partial x$, the horizontal pressure gradient therefore becomes the small residual between two near-cancelling terms, leading to a relatively large error in the horizontal pressure gradient. In Mesinger [1982] and Mellor et al. [1994] it is recognised that this leads to a further hydrostatic inconsistency issue, whereby numerical convergence is only assured if limitations are imposed upon both the horizontal and vertical mesh resolutions. A popular method for reducing this problem is to subtract a global mean density profile [Gary, 1973, Mellor et al., 1994] leading to a decrease in the error for systems with density primarily dependent upon the vertical coordinate. Other mitigation methods include the subtraction of a local density profile [Chu and Fan, 2003], computation of the pressure gradient force on a z-coordinate mesh [Beckmann and Haidvogel, 1993], and the use of higher order methods [Chu and Fan, 1997, McCalpin, 2005]. For a more detailed overview of methods used to mitigate these pressure gradient errors, see Berntsen and Thiem [2007].

A similar issue has been recognised in unstructured mesh finite element modelling. In Ford et al. [2004a] and Ford et al. [2004b] it is identified that pressure gradient errors lead to a spurious vertical motion in the presence of steep topography. The issue is addressed by decomposing the pressure into a buoyancy filtering hydrostatic component and a residual component, and then solving for the hydrostatic component diagnostically. This is an unstructured mesh generalisation of the pressure decomposition method described in Marshall et al. [1997a] and Marshall et al. [1997b]. Piggott et al. [2008b] extend this approach, decomposing the pressure into a combined hydrostatic-geostrophic component and a residual, and hence further address similar pressure gradient errors in a near geostrophically balanced flow.

In this chapter we address the issue of geostrophic and hydrostatic balance representation in unstructured mesh finite element ocean modelling. This is tested using simulations on fixed, non-adaptive, meshes. Additional physical balance errors arising from the application of dynamic mesh

adaptivity will be described in chapter 3, and the application of dynamic mesh adaptivity to the thermally driven rotating annulus will be described in chapter 4.

In this chapter the physical and mathematical basis for balance pressure decomposition methods will be provided, the accuracy of the approach quantified, and the utility for geophysical fluid dynamics modelling demonstrated. In section 2.1.1 the origin of balance errors will be described by considering a Helmholtz decomposition of terms in the momentum equation. This will be extended in section 2.1.2 to consider the origin of numerical balance errors in a continuous Galerkin finite element discretisation of the Navier-Stokes equations using the fractional timestep pressure projection method. This discussion will be used to motivate balanced pressure decomposition methods in section 2.1.3, which enable the accuracy of physical balance representation to be increased. The utility of the balanced pressure decomposition method will be quantified in section 2.2.2, using a fully non-linear numerical benchmark in exact geostrophic balance, and in section 2.3, via comparison of simulations of the thermally driven rotating annulus against experimental observations.

2.1. Origin of physical balance errors

In this section we discuss hydrostatic and geostrophic balance in the context of the Helmholtz decomposition of the incompressible Navier-Stokes equations. In section 2.1.1 the source of balance errors is identified via a Helmholtz decomposition of the Coriolis and buoyancy accelerations in the momentum equation. In section 2.1.2 the discretised finite element approximation to the continuous system is presented, and the origin of numerical imbalance errors discussed. This discussion is used to motivate balanced pressure decomposition methods in section 2.1.3, whereby the accuracy of balance representation is increased via the use of an additional diagnostic solve for the pressure associated with buoyancy and Coriolis accelerations.

2.1.1. Continuous formulation

Consider the incompressible Navier-Stokes equations for a Boussinesq fluid subject to Dirichlet boundary conditions:

$$\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u} \cdot \nabla \boldsymbol{u} + \boldsymbol{f} = -\nabla p + \boldsymbol{b} + \boldsymbol{F}, \tag{2.3a}$$

$$\nabla \cdot \boldsymbol{u} = 0, \tag{2.3b}$$

$$\boldsymbol{u} = \boldsymbol{u}_D \text{ on } \partial\Omega, \tag{2.3c}$$

where $\boldsymbol{u}$ is the Eulerian velocity, $\boldsymbol{u}_D$ is the value of velocity on the boundary $\partial\Omega$ bounding the domain Ω, $\boldsymbol{f}$ is the Coriolis acceleration, $\boldsymbol{b}$ is the buoyancy acceleration, $\boldsymbol{F}$ contains all remaining forcing terms (including any viscous dissipation) and p is the pressure (divided by a reference density ρ_0). We limit ourselves to the case of boundary conditions satisfying no-normal flow, $\boldsymbol{u} \cdot \hat{\boldsymbol{n}} = 0$ on $\partial\Omega$, where $\hat{\boldsymbol{n}}$ is a unit normal outward on $\partial\Omega$. Combining all forcing terms, equation (2.3a) can be re-written:

$$\frac{\partial \boldsymbol{u}}{\partial t} + \nabla p = \boldsymbol{G}, \tag{2.4}$$

where $\boldsymbol{G} = \boldsymbol{b} + \boldsymbol{F} - \boldsymbol{u} \cdot \nabla \boldsymbol{u} - \boldsymbol{f}$. We now make use of the Helmholtz decomposition [Weyl, 1940, Ladyzhenskaya, 1969, Denaro, 2003]: for a vector field $\boldsymbol{G} \in L^2(\Omega)$ with one boundary condition of:

$$\hat{\boldsymbol{n}} \cdot \boldsymbol{G} = G_{\partial\Omega,n} \text{ on } \partial\Omega, \qquad \hat{\boldsymbol{n}} \times \boldsymbol{G} = G_{\partial\Omega,t} \text{ on } \partial\Omega, \tag{2.5}$$

there exists a unique and orthogonal decomposition:

$$\boldsymbol{G} = \nabla\Phi + \nabla \times \boldsymbol{A} + \boldsymbol{H}, \tag{2.6}$$

for some scalar potential $\Phi \in H^1(\Omega)$, vector potential $\boldsymbol{A} \in H^1(\Omega)$ and harmonic residual $\boldsymbol{H} \in L^2(\Omega)$. $\nabla\Phi$ is irrotational, and is hereafter referred to as the *conservative* component of $\boldsymbol{G}$. $\nabla \times \boldsymbol{A}$ is solenoidal, and is hereafter referred to as the *rotational* component of $\boldsymbol{G}$. The harmonic residual $\boldsymbol{H}$ is both irrotational and solenoidal, and its presence is dependent upon the topology of Ω. For the remainder of this thesis we denote the conservative component $(\nabla\Phi)_{\boldsymbol{V}}$ of a vector $\boldsymbol{V}$ as $\boldsymbol{V_C}$, and

the solenoidal residual component $(\nabla \times \boldsymbol{A} + \boldsymbol{H})_{\boldsymbol{V}}$ as $\boldsymbol{V_R}$. In many case the rotational component and the solenoidal residual component are equivalent, dependent upon the system topology.

Since the velocity must be divergence free (2.3b), the Eulerian acceleration must also be divergence free, and hence from equations (2.3a) and (2.6):

$$\boldsymbol{G_C} = \nabla\Phi = \nabla p, \tag{2.7}$$

and the scalar potential in the Helmholtz decomposition of the forcing terms for incompressible Navier-Stokes can be identified as the dynamical pressure. Equation (2.6) is sufficient to diagnose the pressure field uniquely, up to some physically unimportant additive constant, demonstrating the role of the pressure gradient as a Lagrange multiplier for the momentum equation via which the incompressibility constraint is applied. Restated, the pressure gradient acts, and can only ever act, as a filter by which all conservative components of forcing terms in the momentum equation are exactly cancelled. It is this independence of the dynamics on conservative forcings that motivates a move to vorticity space formulations, in which the pressure gradient acceleration is entirely absent.

Note that the Helmholtz decomposition is unique only when appropriate boundary conditions are supplied. For the forcing terms in the incompressible momentum equation, these boundary conditions can be diagnosed directly from the velocity field boundary conditions. Also note that, while the potentials Φ and $\boldsymbol{A}$ are not unique, as one can arbitrarily add a constant to the former and a gradient term to the latter, any such modifications to the potentials have no physical significance – the values of the components $\nabla\Phi$ and $\nabla \times \boldsymbol{A}$ are themselves unique.

Turning to the issue of balance, equation (2.4) can be re-written:

$$\frac{\partial \boldsymbol{u}}{\partial t} + \nabla p = \boldsymbol{B} + \boldsymbol{F'}, \tag{2.8}$$

where $\boldsymbol{B} = -\boldsymbol{f} + \boldsymbol{b}$ and $\boldsymbol{F'} = -\boldsymbol{u} \cdot \nabla \boldsymbol{u} + \boldsymbol{F}$. Performing separate Helmholtz decompositions for $\boldsymbol{B}$ and $\boldsymbol{F'}$, and now noting that the scalar potential is associated with a pressure:

$$\boldsymbol{B} = \nabla p_b + \nabla \times \boldsymbol{A}_b + \boldsymbol{H}_b, \tag{2.9a}$$

$$\boldsymbol{F}' = \nabla p_r + \nabla \times \boldsymbol{A}_r + \boldsymbol{H}_r. \tag{2.9b}$$

Here ∇p_b filters only the conservative components of the Coriolis and buoyancy accelerations, and ∇p_r filters the conservative components from all remaining accelerations. We refer to p_b as the *balanced pressure* component of pressure, and p_r as the *residual pressure* component. Noting that the dynamics are dependent only upon the rotational components of the forcing terms in the momentum equation, in order to achieve an accurate representation of physical balance in a numerical model we therefore require an accurate representation of $\boldsymbol{B_R}$. Writing out equation (2.9a) in components:

$$B_{R,x} = -f_x - \frac{\partial p_b}{\partial x}, \tag{2.10a}$$

$$B_{R,y} = -f_y - \frac{\partial p_b}{\partial y}, \tag{2.10b}$$

$$B_{R,z} = -\rho g - \frac{\partial p_b}{\partial z}, \tag{2.10c}$$

where $\boldsymbol{b} = -\rho g \hat{\boldsymbol{z}}$, ρ is the fluid density (divided by a reference density ρ_0), and g is the gravitational acceleration. For flows that are close to geostrophic and hydrostatic balance we expect $\left|B_{R,i}\right| \ll \left|B_{C,i}\right|$ for $i \in \{x, y, z\}$, and hence the components of $\boldsymbol{B_R}$ become the small residual between two near-cancelling terms. This is a generalisation of the issue addressed in equation (2.2) for σ-coordinate models: in a numerical discretisation of the incompressible momentum equation there is potential for small errors in $\boldsymbol{B}$ or $\boldsymbol{B_C}$ to lead to significant errors in $\boldsymbol{B_R}$, and hence to lead to significant pressure gradient errors that can pollute the resulting model dynamics.

2.1.2. Discrete formulation

We consider a continuous Galerkin finite element discretisation of the incompressible Navier-Stokes equations (2.3a) and (2.3b) using the fractional timestep pressure projection method [Chorin, 1967, Temam, 1968, Gresho, 1990] and semi-implicit theta timestepping. We denote the finite element approximation of a field v as v^δ, and the nodal values of v^δ as $\tilde{v}$. The discretised momentum equation takes the form:

$$(M + \theta\Delta t(fL + N + A))\frac{\Delta\tilde{\boldsymbol{u}}_*}{\Delta t} = -N\tilde{\boldsymbol{u}}^n - fL\tilde{\boldsymbol{u}}^n + C\tilde{p}^{n-\frac{1}{2}} + \left[\int_\Omega \psi_i \left(\boldsymbol{b}^{\delta,n+\frac{1}{2}} + \boldsymbol{F}\right)\right], \tag{2.11}$$

where ψ_i are the velocity basis functions, $\Delta\tilde{\boldsymbol{u}}_* = \tilde{\boldsymbol{u}}_*^{n+1} - \tilde{\boldsymbol{u}}^n$ are the nodal values for the velocity increment, $\tilde{\boldsymbol{u}}^n$ are the nodal values for the velocity at time level n, and $\tilde{p}^{n-1/2}$ are the nodal values for pressure at the previous intermediate time level $n - 1/2$. $\boldsymbol{b}^{\delta,n+1/2}$ is the buoyancy acceleration at the next intermediate time level $n + 1/2$ solved via some discretisation of the prognostic tracer fields. $\boldsymbol{F}$ are all other forcing terms, which can include viscosity (not presented here). The final term in equation (2.11) is a vector of length equal to the number of velocity basis functions. Δt is the timestep size, and θ is an implicitness parameter in the range $[0, 1]$ with, for example, $\theta = 1$ yielding a first order backward Euler time discretisation and $\theta = 1/2$ yielding a second order Crank-Nicolson time discretisation [Crank and Nicolson, 1947]. For the remainder of this thesis we consider only the $\theta = 1/2$ Crank-Nicolson case.

M and $-C$ are the mass and gradient matrices respectively, N is the advection matrix, L the Coriolis matrix and A is a matrix for all terms in $\boldsymbol{F}$ that are treated implicitly. Expressed in terms of elemental basis functions the system matrices take the form:

$$M = \mathrm{diag}\left(M', M', M'\right), \tag{2.12a}$$

$$N = \mathrm{diag}\left(N', N', N'\right), \tag{2.12b}$$

$$L = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} M', \tag{2.12c}$$

$$C = \left(C^x, C^y, C^z\right)^T. \tag{2.12d}$$

where:

$$(M')_{ij} = \int_\Omega \psi_i \psi_j, \tag{2.13a}$$

$$(N')_{ij} = \int_\Omega \psi_i \bar{\boldsymbol{u}}^{\delta,n} \cdot \nabla \psi_j, \tag{2.13b}$$

$$(C^q)_{ij} = \int_\Omega \frac{\partial \psi_i}{\partial q} \xi_j \quad q \in \{x, y, z\}, \tag{2.13c}$$

and where the ψ_i and ξ_i are the velocity and pressure elemental basis functions respectively. Non-linear advection is treated by successive Picard iterations of the discrete Navier-Stokes equation (2.11), where overbars denote, on Picard iteration m:

$$\bar{\boldsymbol{u}}_m^{\delta,n} = \begin{cases} \boldsymbol{u}^{\delta,n} & \text{for } m = 1 \\ \frac{1}{2}\left(\hat{\boldsymbol{u}}_{m-1}^{\delta,n+1} + \boldsymbol{u}^{\delta,n}\right) & \text{for } m > 1 \end{cases} \tag{2.14}$$

and where $\hat{\boldsymbol{u}}_{m-1}^{\delta,n+1}$ is the velocity computed on the previous iteration (after the pressure projection step, described below). Second order accuracy in time is achieved with two Picard iterations [Ford et al., 2004a], and this is applied for the remainder of this thesis.

Note that in equation (2.11) the pressure from the previous intermediate timestep is used. This does not, in general, filter all conservative components of forcing terms in the new timestep, and hence the resulting velocity $\tilde{\boldsymbol{u}}_*^{n+1} = \tilde{\boldsymbol{u}}^n + \Delta\tilde{\boldsymbol{u}}_*$ is not discrete divergence free. Continuity is restored by projecting the velocity onto the divergence free space via the use of a Galerkin projection of the Helmholtz decomposition of the intermediate divergent velocity $\tilde{\boldsymbol{u}}_*^{n+1}$ together with the discretised continuity equation:

$$M\tilde{\boldsymbol{u}}^{n+1} = M\tilde{\boldsymbol{u}}_*^{n+1} + \Delta t C \Delta\tilde{p}, \tag{2.15a}$$

$$C^T \tilde{\boldsymbol{u}}^{n+1} = 0, \tag{2.15b}$$

where $\tilde{\boldsymbol{u}}^{n+1}$ are the nodal values of the non-divergent velocity at time level $n+1$, and $\Delta\tilde{p} = \tilde{p}^{n+1/2} - \tilde{p}^{n-1/2}$ are the nodal values of the pressure increment. Pre-multiplying equation (2.15a) by $C^T M^{-1}$ and applying the discrete continuity condition (2.15b) yields a diagnostic Poisson equation for the pressure increment:

$$C^T M^{-1} C \Delta\tilde{p} = -\frac{1}{\Delta t} C^T \tilde{\boldsymbol{u}}_*^{n+1}. \tag{2.16}$$

As in the continuous case, we see that the Helmholtz decomposition of the forcing terms in the momentum equation, here expressed as an intermediate divergent velocity, is sufficient to diagnose a unique pressure field increment.

The Dirichlet boundary condition (2.3c) can be imposed in either the strong (node-wise) sense via removal of the basis functions associated with boundary condition nodes, or in the weak (integral) sense by integrating the advection term in the momentum equation (2.11) and the divergence operator in the continuity equation (2.15b) by parts:

$$(N')_{ij}\,\tilde{u}_j^q = \int_{\partial\Omega} \psi_i \bar{\boldsymbol{u}}^\delta \cdot \hat{\boldsymbol{n}} \tilde{u}_D^q - \int_\Omega \nabla\psi_i \cdot \bar{\boldsymbol{u}}^\delta \psi_j \tilde{u}_j^q \quad q \in \{x, y, z\}, \tag{2.17a}$$

$$\left(C^T\right)_{ij}^q \tilde{u}_j^q = \int_{\partial\Omega} \xi_i \tilde{u}_D^q \hat{n}^q - \int_\Omega \frac{\partial \xi_i}{\partial q} \psi_i \tilde{u}_j^q \quad q \in \{x, y, z\}, \tag{2.17b}$$

Here advection is treated using a centred Galerkin discretisation. Greater stability can be achieved via upwinding schemes [Brooks and Hughes, 1982, Hughes et al., 1986, Hughes, 1987a], but these methods are not used in simulations described in this thesis. For a more complete description of the discretisation of the continuous Galerkin finite element discretisation of the Navier-Stokes equations the reader is referred to Ford et al. [2004a], Ford et al. [2004b] and Piggott et al. [2008a].

It remains to specify the choice of elemental basis functions for velocity and pressure: the velocity-pressure finite element pair. This choice is limited by the need to avoid spurious computational modes – in particular the well-known Ladyzhenskaya-Babuška-Brezzi (LLB) stability

criterion must be satisfied if spurious pressure modes are to be avoided [Ladyzhenskaya, 1969, Babuška, 1971, Brezzi, 1974]. A necessary but not sufficient condition for LBB stability is that there are more degrees of freedom for velocity than for pressure. We consider four choices of velocity and pressure basis functions, each defined in terms of Lagrange interpolating polynomials over simplex elements: piecewise linear continuous basis functions for velocity and pressure, P_1P_1; piecewise quadratic continuous basis functions for velocity and piecewise linear continuous basis functions for pressure, P_2P_1; piecewise linear discontinuous basis functions for velocity and piecewise linear continuous basis functions for pressure, $P_{1\mathrm{DG}}P_1$; piecewise linear discontinuous basis functions for velocity and piecewise quadratic continuous basis functions for pressure, $P_{1\mathrm{DG}}P_2$.

The P_1P_1 element is discussed in the context of unstructured mesh ocean modelling in Pain et al. [2005] and Piggott et al. [2008a]. Since in this case the inverse mass matrix M^{-1} is not sparse, in order for assembly of the discrete Laplacian matrix $C^TM^{-1}C$ in equation (2.16) to be practical the mass matrix must be diagonalised, for example by replacing the mass matrix M' with the row-summed lumped mass matrix:

$$(M'_L)_{ij} = \delta_{ij} \sum_i \int_\Omega \psi_i \psi_j. \tag{2.18}$$

The P_1P_1 element pair also does not satisfy the LBB stability criterion, and hence pressure modes must be filtered from the system. Following Piggott et al. [2008a] a fourth order diffusion term is introduced into the continuity equation (2.15b):

$$C^T\tilde{\boldsymbol{u}}^{n+1} = -\Delta t K \tilde{p}^{n+\frac{1}{2}}, \tag{2.19}$$

where the stabilisation matrix K takes the form:

$$K = Q^T M_L^{-1} Q, \qquad Q_{ij} = \frac{1}{2} \int_\Omega (\nabla \xi_i)^T h \nabla \xi_j, \tag{2.20}$$

and where h is an anisotropic mesh length scale. Here we take h to be the edge length tensor derived from the polar decomposition of the Jacobian mapping an element to the ideal simplex [Formaggia

and Perotto, 2001, Micheletti and Perotto, 2006]. This leads to the stabilised diagnostic Poisson equation for pressure:

$$\left(C^T M_L^{-1} C + K\right) \Delta\tilde{p}_* = -\frac{1}{\Delta t} C^T \tilde{\boldsymbol{u}}_*^{n+1} - K\tilde{p}_*^{n+\frac{1}{2}}, \tag{2.21}$$

where $\tilde{p}^{n+\frac{1}{2}} = \tilde{p}_*^{n+\frac{1}{2}} + \Delta\tilde{p}_*$, and where $\tilde{p}_*^{n+\frac{1}{2}}$ is the pressure from the previous Picard iteration (or the previous timestep on the first iteration). The addition of the stabilisation term into the continuity equation (2.19) preserves global mass conservation [Piggott et al., 2008a], but this is now an inconsistent discretisation of the original continuous system.

The P_2P_1 element pair is LBB stable [Hughes, 1987b] and hence no pressure filtering is required. As in the P_1P_1 case, in order for assembly of the discrete Laplacian matrix to be practical the mass matrix must be diagonalised. A row-summed lumped mass matrix as defined by equation (2.18) is unsuitable for this element pair as it is singular for piecewise quadratic velocity basis functions ψ_i [Zienkiewicz et al., 2005]. Instead, we follow Wilson [2009] and replace the mass matrix with the row-summed P_1 mass matrix assembled on a mesh composed of the sub-division of each quadratic element into (in two dimensions) four triangles (i.e. the P_1 iso P_2 row-summed lumped mass matrix). This sub-division procedure is somewhat ambiguous for tetrahedral elements (C. Wilson, personal communication), and hence the three-dimensional case is not considered.

The mixed discontinuous-continuous $P_{1\mathrm{DG}}P_2$ element pair is an LBB stable discretisation introduced in Cotter et al. [2009b]. The discontinuous velocity field is discretised using a discontinuous Galerkin method, and the discrete momentum equation (2.11) is modified to include integrals over interior element interfaces. See Ham and Cotter [2011] for a complete description of the $P_{1\mathrm{DG}}P_2$ discretisation of the incompressible Navier-Stokes equations. In $P_{1\mathrm{DG}}P_2$ simulations described in this thesis, interior element interface fluxes for advection are treated using the upwind flux, and viscosity is treated using the Bassi-Rebay scheme [Bassi and Rebay, 1997]. For this element pair the discontinuous velocity mass matrix is block diagonal, and hence the full consistent mass matrix can be used without any diagonalisation procedure. The $P_{1\mathrm{DG}}P_2$ element pair is tested in Cotter et al. [2009a] in the context of shallow-water ocean modelling, and is shown to have the property that geostrophically balanced states with a constant stream function on the boundary are exactly steady solutions of the discretised linear shallow-water equations on an f-plane. This element pair is further tested in Comblen et al. [2009] against other low order discontinuous finite element pairs,

and is shown to be the most accurate choice, for a given number of mesh elements, amongst those tested.

The $P_{1\mathrm{DG}}P_1$ element pair is formulated in an identical manner to the $P_{1\mathrm{DG}}P_2$ element pair, with the only difference being the lower order piecewise linear continuous representation of the pressure field. This element pair is LBB stable, but can require a high viscosity in order for optimal numerical convergence rates to be achieved [Comblen et al., 2009].

The P_1P_1 element pair is analogous to the unstaggered Arakawa and Lamb [1977] A-grid finite difference formulation. The $P_{1\mathrm{DG}}P_2$ element pair is a second order counterpart to the first order P_0P_1 finite element pair, which is itself analogous to the staggered Arakawa and Lamb [1977] B-grid finite difference formulation (see, for example, Williams and Zienkiewicz [1981]).

We now consider sources of numerical balance errors in the finite element discretisation of the incompressible Navier-Stokes equations. Limiting ourselves to consideration of geostrophic balance (hydrostatic balance can be considered using a similar argument), we note from equations (2.11) and (2.15b) that in order for geostrophically balanced states to be exactly steady solutions of the linearised system it is required that:

$$\forall \boldsymbol{u}^\delta \in U, p^\delta \in P, \quad fL\tilde{\boldsymbol{u}} = C\tilde{p} \Rightarrow C^T\tilde{\boldsymbol{u}} = 0, \tag{2.22}$$

where U and P are the discrete velocity and pressure spaces respectively. In Cotter et al. [2009a] and Cotter and Ham [2011] it is shown that, with a no-normal-flow boundary condition and for a velocity field associated with a constant stream function value on the boundary, a sufficient condition for this is that the following two embedding conditions be satisfied:

$$\forall \boldsymbol{u}^\delta \in U,\ f\hat{\boldsymbol{z}} \times \boldsymbol{u}^\delta \in U, \tag{2.23a}$$

$$\forall p^\delta \in P,\ \nabla p^\delta \in U. \tag{2.23b}$$

On an f-plane, condition (2.23a) states that the skew operator $\hat{\boldsymbol{z}} \times \boldsymbol{u}^\delta$ maps the velocity space into itself. Condition (2.23b) states that the gradient of the pressure space maps into the velocity space. The $P_{1\mathrm{DG}}P_2$ element pair satisfies these two embedding conditions [Cotter et al., 2009a,b, Cotter

and Ham, 2011], resulting in exactly steady geostrophically balanced states (with a constant stream function on the boundary) in the linearised system on an f-plane.

For systems close to geostrophic balance, or for which the optimal balance properties (2.23a) and (2.23b) do not hold, the accuracy of balance representation is limited by the accuracy of the diagnostic pressure solve (2.16). Considering the $\theta = 0$ explicit time-stepping case, and neglecting buoyancy, advection, and all forcing terms $\boldsymbol{F}$, the discretised momentum equation becomes:

$$M\frac{\boldsymbol{\Delta\tilde{u}}_*}{\Delta t} = -fL\tilde{\boldsymbol{u}}^n + C\tilde{p}^{n-\frac{1}{2}}. \tag{2.24}$$

Substituting this into the pressure correction step (2.15a), using $C^T\tilde{\boldsymbol{u}}^n = 0$, and omitting any pressure stabilisation gives:

$$M\boldsymbol{\Delta\tilde{u}} = -\Delta t f L\tilde{\boldsymbol{u}}^n + \Delta t C\tilde{p}^{n+\frac{1}{2}}, \tag{2.25}$$

which, expressed in components, becomes:

$$M\Delta\tilde{u}_x = \Delta t f M\tilde{u}_y^n + \Delta t C_x\tilde{p}^{n+\frac{1}{2}}, \tag{2.26a}$$

$$M\Delta\tilde{u}_y = -\Delta t f M\tilde{u}_x^n + \Delta t C_y\tilde{p}^{n+\frac{1}{2}}. \tag{2.26b}$$

For a state close to geostrophic balance, $M\Delta\tilde{u}_x$ and $M\Delta\tilde{u}_y$ are the small residual between two near-cancelling terms. Hence it is essential that a high accuracy solution for $\tilde{p}^{n+1/2}$ be obtained if significant pressure gradient errors are to be avoided.

2.1.3. Balanced pressure decomposition methods

Balanced pressure decomposition methods are based upon a separation of pressure into a component associated with the Coriolis and buoyancy accelerations, and a residual component, as per equations (2.9a) and (2.9b). We consider this in a slightly altered form to that presented previously

in Ford et al. [2004a], Pain et al. [2005] and Piggott et al. [2008b], in terms of a direct modification to the pressure projection equation (2.15a):

$$M\tilde{\boldsymbol{u}}^{n+1} = M\tilde{\boldsymbol{u}}_*^{n+1} + \Delta t \left(C\Delta\tilde{p}_r + C_b\tilde{p}_b\right), \tag{2.27}$$

where $p_b^\delta \in P_b$ is some diagnostic solution for the balanced pressure, $p_r^\delta \in P$ is the residual pressure component enforcing the incompressibility constraint (2.15b), and $\tilde{p}_b$ and $\tilde{p}_r$ are the nodal values for the balanced and residual pressures respectively. $-C_b$ is the balanced pressure gradient matrix:

$$C_b = \left(C_b^x, C_b^y, C_b^z\right)^T. \tag{2.28}$$

where:

$$\left(C_b^q\right)_{ij} = -\int_\Omega \psi_i \frac{\partial \chi_j}{\partial q} \quad q \in \{x, y, z\}, \tag{2.29}$$

and where the χ_i are the basis functions used in the discrete representation of the balanced pressure. This leads to a modified diagnostic Poisson solve for the residual pressure (omitting any pressure stabilisation):

$$C^T M^{-1} C\Delta\tilde{p}_r = -\frac{1}{\Delta t} C^T \tilde{\boldsymbol{u}}_*^{n+1} - C^T M^{-1} C_b \tilde{p}_b. \tag{2.30}$$

Equation (2.27) is equivalent to a Galerkin projection of the Helmholtz decomposition of $(\tilde{\boldsymbol{u}}_*^{n+1}/\Delta t + M^{-1}C_b\tilde{p}_b)$, where $-M^{-1}C_b\tilde{p}_b$ is the Galerkin projection of the balanced pressure gradient onto the velocity space. Hence the residual pressure still acts to preserve a discrete divergence free $\tilde{\boldsymbol{u}}^{n+1}$.

Alternatively, this modification of the pressure projection equation is directly equivalent to the introduction of $C_b\tilde{p}_b$ as an explicit right-hand-side term in the momentum equation (2.11):

$$(M + \theta\Delta t(fL + N + A))\frac{\Delta\tilde{\boldsymbol{u}}_*}{\Delta t} = -N\bar{\boldsymbol{u}}^n - fL\tilde{\boldsymbol{u}}^n + C\tilde{p}_r^{n-\frac{1}{2}} + C_b\tilde{p}_b + \left[\int_\Omega \psi_i \left(\boldsymbol{b}^{\delta,n+\frac{1}{2}} + \boldsymbol{F}\right)\right]. \quad (2.31)$$

Hence the balance pressure decomposition appears as an additional conservative forcing term $C_b\tilde{p}_b$ in the momentum equation. The residual pressure $\tilde{p}_r$ acts to enforce incompressibility, while the balanced pressure $\tilde{p}_b$ can be chosen so as to given a more accurate discretisation of the solenoidal component of the Coriolis and buoyancy accelerations.

For an explicit $\theta = 0$ time-stepping scheme, if $\tilde{p}_b$ is solved for using the Galerkin projection of the Helmholtz decomposition of the buoyancy and Coriolis accelerations with $P_b = P$, then equation (2.27) reduces to the original projection equation (2.15a). If a solution for $\tilde{p}_b$ is obtained using a method that is more accurate than the Galerkin projection of the Helmholtz decomposition of the buoyancy and Coriolis accelerations with $P_b = P$, then the overall accuracy of the Helmholtz decomposition of the intermediate divergent velocity will be increased. In particular, for the P_1P_1 element pair, if a consistent second order accurate solution for $\tilde{p}_b$ can be found then this can correct for the inconsistency introduced by pressure stabilisation methods.

In Ford et al. [2004a] and Piggott et al. [2008b] a balanced pressure decomposition of this form is used for just the hydrostatic (buoyancy acceleration filtering) part of pressure. In this case the Helmholtz decomposition of the buoyancy acceleration becomes, in the continuous space:

$$\boldsymbol{b_R} = \boldsymbol{b} - \nabla p_h, \quad (2.32)$$

where p_h is the hydrostatic pressure and $\nabla \cdot \boldsymbol{b_R} = 0$. By the fundamental theorem of calculus, and assuming $\boldsymbol{b}$ is in the z-direction:

$$\exists p_h \text{ s.t. } \frac{\partial p_h}{\partial z} = -\rho g. \quad (2.33)$$

Computing p_h using equation (2.33), equation (2.32) becomes:

$$\begin{aligned}\boldsymbol{b_R} &= \boldsymbol{b} - \frac{\partial p_h}{\partial x}\hat{\boldsymbol{x}} - \frac{\partial p_h}{\partial y}\hat{\boldsymbol{y}} - \frac{\partial p_h}{\partial z}\hat{\boldsymbol{z}} \\ &= -\frac{\partial p_h}{\partial x}\hat{\boldsymbol{x}} - \frac{\partial p_h}{\partial y}\hat{\boldsymbol{y}}. \end{aligned} \tag{2.34}$$

Noting that it is only the non-divergent residual of forcing terms in the momentum equation that are of dynamical importance, $\boldsymbol{b}$ can be replaced with $\boldsymbol{b_R}$ in the momentum equation.

If a no-normal-flow boundary condition is specified at lateral boundaries, then we require, on the lateral boundary $\partial\Omega_l$, zero normal acceleration of fluid parcels by $\boldsymbol{b_R}$:

$$\begin{aligned}\boldsymbol{b_R}\cdot\hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega_l &\Leftrightarrow \left(\frac{\partial p_h}{\partial x}, \frac{\partial p_h}{\partial y}, 0\right)^T \cdot \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega_l \\ &\Leftrightarrow (\hat{\boldsymbol{n}}\cdot\nabla_H)\, p_h = 0 \text{ on } \partial\Omega_l \\ &\Leftrightarrow (\hat{\boldsymbol{n}}\cdot\nabla_H) \int \frac{\partial p_h}{\partial z} dz = 0 \text{ on } \partial\Omega_l \\ &\Leftrightarrow (\hat{\boldsymbol{n}}\cdot\nabla_H) \int \rho g dz = 0 \text{ on } \partial\Omega_l, \end{aligned} \tag{2.35}$$

where $\nabla_H = (\partial/\partial x, \partial/\partial y, 0)^T$ is the horizontal gradient operator. A sufficient (but not necessary) condition for this to be satisfied is $(\hat{\boldsymbol{n}}\cdot\nabla_H)\rho = 0$ on $\partial\Omega_l$, i.e. zero normal horizontal density gradients on lateral boundaries. If the boundary condition (2.35) is not satisfied then the residual pressure must compensate for the inconsistency, leading to the potential for significant hydrostatic pressure gradient errors in the lateral boundary region. In particular, this approach is not in general suitable for systems with lateral buoyancy forcing, such as the thermally driven rotating annulus.

Piggott et al. [2008b] suggest two discontinuous Galerkin solvers for the hydrostatic pressure equation (2.33), one solving for the hydrostatic pressure p_h, and one solving directly for the horizontal gradient of the hydrostatic pressure $\nabla_H p_h$.

In this chapter we consider the extension of this method to a decomposition of pressure into a combined hydrostatic-geostrophic component and a residual [Piggott et al., 2008b], as per equations (2.9a) and (2.9b). Re-writing equation (2.9a):

$$\boldsymbol{B_R} = \boldsymbol{B} - \nabla p_b, \tag{2.36}$$

where $\nabla \cdot \boldsymbol{B_R} = 0$. Taking the divergence yields:

$$0 = \nabla \cdot \boldsymbol{B} - \nabla^2 p_b. \tag{2.37}$$

This equation can be discretised to gain a diagnostic solution for p_b, suitable for preconditioning the Helmholtz decomposition of the intermediate divergent velocity (2.27). Multiplying by a test function ϕ, integrating over the domain Ω and integrating by parts yields the weak form:

$$\int_\Omega \nabla\phi \cdot \nabla p_b = -\int_{\partial\Omega} \phi \left[\boldsymbol{B} - \nabla p_b\right] \cdot \hat{\boldsymbol{n}} + \int_\Omega \nabla\phi \cdot \boldsymbol{B}, \quad \forall\phi, \tag{2.38}$$

where $\partial\Omega$ bounds Ω and $\hat{\boldsymbol{n}}$ is a unit normal outward on $\partial\Omega$. If a no-normal-flow boundary condition is supplied on $\partial\Omega$, then we require zero normal acceleration of fluid parcels by $\boldsymbol{B}_R$, corresponding to the natural boundary condition $\boldsymbol{B_R} \cdot \hat{\boldsymbol{n}} = (\boldsymbol{B} - \nabla p_b) \cdot \hat{\boldsymbol{n}} = 0$ on $\partial\Omega$. Hence the weak form equation becomes:

$$\int_\Omega \nabla\phi \cdot \nabla p_b = \int_\Omega \nabla\phi \cdot \boldsymbol{B}, \quad \forall\phi. \tag{2.39}$$

Restricting ϕ and p_b to the C^0 continuous finite-dimensional space P_b completes the continuous Galerkin discretisation:

$$\int_\Omega \nabla\chi_i \cdot \nabla p_b^\delta = \int_\Omega \nabla\chi_i \cdot \boldsymbol{B}^\delta. \tag{2.40}$$

In this approach p_b^δ is obtained by a Galerkin projection of the continuous Poisson equation (2.37), rather than a Galerkin projection of the continuous Helmholtz decomposition (2.36). Hence, since a discrete incompressibility constraint $C_b^T \tilde{\boldsymbol{B}}_{\boldsymbol{R}} = 0$ is not imposed, LBB stability constraints do not apply in the choice of space P_b. This, for example, enables the basis functions for p_b^δ to be chosen to be of higher order that those for p_r^δ, enabling a more accurate discretisation of the Helmholtz decomposition (2.27).

In applying this balanced pressure decomposition to the Picard iterated theta-stepped momentum equation (2.11), we choose:

$$\boldsymbol{B}^{\delta} = -\boldsymbol{f}_{*}^{\delta} + \boldsymbol{b}^{\delta,n+\frac{1}{2}}, \tag{2.41}$$

where on Picard iteration m:

$$M\tilde{\boldsymbol{f}}_{*} = \begin{cases} fL\tilde{\boldsymbol{u}}^{n} & \text{for } m = 1 \\ fL\left[(1-\gamma)\,\tilde{\boldsymbol{u}}^{n} + \gamma\tilde{\boldsymbol{u}}_{m-1}^{n+1}\right] & \text{for } m > 1 \end{cases} \tag{2.42}$$

where $\tilde{\boldsymbol{u}}_{m-1}^{n+1}$ is the velocity field computed on the previous Picard iteration. This formulation allows the balanced pressure gradient term to be assembled directly in the momentum equation as per equation (2.31). γ is in the range $[0, 1]$ and represents a compromise between the more explicit use of the velocity field from the previous time level n, and the use of the current "best guess" for the velocity field at the next time level $n + 1$. If all terms other than the Coriolis acceleration are neglected in the momentum equation (2.11), and if the pressure correction step is neglected, then for $m > 1$ a value of $\gamma = \frac{1}{2}$ yields, for $\theta = \frac{1}{2}$ Crank-Nicolson time-stepping:

$$\tilde{\boldsymbol{f}}_{*} = -\frac{\Delta\tilde{\boldsymbol{u}}_{*}}{\Delta t} = (M + \tfrac{1}{2}\Delta t(L))^{-1}\,(fL\tilde{\boldsymbol{u}}^{n})\,. \tag{2.43}$$

i.e. for $\gamma = \frac{1}{2}$ and $m > 1$, the Coriolis acceleration as defined by equation (2.42) is, neglecting all other forcing terms and the pressure correction step, equivalent to the Crank-Nicolson integrated Coriolis acceleration. For the remainder of this thesis a value of $\gamma = \frac{1}{2}$ is applied.

Using this approach one can obtain a more accurate discretisation of the Helmholtz decomposition of the buoyancy and Coriolis accelerations, expressed in terms of a decomposition of an intermediate divergent velocity. In particular, this balanced pressure decomposition method results in a consistent discretisation of the Helmholtz decomposition of the buoyancy and Coriolis accelerations, even where pressure stabilisation methods are applied.

2.2. Balanced pressure decomposition accuracy

In this section the combined hydrostatic-geostrophic pressure decomposition method detailed in the previous section and defined by equation (2.37) is tested. In section 2.2.1 we demonstrate that, for the stabilised P_1P_1 discretisation of the incompressible Navier-Stokes equations, this leads to an increase in the accuracy of a Helmholtz decomposition of the Coriolis acceleration. In section 2.2.2 a fully non-linear test case is used to quantify the numerical geostrophic balance error when using the stabilised P_1P_1, P_2P_1, $P_{1\mathrm{DG}}P_1$ and $P_{1\mathrm{DG}}P_2$ discretisations, and to demonstrate that the balanced pressure decomposition method leads to a decrease in the discrete geostrophic balance error.

2.2.1. Helmholtz decomposition accuracy

The diagnostic equation (2.40) for the balanced pressure was implemented as part of Fluidity-ICOM with the resulting linear system solved with preconditioned conjugate gradients using the PETSc library [Balay et al., 1997, 2008, 2009]. The discretisation was implemented in a generic fashion, enabling the use of arbitrary order Lagrange interpolating basis functions for p_b.

The implementation was verified using structured triangular meshes in a 2D square domain $0 \leq x \leq 1$ and $0 \leq y \leq 1$ with a continuous solution $p_b = \frac{1}{\pi}\sin\pi x \cos\pi y$ and with $\boldsymbol{B} = \boldsymbol{B_C} = \nabla p_b$. The errors are shown in figure 2.1. For P_N elements the expected order $n+1$ convergence in the L_2 norm is observed for all element types tested, up to the errors associated with the linear solver tolerances.

The accuracy of the stabilised P_1P_1 Helmholtz decomposition of the Coriolis acceleration was tested with and without application of the balanced pressure decomposition method. The test was conducted using pseudo-isotropic unstructured triangular meshes (generated using Gmsh [Geuzaine and Remacle, 2009]) in a 2D doubly periodic square domain $0 \leq x \leq 1$ and $0 \leq y \leq 1$ using a velocity field of the form:

$$\begin{aligned}\boldsymbol{u} = &\left[-10\pi\sin(10\pi x)\cos(10\pi y) + 10\pi\cos(10\pi x)\right]\hat{\boldsymbol{x}} \\ &+ 10\pi\cos(10\pi x)\sin(10\pi y)\,\hat{\boldsymbol{y}},\end{aligned} \tag{2.44}$$

and with $f = 1$. This corresponds to a Helmholtz decomposition of the Coriolis acceleration of:

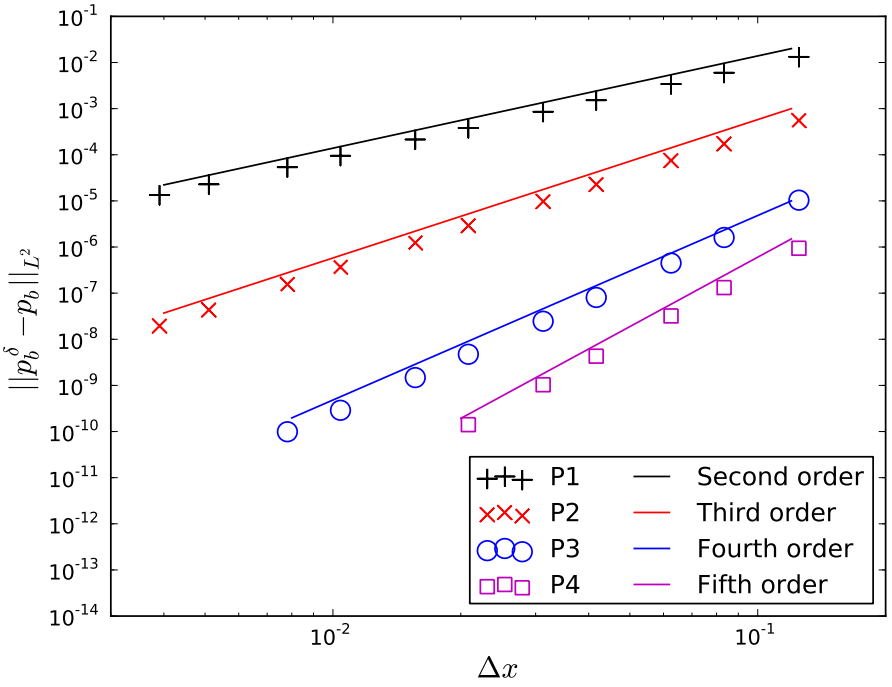

Figure 2.1.: Convergence test for the implementation of the balanced pressure solver (2.40) using P_1, P_2, P_3 and P_4 elements. The tests were conducted on structured triangular meshes with $0 \leq x \leq 1$ and $0 \leq y \leq 1$, and of progressively increasing resolution, with continuous solution $p_b = \frac{1}{\pi} \sin \pi x \cos \pi y$ and with $\boldsymbol{B} = \boldsymbol{B_C} = \nabla p_b$. All element types exhibit the expected order of convergence. For the P_3 and P_4 cases no further convergence is observed below an L_2 error of 10^{-10} (not shown), corresponding to the error associated with the tolerance of the linear solvers.

$$\begin{aligned} \Phi &= \sin(10\pi x) \sin(10\pi y), \\ \boldsymbol{B_R} &= -10\pi \cos(10\pi x)\, \hat{\boldsymbol{y}}. \end{aligned} \tag{2.45}$$

The errors in the conservative component of the Coriolis acceleration are shown in figure 2.2 for a stabilised P_1P_1 discretisation with no balanced pressure decomposition, and for a stabilised P_1P_1 discretisation with a P_1 and P_2 balanced pressure decomposition. With no balanced pressure decomposition, and at lower rotation rates, a convergence of order 1.2 is observed. At higher resolutions the order of convergence decreases significantly to 0.3. The introduction of a P_1 balanced pressure maintains an order of convergence of 1.1 at all resolutions. Increasing the balanced pressure element order from P_1 to P_2 leads to, on average, a factor 1.3 increase in the accuracy of the solution, but has little influence on the overall order of convergence. We conclude that the inconsistency in the discretisation of the Helmholtz decomposition of the Coriolis acceleration, introduced by the use of stabilisation in the discrete continuity equation (2.19), acts to decrease

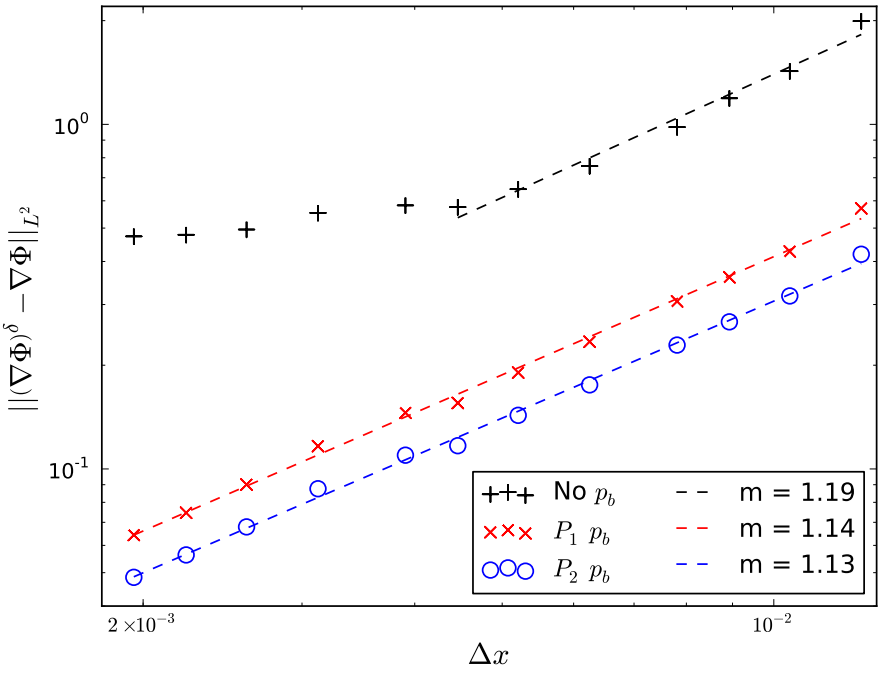

Figure 2.2.: L_2 errors in the conservative component of the Coriolis acceleration associated with the velocity field (2.44) in a 2D doubly periodic square domain $0 \leq x \leq 1$ and $0 \leq y \leq 1$, for a stabilised P_1P_1 discretisation with no balanced pressure decomposition, and for a stabilised P_1P_1 discretisation with P_1 and P_2 balanced pressure decomposition. Here the discrete scalar potential gradient is computed via a lumped mass Galerkin projection, as in the P_1P_1 Galerkin discretisation of the Helmholtz decomposition: $(\nabla\Phi)^\delta = \sum_i \psi_i M_L^{-1}\left(-C\tilde{\Phi} - C_b\tilde{\Phi}_b\right)$.

the order of convergence of the method. The use of a balanced pressure decomposition is able to correct for this inconsistency, leading to a higher order of convergence.

2.2.2. Geostrophic balance accuracy

In a discretisation of the incompressible Navier-Stokes equations, any imbalance errors injected into the simulation are combined with errors from a range of other sources. This complicates the task of numerical balance error quantification. Here, we present a mechanism by which an indication of the numerical geostrophic balance error can be computed in specific special cases, while making no further approximations to the discretised momentum equation (2.11).

It is noted in Hide [1998] that the domain topology can impose important constraints upon the dynamics in a rotating system. In particular, it is noted that the dynamics are entirely independent of system rotation if the following conditions are met:

1. The fluid is barotropic;

2. The domain is simply connected;

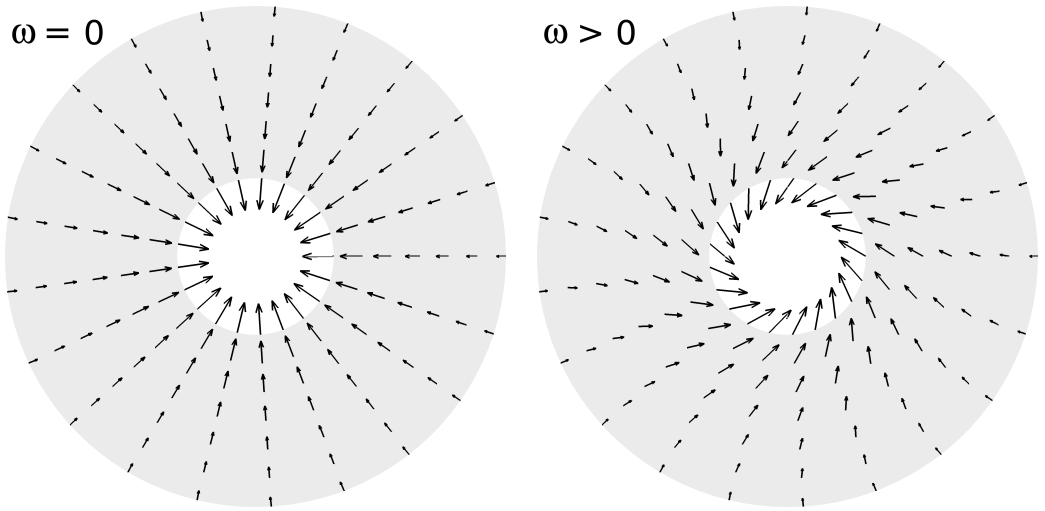

Figure 2.3.: Flow in an annular channel at Reynolds number 5 with an inflow boundary condition at the outer wall and a natural outflow boundary condition at the inner wall. Computed using a stabilised P_1P_1 discretisation of the Navier-Stokes equations with a P_2 balanced pressure decomposition. Left: Non-rotating case. Right: Rossby number $Ro = 1$ case. The flow is dependent on the rotation magnitude.

3. The f-plane approximation is used.

This result is a direct consequence of the application of Stokes' theorem to the two-dimensional vorticity equation [Hide, 1998]. In any system for which these three conditions are satisfied a pressure field is established that exactly cancels the Coriolis acceleration from the momentum equation – the Coriolis acceleration is guaranteed, in the continuous space, to be purely conservative and therefore of no dynamical significance.

As an example, consider a two dimensional annular channel with an inflow boundary condition at the outer wall and a free outflow boundary condition at the inner wall. In the the non-rotating case the fluid flows radially across the channel. In the rotating case an azimuthal Coriolis force acts on the flow. Since the pressure must be single valued, and since the domain is multiply connected, this azimuthal Coriolis force cannot be balanced by any pressure gradient force. An azimuthal component in the solenoidal residual of the Coriolis acceleration is established, and the flow is deflected to form a spiral as shown in figure 2.3. If a full radial barrier is placed across the channel an azimuthal pressure gradient is now permitted. All three conditions above are satisfied, and the flow is radial across the channel in both the non-rotating and rotating cases, as shown in figure 2.4. In the open rotating case angular momentum conservation requires the establishment of an azimuthal velocity component, while in the blocked rotating case angular momentum can be dissipated non-locally to the radial barrier via the pressure gradient force. Alternatively, viewed in vorticity space, the boundary conditions for the vorticity equation are independent of the magnitude of rotation in the simply connected case, and are dependent upon the magnitude of rotation in the multiply connected case.

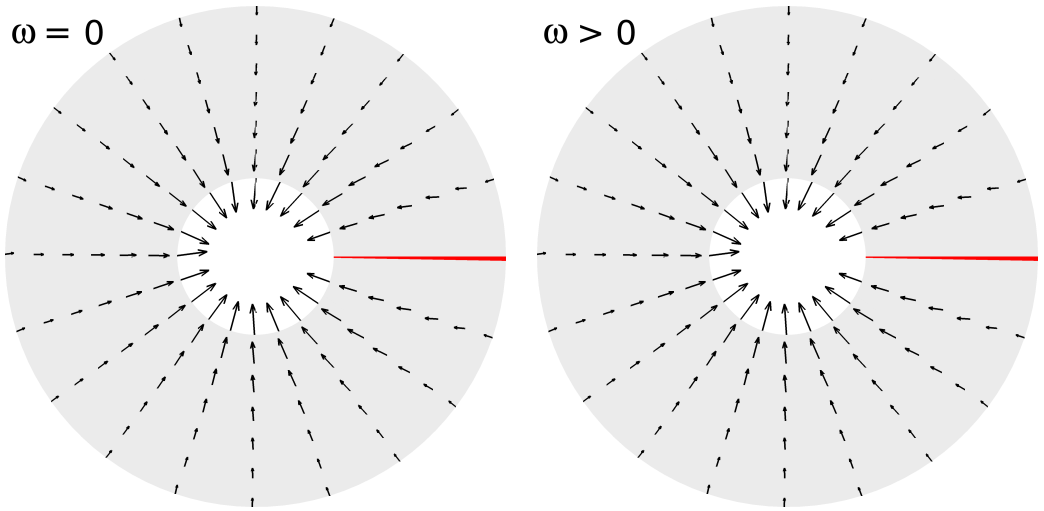

Figure 2.4.: Flow in an annular channel with a full radial barrier at Reynolds number 5 with an inflow boundary condition at the outer wall and a natural outflow boundary condition at the inner wall. Computed using a stabilised P_1P_1 discretisation of the Navier-Stokes equations with a P_2 balanced pressure decomposition. Left: Non-rotating case. Right: Rossby number $Ro = 1$ case. The flow is independent of rotation.

Hence if one chooses a configuration satisfying these three properties with some final steady state solution, the difference between the rotating and non-rotating cases gives a measure the truncation errors in the Coriolis acceleration, combined with numerical imbalance error. Applying this test with and without balanced pressure decomposition therefore gives a measure of the reduction of the numerical geostrophic balance error by the balanced pressure decomposition method.

The configuration chosen for geostrophic balance error quantification was the lid driven cavity benchmark [Botella and Peyret, 1998, Bruneau and Saad, 2006]. This test consists of a unit domain with a tangential forcing on the upper boundary and no-slip boundary conditions on all other boundaries. A regularised boundary condition $\boldsymbol{u} = 16x^2 (x-1)^2 \hat{\boldsymbol{x}}$ [Shen, 1991] was used on the upper boundary to avoid pressure singularities. A Reynolds number 1000 configuration was used for which a single steady state solution exists, shown in figure 2.5. Pseudo-isotropic unstructured triangular meshes of progressively increasing resolution were generated using Gmsh [Geuzaine and Remacle, 2009]. Tests were conducted using the stabilised P_1P_1, the P_2P_1, the $P_{1\mathrm{DG}}P_1$, and the $P_{1\mathrm{DG}}P_2$ element pairs, without any balanced pressure decomposition, and with a P_1 and P_2 balanced pressure decomposition applied. In addition, the $P_{1\mathrm{DG}}P_2$ element pair was tested with a P_3 balanced pressure decomposition applied. The boundary conditions were applied in the strong sense for the P_1P_1 and P_2P_1 cases, and in the weak sense for the $P_{1\mathrm{DG}}P_1$ and $P_{1\mathrm{DG}}P_2$ cases. For each discretisation and mesh resolution, non-rotating and rotating cases were simulated. The rotation rate used corresponded to Rossby number $Ro = 0.05$ with a rotation vector normal to the plane of the domain. The model was integrated with a maximum Courant number of 0.1 [Courant

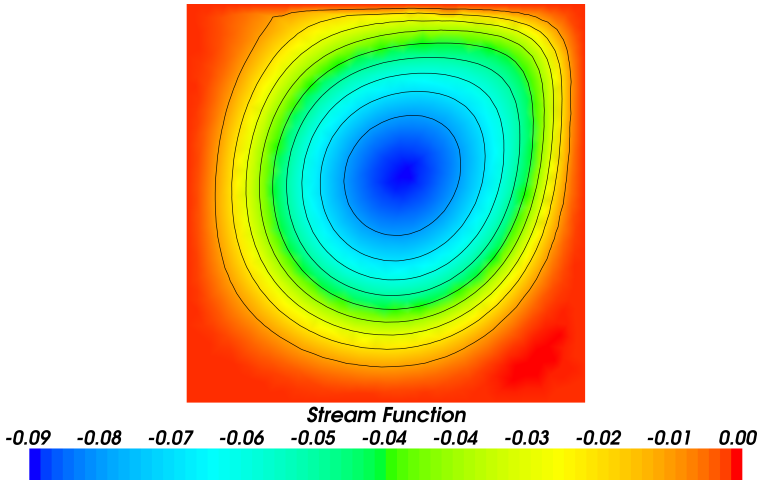

Figure 2.5.: Steady state stream function for a stabilised P_1P_1 simulation of the lid driven cavity at Reynolds number 1000. Here a negative stream function indicates clockwise rotation.

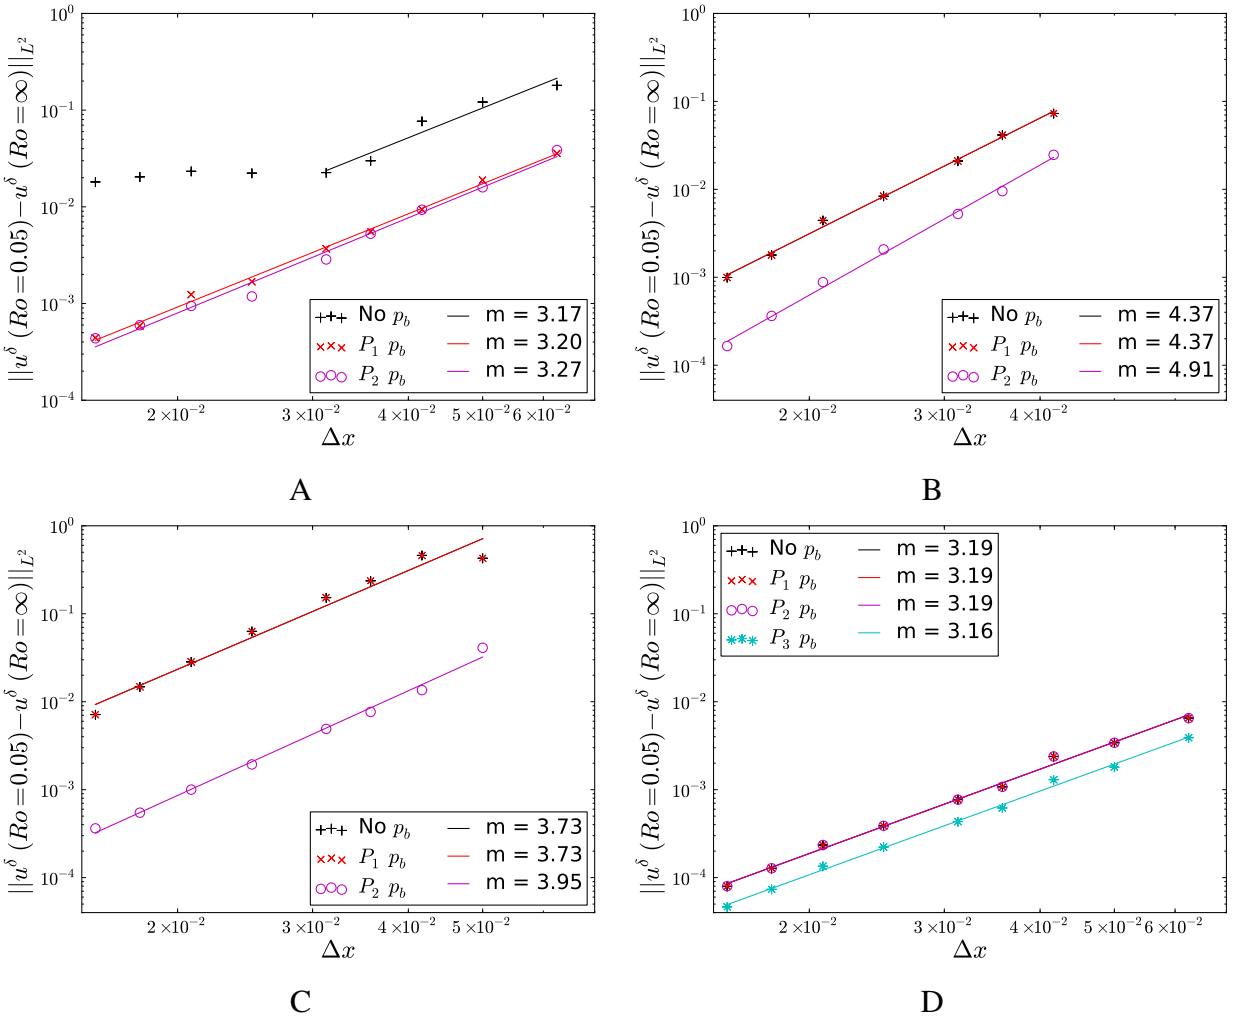

Figure 2.6.: Numerical imbalance convergence tests for the rotating lid driven cavity, conducted with and without balanced pressure decomposition. The L_2 difference between rotating and non-rotating cases is shown for a range of pseudo-isotropic unstructured triangular mesh resolutions. A: Stabilised P_1P_1. B: P_2P_1. C: $P_{1\mathrm{DG}}P_1$. D: $P_{1\mathrm{DG}}P_2$.

et al., 1928] until a steady state was attained, defined as the first timestep for which the maximum change in the velocity and pressure fields was $< 10^{-9}$.

The L_2 difference between rotating and non-rotating cases are shown in figure 2.6. An example of the difference between the rotating and non-rotating case, for the stabilised P_1P_1 element pair, is shown in figure 2.7. For the stabilised P_1P_1 configuration with no balanced pressure decomposition, a convergence between the rotating and non-rotating cases of order 3.2 is observed at low resolution. At higher resolution a much slower convergence rate of 0.5 is observed. This demonstrates the inconsistency introduced into the numerical discretisation by the fourth order stabilisation term in the continuity equation (2.19), previously seen in figure 2.2. Convergence is restored when the stabilisation term is removed (not shown), although a mesh level noise in the pressure field, characteristic of LBB instability, is observed. With the application of a P_1 balanced pressure decomposition to the stabilised P_1P_1 configuration, an order 3.2 convergence is observed over all resolutions. Increasing the balanced pressure order to P_2 decreases the L_2 difference between the rotating and non-rotating cases by an average factor of 1.1, and does not significantly affect the order of convergence. Hence it can be seen that the application of a balanced pressure decomposition method can guarantee a consistent discretisation of the Helmholtz decomposition of the Coriolis acceleration, and hence an accurate representation of geostrophic balance, when stabilisation methods are applied.

For the P_2P_1 configuration with no balanced pressure decomposition a convergence between the rotating and non-rotating cases of order 4.4 is observed. A P_1 balanced pressure decomposition does not alter the solution. Increasing the balanced pressure order to P_2 decreases the L_2 difference between the rotating and non-rotating cases by an average factor of 4.5, but leads to only a small change in the order of convergence to 4.9.

For the $P_{1\mathrm{DG}}P_1$ configuration with no balanced pressure decomposition there is a convergence between the rotating and non-rotating cases of order 3.7. A P_1 balanced pressure decomposition does not alter the solution, while a P_2 balanced pressure decomposition decreases the L_2 difference between the rotating and non-rotating cases by a substantial average factor of 26.8, but again leads to only a small change in the order of convergence to 4.0.

For the $P_{1\mathrm{DG}}P_2$ configuration with no balanced pressure decomposition there is a convergence between the rotating and non-rotating cases of order 3.2. The introduction of either a P_1 or a P_2 balanced pressure decomposition does not alter the solution. For the $P_{1\mathrm{DG}}P_2$ element pair, the discrete Laplacian matrix in the pressure projection equation (2.15a) and the P_2 balanced pressure decomposition equation (2.40) are, neglecting differing implicit boundary conditions, identical [Cotter et al., 2009a]. Hence a balanced pressure decomposition of up to order two offers no additional geostrophic balance accuracy. A P_3 balanced pressure decomposition yields a modest

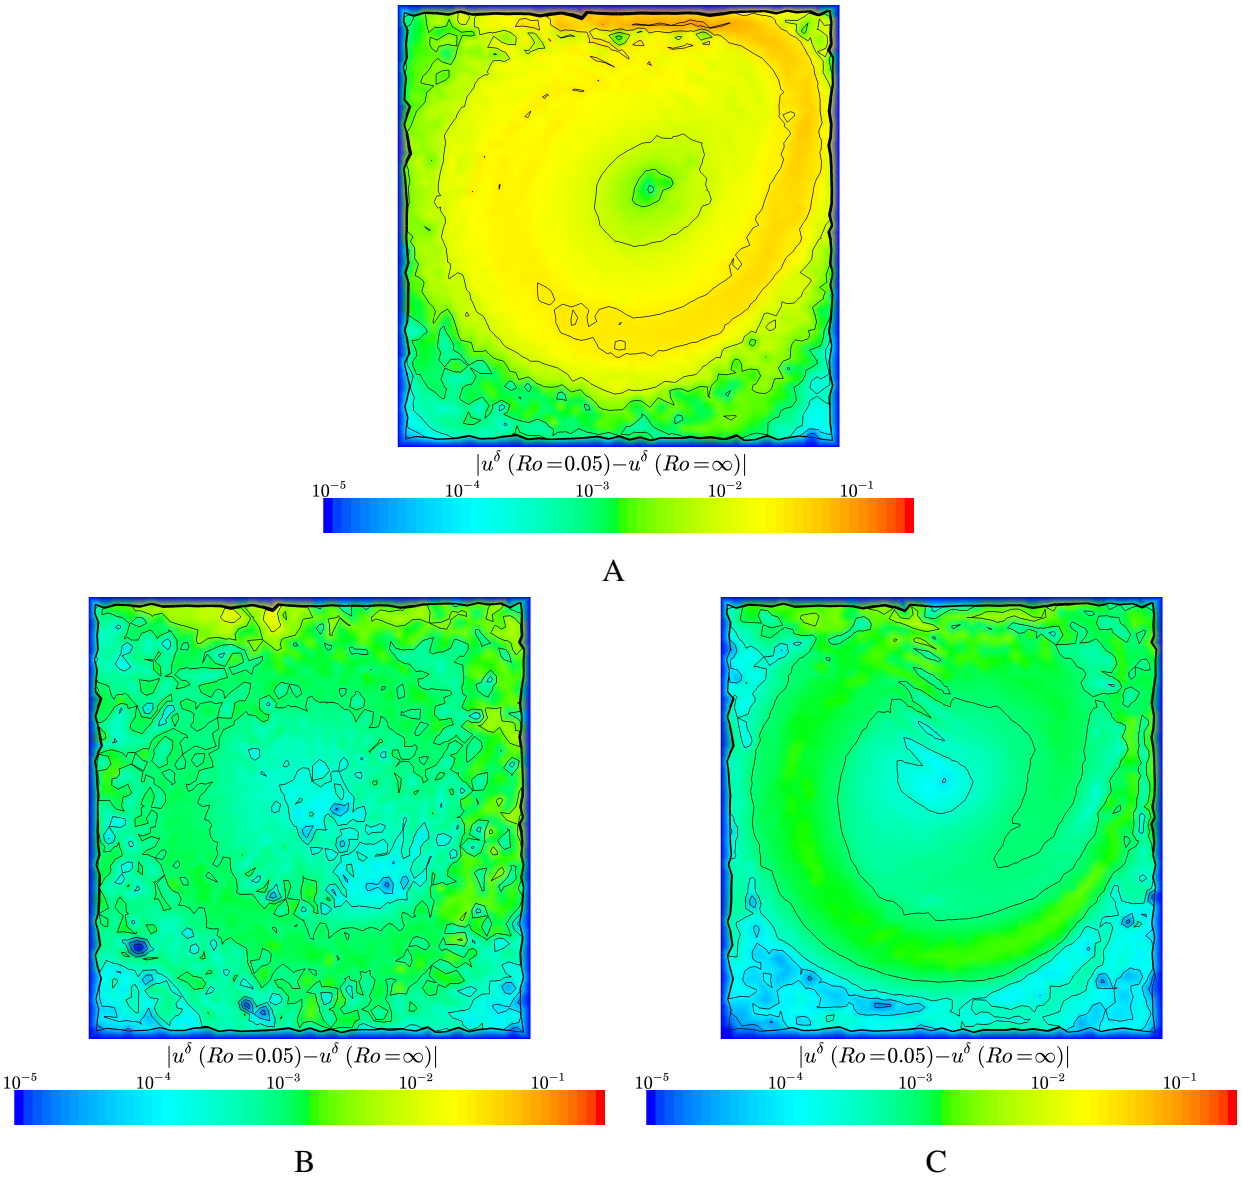

Figure 2.7.: Difference between the rotating and non-rotating cases for the numerical imbalance convergence tests with the rotating lid driven cavity. Shown for resolution $\Delta x = 1/40$, with a logarithmic colour scale. A: Stabilised P_1P_1. B: Stabilised P_1P_1 with P_1 balanced pressure. C: Stabilised P_1P_1 with P_2 balanced pressure.

factor 1.8 decrease in the difference between the rotating and non-rotating cases, with no significant change to the order of convergence. This demonstrates the particularly accurate representation of geostrophic balance using this element pair.

Hence it is concluded that, where the method used to compute the balanced pressure is more accurate than the method used to compute the residual pressure, a balanced pressure decomposition ensures a more accurate representation of geostrophic balance. For the stabilised P_1P_1 case a P_1 balanced pressure decomposition is able to ensure a consistent discretisation of geostrophic balance. For the P_2P_1 and $P_{1\mathrm{DG}}P_1$ case a P_1 balanced pressure decomposition has no effect, while a P_2 balanced pressure decomposition leads to a modest increase in the accuracy of geostrophic balance representation. For the $P_{1\mathrm{DG}}P_2$ case a balanced pressure decomposition of up to order two offers no additional geostrophic balance accuracy.

2.3. Application to the thermally driven rotating annulus

In this section unstructured mesh finite element simulations of the thermally driven rotating annulus are considered. Section 2.3.1 considers simulations using the stabilised P_1P_1 finite element pair, with and without the application of a balanced pressure decomposition. It is demonstrated that the balanced pressure decomposition method leads to a significantly more physical representation of the system. This improved accuracy is quantified via comparison with laboratory measurements. Section 2.3.2 considers simulations using the $P_{1\mathrm{DG}}P_2$ finite element pair. The $P_{1\mathrm{DG}}P_2$ element pair is known to have excellent geostrophic balance properties, and hence simulations using this element pair, with no balance pressure decomposition, are considered, and compared against the stabilised P_1P_1 simulations.

This section specifically tests only the application of balanced pressure decomposition methods, and the application of the optimally balanced $P_{1\mathrm{DG}}P_2$ finite element pair. A more complete comparison with laboratory data will be deferred until chapter 4.

2.3.1. Stabilised P_1P_1 finite element pair

The configuration of Read [2003] was used for this validation, with fluid parameters as given in table 2.1 and with a rigid in-contact lid at the upper boundary. The fluid diffusivity, viscosity, and expansion coefficient are derived from the data in Fowlis and Rossby [1964]. In Fowlis and Rossby

Parameter	Symbol	Value
Minor radius	a	2.5 cm
Major radius	b	8.0 cm
Depth	d	14.0 cm
Kinematic viscosity	ν	2.1×10^{-2} cm^2 s^{-1}
Thermal diffusivity	κ	1.3×10^{-3} cm^2 s^{-1}
Volumetric expansion coefficient	α	3.3×10^{-4} K^{-1}
Gravitational acceleration	g	981 cm s^{-2}
Inner wall temperature	T_A	18 °C
Outer wall temperature	T_B	22 °C
Rate of rotation	ω	1 rad s^{-1}

Table 2.1.: System parameters used for the thermally driven rotating annulus validation. The parameters are as used in Read [2003].

[1964] the fluid parameters are quoted as a function of the mean fluid density and temperature. A fourth order polynomial fit has previously been used to interpolate the fluid parameters from this data set. It was found during the course of this research that this interpolation method leads to a high interpolation error for the viscosity – a more accurate interpolation using linear interpolation of a Delauney triangulation of the data yielded a viscosity of 1.77×10^{-2} cm^2 s^{-1}, indicating a previous error of $\sim 20\%$. All simulations described in this section (and in chapter 4) use the higher viscosity of 2.1×10^{-2} cm^2 s^{-1}, as quoted in Read [2003]. Hence there is some discrepancy between the laboratory fluid viscosity and the simulation viscosity. The influence of this discrepancy will be discussed in section 2.3.4.

Laboratory data are available for this system for a range of rotation rates and for a number of annulus regimes [Hignett et al., 1985, Hignett, 1985, Read, 2003]. These data contain measurements of the system heat transport, here defined in terms of a non-dimensional Nusselt number: the heat transport of the system relative to the transport of a solid body of equal conductivity. In terms of the diffusive flow through the tank inner wall, $\partial\Omega_I$, this takes the form:

$$Nu = \frac{\ln b/a}{2\pi d \Delta T} \int_{\partial\Omega_I} \nabla T \cdot \hat{\boldsymbol{n}}. \tag{2.46}$$

Note that, as defined here, a *negative* Nusselt number corresponds to flow *out* of the domain. Hence the inner wall Nusselt number is negative. In addition, a thermocouple array placed at mid-radius and mid-height was used for direct in-situ measurement of the fluid temperature. More complete background information concerning the laboratory data is provided in section 4.1.2. The

calculation of system Nusselt numbers for unstructured mesh simulations is described in appendix A.2, and the algorithm used for virtual thermocouple arrays with unstructured meshes is described in appendix A.3

The system was simulated using Fluidity-ICOM using the stabilised P_1P_1 element pair, with and without the application of a balanced pressure decomposition. The simulations were conducted in parallel on 8 MPI processes using the Oxford Supercomputing Centre "redqueen" machine. Due to limitations inherent in the configuration of the Fluidity-ICOM simulations a linear equation of state was applied, a temperature independent diffusivity and viscosity used, and the centrifugal acceleration neglected. Note that more recent versions of Fluidity-ICOM permit temperature dependent parameters and the introduction of arbitrary source terms in the momentum equation (which can be used to introduce the centrifugal force). The prognostic temperature equation was solved via a control volume finite element method using the Sweby flux limiter [Sweby, 1984, Wilson, 2009]. No-slip boundary conditions for velocity were applied on all surfaces in the strong sense, and the thermal forcing applied on the inner and outer annulus walls via a strong boundary condition for temperature. A no-flux boundary condition for temperature was applied on the base and lid via a naturally imposed Neumann boundary condition. The resulting linear systems were solved via preconditioned conjugate gradients using the PETSc library [Balay et al., 1997, 2008, 2009, Saad and Schultz, 1986], with a smoothed aggregation algebraic multigrid method applied for the pressure solvers for both the residual and balanced pressure components [Vaněk et al., 1996, Kramer et al., 2010]. Some simulations also used the BoomerAMG multigrid method [Henson and Yang, 2002], via the hypre library [Falgout and Yang, 2009], for the pressure solvers.

The computational meshes used for these simulations were based upon the mesh used by MORALS (see section 1.2.4). MORALS makes use of a structured mesh with resolution stretched to resolve the side-wall (thermal and Ekman) and base and lid (Ekman) boundary layers, via a hyperbolic tangent stretch [Farnell and Plumb, 1975]. The mesh used for a given MORALS simulation is therefore dependent upon the system parameters. For a given structured MORALS mesh, a structured tetrahedral mesh was generated by dividing each cubic region into six tetrahedra [Tanizume et al., 1990]. A fully unstructured mesh was then generated using the Fluidity-ICOM mesh adaptivity library [Pain et al., 2001], with a mesh quality metric derived from the structured tetrahedral mesh via a polar decomposition of the Jacobian mapping elements to the ideal tetrahedron [Formaggia and Perotto, 2001, Micheletti and Perotto, 2006]. This generated fully unstructured meshes of locally equal resolution to a given MORALS mesh, as demonstrated in figure 2.8.

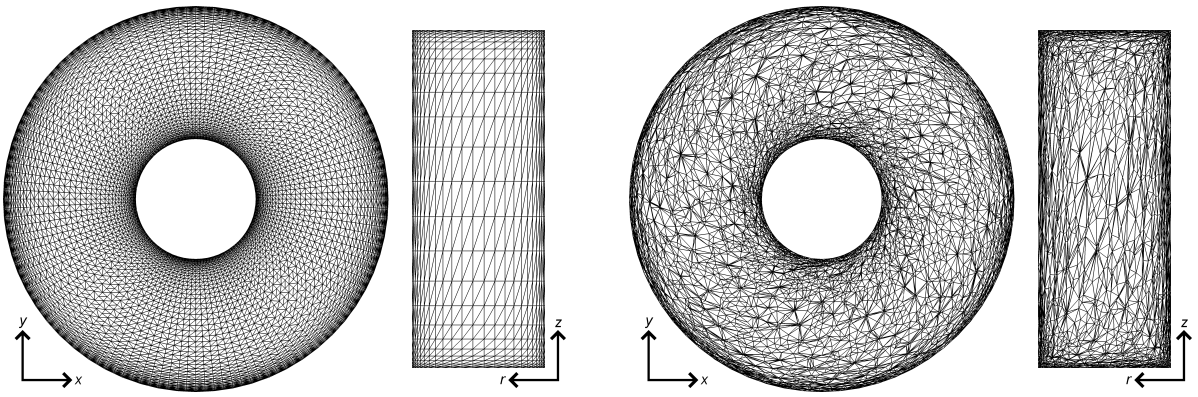

Figure 2.8.: Left: Vertical and horizontal sections through a structured tetrahedral mesh of resolution $24 \times 24 \times 64$ in the radial, vertical and azimuthal dimensions respectively, with nodes concentrated towards the boundaries via a boundary layer resolving hyperbolic tangent stretch as per Farnell and Plumb [1975] and as used by MORALS. Note that MORALS uses an Arakawa and Lamb [1977] C-grid formulation. The mesh shown here represents a structured tetrahedralisation of nodal points used by MORALS. Right: An unstructured tetrahedral mesh of locally equal resolution to the structured mesh, generated using the Fluidity-ICOM adaptivity library. Note that the meshes shown here are slices though the tetrahedral mesh. Hence the quality of the resulting triangulation is not necessarily an indication of the quality of the full 3D tetrahedralisation. The structured mesh has 36864 degrees of freedom, 203136 elements, and a maximum anisotropy (maximum-to-minimum element edge length ratio) of 43. The unstructured mesh has 37178 degrees of freedom, 195561 elements and a maximum anisotropy of 110.

Simulations were conducted at a rotation rate $\omega = 1$ rad s^{-1} using a mesh derived from a MORALS mesh with resolution $24 \times 24 \times 64$ nodes in the radial, vertical and azimuthal dimensions respectively. In order to maximise model efficiency, the timestep size was adapted after every timestep to target a maximum Courant number of unity. While the Crank-Nicolson time discretisation is formally stable [Gresho and Sani, 2000], it was found that a Courant number of order unity was required to achieve stable simulations. It is possible that this results from the explicit form of the balanced pressure decomposition time discretisation defined by equation 2.41. The model was initialised with zero velocity and with a linear stratification, and integrated for a total of 4000 s of simulated laboratory time. For further comparison, a simulation was conducted using the MORALS code at a resolution of $32 \times 32 \times 128$ in the radial, vertical and azimuthal dimensions respectively. The MORALS simulation was conducted with a linear equation of state, with temperature independent diffusivity and viscosity, and initialised with zero velocity and with a linear stratification [1].

The resulting mid-height temperature fields are shown in figure 2.9. Corresponding Hovmöller plots derived from the mid-height and mid-radius temperature fields are shown in figure 2.10, together with the corresponding laboratory observations. The Fluidity-ICOM simulation with no balanced pressure decomposition bears little similarity to the corresponding MORALS simulation or laboratory observations. In particular, with no balanced pressure decomposition, the interior

[1] MORALS data courtesy of Dr. R. M. B. Young, Department of Physics, University of Oxford.

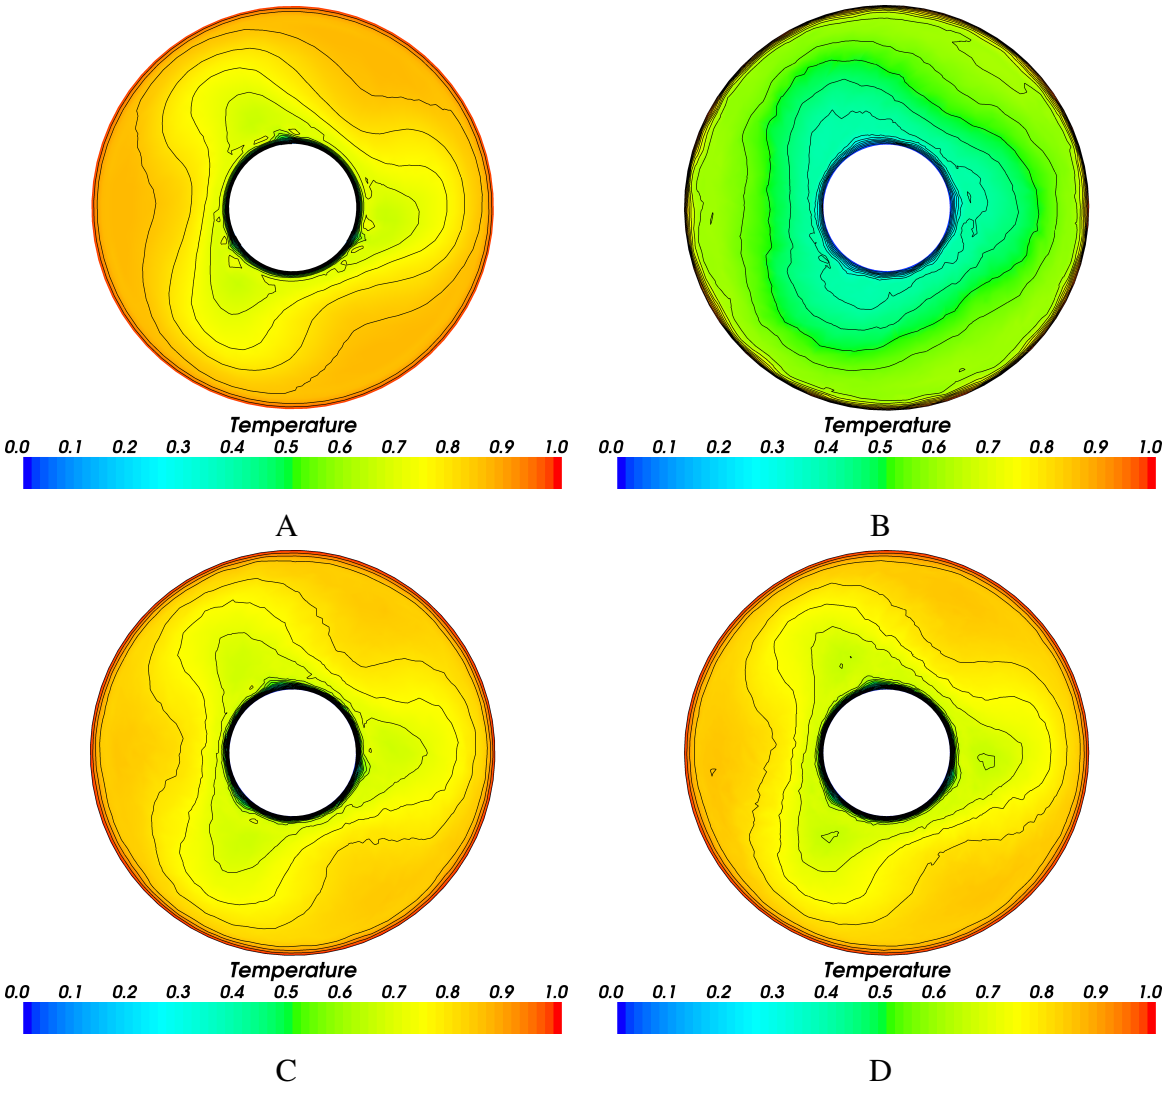

Figure 2.9.: Mid-height temperature field for numerical simulations of the thermally driven rotating annulus with parameters as given in table 2.1, after 4000 s of simulated laboratory time. A: MORALS simulation with $32 \times 32 \times 128$ nodes in the radial, vertical and azimuthal dimensions respectively. B: Unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with no balanced pressure decomposition. C: Unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with a P_1 balanced pressure decomposition. D: Unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with a P_2 balanced pressure decomposition. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$.

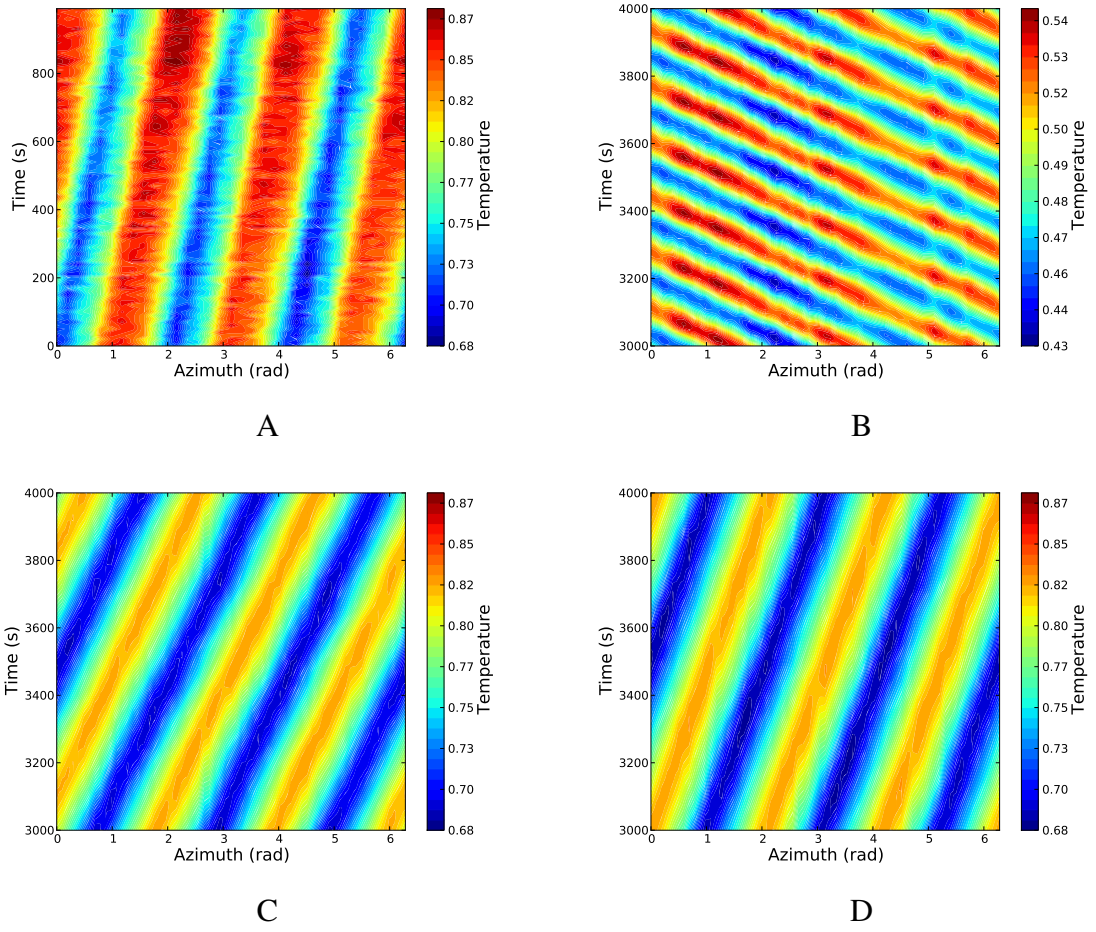

Figure 2.10.: Hovmöller plots derived from the mid-radius and mid-height temperature field in thermally driven rotating annulus experiments and Fluidity-ICOM simulations. A: Laboratory observations from an experiment with parameters as given in table 2.1, obtained from an in-situ thermocouple array. B: Corresponding measurements from an unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with no balanced pressure decomposition. C: Measurements from an unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with a P_1 balanced pressure decomposition. D: Measurements from an unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with a P_2 balanced pressure decomposition. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$. A, C and D are shown on a common colour scale.

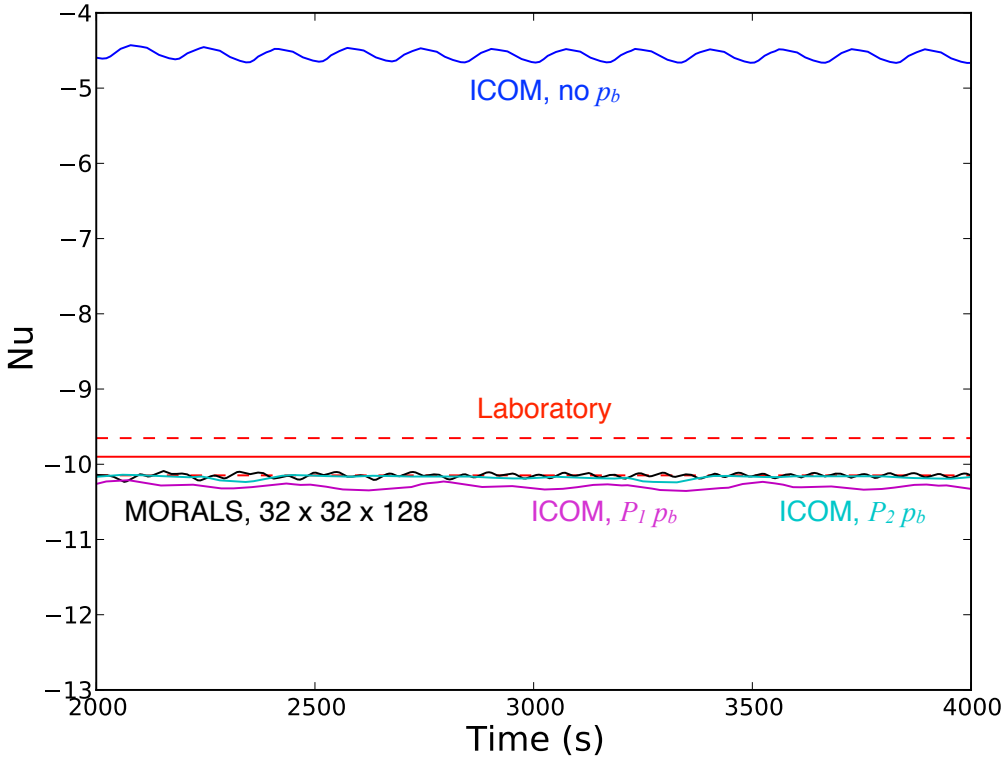

Figure 2.11.: Measurements of the system heat transport from numerical simulations of the thermally driven rotating annulus with parameters as given in table 2.1. Laboratory measurements, in red, are displayed with the accompanying error [Read, 2003]. Three unstructured mesh Fluidity-ICOM simulations are displayed: a simulation using no balanced pressure decomposition, using a P_1 balanced pressure decomposition, and using a P_2 balanced pressure decomposition.

baroclinic wave is very weak as compared against the laboratory observations. When a P_1 or P_2 balanced pressure decomposition is applied a clear mode three baroclinic wave is observed, of qualitative similarity to the laboratory observations. Measurements of the system heat transport are shown in figure 2.11. The Fluidity-ICOM simulation with no balanced pressure decomposition has a system heat transport that is significantly below observed values. By contrast, the Fluidity-ICOM simulations with a P_1 and P_2 balanced pressure decomposition show a much closer agreement with experiment.

For the P_1 and P_2 balanced pressure cases, the phase speed of the baroclinic wave differs significantly from that observed in experiment. The difference in phase speed is attributed to the absence of effects that act to break the symmetry of the system. Such effects include the use of temperature independent viscosity [Hide and Mason, 1970] and diffusivity, and the removal of the centrifugal force. In addition, the presence of an in-situ thermocouple array may lead to further discrepancies [James et al., 1981, Hignett et al., 1985]. Note that the use of a higher kinematic viscosity than present in experiment does not lead to a significant improvement in the simulated drift rate (see section 2.3.4, and particularly figure 2.19).

The balanced pressure decomposition can be applied to hydrostatic and geostrophic balance separately. Replacing $\boldsymbol{B}^{\delta}$ with $\boldsymbol{b}^{\delta}$ in equation (2.40) gives:

$$\int_{\Omega} \nabla \chi_i \cdot \nabla p_h^{\delta} = \int_{\Omega} \nabla \chi_i \cdot \boldsymbol{b}^{\delta}, \tag{2.47}$$

where p_h^{δ} denotes a "hydrostatic pressure". Similarly, replacing $\boldsymbol{B}^{\delta}$ with $-\boldsymbol{f}^{\delta}$ in equation (2.40) gives:

$$\int_{\Omega} \nabla \chi_i \cdot \nabla p_g^{\delta} = -\int_{\Omega} \nabla \chi_i \cdot \boldsymbol{f}^{\delta}, \tag{2.48}$$

where p_g^{δ} denotes a "geostrophic pressure". Further Fluidity-ICOM simulations were conducted using only a hydrostatic pressure decomposition and only a geostrophic pressure decomposition. The resulting spatial heat flux through the tank inner wall is shown in figure 2.12, and the integrated heat flow in figure 2.13. As previously noted, when no balanced pressure decomposition is applied the system heat transport is significantly below the observed value. When a hydrostatic pressure decomposition is used the system heat transport increases significantly. At the inner wall the additional heat transport is seen to manifest itself as a predominantly axisymmetric flux concentrated towards the tank lid, and is identified as an increased Ekman layer flux. When a geostrophic pressure decomposition is used, a somewhat more modest increase in heat transport is observed. In this case, at the inner wall a distinct mode three pattern in the heat flux is observed, and hence the additional heat transport is identified as an increased advective transport via the establishment of a fully developed baroclinic wave. When a combined hydrostatic and geostrophic pressure decomposition is applied, both an Ekman flux and a baroclinic eddy flux is observed, and the integrated heat transport approaches the laboratory value. The simulation using only a hydrostatic pressure decomposition actually has an increased heat transport relative to that using both a geostrophic and hydrostatic pressure decomposition. Conclusions regarding this overshoot are difficult to draw due to its non-physical origin. Hence, in this case, it is concluded that an accurate representation of hydrostatic balance representation is required for a physical representation of the Ekman layer dynamics, while an accurate representation of geostrophic balance is required for the establishment of baroclinic waves. The balanced pressure decomposition method is able to increase the accuracy of hydrostatic and geostrophic balance representation, and is required for the stabilised P_1P_1 Fluidity-ICOM simulations to capture these features.

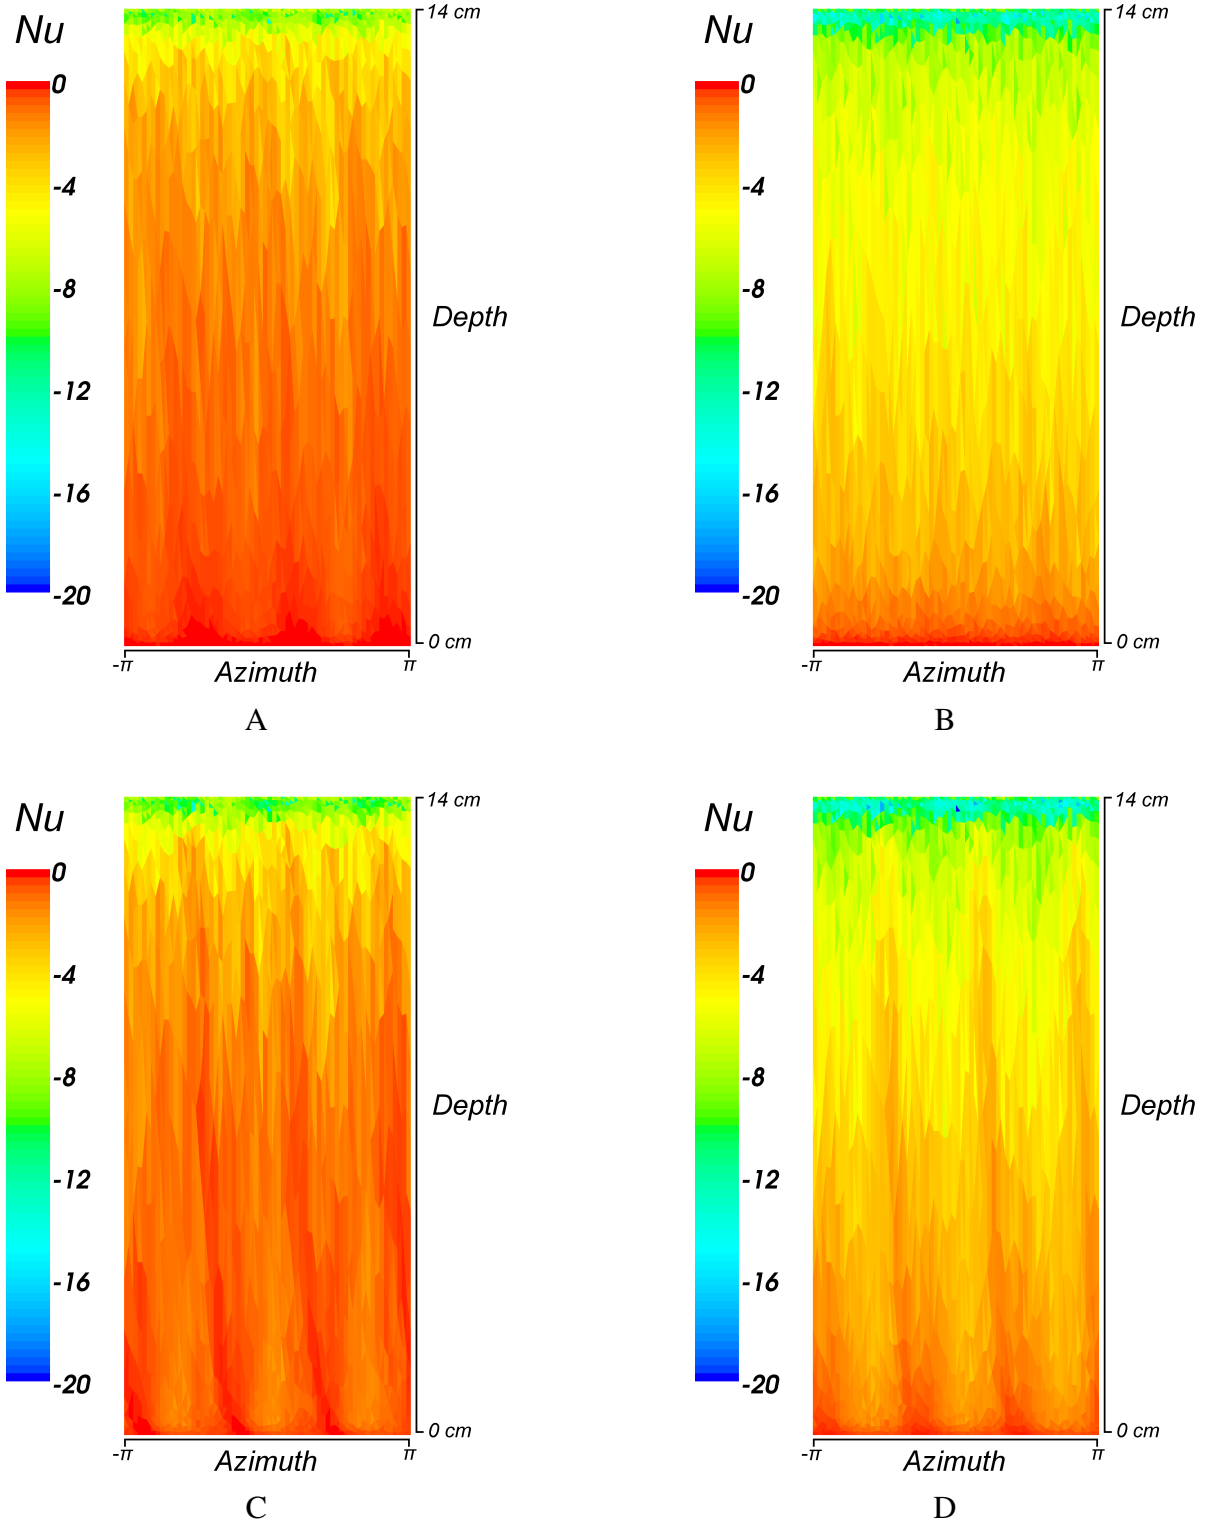

Figure 2.12.: Heat flux through the tank inner wall, expressed as a non-dimensional Nusselt number $Nu = [(\ln b/a) / (2\pi d \Delta T)]\,[2\pi a d\,(\nabla T \cdot \hat{\boldsymbol{n}})]$, for Fluidity-ICOM simulations of the thermally driven rotating annulus with parameters as given in table 2.1, after 4000 s of simulated laboratory time. A: Unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with no balanced pressure decomposition. B: Unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with a P_2 hydrostatic pressure decomposition. C: Unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with a P_2 geostrophic pressure decomposition. D: Unstructured mesh stabilised P_1P_1 Fluidity-ICOM simulation with a P_2 balanced pressure decomposition, combining the hydrostatic and geostrophic pressure decompositions.

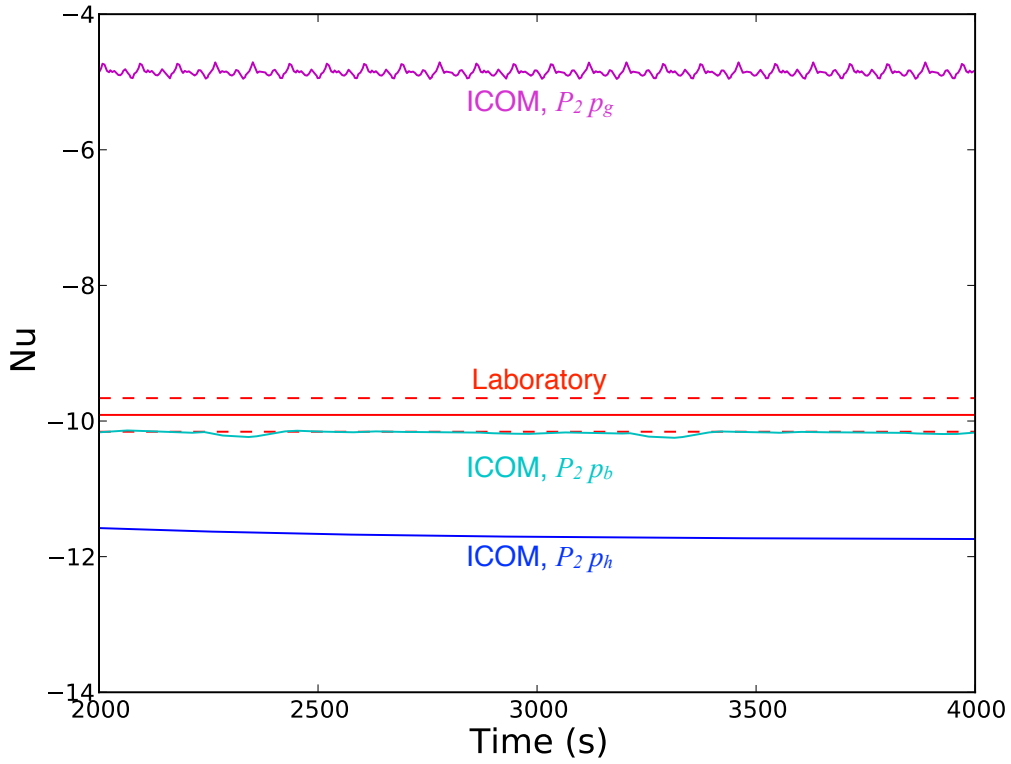

Figure 2.13.: Measurements of the system heat transport from numerical simulations of the thermally driven rotating annulus with parameters as given in table 2.1. Laboratory measurements, in red, are displayed with the accompanying error [Read, 2003]. Three unstructured mesh Fluidity-ICOM simulations are displayed: a simulation using only a hydrostatic pressure (p_h) decomposition, only a geostrophic pressure decomposition (p_g), and both a hydrostatic and geostrophic pressure decomposition (p_b).

For this system, the absence of any balanced pressure decomposition was observed to lead to a poor representation of the thermally driven rotating annulus dynamics. The introduction of a balanced pressure decomposition, using P_1 or P_2 elements, was able to restore the primary features of the flow and lead to a more favourable comparison with experiment.

For comparison, figure 2.14 shows the heat transport of Fluidity-ICOM simulations conducted on a fully structured mesh (with stretching as in MORALS) with $24 \times 24 \times 64$ nodes in the radial, vertical and azimuthal dimensions respectively, with and without balanced pressure decomposition. As in the unstructured case, the system heat transport is poorly represented without any balanced pressure decomposition, while a P_1 balanced pressure restores the system heat transport towards the experimental value. Hence the pressure gradient errors associated with a stabilised P_1P_1 discretisation of the Navier-Stokes equations, which act to decrease the accuracy of hydrostatic and geostrophic balance, are an inherent limitation of the discretisation, and not due to the use of unstructured meshes. However, once a balanced pressure decomposition is applied, the method is robust on both structured and unstructured meshes.

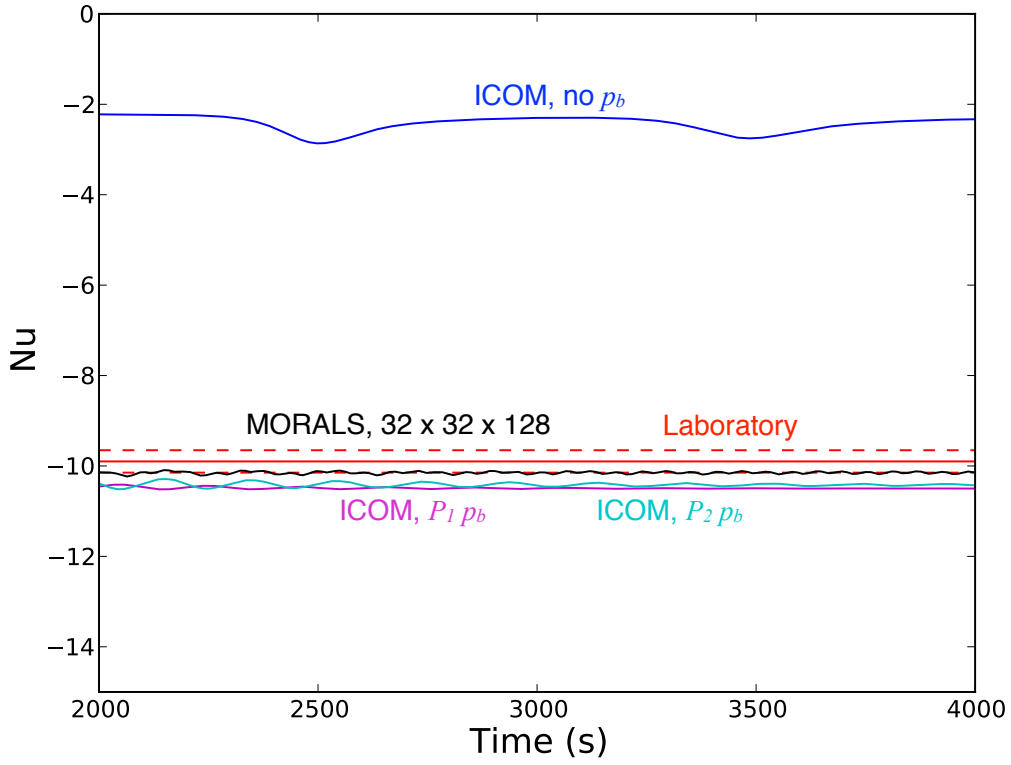

Figure 2.14.: Measurements of the system heat transport from numerical simulations of the thermally driven rotating annulus with parameters as given in table 2.1. Laboratory measurements, in red, are displayed with the accompanying error [Read, 2003]. Three structured mesh Fluidity-ICOM simulations are displayed: a simulation using no balanced pressure decomposition, using a P_1 balanced pressure decomposition, and using a P_2 balanced pressure decomposition.

2.3.2. $P_{1DG}P_2$ finite element pair

As previously noted in section 2.1.2, the $P_{1DG}P_2$ element pair is known to have excellent geostrophic balance representation. This was demonstrated in section 2.2.2 where it was shown that, for this element pair, the use of either a P_1 or P_2 balanced pressure decomposition has no influence on the accuracy of geostrophic balance representation. A P_3 balanced pressure decomposition led to only a very modest improvement.

Here, we test the application of the $P_{1DG}P_2$ finite element pair to simulations of the thermally driven rotating annulus. The momentum equation discretisation is as described in section 2.1.2 and Ham and Cotter [2011]. The prognostic temperature equation was discretised using a discontinuous Galerkin formulation. For both the momentum and temperature discretisations, the interior element fluxes for advection were treated using the upwind flux, and diffusivity terms treated using the Bassi-Rebay scheme [Bassi and Rebay, 1997]. The no-slip velocity and thermal forcing temperature boundary conditions were applied in the strong sense. The resulting linear systems were solved with preconditioned conjugate gradients using the PETSc library [Balay et al., 1997, 2008, 2009] and hypre library [Falgout and Yang, 2009], with the BoomerAMG multigrid solver applied

Mesh resolution	Mesh nodes		
	P_1 (coordinate)	P_{1DG} (velocity, temperature)	P_2 (pressure)
$12 \times 12 \times 32$	4608	92928	33856
$15 \times 15 \times 40$	9000	188160	67280
$18 \times 18 \times 50$	16200	346800	122500
$24 \times 24 \times 64$	36864	–	–

Table 2.2.: Resolutions of the meshes used for $P_{1DG}P_2$ simulations of the thermally driven rotating annulus, including all strong boundary condition nodes. The mesh resolutions are given in terms of the number of nodes in the radial, vertical and azimuthal dimensions respectively.

for the pressure solve [Henson and Yang, 2002]. The model was integrated with the timestep size adapted after every timestep to target a maximum Courant number of unity.

The simulations were conducted on a series of structured MORALS meshes of varying resolutions, using the Oxford Supercomputing Centre "redqueen" machine. The mesh resolutions used for these simulations, together with the degrees of freedom for the simulation prognostic fields, are shown in table 2.2. Note that the number of velocity degrees of freedom exceeds the number of pressure degrees of freedom, as is necessary for an LBB stable element pair.

Figure 2.15 shows an isopycnal, figure 2.16 shows the mid-height temperature field, and figure 2.17 shows a Hovmöller plot derived from the mid-height and mid-radius temperature field, for the $P_{1DG}P_2$ simulation with $15 \times 15 \times 40$ coordinate mesh nodes in the radial, vertical, and azimuthal dimensions respectively. A fully developed mode 3 baroclinic wave is observed, as expected for these system parameters. Under- and over-shoot errors are also observed, leading to a temperature range that exceeds the bounds imposed by the thermal forcing. These overshoots are primarily located in the inner and outer wall thermal boundary layer regions – i.e. in the regions of strongest temperature gradients. This numerical artefact could be removed via an appropriate slope limiter (not available in Fluidity-ICOM at the time these simulations were conducted), but here we note that for the lowest resolution simulation these overshoots occurred in only 0.023% of the domain, and for the highest resolution simulation occurred in only 0.007% of the domain.

Figure 2.18 shows the heat transport for these simulations, with laboratory data and the results of a stabilised P_1P_1 simulation for comparison. The stabilised P_1P_1 simulation was conducted on a structured MORALS mesh with resolution $24 \times 24 \times 64$ nodes in the radial, vertical, and azimuthal dimensions respectively, and using a P_2 balanced pressure decomposition. The $P_{1DG}P_2$ simulations all have a physically credible heat transport, with the heat transport approaching the laboratory value as the mesh resolution is increased. Hence the $P_{1DG}P_2$ finite element pair, with a P_{1DG}

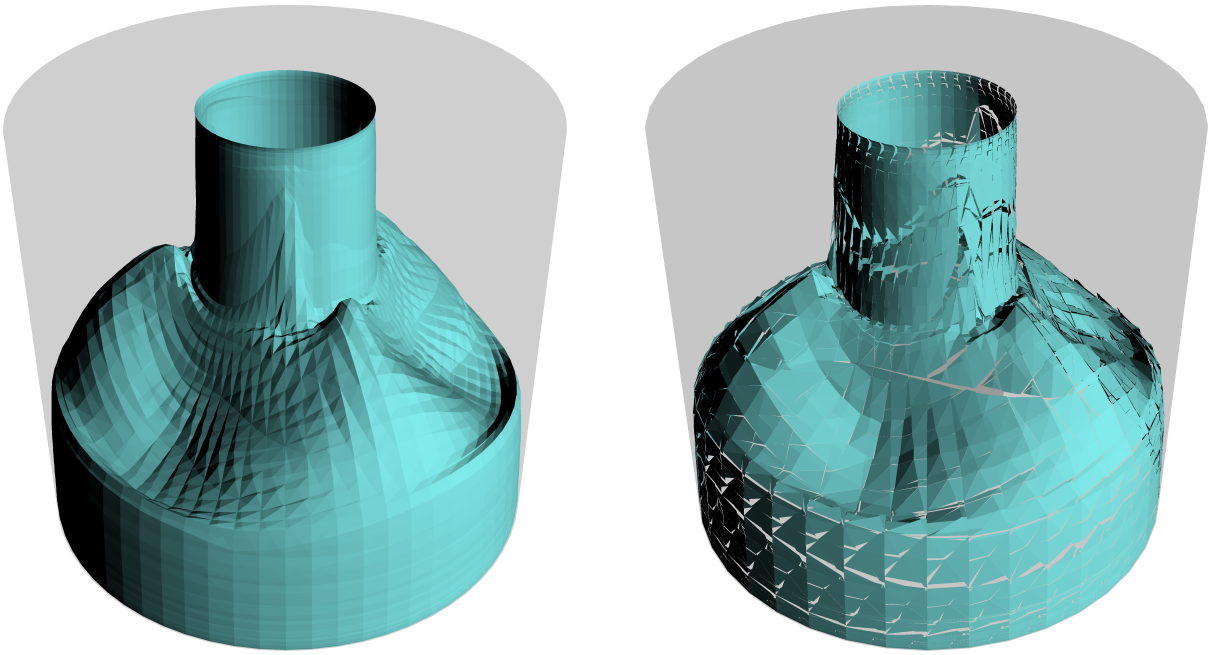

Figure 2.15.: $T = T_A + \frac{3}{4}\Delta T$ isopycnal for simulations of the thermally driven rotating annulus. Left: Stabilised P_1P_1 simulation with a P_2 balanced pressure decomposition, with $24\times24\times64$ nodes in the radial, vertical, and azimuthal dimensions respectively. Right: $P_{1\mathrm{DG}}P_2$ simulation, with $15\times15\times40$ coordinate mesh nodes in the radial, vertical, and azimuthal dimensions respectively. With this discretisation the temperature field is discontinuous, and hence the isopycnal is discontinuous.

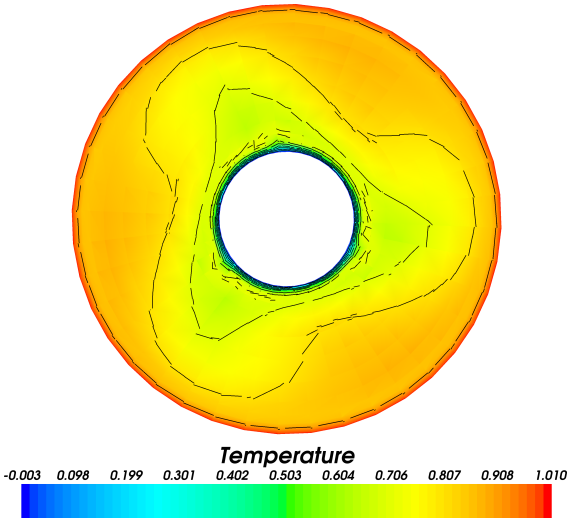

Figure 2.16.: Mid-height temperature field for a $P_{1\mathrm{DG}}P_2$ simulation of the thermally driven rotating annulus, with $15 \times 15 \times 40$ coordinate mesh nodes in the radial, vertical, and azimuthal dimensions respectively. Parameters are as given in table 2.1. The contour lines are discontinuous as a result of the broken $P_{1\mathrm{DG}}$ space used for temperature. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$.

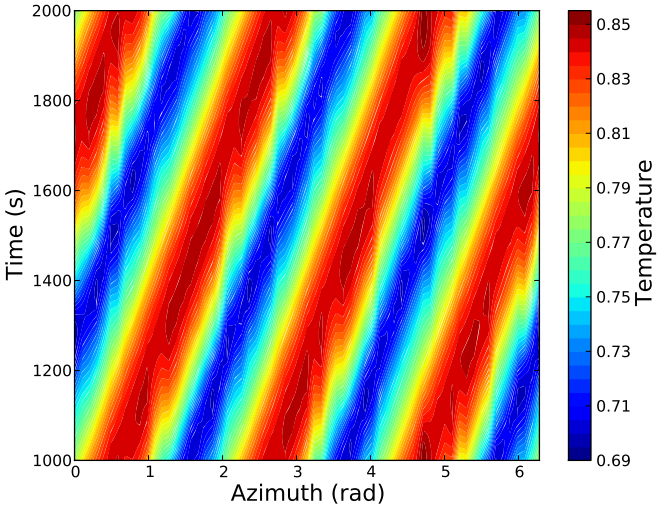

Figure 2.17.: Hovmöller plot derived from the mid-radius and mid-height temperature field for a $P_{1\mathrm{DG}}P_2$ simulation of the thermally driven rotating annulus, with $15 \times 15 \times 40$ coordinate mesh nodes in the radial, vertical, and azimuthal dimensions respectively. Parameters are as given in table 2.1. Note that this diagram is produced from point-wise interpolation of the discontinuous temperature field, which is not in general well-defined at element interfaces for the discontinuous temperature field. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$.

temperature discretisation, when applied to simulations of the thermally driven rotating annulus, is able to achieve a representation of geostrophic and hydrostatic balance of sufficient accuracy to yield a physically realistic system heat transport. This is achieved without the application of any balanced pressure decomposition method.

The stabilised P_1P_1 (with a P_2 balanced pressure decomposition) simulation at resolution $24 \times 24 \times 64$ has a heat transport that lies between that of the $P_{1\mathrm{DG}}P_2$ simulations of resolution $12 \times 12 \times 32$ and $15 \times 15 \times 40$ coordinate mesh nodes in the radial, vertical, and azimuthal dimensions respectively. Note, however, that these $P_{1\mathrm{DG}}P_2$ simulations each have considerably more velocity and temperature degrees of freedom than the stabilised P_1P_1 simulation, as seen in table 2.2 (comparing the number of P_1 mesh nodes with the number of $P_{1\mathrm{DG}}$ mesh nodes). Hence in this test, and for this diagnostic, the $P_{1\mathrm{DG}}P_2$ simulations require a greater number of velocity and temperature degrees of freedom than a corresponding stabilised P_1P_1 Fluidity-ICOM simulation (with a P_2 balanced pressure decomposition) of comparable accuracy. In terms of computational cost, the $15 \times 15 \times 40$ $P_{1\mathrm{DG}}P_2$ simulations were found to be $\sim 1 - 3$ times slower than a typical $24 \times 24 \times 64$ stabilised P_1P_1 simulation with P_2 balanced pressure, although this figure is expected to be sensitive to the details of the linear solver configurations. In particular, the $P_{1\mathrm{DG}}P_2$ simulations used the BoomerAMG multigrid method for the pressure solve, while the stabilised P_1P_1 simulations with P_2 balance pressure primarily used the multigrid method of Kramer et al. [2010].

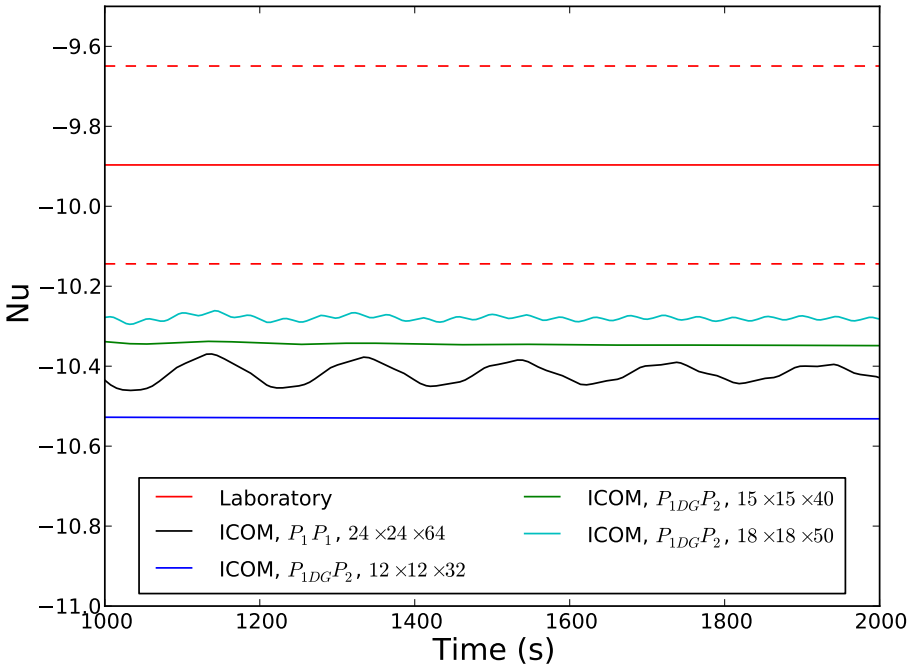

Figure 2.18.: Measurements of the system heat transport for $P_{1\mathrm{DG}}P_2$ numerical simulations of the thermally driven rotating annulus with parameters as given in table 2.1 and for mesh resolutions as given in table 2.16. Laboratory measurements, in red, are displayed with the accompanying error [Read, 2003]. A stabilised P_1P_1 Fluidity-ICOM simulation, with a P_2 balanced pressure decomposition and with resolution 24 × 24×64 nodes in the radial, vertical, and azimuthal dimensions respectively, is shown in black (previously shown in figure 2.14).

2.3.3. Computational cost

Benchmarks were conducted to measure the computational cost of the diagnostic solve for the balanced component of pressure, as given by equation (2.40). The system was initialised using the final $t = 4000$ s output from the $\omega = 1$ rad s^{-1} unstructured mesh annulus simulation described in section 2.3. The initial simulation was conducted using a P_2 balanced pressure decomposition. The model timestep was fixed, and the simulation integrated for 100 timesteps of 0.259 s, corresponding to a maximum initial Courant number of unity. Fluidity-ICOM was compiled using version 10.1 of the Intel Fortran compiler and version 4.4 of the GNU C and C++ compilers. The profiling was conducted in serial on a 64-bit Linux workstation with a 2.66GHz Intel 6700 dual core processor and 2GB of system RAM, and using version 3.8 of the libhoard fast memory allocator [Berger et al., 2000]. The MPI routine mpi_wtime[2] was used to measure the time taken to assemble the sparse matrices for the balanced pressure solve, the time spent in the multigrid solvers, and the total simulation runtime. Note that the total runtime includes all input/output and all run-time

[2]Note that, although the MPI routine mpi_wtime was used to measure the model wall time, these benchmarks were conducted in serial.

diagnostics, the cost of which increase with the introduction of a balanced pressure decomposition. The measurements are given in table 2.3.

Table 2.3.: Benchmark timings for 100 timesteps of a simulation of the thermally driven rotating annulus using Fluidity-ICOM with parameters as given in 2.1. All timings are given in seconds.

Statistic	No p_b	P_1 p_b	P_2 p_b
Total runtime	2667.9	2880.0	3612.4
Total p_b solution time	–	141.4	803.7
p_b linear system assembly	–	74.2	194.3
p_b linear solve	–	67.0	606.9

Using a P_1 balanced pressure decomposition increased the simulation runtime by 8.0%, with the sparse matrix assembly and multigrid solver accounting for 2.8% and 2.5% of the increase respectively. Using a P_2 balanced pressure decomposition increased the simulation runtime by 35.4%, with the sparse matrix assembly and multigrid solver accounting for 7.3% and 22.7% of the increase respectively. Clearly the P_2 balanced pressure solver has a significant computational cost, while the P_1 balanced pressure solver has a much more modest computational cost. Noting that a P_1 balanced pressure decomposition enables significantly more physical results to be obtained, with the more accurate P_2 balanced pressure decomposition yielding only marginal improvements, it is concluded that the P_1 balanced pressure decomposition, for this system, represents a reasonable compromise between the need for physical balance accuracy and the need for computational efficiency.

Note that, in the case of a linearly stratified flow, a P_2 balanced pressure decomposition yields an exact representation of hydrostatic balance. Hence there are cases where a P_2 balanced pressure decomposition may be expected to yield superior results to a P_1 balanced pressure decomposition, and hence there may be cases where the increased cost of the P_2 balanced pressure decomposition is justified.

2.3.4. Model viscosity

It was previously noted that the viscosity used for these simulations is too high, resulting from an inaccurate interpolation of the viscosity parameter from the data in Fowlis and Rossby [1964]. Linear interpolation of a Delauney triangulation of the data yielded a lower viscosity of 1.77×10^{-2} cm^2 s^{-1}and a marginally lower diffusivity of 1.29×10^{-2} cm^2 s^{-1}. A P_1P_1 Fluidity-ICOM simulation using these corrected parameters was conducted using a P_2 balanced pressure decomposi-

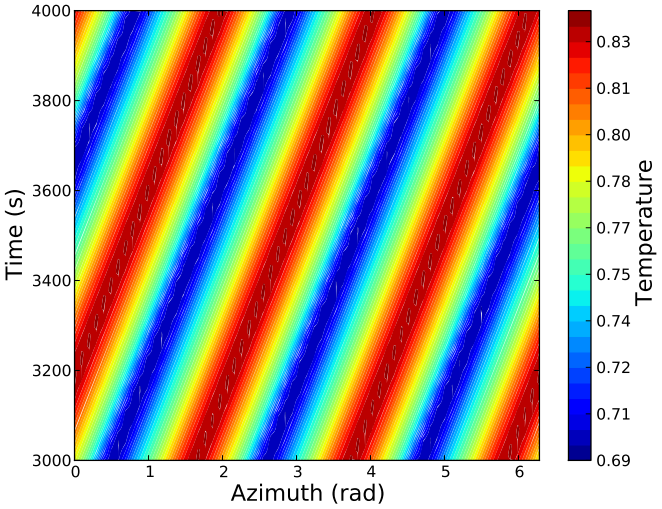

Figure 2.19.: Hovmöller plots derived from the mid-radius and mid-height temperature field in thermally driven rotating annulus Fluidity-ICOM simulation with a viscosity of 1.77×10^{-2} cm^2 s^{-1}and a diffusivity of 1.29×10^{-2} cm^2 s^{-1}. Other parameters are as given in table 2.1. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$.

tion, and simulated on a structured mesh derived from a MORALS mesh with resolution 24×24×64 nodes in the radial, vertical and azimuthal dimensions respectively. A Hovmöller plot derived from the mid-height and mid-radius temperature field is shown in figure 2.19, and the heat flux is shown in figure 2.20. The heat flux is significantly increased in this simulation, and agrees less well with the laboratory observations than the simulation conducted with the higher viscosity (and diffusivity). It is therefore possible that the very close agreement between experiment and the numerical simulations at the higher viscosity (and diffusivity) is somewhat fortuitous. Note, however, that similarly large displacements in phase space have previously been attributed to the presence of the in-situ thermocouple array in the experiments [Hignett et al., 1985].

2.4. Summary

We have discussed the source of pressure gradient errors in unstructured mesh finite element modelling. Since, for geophysical systems, the flows are typically in geostrophic balance and hydrostatic balance to leading order, the dynamically unimportant conservative components of the buoyancy and Coriolis accelerations are much larger than the dynamically important non-divergent residual components. The pressure associated with the buoyancy and Coriolis accelerations must therefore be solved for to a high degree of accuracy if pressure gradient errors, of significant magnitude when compared to these non-divergent residual components, are to be avoided.

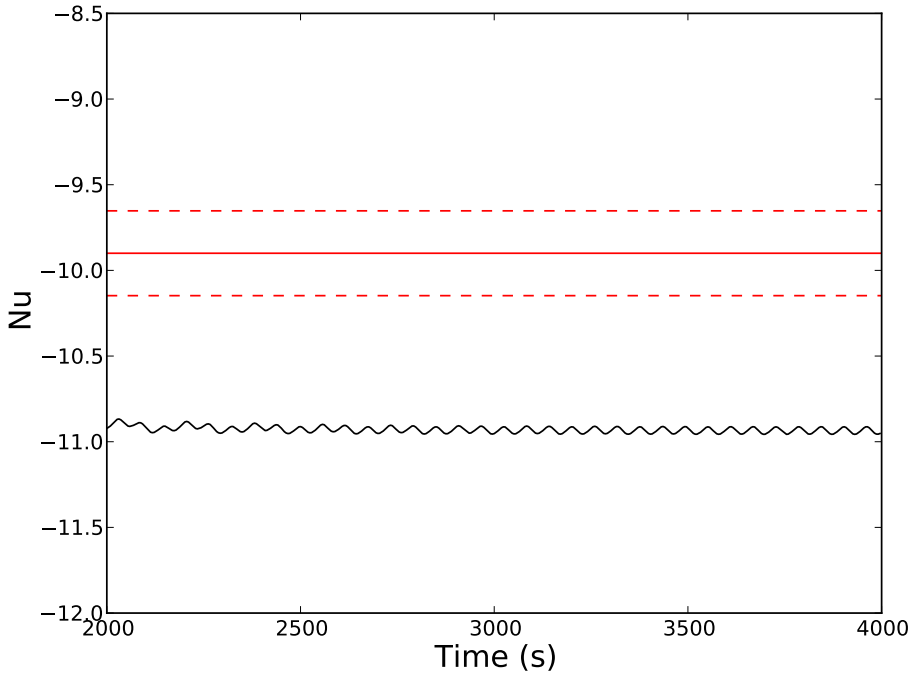

Figure 2.20.: Measurements of the system heat transport from numerical simulations of the thermally driven rotating annulus with a viscosity of 1.77×10^{-2} cm^2 s^{-1}and a diffusivity of 1.29×10^{-2} cm^2 s^{-1}. Other parameters are as given in table 2.1. Laboratory measurements, in red, are displayed with the accompanying error [Read, 2003], and the Fluidity-ICOM simulation (using a P_2 balanced pressure decomposition) is shown in black.

We have further discussed the fractional time-step pressure projection method, commonly used to solve the Navier-Stokes equations. In particular, it was identified that stabilisation methods for LBB unstable elements have the potential to reduce the order of accuracy of the pressure solve, and hence reduce the accuracy of physical balance representation. A balanced pressure decomposition method has been presented. Physically, this may be interpreted as a division of the dynamical pressure into a component associated with the hydrostatic and geostrophic pressure, and a residual. More precisely, the Helmholtz decomposition used to project the intermediate divergent velocity onto the solenoidal space is preconditioned via the addition of an extra pressure gradient term. The balanced pressure associated with this additional term can be solved for using methods that are free from LBB stability constraints, while the residual pressure filters conservative components from all remaining forcings and enforces incompressibility.

We have presented several numerical examples of this approach. Noting that, in the continuous space, simply connected barotropic flows on an f-plane are independent of rotation, we have quantified the increase in geostrophic balance accuracy obtained via the use of a balanced pressure decomposition. At higher resolutions, stabilised P_1P_1 simulations demonstrated slow convergence between rotating and non-rotating cases. Application of a P_1 balanced pressure decomposition ensured consistency and maintained a higher rate of convergence. $P_{1\mathrm{DG}}P_2$ simulations were unaf-

fected by the use of either a P_1 or P_2 balanced pressure decomposition, illustrating the excellent geostrophic balance properties of this element pair. We have further demonstrated the utility of the balanced pressure decomposition method for stabilised P_1P_1 unstructured mesh simulations of the thermally driven rotating annulus. For this system the application of a balanced pressure decomposition method was found to be crucial for the system heat transport to be well represented. In particular, hydrostatic balance accuracy was required for a physical representation of the Ekman layer heat transport, and geostrophic balance accuracy was required for a good representation of the interior baroclinic waves. Simulations of the thermally driven rotating annulus using the $P_{1\mathrm{DG}}P_2$ finite element pair required no balance pressure decomposition in order for a physically realistic system heat transport to be obtained. The simulations required a greater number of degrees of freedom, relative to the stabilised P_1P_1 simulations with a P_2 balanced pressure, in order to achieve a comparable simulated heat transport accuracy.

During the course of this research a discrepancy was identified between the annulus fluid parameters previously given in the literature and the laboratory fluid parameters, arising from an inaccurate interpolation of the parameters. In particular, this led to a significant error in the quoted viscosity – a more accurate interpolation method yielded a kinematic viscosity of 1.77×10^{-2} cm^2 s^{-1}, versus the quoted 2.1×10^{-2} cm^2 s^{-1}. This was found to have a significant influence on the system heat transport.

In this chapter we have focused on the issue of physical balance accuracy associated with pressure gradient errors inherent in the discretisation of the incompressible Navier-Stokes equations. Pressure gradient errors of this kind are inherent in the numerical method, and are not due to the choice of computational mesh. Hence, while these methods are robust for meshes that are unstructured in all three dimensions, they are also applicable when using fully structured meshes, as demonstrated in the simulations of the thermally driven rotating annulus in section 2.3.

Hydrostatic and geostrophic pressure gradient errors are a particularly challenging issue for vertically unstructured simulations in large aspect ratio domains, such as for the global atmosphere or ocean. While the Coriolis acceleration (applying the traditional approximation) is confined to the horizontal directions, the divergence free residual of the Coriolis acceleration has a component in the vertical direction. Since, for the global atmosphere or ocean, the velocity scales in the horizontal and vertical directions are separated by several orders of magnitude, there is potential for meshes that are unstructured in the vertical to be highly sensitive to pressure gradient errors in the vertical direction, originating from the pressure associated with the horizontal Coriolis acceleration. This is separate from the issue with non-hydrostatic pressure solvers on unstructured meshes

addressed in Kramer et al. [2010], where the horizontal and vertical velocity scales are shown to lead to an ill conditioned discrete Laplacian matrix, requiring careful treatment for efficient solution of the linear system.

In the following chapter, the application of dynamic mesh adaptivity to numerical simulations of geophysical systems will be considered. In particular, the potential for imbalance injection by the mesh adaption procedure itself will be addressed.

Chapter 3.

Balance errors in mesh adaptive ocean modelling

This chapter is based and expands upon Maddison et al. [2011a] and Farrell and Maddison [2011].

When applying dynamic mesh adaptivity, once an optimised mesh has been generated via some optimisation procedure, controlled via some metric derived from the simulation fields, an interpolant is required to transfer the fields from the old donor mesh to the new target mesh. This interpolant need not preserve geostrophic balance accurately: a state close to geostrophic balance on the old donor mesh need not remain close to balance after interpolation onto the new target mesh. In particular, the *hr*-adaptivity scheme implemented in Fluidity-ICOM leads to target meshes that are unrelated to the original donor mesh, other than that each is some covering simplex partitioning of the same original domain. This presents a particularly challenging interpolation problem.

In this previous chapter, the issue with accurate physical balance representation in a numerical discretisation of the Navier-Stokes equations was described. In this chapter, this is extended to consider accurate balance preservation when applying dynamic mesh adaptivity. In section 3.1 a method for the L_2 optimal (Galerkin) projection of solution fields between meshes is described. In section 3.2, in the context of a $P_{1\mathrm{DG}}P_2$ discretisation of the linearised shallow-water equations, an interpolant is presented that guarantees that, for the linearised equations on an f-plane, a steady and balanced state on a donor mesh remains steady and balanced after interpolation onto an arbitrary target mesh. Several numerical examples using this interpolant are provided in section 3.3. Finally,

in section 3.4, the interpolant is tested with other finite element discretisations, and the application to Navier-Stokes is considered in section 3.5.

The mesh optimisation procedure utilised by Fluidity-ICOM, to be described in more detail in section 4.2, performs local element operations and nodal perturbations. Hence while in general the donor and target meshes need not be related, Fluidity-ICOM, does not discard the donor mesh and generate an entirely new target mesh in each mesh adapt. Hence many of the interpolation tests described in this chapter (and also later, in section 4.2.3), represent worst-case stress tests of interpolants.

3.1. Galerkin projection

Consider the interpolation of a discrete field q^δ with nodal values $\tilde{q}$ and elemental basis functions ϕ_i, from a donor mesh A to a target mesh B, where the meshes A and B cover the same domain Ω. Let $(\ldots)^A$ and $(\ldots)^B$ denote "on donor" and "on target" respectively – for example, ϕ_i^A are the basis functions for $\left(q^\delta\right)^A$ on the donor mesh A. We consider three interpolants: collocation, Galerkin projection, and a first order conservative interpolant described in Grandy [1999] ("Grandy interpolation").

Collocation (sometimes referred to as *consistent interpolation*) is the evaluation of the donor field at the nodal locations of the target mesh. This evaluation is well-defined everywhere within the domain Ω provided q^δ is C^0 continuous. For piecewise linear fields collocation is second order accurate. In addition, for piecewise linear fields collocation results in a bounded projection – the range of $\left(q^\delta\right)^B$ lies within the range of $\left(q^\delta\right)^A$ – but, in general, the field bounds are eroded by the interpolation [Farrell, 2009]. Collocation is not well-defined for discontinuous fields – ad-hoc averaging is required at elemental boundaries for fields represented using $P_{N\mathrm{DG}}$ basis functions [Farrell, 2009].

Galerkin projection is the L_2 optimal projection of the donor field, minimising $\left\|\left(q^\delta\right)^B - \left(q^\delta\right)^A\right\|_{L_2}$. This takes the form [George and Borouchaki, 1998]:

$$\int_\Omega \phi_i^B \left(q^\delta\right)^B = \int_\Omega \phi_i^B \left(q^\delta\right)^A, \tag{3.1}$$

where ϕ_i^B and ϕ_i^A are the target and donor basis functions for $\left(q^\delta\right)^B$ and $\left(q^\delta\right)^A$ respectively. Re-written:

$$M^B \tilde{q}^A = M^{BA} \tilde{q}^A, \tag{3.2}$$

where:

$$\left(M^B\right)_{ij} = \int_\Omega \phi_i^B \phi_j^B, \qquad \left(M^{BA}\right)_{ij} = \int_\Omega \phi_i^B \phi_j^A. \tag{3.3}$$

M^B is the target mesh mass matrix, and M^{BA} is a "mixed" mass matrix, consisting of the inner products $\phi_i^B \phi_j^A$ between the target and donor basis functions. For piecewise linear fields Galerkin projection between meshes is second order accurate. With Galerkin projection there are no restrictions on the continuity of the interpolated field. Hence fields represented using discontinuous basis functions can be interpolated using this method. It follows from equation (3.2) that Galerkin projection between meshes conserves the integral of the field. However, it is not bounded – the range of $\left(q^\delta\right)^B$ can exceed the range of $\left(q^\delta\right)^A$ [Farrell, 2009].

It is challenging to assemble the mixed mass matrix as, *within* each donor or target element, the inner products are piecewise polynomials, which are challenging to integrate numerically. The direct numerical quadrature of these inner products is attempted in Farrell and Maddison [2011] using the adaptive quadrature package described in Cools and Haegemans [2003], but without success. One method for computing these inner products is via the construction of an intermediate "supermesh" – the mesh of intersection between the donor and target meshes, as demonstrated in figure 3.1. For each element in the supermesh the inner products $\phi_i^B \phi_j^A$ are (non-piecewise) polynomials, enabling the use of conventional polynomial quadrature rules and therefore enabling the assembly of the mixed mass matrix. Since, by construction, both the donor and target mesh basis functions are supported on the supermesh, both the donor field $\left(q^\delta\right)^A$ and the target field $\left(q^\delta\right)^B$ can be represented on the supermesh with zero interpolation error.

Galerkin projection via mesh intersection is described in George and Borouchaki [1998], but without practical applications. Assembly of the mixed mass matrix in two and three dimensions using approximate numerical quadrature is described in El Hraiech et al. [2005]. Assembly of the mixed mass matrix by mesh intersection is described in Heinstein and Laursen [2003], Jiao and

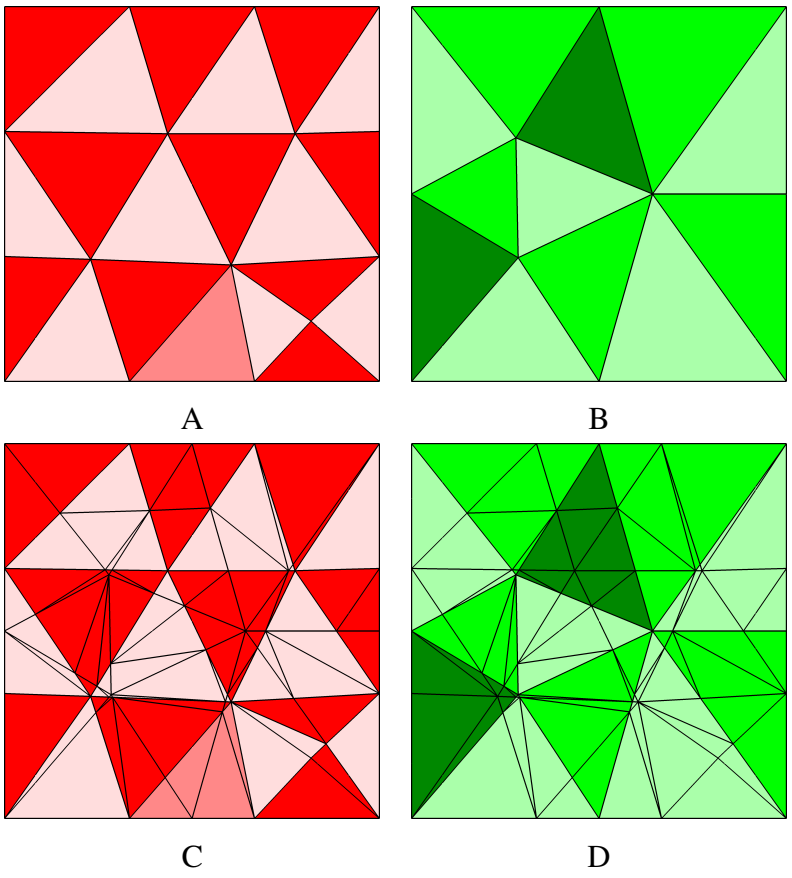

Figure 3.1.: Two triangular meshes A and B, and the mesh of their intersection C and D. In C the elements of mesh A are highlighted, and in D the elements of mesh B are highlighted.

Heath [2004a] and Jiao and Heath [2004b] for two dimensional meshes. In Farrell et al. [2009] interpolation using Galerkin projection is performed by exploiting an equivalence between mesh intersection and constrained Delauney triangulation. In addition, in Farrell et al. [2009], a bounded minimally diffusive interpolant is described, applying selective diffusion to bound the result of a Galerkin projection.

In Farrell and Maddison [2011] interpolation using Galerkin projection is applied in two and three dimensions via local assembly of equation (3.2). In this approach, each intersection between donor and target elements is constructed individually, and the corresponding contributions to the right-hand-size of equation (3.2) integrated separately, without assembling the full mixed mass matrix M^{BA} and without constructing the supermesh in its entirety. The problem of supermesh construction then reduces to the problem of the intersection between linear simplices. In Farrell and Maddison [2011] intersecting donor-target element pairs are identified using an advancing front algorithm, and each intersection is performed using a repeated half-space clipping algorithm as described in Eberly [2001] and demonstrated in figure 3.2. This approach is efficient, as only a small section of the supermesh is constructed at any one time. In addition, it can be extended to

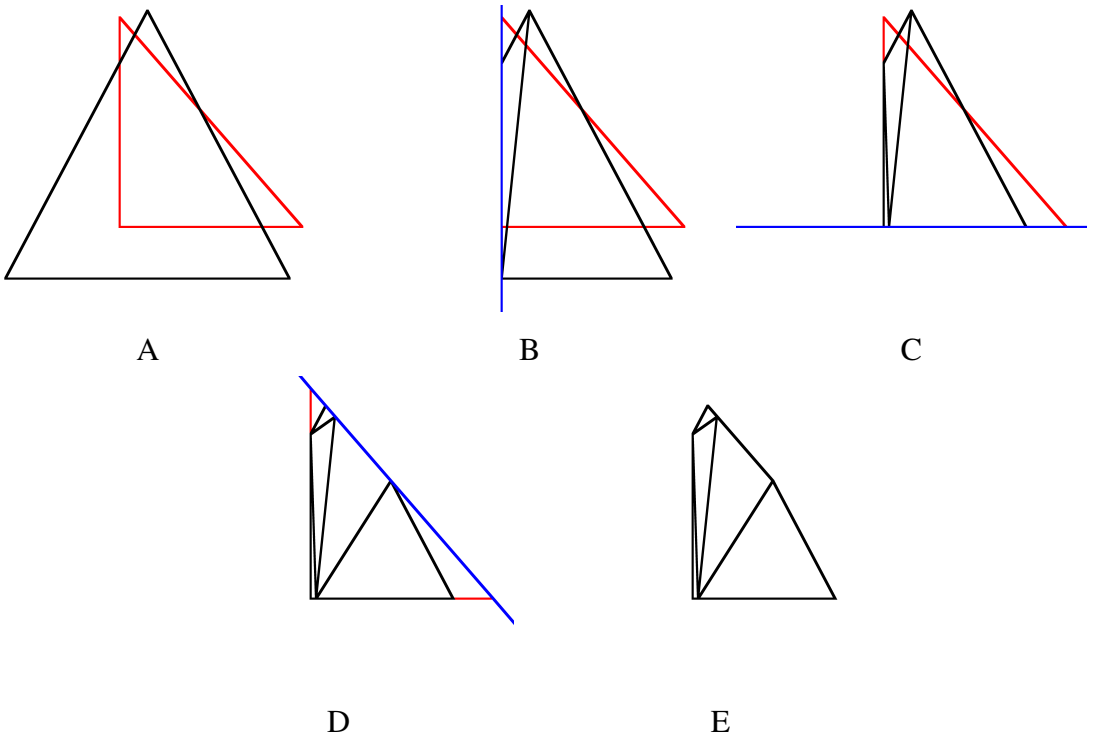

Figure 3.2.: An example of simplex intersection via repeated half-space clipping, following the algorithm of Eberly [2001] (from figure 5 of Farrell and Maddison [2011]). A: Two triangles to be intersected. B–D: The black triangle is successively clipped using each edge of the red triangle. In each clip the triangles from the previous stage are individually subdivided to form a new triangular mesh. Hence the intersection procedure is composed of a number intersections of triangles with half-spaces. E: The resulting mesh of the intersection. The contribution of this single intersection to the right-hand-side of equation (3.2) can be computed using this intersection mesh, and then the intersection mesh can be discarded.

three-dimensional tetrahedral meshes. For a full description of the Galerkin projection of fields via local mesh intersection, see Farrell [2009] and Farrell and Maddison [2011].

Grandy interpolation is a first order interpolant suggested in Grandy [1999]. It is the mesh-to-mesh Galerkin projection of the Galerkin projection of q^δ onto a P_0 function space on the donor mesh, or equivalently:

$$\int_\Omega \phi_i^B \left(q^\delta\right)^B \int_\Omega \psi_i^B \left(q^\delta\right)^A, \tag{3.4}$$

where ψ_i^B are the P_0 basis functions on the target mesh B. While this interpolant is conservative, the first order accuracy leads to a high level of numerical diffusion. In Grandy [1999] this interpolant is suggested only for (quote) *"single-use applications"*, and not for repeated use within a simulation. However, this is an early example of conservative interpolation of simulation fields, with application to discontinuous methods.

These interpolants were compared using a simple 1D test case. An initial structured mesh $0 \leq x \leq 1$ with 256 nodes was created, and a P_1 field q^δ initialised with a Gaussian profile, with $0 \leq q^\delta \leq 1$ about a mean of $x = 0.5$ and with a standard deviation of 8 times the initial mesh spacing. One hundred further 256 node meshes were generated with internal mesh nodes distributed randomly. q^δ was interpolated successively between each of these meshes using collocation, Grandy interpolation, Galerkin projection, Galerkin projection using a lumped target mesh mass matrix M^B (lumped according to equation (2.18)), and Galerkin projection combined with the diffusive bounding algorithm described in Farrell et al. [2009]. The maximum value, integral, and variance of q^δ are shown in figure 3.3.

All of Grandy interpolation, Galerkin projection, lumped mass Galerkin projection and Galerkin projection with the diffusive bounding method described in Farrell et al. [2009] preserve the field integral. When applying collocation, after 100 interpolations the maximum value of q^δ falls to 63.4% of its initial value. The integral of q^δ decreases to 99.1% of its initial value, and the variance increases by a factor of 2.4. In this example collocation is found to seriously erode the field bounds and to result in a high level of numerical diffusion. Note also that collocation is not purely dissipative – the integral of q^δ at times exceeds its initial value. When applying Grandy interpolation, after 100 interpolations the maximum value of q^δ falls to 54.9% of its initial value, and the variance increases by a factor of 3.1. The high level of numerical diffusion resulting from the first order accuracy of this method is clear. When interpolating using Galerkin projection with a

A

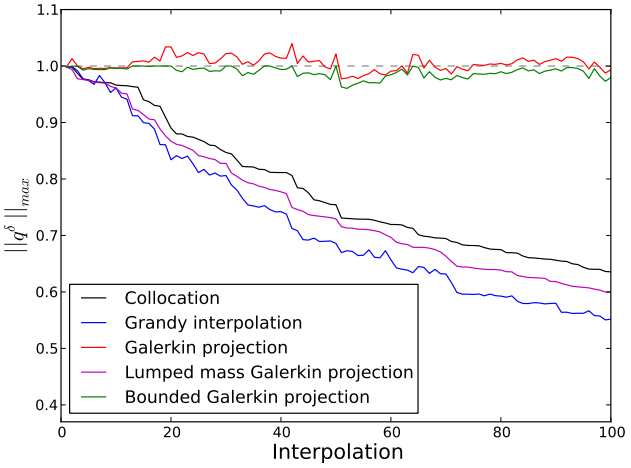

B

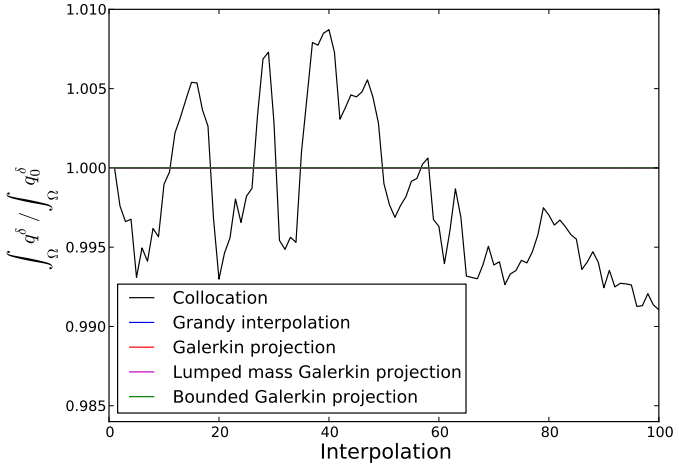

C

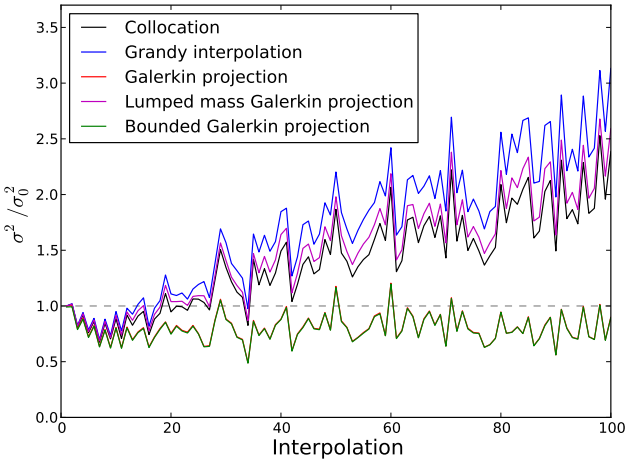

Figure 3.3.: Repeated interpolation of a 1D Gaussian onto random target meshes using collocation, Grandy interpolation, Galerkin projection, Galerkin projection with a lumped target mesh mass matrix, and Galerkin projection together with the diffusive bounding algorithm of Farrell et al. [2009]. A: The maximum value of the interpolated field. B: The integral of the interpolated field, normalised by the initial field integral. All tested interpolants other than collocation conserve the field integral. C: The variance of the interpolated Gaussian, $\sigma^2 = \sum_i q_i(x_i - \mu)^2 / \sum_i q_i$ where $\mu = \sum_i q_i x_i / \sum_i q_i$, normalised by the initial variance. Note that Galerkin projection and Galerkin projection combined with the diffusive bounding algorithm are almost coincident.

lumped target mesh mass matrix, after 100 interpolations the maximum value of q^δ falls to 59.6% of its initial value, and the variance increases by a factor of 2.6. This demonstrates the additional numerical diffusion associated with mass lumping [Zienkiewicz et al., 2005]. When interpolating using Galerkin projection, after 100 interpolations the maximum value of q^δ increases by 0.02% of its initial value. The maximum value of q^δ reaches a peak value of 4.0% above its initial value, demonstrating the lack of boundedness of this approach. The variance decreases to 0.9 times its initial value. Applying the diffusive bounding method of Farrell et al. [2009] has only a very small affect on the variance, and successfully prevents the interpolated field exceeding its initial bounds.

Hence interpolation using Galerkin projection has a significantly lower associated numerical diffusion than either collocation or Grandy interpolation. Galerkin projection with a lumped target mesh mass matrix is associated with an unacceptable level of numerical diffusion. Galerkin projection, combined with the diffusive bounding algorithm of Farrell et al. [2009], results in an interpolant that is bounded, conservative, and has a low level of numerical diffusion.

3.2. Geostrophic balance preserving interpolants

The preservation of geostrophic balance in mesh-to-mesh interpolation is now considered, in the context of linearised shallow-water ocean modelling. In section 3.2.1 the finite element discretisation of the linearised shallow-water equations using the $P_{1\mathrm{DG}}P_2$ finite element pair is outlined. A set of properties for geostrophic balance preserving interpolants is presented in section 3.2.2 for which, for the linearised system on an f-plane, an initially steady and geostrophically balanced state remains steady and balanced after interpolation onto an arbitrary target mesh. An interpolant satisfying these properties is given in section 3.2.3.

3.2.1. Discretised shallow-water equations

The linearised shallow-water equations with free slip boundary conditions are:

$$\frac{\partial \boldsymbol{u}}{\partial t} + f\hat{\boldsymbol{z}} \times \boldsymbol{u} + g\nabla\eta = 0, \tag{3.5a}$$

$$\frac{\partial \eta}{\partial t} + H\nabla \cdot \boldsymbol{u} = 0, \tag{3.5b}$$

$$\boldsymbol{u} \cdot \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega, \tag{3.5c}$$

where $\boldsymbol{u}$ is the (horizontal) velocity, H is the mean layer thickness, η is the deviation of the layer thickness from H, f is the Coriolis parameter, g is the gravitational acceleration and $\hat{\boldsymbol{n}}$ is a unit normal outward on the boundary $\partial\Omega$. From this one may define two non-dimensional parameters: the Rossby number $Ro = U/fD$ and the Froude number $Fr = \sqrt{U/gH}$, where U and D are characteristic flow speeds and spatial scales respectively.

Multiplying equations (3.5a) and (3.5b) by test functions $\boldsymbol{w}$ and ϕ respectively, integrating over the domain Ω, integrating by parts and applying the free slip boundary condition yields the weak form:

$$\int_\Omega \boldsymbol{w} \cdot \frac{\partial \boldsymbol{u}}{\partial t} + f\boldsymbol{w} \cdot (\hat{\boldsymbol{z}} \times \boldsymbol{u}) + g\boldsymbol{w} \cdot \nabla\eta = 0 \quad \forall \boldsymbol{w}, \tag{3.6a}$$

$$\int_\Omega \phi \frac{\partial \eta}{\partial t} - H \int_\Omega \nabla\phi \cdot \boldsymbol{u} = 0 \quad \forall \phi. \tag{3.6b}$$

Choosing some simplex covering partition of Ω (the mesh), restricting $\boldsymbol{w}$ and $\boldsymbol{u}$ to be piecewise linear discontinuous and restricting ϕ and η to be piecewise quadratic C^0 continuous completes the $P_{1\mathrm{DG}}P_2$ spatial discretisation:

$$\int_\Omega \boldsymbol{w}^\delta \frac{\mathrm{d}\boldsymbol{u}^\delta}{\mathrm{d}t} + f\boldsymbol{w}^\delta \cdot \left(\hat{\boldsymbol{z}} \times \boldsymbol{u}^\delta\right) + g\boldsymbol{w}^\delta \cdot \nabla\eta^\delta = 0 \quad \forall \boldsymbol{w}^\delta, \tag{3.7a}$$

$$\int_\Omega \phi^\delta \frac{\mathrm{d}\eta^\delta}{\mathrm{d}t} - H \int_\Omega \nabla\phi^\delta \cdot \boldsymbol{u}^\delta = 0 \quad \forall \phi^\delta, \tag{3.7b}$$

where v^δ denotes the finite element approximation for a field v. Introducing basis function expansions of $\boldsymbol{w}^\delta$, $\boldsymbol{u}^\delta$, ϕ^δ and η^δ, this can be re-expressed as:

$$\frac{\mathrm{d}}{\mathrm{d}t} M_1 \tilde{\boldsymbol{u}} + fL\tilde{\boldsymbol{u}} - gC\tilde{\eta} = 0, \tag{3.8a}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} M_2 \tilde{\eta} + H C^T \tilde{\boldsymbol{u}} = 0, \tag{3.8b}$$

where $\tilde{\boldsymbol{u}}$ and $\tilde{\eta}$ are the nodal values for velocity and layer thickness respectively and:

$$\begin{aligned} M_1 &= \mathrm{diag}\left(M_1', M_1'\right), & M_2 &= \mathrm{diag}\left(M_2', M_2'\right), \\ L &= \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} M_1', & C^T &= \left(C^x, C^y\right). \end{aligned} \tag{3.9}$$

M_1' and M_2' are the velocity space and layer thickness space mass matrices respectively, and C^T is the discrete divergence matrix:

$$\left(M_1'\right)_{ij} = \int_\Omega \psi_i \psi_j, \qquad \left(M_2'\right)_{ij} = \int_\Omega \xi_i \xi_j,$$

$$\left(C^q\right)_{ij} = -\int_\Omega \psi_i \frac{\partial \xi_j}{\partial q} \quad q \in \{x, y\}, \tag{3.10}$$

where the ψ_i and ξ_i are the $P_{1\mathrm{DG}}$ and P_2 elemental basis functions respectively. Here M_1' corresponds to M' in equation (2.13a). C^q corresponds to equation (2.13c), where the divergence operator has been integrated by parts as per equation (2.17b) and the boundary condition applied. Choosing some time discretisation allows equations (3.8a) and (3.8b) to be integrated on a computer.

3.2.2. Geostrophic balance preserving interpolants

Consider interpolation between a donor mesh A and a target mesh B. As before, let $(\ldots)^A$ and $(\ldots)^B$ denote "on donor" and "on target" respectively – for example $(C^T)^A$ and $(C^T)^B$ are the divergence matrices, as per (3.10), assembled on the donor and target meshes respectively. Consider an interpolation procedure as follows:

1. Perform a Helmholtz decomposition of the Coriolis acceleration $\tilde{\boldsymbol{F}}_*^A = f\hat{\boldsymbol{z}} \times \tilde{\boldsymbol{u}}^A$ on the donor mesh:

$$M_1^A \tilde{\boldsymbol{F}}^A = M_1^A \tilde{\boldsymbol{F}}_*^A + C^A \tilde{\Phi}^A, \tag{3.11}$$

   for some scalar potential $\tilde{\Phi}^A$ and discrete divergence free $\tilde{\boldsymbol{F}}^A$.

2. Interpolate $\tilde{\boldsymbol{F}}^A$, $\tilde{\Phi}^A$ and $\tilde{\eta}^A$ from the donor to the target to form $\tilde{\boldsymbol{F}}^B$, $\tilde{\Phi}^B$ and $\tilde{\eta}^B$, using interpolants with the following properties:

$$\left(M_1^A\right)^{-1} C^A \tilde{\Phi}^A \times \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega \Rightarrow \left(M_1^B\right)^{-1} C^B \tilde{\Phi}^B \times \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega, \tag{3.12a}$$

$$\tilde{\boldsymbol{F}}^A = 0 \Rightarrow \tilde{\boldsymbol{F}}^B = 0, \tag{3.12b}$$

$$C^A \tilde{\Phi}^A = g C^A \tilde{\eta}^A \Rightarrow C^B \tilde{\Phi}^B = g C^B \tilde{\eta}^B. \tag{3.12c}$$

3. Recompose $\tilde{\boldsymbol{F}}_*^B$ from $\tilde{\Phi}^B$ and $\tilde{\boldsymbol{F}}^B$:

$$M_1^B \tilde{\boldsymbol{F}}_*^B = M_1^B \tilde{\boldsymbol{F}}^B - C^B \tilde{\Phi}^B. \tag{3.13}$$

Note that here, and for the remainder of this chapter, $\tilde{\boldsymbol{F}}_*$ denotes the Coriolis acceleration, and $\tilde{\boldsymbol{F}}$ denotes the solenoidal component of the Coriolis acceleration.

Property (3.12b) states the somewhat trivial requirement that zero is preserved by the interpolant for $\tilde{\boldsymbol{F}}$. Property (3.12c) couples the velocity and layer thickness interpolation, and can be achieved if $\tilde{\Phi}$ and $\tilde{\eta}$ use the same interpolant and that interpolant is scale invariant. Property (3.12a) asserts that the interpolant for $\tilde{\Phi}$ preserves zero tangential gradients on the domain boundary. This is required to avoid the injection of mesh scale divergence in the velocity field at the boundary, corresponding to generation of mesh scale boundary Kelvin waves by the interpolation.

We now proceed to prove that this interpolant, when applied to a $P_{1\mathrm{DG}}P_2$ discretisation of the linearised shallow-water equations on an f-plane, guarantees that a steady and geostrophically balanced state on the donor mesh results in a state that is steady and balanced on the target mesh. By definition, for a geostrophically balanced state on the donor:

$$fL^A\tilde{\boldsymbol{u}}^A = gC^A\tilde{\eta}^A. \tag{3.14}$$

On an f-plane this can be re-written:

$$M_1^A\tilde{\boldsymbol{F}}_*^A = gC^A\tilde{\eta}^A. \tag{3.15}$$

Hence by equation (3.11) $\tilde{\boldsymbol{F}}^A = 0$ and $C^A\tilde{\Phi}^A = gC^A\tilde{\eta}^A$, and hence by properties (3.12b) and (3.12c) $\tilde{\boldsymbol{F}}^B = 0$ and $C^B\tilde{\Phi}^B = gC^B\tilde{\eta}^B$. Hence, by equation (3.13), on the target:

$$fL^B\tilde{\boldsymbol{u}}^B = gC^B\tilde{\eta}^B. \tag{3.16}$$

Also, since $\tilde{\boldsymbol{F}}_*$ is perpendicular to $\tilde{\boldsymbol{u}}$:

$$\begin{aligned}
\tilde{\boldsymbol{u}}^A \cdot \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega &\Rightarrow \tilde{\boldsymbol{F}}_*^A \times \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega \\
&\Rightarrow (M_1^A)^{-1}C^A\tilde{\Phi}^A \times \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega \qquad \text{by (3.14)} \\
&\Rightarrow (M_1^B)^{-1}C^B\tilde{\Phi}^B \times \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega \qquad \text{by (3.12a)} \\
&\Rightarrow \tilde{\boldsymbol{F}}_*^B \times \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega \\
&\Rightarrow \tilde{\boldsymbol{u}}^B \cdot \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega.
\end{aligned} \tag{3.17}$$

From Cotter et al. [2009a], if using the $P_{1\mathrm{DG}}P_2$ element pair on an f-plane:

$$\begin{aligned}
fL\tilde{\boldsymbol{u}} - gC\tilde{\eta} = 0 \text{ and } \tilde{\boldsymbol{u}} \cdot \hat{\boldsymbol{n}} = 0 \text{ on } \partial\Omega \\
\Leftrightarrow \frac{\partial\tilde{\boldsymbol{u}}}{\partial t} = 0 \text{ and } \frac{\partial\tilde{\eta}}{\partial t} = 0.
\end{aligned} \tag{3.18}$$

Hence by (3.16), (3.17) and (3.18) the solution on the target mesh is geostrophically balanced and exactly steady.

3.2.3. Implementation and boundary conditions

The Helmholtz decomposition of the Coriolis acceleration on the donor mesh is equivalent to the pressure projection method described in section 2.1.2. Multiplying equation (3.11) by $(C^T)^A(M_1^A)^{-1}$ and using $C^T\tilde{\boldsymbol{F}}^A = 0$ leads to the elliptic equation:

$$(C^T)^A(M_1^A)^{-1}C^A\tilde{\Phi}^A = -(C^T)^A\tilde{\boldsymbol{F}}_*^A. \tag{3.19}$$

As mentioned in section 2.1.2, here the consistent mass matrix can be used as the $P_{1\mathrm{DG}}$ mass matrix M_1^A is block diagonal, and hence the Laplacian matrix $(C^T)^A(M_1^A)^{-1}C^A$ is sparse. From this $\tilde{\Phi}^A$ and $\tilde{\boldsymbol{F}}^A$ can be determined. Following interpolation of the Helmholtz decomposition $\tilde{\boldsymbol{F}}_*^B$ can be diagnosed directly from $\tilde{\Phi}^B$ and $\tilde{\boldsymbol{F}}^B$ via equation (3.13). Therefore, the key step in forming a geostrophic balance preserving interpolant is to choose interpolants for $\tilde{\Phi}$, $\tilde{\eta}$ and $\tilde{\boldsymbol{F}}$ such that the properties (3.12a), (3.12b) and (3.12c) are satisfied.

One simple approach is to apply a Galerkin projection of $\tilde{\boldsymbol{F}}$ from the donor mesh to the target mesh, as described in section 3.1, and to interpolate $\tilde{\Phi}$ and $\tilde{\eta}$ using collocation. Since Galerkin projection and collocation are linear, properties (3.12b) and (3.12c) are satisfied. Collocation also preserves constant boundary values, and hence property (3.12a) is satisfied. However, as noted in section 3.1, collocation can (at least for piecewise linear fields) erode solutions bounds, and can have a high associated numerical diffusion.

A more accurate approach is to interpolate $\tilde{\Phi}$ and $\tilde{\eta}$ using Galerkin projection. This does not in general satisfy property (3.12a), although this issue can be resolved by using a further decomposition of $\tilde{\Phi}$, with an equivalent decomposition of $\tilde{\eta}$ in order to satisfy property (3.12c). Assuming Ω is simply connected, $\tilde{\Phi}^A$ can be decomposed into (C. Cotter, personal communication):

$$\tilde{\Phi}^A = \tilde{\Phi}_C^A + \tilde{\Phi}_R^A, \tag{3.20}$$

where $\tilde{\Phi}_C^A$ is equal to some constant c on the boundary, and $\tilde{\Phi}_R^A$ is some residual. $\tilde{\Phi}_C^A$ can be re-expressed:

$$\tilde{\Phi}_C^A = \tilde{\Phi}_0^A + c\tilde{\Phi}_1. \tag{3.21}$$

Here $\tilde{\Phi}_0$ and $\tilde{\Phi}_1$ satisfy:

$$N_0^A \tilde{\Phi}_0^A = N^A \tilde{\Phi}^A, \tag{3.22a}$$

$$N_1^A \tilde{\Phi}_1^A = 0, \tag{3.22b}$$

where N^A is the discrete Laplacian matrix $N^A = (C^T)^A (M_1^A)^{-1} C^A$, N_0^A is N^A with a Dirichlet boundary condition of zero on $\partial\Omega$ and N_1^A is N^A with a Dirichlet boundary condition of one on $\partial\Omega$. Minimising $\left\|\tilde{\Phi}_R^A\right\|_{L_2}$ subject to (3.20) then yields a unique value for c (C. Cotter, personal communication):

$$c = \frac{\left\langle \tilde{\Phi}_1^A, \tilde{\Phi}^A - \tilde{\Phi}_0^A \right\rangle_{L_2}}{\left\langle \tilde{\Phi}_1^A, \tilde{\Phi}_1^A \right\rangle_{L_2}}. \tag{3.23}$$

This choice of c has the property that if $\tilde{\Phi}^A$ is constant on $\partial\Omega$, then $\tilde{\Phi}_R^A = 0$. Applying a Galerkin projection of $\tilde{\Phi}_C^A$ and $\tilde{\Phi}_R^A$ from the donor mesh to the target mesh, with a Dirichlet boundary condition of c on $\partial\Omega$ for $\tilde{\Phi}_C^A$, therefore guarantees that the boundary property (3.12a) is satisfied.

Note that using the mass matrix in place of the Laplacian matrix, $N^A = M_2^A$, is not suitable here, as this results in non-smooth $\tilde{\Phi}_C^A$ and $\tilde{\Phi}_R^A$, with strong gradients close to the boundary which can generate significant noise in the donor-to-target Galerkin projection.

The full geostrophic balance preserving interpolation procedure is therefore:

1. Compute $\tilde{\boldsymbol{F}}_*^A$ from $\tilde{\boldsymbol{u}}^A$.
2. Solve equation (3.19) for $\tilde{\Phi}^A$ and compute $\tilde{\boldsymbol{F}}^A$ using equation (3.11).
3. Solve equations (3.22a) and (3.22b), with $N^A = (C^T)^A (M_1^A)^{-1} C^A$, for $\tilde{\Phi}_0^A$ and $\tilde{\Phi}_1^A$, and compute $\tilde{\Phi}_C^A$ and $\tilde{\Phi}_R^A$. Perform a similar decomposition for $\tilde{\eta}^A$ to form $\tilde{\eta}_C^A$ and $\tilde{\eta}_R^A$.

4. Apply a Galerkin projection from the donor mesh to the target mesh of $\tilde{\Phi}_C^A$, $\tilde{\Phi}_R^A$, $\tilde{\eta}_C^A$, $\tilde{\eta}_R^A$ and $\tilde{\boldsymbol{F}}^A$, with Dirichlet boundary conditions for $\tilde{\Phi}_C^A$ and $\tilde{\eta}_C^A$ as determined from equation (3.23), to form $\tilde{\Phi}_C^B$, $\tilde{\Phi}_R^B$, $\tilde{\eta}_C^B$, $\tilde{\eta}_R^B$ and $\tilde{\boldsymbol{F}}^B$.

5. Compute $\tilde{\Phi}^B$ from $\tilde{\Phi}_C^B$ and $\tilde{\Phi}_R^B$ using equation (3.20). Similarly compute $\tilde{\eta}^B$ from $\tilde{\eta}_C^B$ and $\tilde{\eta}_R^B$.

6. Compute $\tilde{\boldsymbol{F}}_*^B$ using equation (3.13).

7. Compute $\tilde{\boldsymbol{u}}^B$ from $\tilde{\boldsymbol{F}}_*^B$.

3.3. Geostrophic balance preserving interpolation examples

In this section several numerical examples of geostrophic balance preservation using the interpolation procedure presented above are given. In section 3.3.1 it is demonstrated that the geostrophic balance preserving interpolant ensures that a steady and balanced state remains steady and balanced after interpolation. In section 3.3.2 a state close to geostrophic balance is considered, and it is shown that the geostrophic balance preserving interpolant avoids imbalance injection. The interpolant is applied to a Kelvin wave in section 3.3.3, the accuracy of the interpolant in the L_2 norm is quantified in section 3.3.4, and an indication of the computation cost is given in section 3.3.5.

3.3.1. Preservation of balance

The $P_{1\mathrm{DG}}P_2$ linearised shallow-water equations (3.8a) and (3.8b) on an f-plane were discretised in time using Crank-Nicolson finite differencing [Crank and Nicolson, 1947], and the linear systems solved with preconditioned conjugate gradients using the PETSc library [Balay et al., 1997, 2008, 2009]. Further details of the discretisation are given in Cotter et al. [2009a].

In order to test for imbalance injection by mesh-to-mesh interpolation, two pseudo-isotropic circular meshes A and B were generated using Gmsh [Geuzaine and Remacle, 2009] and the ani2d mesh optimisation library [Vasilevskii and Lipnikov, 1999], with mesh B of one half the resolution of mesh A, as shown in figure 3.4. Meshes A and B have 2447 and 557 nodes respectively. Following the balance preservation test of Le Roux et al. [1998] and Cotter et al. [2009b], the system was initialised on mesh A with a Gaussian profile for layer thickness, shown in figure 3.5, and with a velocity field initialised so as to be in discrete geostrophic balance with that profile as per equation (3.14). The solution was then interpolated backwards and forwards between meshes A and B at

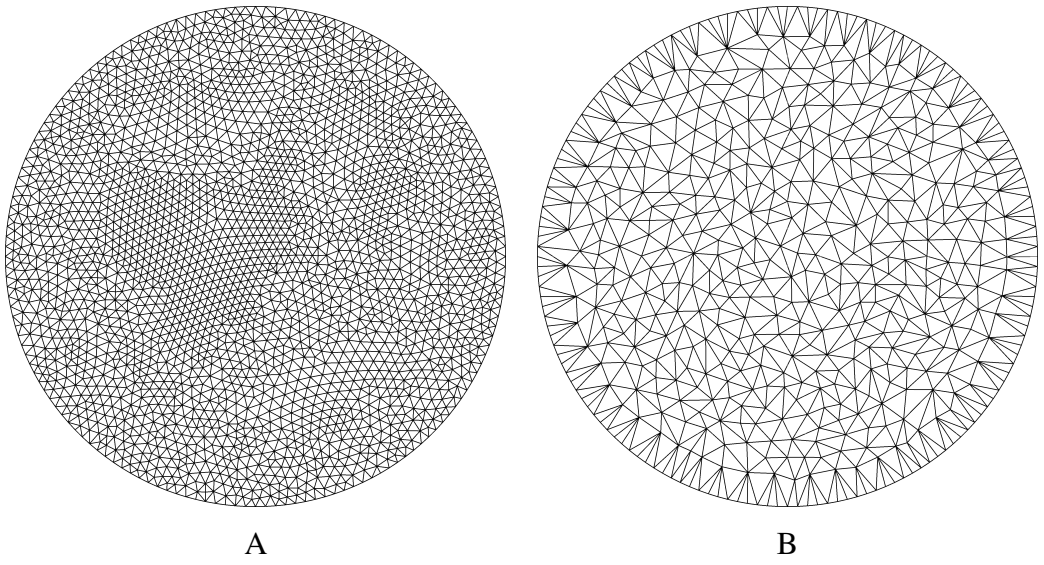

Figure 3.4.: Pseudo-isotropic meshes used to test for imbalance injection by mesh-to-mesh interpolation. Mesh B has one half the resolution of mesh A.

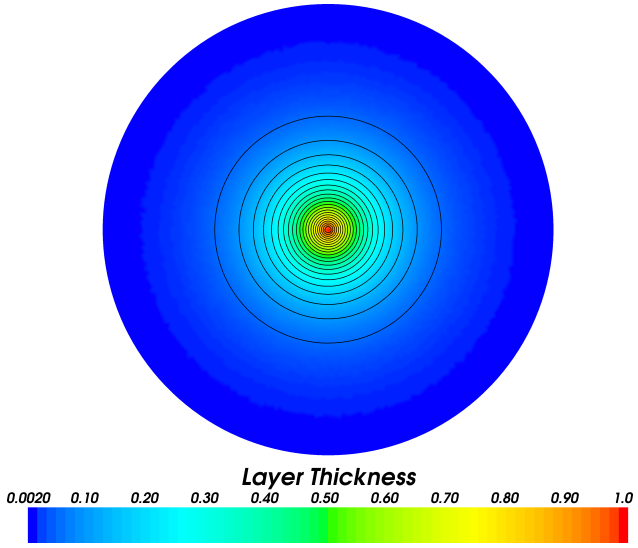

Figure 3.5.: Gaussian profile layer thickness used as an initial condition for the geostrophic balance preservation test.

ten timestep intervals. Since geostrophically balanced states with a constant stream function on the boundary are known to be exactly steady when using the $P_{1\mathrm{DG}}P_2$ element pair [Cotter et al., 2009a], any transience observed in the simulation is purely due to imbalance injection by the interpolation procedure.

The model was integrated for 210 timesteps of $6 \times 10^{-4}(D/U)$ at Rossby number 0.06 and Froude number 0.07 for a total of 20 interpolations between meshes A and B. Three interpolants were tested: Grandy interpolation for velocity and layer thickness, Galerkin projection for velocity and layer thickness, and the geostrophic balance preserving interpolant presented in the previous section. Collocation was not tested, as this is not well-defined at element boundaries for the discontinuous velocity field.

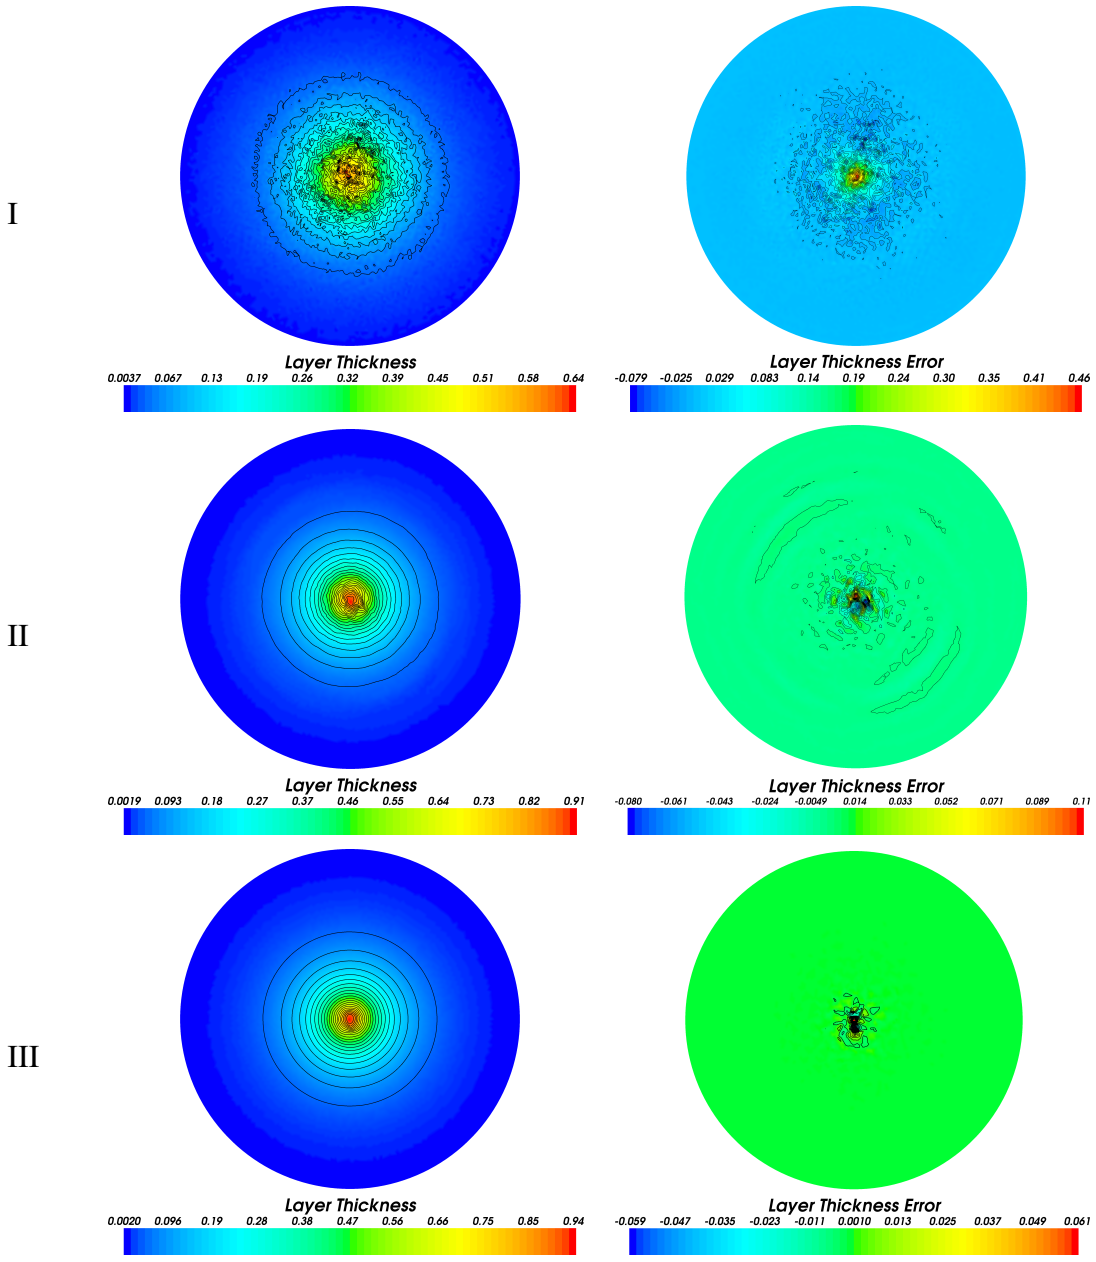

Figure 3.6.: The final solution of the geostrophic balance preservation test after 20 repeated interpolations backwards and forwards between the pseudo-isotropic meshes A and B in figure 3.4. Left: Final layer thickness. Right: Change in layer thickness from the initial condition in figure 3.5. I: Grandy interpolation. II: Galerkin projection. III: Helmholtz decomposed geostrophic balance preserving interpolation.

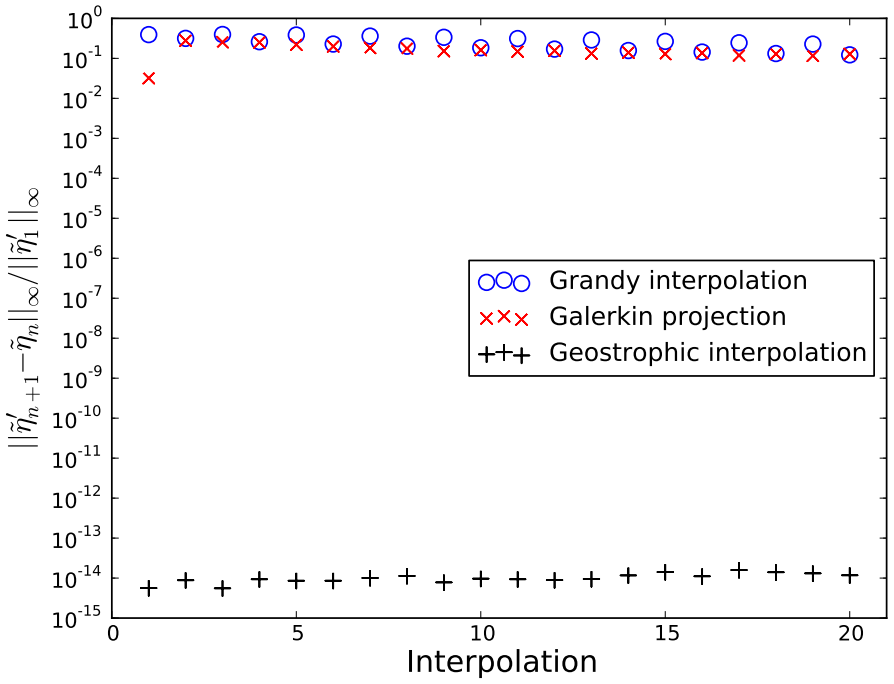

Figure 3.7.: The maximum change in layer thickness nodal values between interpolations backwards and forwards between the pseudo-isotropic meshes A and B in figure 3.4.

The final layer thickness and change in layer thickness from the initial condition are shown in figure 3.6. The maximum change in layer thickness nodal values between each interpolation, $\left\|\tilde{\eta}'_{n+1} - \tilde{\eta}_n\right\|_\infty$ where $\tilde{\eta}'_{n+1}$ are the layer thickness nodal values before interpolation $(n+1)$ and $\tilde{\eta}_n$ are the layer thickness nodal values after interpolation n, is shown in figure 3.7. Note that the maximum change in layer thickness nodal values, $\left\|\tilde{\eta}'_{n+1} - \tilde{\eta}_n\right\|_\infty$, may be less than the maximum change in layer thickness, $\left\|(\eta')^\delta_{n+1} - \eta^\delta_n\right\|_\infty$, due to the use of P_2 basis functions. Grandy interpolation is observed to inject imbalance everywhere after each interpolation, resulting in a severe degradation of the simulation fields after just a single interpolation. Interpolation using Galerkin projection is observed to inject imbalance towards the centre of the domain, near the layer thickness maximum. The resulting gravity waves propagate outwards polluting the global solution, and accumulate after every interpolation. The geostrophic balance preserving interpolant is exactly steady, to within machine precision, after every interpolation, with a change in layer thickness nodal values between interpolations of $\lesssim 10^{-13}$. The residual imbalance between interpolations is attributed to double precision round-off error.

After 20 interpolations the L_2 layer thickness error is 20% (of initial layer thickness L_2 norm) for Grandy interpolation, 2.7% for Galerkin projection and 2.0% for the geostrophic balance preserving interpolant. While Galerkin projection is optimal in the L_2 norm for each interpolation, the imbalance injection and resulting pollution of the solution by gravity waves leads to a reduced ac-

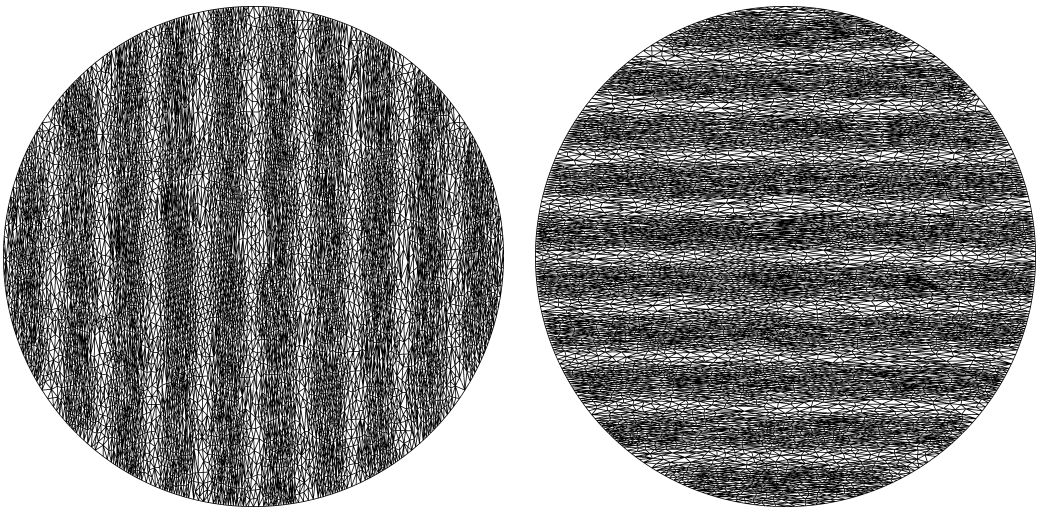

C D

Figure 3.8.: Anisotropic meshes used to test for imbalance injection by mesh-to-mesh interpolation. There is no relationship between meshes C and D, other than that they cover the same domain.

curacy in the L_2 norm of the final model solution with respect to the geostrophic balance preserving interpolant.

To further demonstrate geostrophic balance preservation the test was repeated on two anisotropic circular meshes C and D generated using Gmsh [Geuzaine and Remacle, 2009] and the ani2d mesh optimisation library [Vasilevskii and Lipnikov, 1999] with elements stretched in perpendicular directions as shown in figure 3.8. Meshes C and D have 7986 and 7205 nodes respectively, and a maximum element edge length ratio of ~ 30. The velocity field was initialised to be in discrete geostrophic balance with this layer thickness as before, with interpolations backwards and forwards between the two meshes at 10 timestep intervals for 20 interpolations. Simulations were conducted with fields interpolated using the geostrophic balance preserving interpolant, Galerkin projection, and Grandy interpolation, with the maximum change in layer thickness nodal values between interpolations shown in figure 3.9. When applying the geostrophic balance preserving interpolant the maximum change between interpolations was $\lesssim 10^{-13}$ as before.

Finally, the geostrophic balance preservation the test was repeated using the anisotropic meshes C and D in figure 3.8, with a layer thickness initialised to random values in the interior and a value of zero on the boundary. The velocity field was initialised to be in discrete geostrophic balance with this layer thickness. The model was integrated as before, with interpolations backwards and forwards between the two meshes at 10 timestep intervals for 4 interpolations, using the geostrophic balance preserving interpolant. The maximum change in layer thickness nodal values between interpolations was $\lesssim 10^{-12}$, and hence the solution was observed to be steady to within double precision round-off error.

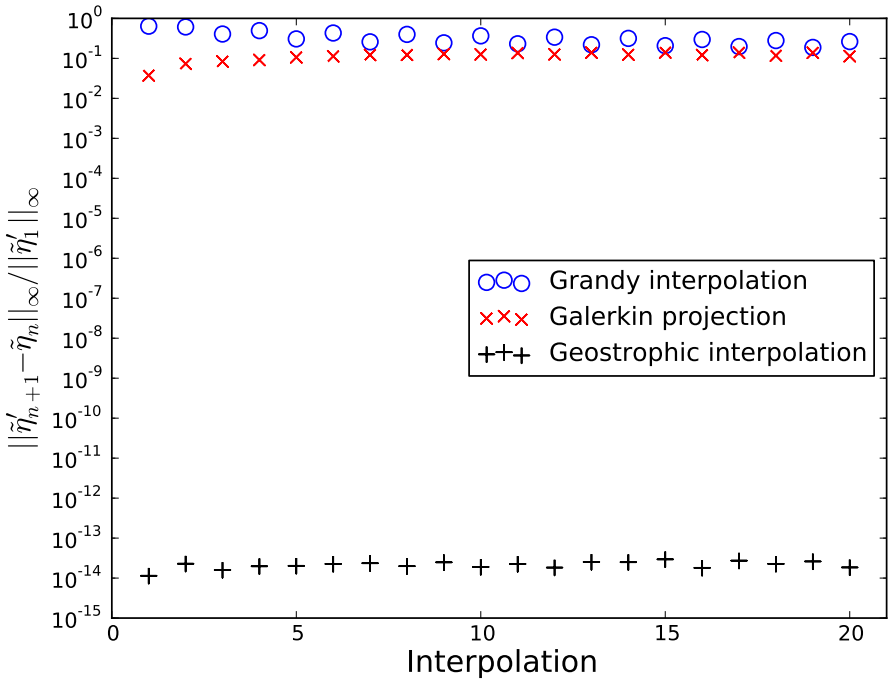

Figure 3.9.: The maximum change in layer thickness nodal values between interpolations backwards and forwards between the anisotropic meshes C and D in figure 3.8.

3.3.2. Nearly balanced states

The utility of the geostrophic balance preserving interpolant was tested for a system close to geostrophic balance. The Gaussian layer thickness profile in figure 3.5 had a perturbation applied of the form:

$$\Delta\tilde{\eta} = \frac{1}{10} X\tilde{\eta}, \tag{3.24}$$

where X is some point-wise random value in the range $[0, 1]$. This perturbation was smoothed using a Helmholtz smoother with a characteristic length scale of $D/8$ (where here D is the domain size) to produce the layer thickness perturbation shown in figure 3.10. The velocity field was initialised to be in discrete geostrophic balance with the unperturbed layer thickness, thereby generating a nearly balanced state.

The system was integrated as before, with interpolations backwards and forwards between the pseudo-isotropic meshes A and B in figure 3.4 at 10 timestep intervals. One can define an "imbalanced layer thickness":

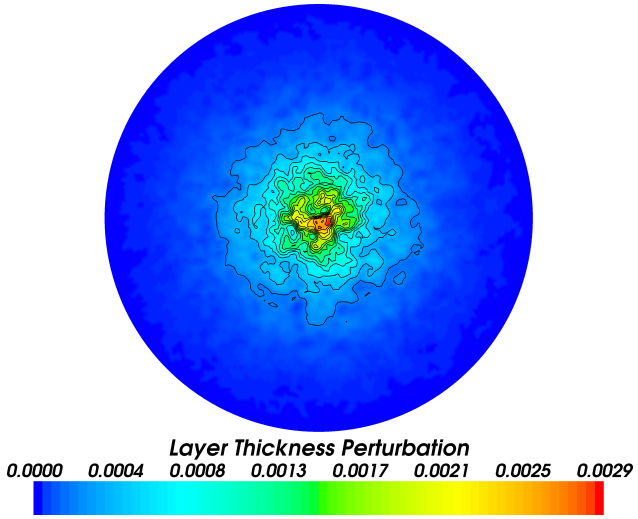

Figure 3.10.: Perturbation applied to the layer thickness in figure 3.5 to test for imbalance injection by interpolation of a nearly balanced state.

$$\tilde{\eta}_{imbal} \coloneqq \tilde{\eta} - \frac{1}{g}\tilde{\Phi}, \tag{3.25}$$

where $\tilde{\Phi}$ is the scalar potential computed from the Helmholtz decomposition of the Coriolis acceleration (with an arbitrary zero point). The final imbalanced layer thickness when interpolating using Galerkin projection and the geostrophic balance preserving interpolant is shown in figure 3.11. When interpolating using Galerkin projection imbalance is observed to be injected near the layer thickness maximum. This additional imbalance dominates over the original layer thickness perturbation after 20 interpolations. When using the geostrophic balance preserving interpolant propagation of the original layer thickness perturbation is observed, with no significant pollution introduced by the interpolation.

Defining a "balanced velocity" $\tilde{\boldsymbol{u}}_{bal}$ where:

$$f L\tilde{\boldsymbol{u}}_{bal} \coloneqq gC\tilde{\eta}, \tag{3.26}$$

and an "imbalanced velocity" $\tilde{\boldsymbol{u}}_{imbal}$:

$$\tilde{\boldsymbol{u}}_{imbal} \coloneqq \tilde{\boldsymbol{u}} - \tilde{\boldsymbol{u}}_{bal}, \tag{3.27}$$

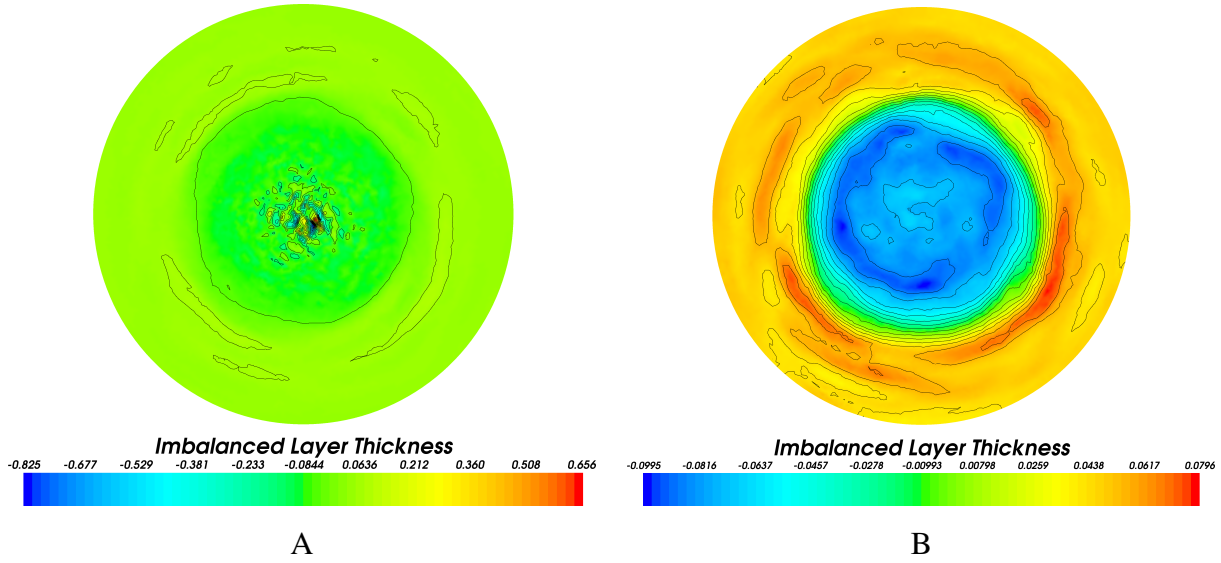

Figure 3.11.: The final imbalanced layer thickness, computed as the difference between the simulation layer thickness and the scaled scalar potential from the Helmholtz decomposition of the Coriolis acceleration. A: Galerkin projection. B: Helmholtz decomposed geostrophic balance preserving interpolation.

allows one to compute an imbalanced kinetic energy:

$$\begin{aligned} T_{imbal} &= \tfrac{1}{2}\left\|\tilde{u}^{\delta}_{imbal}\right\|^2_{L_2} \\ &= \tfrac{1}{2}\left\|\tilde{u}^{\delta} - \frac{g}{f}L^{-1}C\tilde{\eta}\right\|^2_{L_2}. \end{aligned} \tag{3.28}$$

The imbalanced kinetic energies when interpolating using Galerkin projection and the geostrophic balance preserving interpolant are shown in figure 3.12. When interpolating using Galerkin projection the imbalanced kinetic energy, compared before and after each interpolation, is observed to increase by up to a factor of 70, with the imbalanced kinetic energy peaking at 150 times its initial value. When using the geostrophic balance preserving interpolant the imbalanced kinetic energy, compared before and after each interpolation, is observed to increase by at most a factor 1.19. The imbalanced kinetic energy never exceeds its value at simulation start.

3.3.3. Kelvin wave

The interpolant was tested for a Kelvin wave, configured as in in Ham et al. [2005] and Cotter et al. [2009b] with an initial condition:

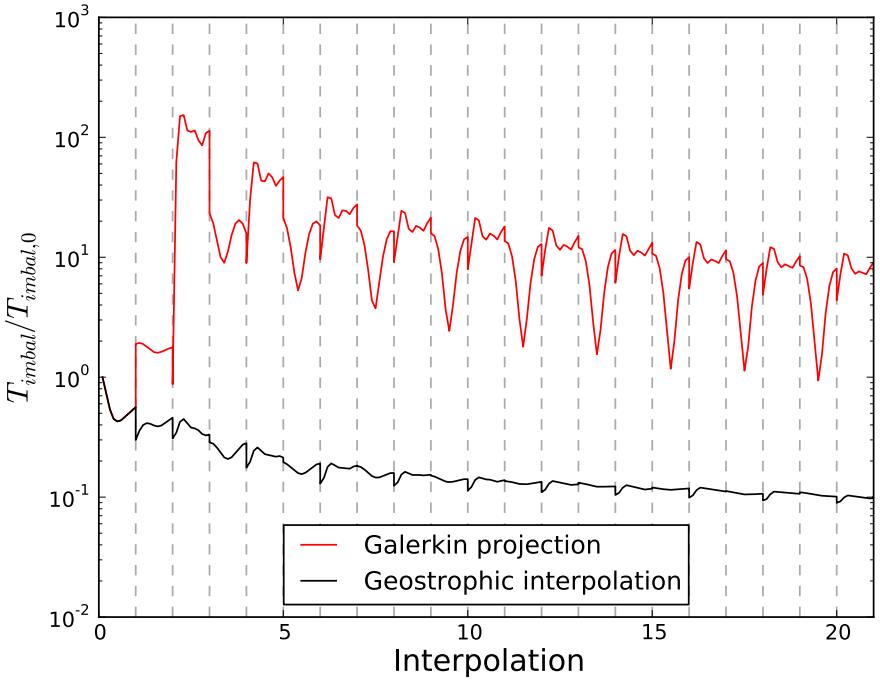

Figure 3.12.: The imbalanced kinetic energy, normalised by the initial imbalanced kinetic energy, for the nearly balanced interpolation test. When interpolating using Galerkin projection large increases in the imbalanced kinetic energy are observed after interpolation, with these increases dominating over the original imbalanced kinetic energy. When using the geostrophic balance preserving interpolant imbalanced kinetic energy injection is significantly reduced.

$$\eta(r,\theta) = \exp\left(\frac{r-r_0}{Ro}\right)\cos\theta, \tag{3.29a}$$

$$u_\theta(r,\theta) = \frac{1}{Fr}\exp\left(\frac{r-r_0}{Ro}\right)\cos\theta, \tag{3.29b}$$

$$u_r = 0, \tag{3.29c}$$

for $Ro = 10$ and $Fr = 1$ in a circular domain of non-dimensional radius $r_0 = 1$. The Kelvin wave is geostrophically balanced in the direction normal to the boundary and imbalanced in the tangential direction. The model was integrated with a timestep of $2\pi \times 10^{-4}(D/U)$ for a total simulated model time of $2\pi(D/U)$, corresponding to the time taken for a single Kelvin wave to perform a circuit of the domain in the limit of large r_0. Two meshes of quasi-uniform resolution with 1473 and 1461 nodes respectively were created using Gmsh [Geuzaine and Remacle, 2009] and the ani2d mesh

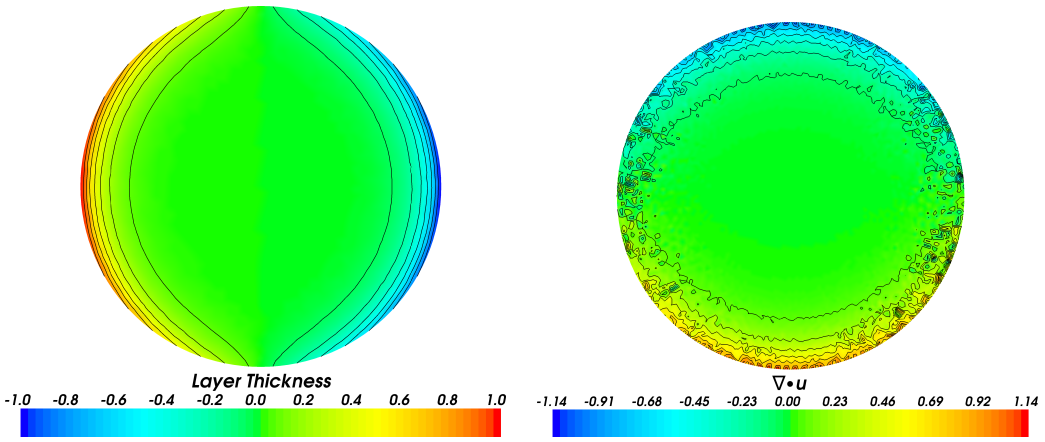

Figure 3.13.: Left: Initial layer thickness used for the Kelvin wave test at $Ro = 10$, $Fr = 1$. Right: Initial velocity divergence, $M_2^{-1}C^T\tilde{\boldsymbol{u}}$. The interpolation of the initial condition (3.29) yields a somewhat noisy initial velocity divergence.

optimisation library [Vasilevskii and Lipnikov, 1999], and the solution interpolated backwards and forwards between these meshes at 10 timestep intervals.

The initial layer thickness and velocity divergence is shown in figure 3.13, and the Helmholtz decomposition of the initial Coriolis acceleration in figure 3.14. The final solutions when interpolating using Galerkin projection and the geostrophic balance preserving interpolant are shown in figure 3.15. Relatively little difference is observed in the final layer thickness field between these simulations. However, when interpolating using Galerkin projection, noise is observed in the velocity divergence field, originating at the boundary. This noise is significantly reduced when using the geostrophic balance preserving interpolant.

A discretisation of the linearised shallow-water equations conserves energy if the layer thickness gradient matrix is, after multiplication by some diagonal matrix, equal to the transpose of the velocity divergence matrix [Ham et al., 2007], and if the implicit midpoint rule is used for timestepping [Leimkuhler and Reich, 2004]. Hence the $P_{1\mathrm{DG}}P_2$ spatial discretisation of the linearised shallow-water equations as presented here conserves the total energy. The kinetic, potential, and total energy of the system when interpolating using Galerkin projection, the geostrophic balance preserving interpolant, and when using a single fixed computational mesh, are shown in figure 3.16. The fixed mesh simulation is observed to conserve the total energy to within one part in 10^4, with the relatively high error attributed to the tolerances used for the linear solvers. The use of direct solvers, combined with more precision robust calculation of the energy diagnostics, is expected to decrease this error. When interpolating between meshes using Galerkin projection a systematic dissipation of both kinetic and potential energy is observed, leading to a decrease in the total system energy of 1.3% after 1000 interpolations, at the end of the simulation. When using the geostrophic balance preserving interpolant a slight increase the potential energy is observed,

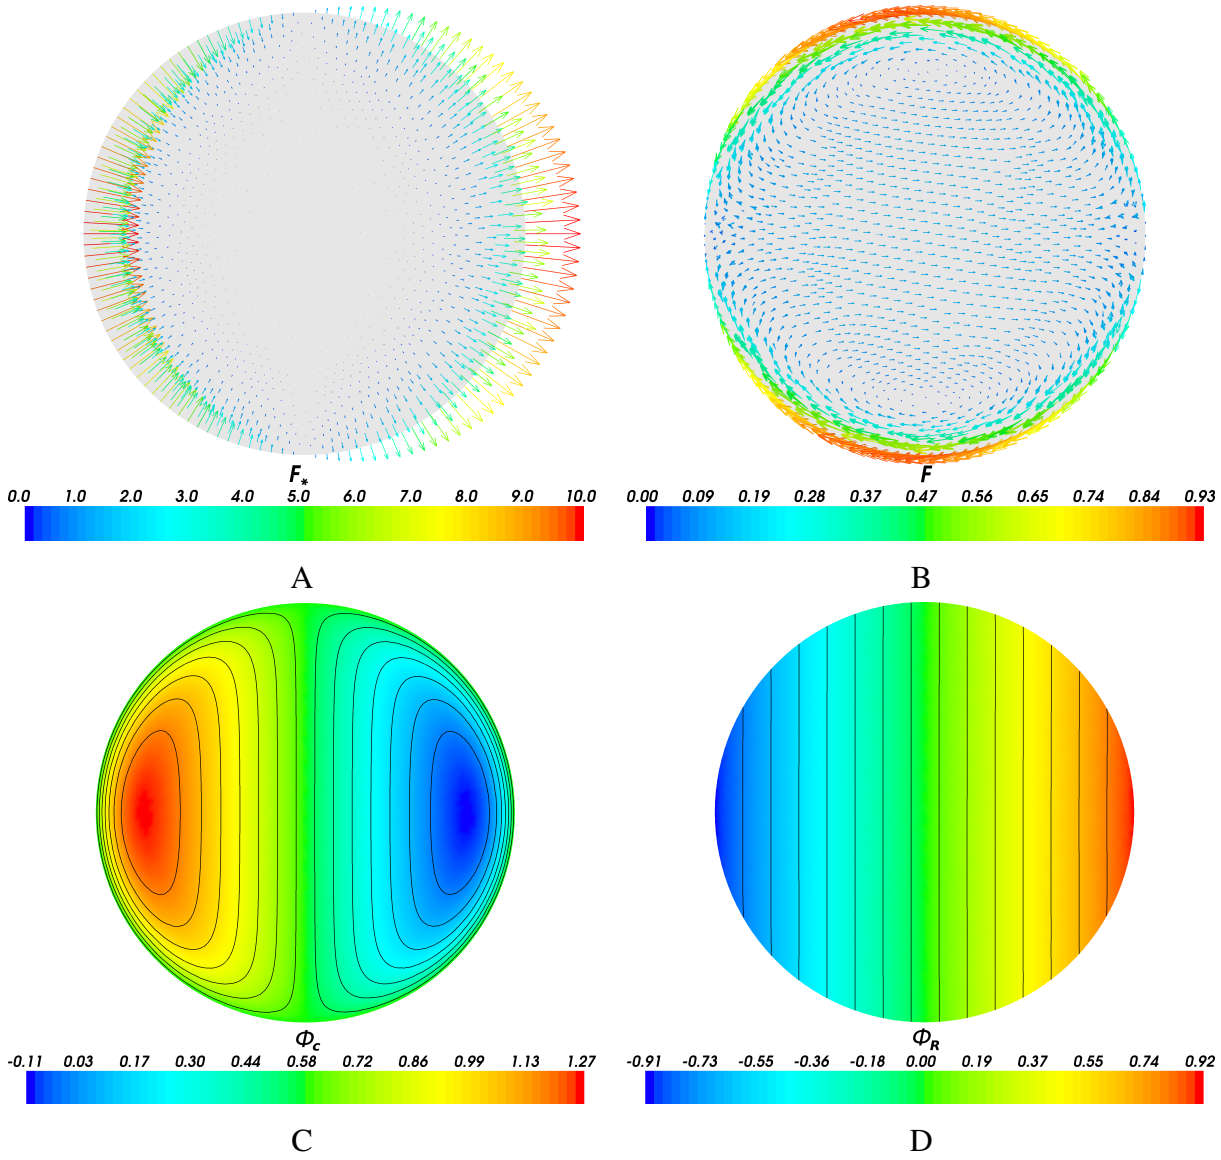

Figure 3.14.: Helmholtz decomposition of initial Coriolis acceleration for the Kelvin wave test at $Ro = 10$, $Fr = 1$. A: Coriolis acceleration, $\boldsymbol{F}_*$. B: Non-divergent residual, $\boldsymbol{F}$. C: Scalar potential with a constant boundary value, Φ_c. D: Scalar potential residual, Φ_R.

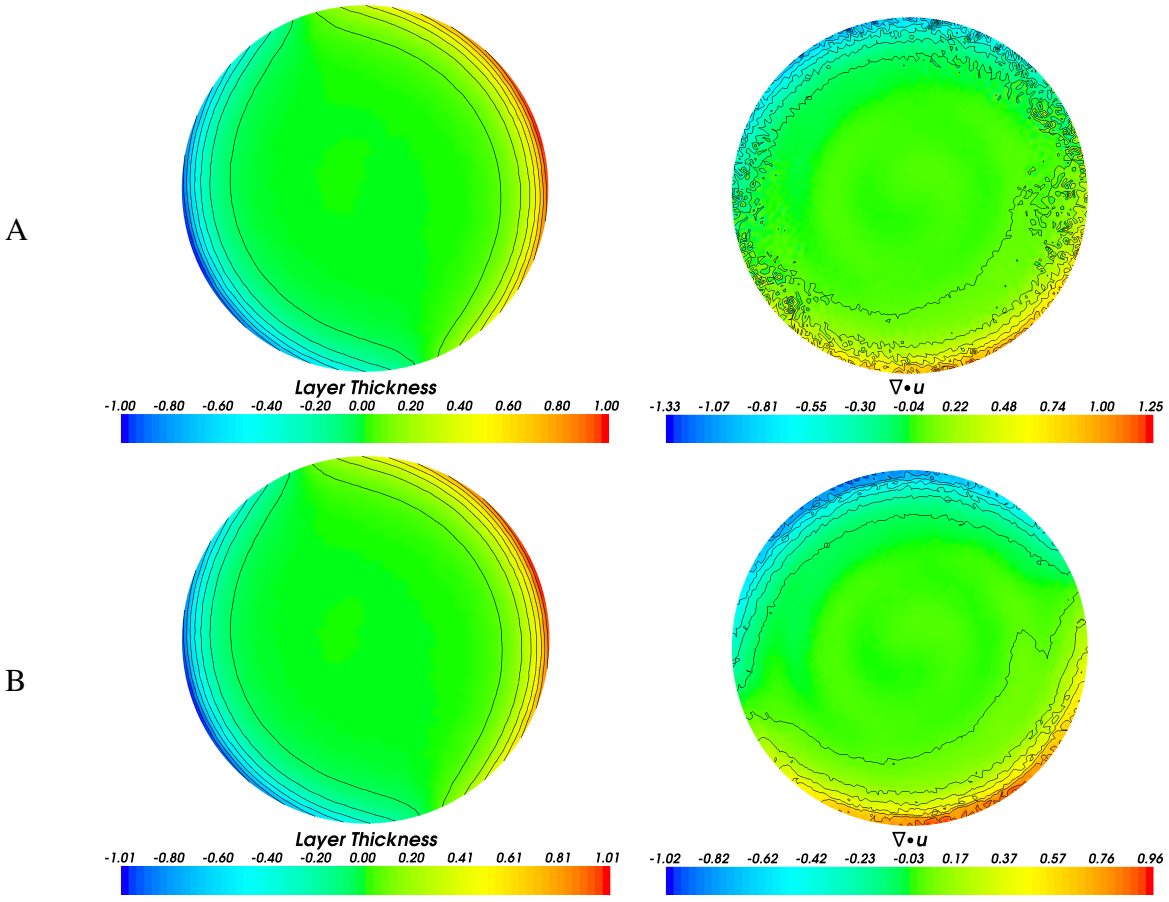

Figure 3.15.: The final solution of the Kelvin wave test at $Ro = 10$, $Fr = 1$. Left: Final layer thickness. Right: Final velocity divergence, $M_2^{-1}C^T\tilde{\boldsymbol{u}}$. A: Galerkin projection. B: Helmholtz decomposed geostrophic balance preserving interpolation.

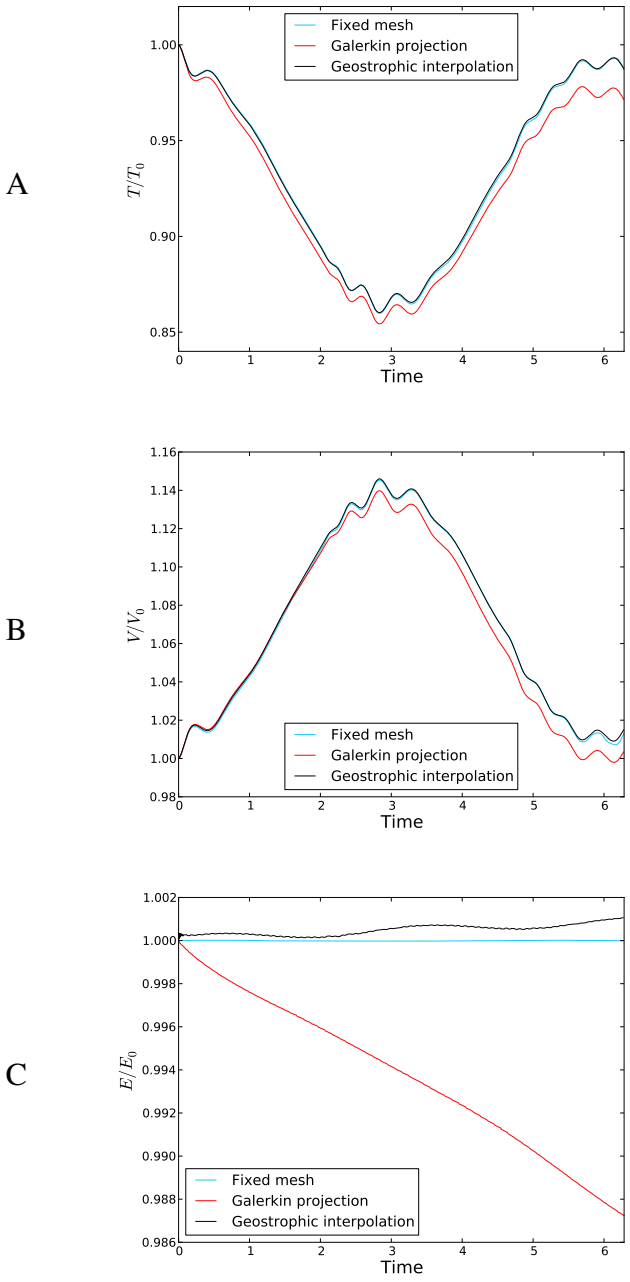

Figure 3.16.: The energy of the Kelvin wave test when interpolating between meshes using Galerkin projection and the geostrophic balance preserving interpolant, and when using a single fixed computational mesh. A: Kinetic energy. B: Potential energy. C: Total energy.

leading to an increase in the total system energy of 0.11% after 1000 interpolations. While the geostrophic balance preserving interpolant is not energy conserving, the change in system energy is, for this test, more than an order of magnitude smaller than that observed when interpolating using Galerkin projection.

In further testing it was found that highly anisotropic elements intersecting the domain boundary led to very poor results when using the geostrophic balance preserving interpolant. This is likely due to significant interpolation errors in the projection of $\tilde{\Phi}$ in this region, possibly as a

Donor mesh resolution ($x \times y$)	Target mesh resolution($x \times y$)
24×26	26×24
32×35	35×32
48×52	52×48
64×70	70×64
96×105	105×96
128×130	130×128
192×215	215×196

Table 3.1.: Mesh resolutions used for the geostrophic balance preserving interpolant convergence test. A mesh resolution of $N \times M$ denotes a division in the x-direction into N sections, a division in the y-direction into M sections, and the division of each resulting quadrilateral into two triangles.

result of the Dirichlet boundary condition for $\tilde{\Phi}_c$ in the mesh-to-mesh Galerkin projection, which pollutes the interpolated Coriolis acceleration. The magnitude of the possible error was not investigated in detail, although the issue was successfully avoided by imposing constraints on the maximum element size for elements directly on the boundary.

3.3.4. Accuracy

Since the geostrophic balance preserving interpolant is composed of a Galerkin discretisation of the Helmholtz decomposition of the Coriolis acceleration followed by a mesh-to-mesh Galerkin projection of the decomposition, when using the $P_{1\mathrm{DG}}P_2$ element pair the interpolant is expected to be second order accurate for velocity and third order accurate for layer thickness.

A series of structured triangular mesh pairs for a 2D unit square $-0.5 \leq x \leq 0.5$ and $-0.5 \leq y \leq 0.5$ were generated with resolutions in the x- and y-directions as given in table 3.1. A layer thickness of the form:

$$\eta = \sin(10\pi x)\sin(10\pi y), \tag{3.30}$$

and a velocity of the form:

$$\boldsymbol{u} = -2\pi\hat{\boldsymbol{z}} \times \nabla\eta + \sin(2x)\hat{\boldsymbol{x}} + \sin(2y)\hat{\boldsymbol{y}}, \tag{3.31}$$

were interpolated between the meshes in each pair using the geostrophic balance preserving interpolant. The first term in (3.31) corresponds to a flow that is, for $f/g = -2\pi$, in geostrophic balance with the layer thickness (3.3.4). The remaining terms correspond to an imbalanced layer thickness. In order to test for additional error introduced by the scalar potential $\tilde{\Phi}$ and layer thickness decomposition, as per equation (3.20), tests were conducted for a doubly periodic and for a bounded domain. The L_2 errors $\left\|\eta_B^\delta - \eta_A^\delta\right\|_{L_2}$, $\left\|u_{x,B}^\delta - u_{x,A}^\delta\right\|_{L_2}$ and $\left\|u_{y,B}^\delta - u_{y,A}^\delta\right\|_{L_2}$, were computed explicitly via supermesh construction, as described in Farrell [2009]. For comparison the fields were also interpolated using Galerkin projection, giving a measure of the quality of the geostrophic balance preserving interpolant relative to the projection that is optimal in the L_2 norm.

The resulting errors are shown for the doubly periodic domain in figure 3.17 and for the bounded domain in figure 3.18. The geostrophic balance preserving interpolant is observed to be second order accurate for velocity and third order accurate for layer thickness, as expected. For the doubly periodic domain the average L_2 norm error for the geostrophic balance preserving interpolant is observed to be 1.37 times the optimal value for velocity, and (since no layer thickness decomposition is applied in this case) optimal for layer thickness. For the bounded domain the error in velocity is not significantly changed, and the error in layer thickness is increased to 1.005 times the optimal value, indicating that the decomposition of the scalar potential $\tilde{\Phi}_A$ and layer thickness $\tilde{\eta}_A$ introduces only a small additional error. For comparison, in Farrell [2009] collocation is found to give, for a field $\sin x + \cos x$, an L_2 error that is $\sim 2 - 2.5$ times the optimal value for piecewise linear elements, and ~ 1.1 times the optimal value for piecewise quadratic elements.

3.3.5. Computational cost

Since the diagnostic solve for the scalar potential (3.19) is equivalent to the pressure projection method, it has a relatively high computational cost. Similarly, the solves for the decomposed scalar potential and layer thickness components $\tilde{\Phi}_0^A$, η_0^A, and $\tilde{\Phi}_1^A$ are also expected to be associated with relatively high computational costs. For each of these (four) linear solves, the cost is expected to be of order the cost of one pressure solve when applying the pressure projection method, and hence of order the cost of one Picard iteration in a solution of the Navier-Stokes equations.

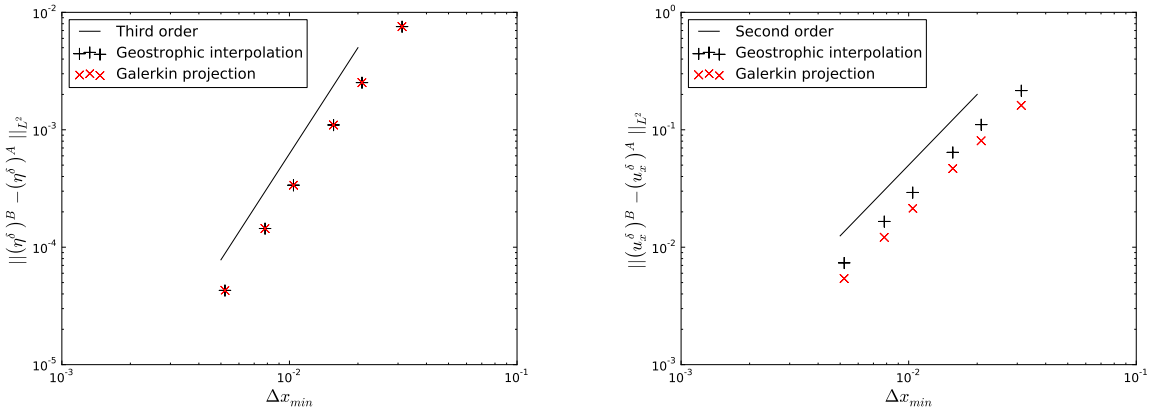

Figure 3.17.: Convergence test for the geostrophic balance preserving interpolant in a doubly periodic domain, with Galerkin projection for comparison. Left: L_2 error in the layer thickness. Right: L_2 error in the x-component of velocity. The error in the y-component of velocity is similar.

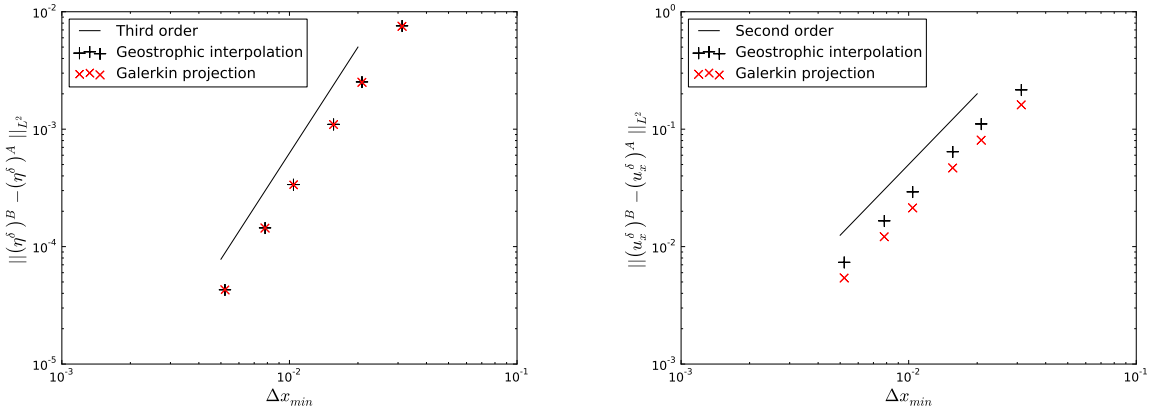

Figure 3.18.: Convergence test for the geostrophic balance preserving interpolant in a bounded domain, with Galerkin projection for comparison. Left: L_2 error in the layer thickness. Right: L_2 error in the x-component of velocity. The error in the y-component of velocity is similar.

The computational cost of the bounded domain test used in section 3.3.4 was measured. The MPI routine mpi_wtime[1] was used to measure the total time spent performing the geostrophic balance preserving interpolation. This measurement includes the time spent performing the Helmholtz decomposition of the Coriolis acceleration on the donor mesh, the time spent decomposing the scalar potential and layer thickness on the donor mesh, the time spent in the mesh-to-mesh Galerkin projection of the decomposed Coriolis acceleration and layer thickness, and the time spent recomposing the Coriolis acceleration and the layer thickness on the target mesh. The linear systems were solved via preconditioned conjugate gradients using the PETSc library [Balay et al., 1997, 2008, 2009], with a smoothed aggregation algebraic multigrid method applied for the scalar potential $\tilde{\Phi}^A$ solve, and for the solves for the decomposed scalar potential and layer thickness components $\tilde{\Phi}^A_0$, η^A_0 and $\tilde{\Phi}^A_1$ [Vaněk et al., 1996, Kramer et al., 2010]. Fluidity-ICOM was compiled using version 11.1 of the Intel Fortran compiler and version 4.4 of the GNU C and C++ compilers. The profiling was conducted in serial on a 64-bit Linux workstation with a 2.66GHz Intel X5355 quad core processor and 4GB of system RAM, and using version 3.8 of the libhoard fast memory allocator [Berger et al., 2000]. For comparison, the computational cost of interpolating the velocity and layer thickness using Galerkin projection was also measured. In all tests, mesh-to-mesh Galerkin projection was performed using the local assembly algorithm of Farrell and Maddison [2011] described in section 3.1. In each case, the interpolation was performed three times and an average taken. The results of these tests are shown in figure 3.19.

The geostrophic balance preserving interpolant is found to scale linearly with problem size. The cost of the geostrophic balance preserving interpolant is found to be 1.9 times the cost of interpolating using Galerkin projection. On average, for the geostrophic balance preserving interpolant, 12.2% of the interpolation time was spent performing the solve for the scalar potential for the Helmholtz decomposition of the Coriolis acceleration on the donor mesh, 15.4% performing the scalar potential and layer thickness decomposition on the donor mesh, 59.4% performing the mesh-to-mesh Galerkin projection of the decomposed Coriolis acceleration and layer thickness, 6.6% performing the recomposition on the target mesh, and 6.4% in other calculations (including the computation of the Coriolis acceleration on the donor mesh). Hence, for this test, the geostrophic balance preserving interpolant is found to approximately double the interpolation cost, relative to interpolating the velocity and layer thickness using Galerkin projection.

Note that this performance analysis is highly susceptible to implementation details. Hence it is likely that the performance of each of these interpolants can be improved – no detailed optimisa-

[1] Note that, although the MPI routine mpi_wtime was used to measure the interpolation time, these benchmarks were conducted in serial.

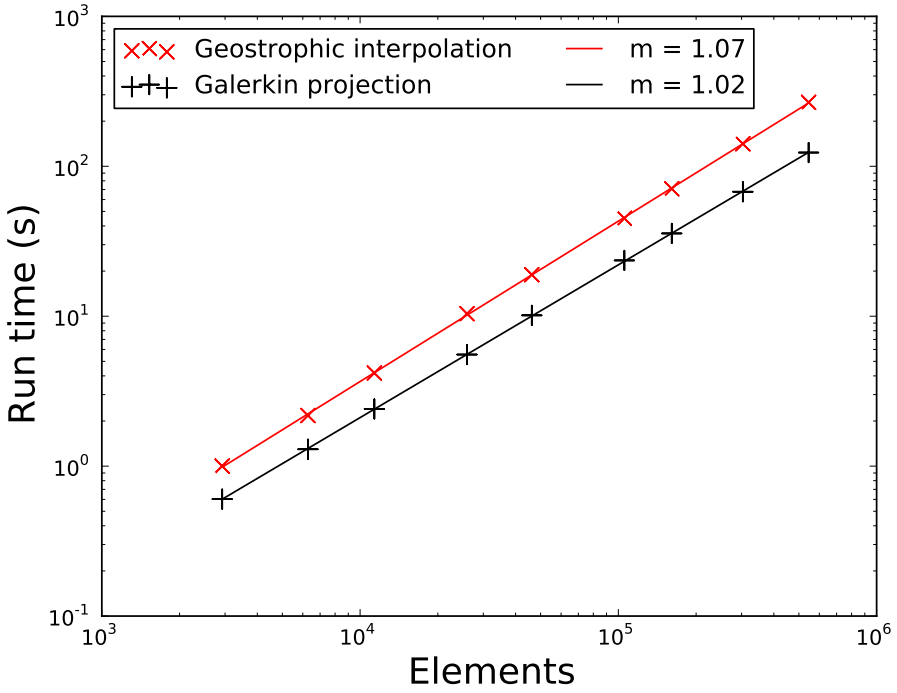

Figure 3.19.: Computational cost of the geostrophic balance preserving interpolant in a bounded domain, with Galerkin projection for comparison.

tions of the code were undertaken. The most important results are therefore that the geostrophic balance preserving interpolant is observed to scale linearly with problem size, and that the additional cost of the decompositions of the Coriolis acceleration and layer thickness are of a comparable magnitude to the cost of the interpolation of the decomposition.

3.4. Application to other finite element pairs

Thus far only the $P_{1\mathrm{DG}}P_2$ element pair has been considered. We now consider the application of the geostrophic balance preserving interpolant to alternative finite element discretisations.

The $P_{1\mathrm{DG}}P_2$ pair is a member of the more general $P_{(N-1)\mathrm{DG}}P_N$ family of element pairs. The first order member of this family, P_0P_1, is discussed in the context of shallow-water ocean modelling in further detail in Williams and Zienkiewicz [1981] and Umgiesser et al. [2004]. All of the element pairs in this family satisfy the embedding conditions (2.23a) and (2.23b), and hence satisfy the optimal balance property (3.18) [Cotter and Ham, 2011]. Hence the geostrophic balance preserving interpolant can be applied to each element pair in this family with, in each case, the property that steady and balanced states remain steady and balanced after interpolation.

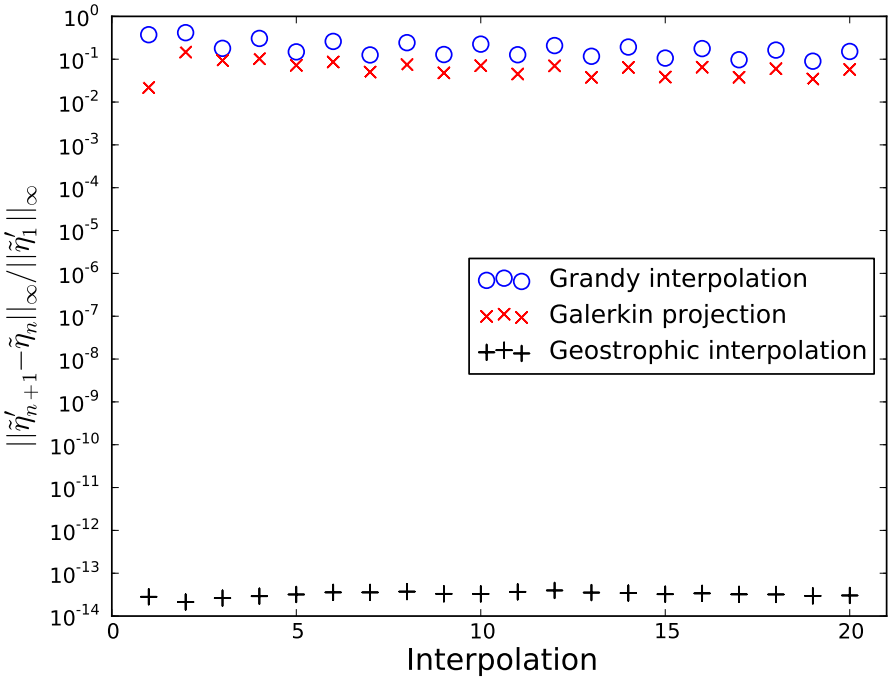

Figure 3.20.: The maximum change in layer thickness nodal values between interpolations backwards and forwards between the pseudo-isotropic meshes A and B in figure 3.4 using a $P_{2DG}P_3$ finite element discretisation of the linearised shallow-water equations.

The geostrophic balance preservation test of section 3.3.1 was repeated for the $P_{2DG}P_3$ element pair using repeated interpolation between meshes A and B in figure 3.4. The model was configured and integrated as before, with interpolations between meshes A and B at ten timestep intervals for a total of 20 interpolations. The test was performed using Grandy interpolation for the velocity and layer thickness, Galerkin projection for the velocity and layer thickness, and the geostrophic balance preserving interpolant. The maximum changes in the layer thickness nodal values between interpolations are shown in figure 3.20. The results are broadly similar to those obtained when using the $P_{1DG}P_2$ element pair. When applying the geostrophic balance preserving interpolant, the maximum change in layer thickness nodal values between interpolations is $\lesssim 10^{-13}$, and hence the simulation is observed to remain exactly steady and balanced between interpolations, to within double precision round-off error.

The P_1P_1 finite element pair satisfies only embedding condition (2.23a), but not condition (2.23b) [Cotter and Ham, 2011]. Hence, in a P_1P_1 discretisation of the linearised shallow-water equations, geostrophically balanced states with a constant stream function on the boundary need not be discrete divergence free, and hence need not be exactly steady solutions of the discrete linearised shallow-water equations. The lack of steady geostrophic balance representation by this element pair is noted in Le Roux et al. [1998] and Le Roux et al. [2007].

There are two issues in applying the geostrophic balance preserving interpolant as presented in section 3.2.3 to the P_1P_1 element pair. First, in order for assembly of the discrete Laplacian matrix $C^T M_1^{-1} C$ in the diagnostic solve for the conservative potential to be practical, the velocity mass matrix M_1 must be diagonalised. For the P_1 velocity basis functions this can be achieved using mass lumping as per equation (2.18). Hence equation (3.11) becomes:

$$(M_1)_L \, \tilde{\boldsymbol{F}} = (M_1)_L \, \tilde{\boldsymbol{F}}_* + C\tilde{\Phi}, \tag{3.32}$$

where $(M_1)_L = \text{diag}\left(M_L', M_L'\right)$. Hence $\tilde{\boldsymbol{F}}_*$ must now be defined via a lumped mass Galerkin projection:

$$(M_1)_L \, \tilde{\boldsymbol{F}}_* = fL\tilde{\boldsymbol{u}}. \tag{3.33}$$

Second, since the P_1P_1 finite element pair is not LBB stable, the resulting diagnostic Poisson solve for the conservative potential is susceptible to spurious modes. Stabilisation methods can be applied, such as the K stabilisation matrix method applied to the fractional timestep pressure projection method in equations (2.19) and (2.21). No stabilisation methods are applied here.

The geostrophic balance preservation test of section 3.3.1 was repeated for the P_1P_1 element pair, using repeated interpolation between meshes A and B in figure 3.4. The model was configured and integrated as before, with interpolations between meshes A and B at ten timestep intervals for a total of 20 interpolations. The test was performed using collocation for the velocity and layer thickness, Galerkin projection for the velocity and layer thickness, and the lumped mass geostrophic balance preserving interpolant as described above. For additional comparison, simulations were conducted on meshes A and B with no mesh-to-mesh interpolation. The final layer thickness and change in layer thickness from the initial condition are shown in figure 3.21. The maximum change in layer thickness nodal values between interpolations is shown in figure 3.22, and the normalised L_2 difference between the layer thickness and its initial condition is shown in figure 3.23.

As expected, the simulations conducted on a fixed mesh, with no mesh-to-mesh interpolation, are not exactly steady, despite the initial conditions being in exact discrete geostrophic balance. The transience appears as a mesh scale noise in the layer thickness. When interpolating backwards

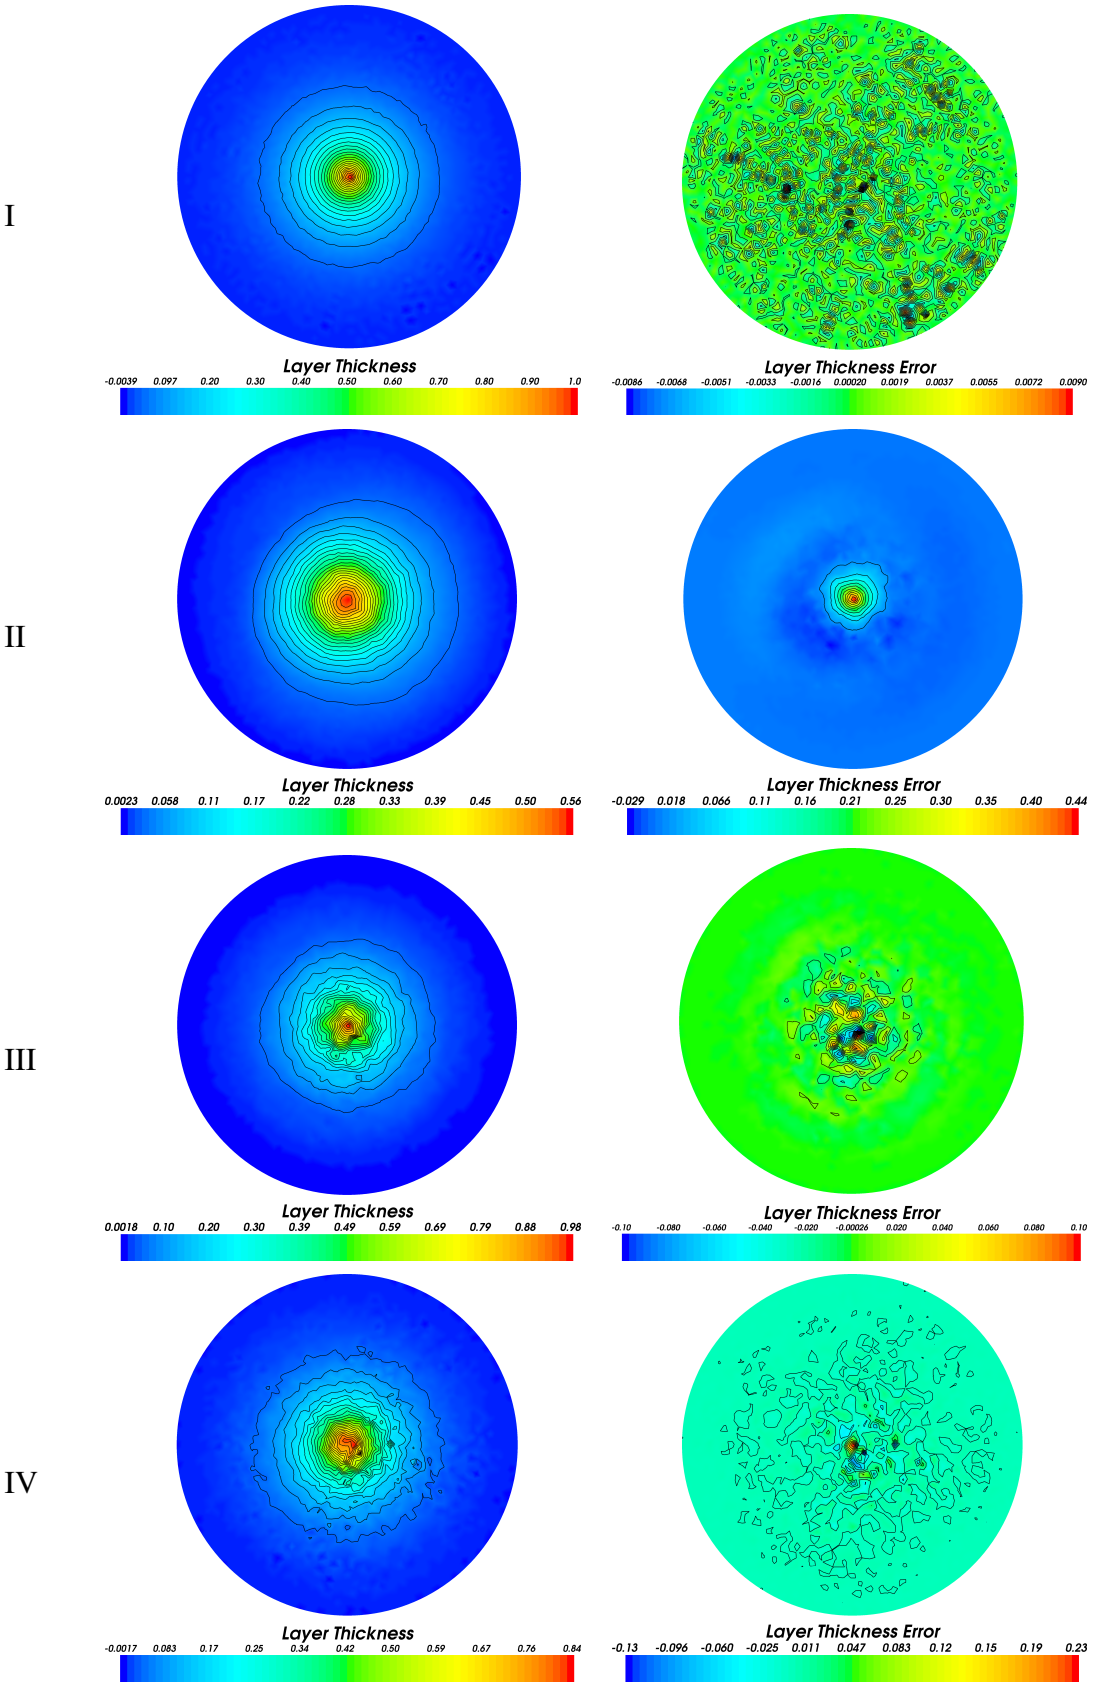

Figure 3.21.: The final solution of the geostrophic balance preservation test using a P_1P_1 finite element discretisation of the shallow-water equations, after 20 repeated interpolations backwards and forwards between the pseudo-isotropic meshes A and B in figure 3.4. Left: Final layer thickness. Right: Change in layer thickness from the initial condition. I: Fixed mesh simulation conducted on mesh A. II: Collocation. III: Galerkin projection. IV: Lumped mass Helmholtz decomposed geostrophic balance preserving interpolation.

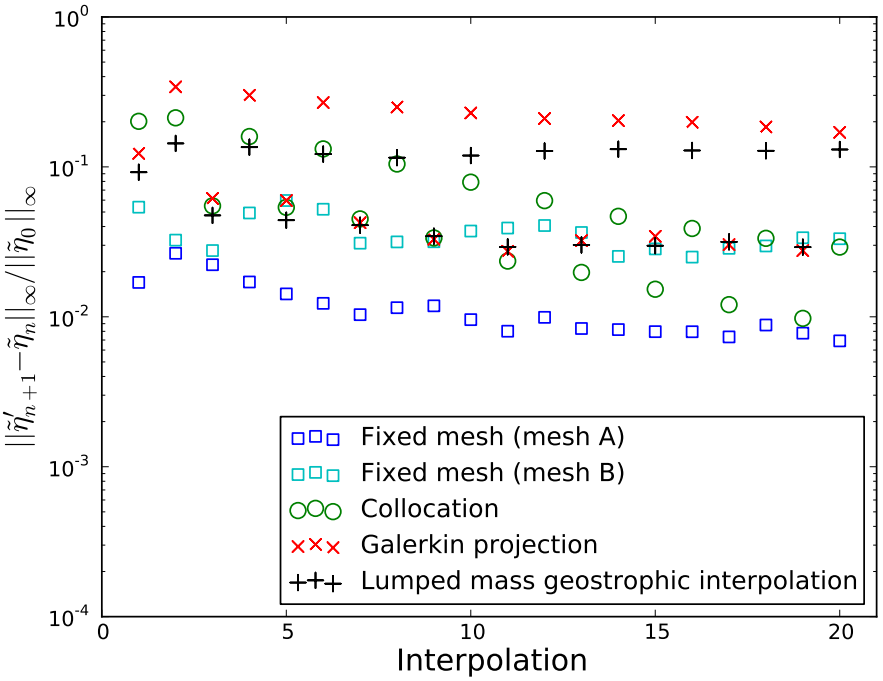

Figure 3.22.: The maximum change in layer thickness between interpolations backwards and forwards between the pseudo-isotropic meshes A and B in figure 3.4 using a P_1P_1 finite element discretisation of the linearised shallow-water equations. Normalised by the initial layer thickness maximum. When applying collocation, the change in layer thickness nodal values decreases in later interpolations, due to the numerical diffusion associated with this interpolant.

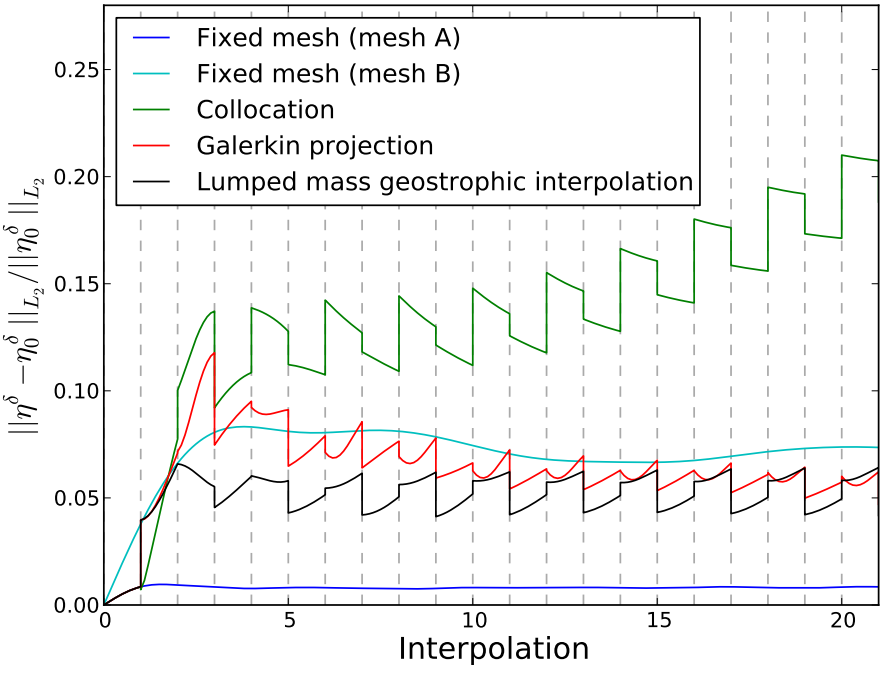

Figure 3.23.: The normalised difference between the layer thickness and its the initial condition for the geostrophic balance preservation test using a P_1P_1 finite element discretisation of the shallow-water equations. Since the simulation conducted on mesh B is also initialised on mesh B, the normalisation factor $\left\|\eta_0^\delta\right\|_{L_2}$ for this test differs marginally (by $\sim 2\%$) from the factor used for the other tests.

and forwards between meshes A and B, the application of collocation is observed to be associated with a high level of numerical diffusion, and consequently a large deviation of the final state from the initial condition. A significant large scale loss of rotational symmetry in the final solution is also observed. By contrast, interpolating using Galerkin projection leads to a relatively small deviation of the final solution from the initial condition, with the normalised L_2 layer thickness change from the initial condition remaining between the values observed for the two fixed mesh simulations. The application of the lumped mass geostrophic balance preserving interpolant results in the lowest level of transience of the three interpolants, as measured by the maximum change in layer thickness nodal values between interpolations, up to and including interpolation 7. The normalised L_2 layer thickness change from the initial condition remains below or comparable to that observed when interpolating using Galerkin projection, although both an increased mesh scale noise and increased diffusion is observed.

Hence, for a P_1P_1 finite element discretisation of the linearised shallow-water equations, the geostrophic balance preserving interpolant is of comparable quality, or is slightly inferior to, interpolation using Galerkin projection.

3.5. Application to Navier-Stokes

The application of the geostrophic balance preserving interpolant to the discrete Navier-Stokes equation is considered, with Dirichlet boundary conditions (2.11), and using the fractional time-step pressure projection method discussed in detail in section (2.1.2). We limit ourselves to the case of strong Dirichlet boundary conditions on all boundaries. In this case, the geostrophic balance preserving interpolant is as for the shallow-water equations, with the layer thickness replaced with the dynamical pressure p, and with three further modifications as follows.

First, it is required that the computed non-divergent Coriolis acceleration, $\tilde{\boldsymbol{F}}$, must be zero for a system in discrete geostrophic balance as defined by the discretisation of the Navier-Stokes equation. The Helmholtz decomposition of the Coriolis acceleration must therefore have an identical discretisation to that used for the Helmholtz decomposition of the intermediate divergent velocity in the pressure correction step of the Navier-Stokes solver (2.15a). In particular, if K matrix stabilisation (2.19) is applied to the pressure solver, this must be applied in the geostrophic balance preserving interpolant. Similarly, if a balanced pressure decomposition is applied in the Navier-Stokes solver, this must be applied in the geostrophic balance preserving interpolant. In the latter case, equation (3.11) becomes:

$$M_1\tilde{\boldsymbol{F}} = M_1\tilde{\boldsymbol{F}}_* + C\tilde{\Phi}_r + C_b\tilde{\Phi}_g, \tag{3.34}$$

where $\tilde{\Phi}_g$ is a "geostrophic scalar potential" component and $\tilde{\Phi}_r$ is a "residual scalar potential" component, where $-C_b$ is as per equation (2.28), and $\tilde{\Phi}_g$ is solved for as per equation (2.48):

$$\int_\Omega \nabla\chi_i \cdot \nabla\Phi_g^\delta = -\int_\Omega \nabla\chi_i \cdot \boldsymbol{F}_*^\delta. \tag{3.35}$$

Second, the application of strong no-normal-flow boundary conditions for velocity means that the scalar potential decomposition as per equations (3.22a) and (3.22b) need not be applied. The application of strong no-normal-flow velocity boundary conditions automatically precludes the generation of mesh scale imbalance injection at the boundary, without any such decomposition.

Finally, in three dimensions the vertical component of velocity can in principle be recovered in the inversion of the Coriolis acceleration for velocity on the target mesh (3.13): it is the difference between the vertical parts of the conservative and residual components of the Coriolis acceleration. For a state close to geostrophic balance these two terms are of comparable magnitude and opposite in sign, and hence the vertical component of velocity is a small residual. Hence it is likely to be more accurate simply to interpolate the vertical component of velocity separately, for example using Galerkin projection.

3.6. Summary

We have presented an interpolation method that, when applied to the $P_{1\mathrm{DG}}P_2$ discretisation of the linearised shallow-water equations on an f-plane, guarantees that steady and geostrophically balanced states on the donor mesh remain steady and geostrophically balanced after interpolation onto an arbitrary target mesh. We have stress tested this balance preserving property with highly anisotropic meshes and randomly initialised balanced states (constrained to satisfy appropriate boundary conditions). We have demonstrated the utility of this interpolant for nearly balanced dynamics, and quantified its accuracy in the L_2 norm.

We have further discussed the application of the interpolant to alternative finite element discretisations of the linearised shallow-water equations. This interpolant can be naturally extended to any of the $P_{(N-1)\mathrm{DG}}P_N$ family of finite element pairs, as demonstrated in geostrophic balance preservation tests using the $P_{2\mathrm{DG}}P_3$ finite element pair. The P_1P_1 discretisation of the linearised shallow-water equations was also considered. For this finite element pair geostrophically balanced states are not exactly steady. Interpolating using Galerkin projection was found to lead to an acceptable deviation from the initial balanced state, of a comparable magnitude to that observed when conducting simulations on fixed meshes. The application of a lumped mass geostrophic balance preserving interpolant leads to at best only a modest improvement, as compared against interpolation using Galerkin projection. The geostrophic balance preserving interpolant presented in this chapter appears to show significant advantages over other interpolants tested for the $P_{1\mathrm{DG}}P_2$ element pair, but shows little benefits over the application of Galerkin projection for the P_1P_1 element pair.

A shortcoming of this approach, at least in the form presented, is that it does not conserve energy. The Helmholtz decomposed interpolation of the Coriolis acceleration does not conserve kinetic energy or potential energy. Despite this, the change in energy when using the geostrophic balance preserving interpolant was found to be more than an order of magnitude smaller than the energy dissipation when interpolating using Galerkin projection. In addition to this, the interpolant does not locally conserve potential vorticity. Geophysical flows are only in geostrophic balance to leading order, and a lack of potential vorticity conservation could, when non-linear advection is included in the fundamental equations, lead to higher order balance loss. Potential vorticity decompositions could be considered where such a conservation is desired [Staquet and Riley, 1989, McIntyre and Norton, 2000], although the benefit of such an approach, bearing in mind that discretisations of the non-linear shallow-water equations are not generally potential vorticity conserving, may be somewhat limited compared to the benefit of leading order balance preservation.

The following chapter will describe dynamic mesh adaptive simulations of the thermally driven rotating annulus. These will be compared against fixed mesh simulations, MORALS simulations, and laboratory data, and strengths and weaknesses of the approach will be identified.

Chapter 4.

Application of dynamic mesh adaptivity to simulations of the thermally driven rotating annulus

In the previous two chapters we have considered issues with the application of unstructured dynamic mesh adaptive numerical methods to geophysical systems. In particular, the issues with representation and maintenance of physical balance were addressed, first for unstructured mesh finite element numerical modelling in chapter 2, and then for mesh-to-mesh interpolation when applying dynamic mesh adaptivity in chapter 3.

In this chapter, the application of dynamic mesh adaptive numerical methods to simulations of the thermally driven rotating annulus is now considered. First, in section 4.1, fixed structured mesh Fluidity-ICOM simulations of the rotating annulus are presented. These are compared against MORALS in section 4.1.1, and against laboratory observations in section 4.1.2. The configuration used for dynamic mesh adaptive simulations of the thermally driven rotating annulus is described in section 4.2. Dynamic mesh adaptive Fluidity-ICOM simulations of the rotating annulus are presented in section 4.3. These are compared against the corresponding fixed mesh simulations and experimental observations.

4.1. Fixed mesh Fluidity-ICOM simulations

The validation of Fluidity-ICOM simulations of the thermally driven rotating annulus proceeds in two stages. In this section simulations on fixed meshes, without application of dynamic mesh

adaptivity, are considered. In section 4.3 dynamic mesh adaptivity is enabled, and the performance compared with the fixed mesh simulations.

The comparison in this section is an extension of the balanced pressure decomposition test described in section 2.3. The configuration uses parameters as given in table 2.1, with all simulations conducted with a thermal forcing of $\Delta T = 4$ K, and at 13 rotation rates ranging from the non-rotating $\omega = 0$ case, to $\omega = 2.764$ rad s^{-1}. These parameters correspond to thermal Rossby numbers of $\Theta \leq 0.08$ and Taylor numbers $Ta \leq 2.5 \times 10^7$ (defined as per equations (1.5) and (1.6)). Note that these simulations were conducted using the higher viscosity, and not the corrected viscosity discussed in section 2.3.4. The comparison discussed in this chapter is therefore somewhat qualified by the use of a kinematic viscosity that is slightly higher than the laboratory value.

The system was simulated with Fluidity-ICOM using the stabilised P_1P_1 finite element pair with a P_2 balanced pressure decomposition, and with the prognostic temperature equation solved via a control volume finite element method. Boundary conditions were applied as before, with strong no-slip boundary conditions for velocity on all boundaries, with the thermal forcing applied via a strong boundary condition for temperature on the inner and outer annulus walls, and with a naturally imposed no-flux boundary condition for temperature on the base and lid. A linear equation of state was applied, a temperature independent diffusivity and viscosity used, and the centrifugal acceleration neglected. The resulting linear systems were solved via preconditioned conjugate gradients using the PETSc library [Balay et al., 1997, 2008, 2009, Saad and Schultz, 1986], with a smoothed aggregation algebraic multigrid method applied for the pressure solvers for both the residual and balanced pressure components [Vaněk et al., 1996, Kramer et al., 2010]. Some simulations also used the BoomerAMG multigrid method [Henson and Yang, 2002], via the hypre library [Falgout and Yang, 2009], for the pressure solvers. Consistent with previous numerical simulations of similar annulus regimes, no sub-grid-scale parameterisation method is applied (see, for example, James et al. [1981]).

It is known that MORALS has a favourable comparison against experiment at lower rotation rates [James et al., 1981]. Hence it is likely that the hyperbolic tangent stretch mesh implemented in MORALS is close to being an optimised mesh for simulation of lower Reynolds number rotating annulus flows. To enable a more realistic assessment of dynamic mesh adaptivity, the fixed mesh simulations were conducted on hyperbolic tangent stretch meshes derived from those used by MORALS, with resolution $24 \times 24 \times 64$ in the radial, vertical, and azimuthal dimensions respectively. For a given structured MORALS mesh, a structured tetrahedral mesh was generated by dividing each cubic region into six tetrahedra [Tanizume et al., 1990], as illustrated in figure 2.8.

The simulations were integrated in parallel on 8 MPI processes for a total of 4000 s of simulated laboratory time, using the Oxford Supercomputing Centre "redqueen" machine. For additional model efficiency, the model timestep size was adapted after every timestep to target a specified maximum Courant number – a maximum Courant number of 0.75 for the $\omega = 2.503$ rad s^{-1} simulation, and unity for all other simulations. For the higher rotation rate configurations very noisy solutions were observed over the early parts of the simulations. The noise exhibited itself as a meridional overturning cell, with rapidly varying magnitude and alternating direction. This was later identified as a form of Courant-Friedrichs-Lewy (CFL) instability. The instability disappeared as the simulations developed. All diagnostics presented in this chapter neglect the unstable portions of the simulations. The reduced Courant number for the $\omega = 2.503$ rad s^{-1} simulation was used in order to reduce the duration of this instability.

Laboratory data are available for this system [Hignett et al., 1985, Hignett, 1985, Read, 2003] enabling quantitative comparison of the simulations against experimental observations (the data set used for comparison in 2.3). For further comparison, MORALS simulations were conducted at resolution $32 \times 32 \times 128$ in the radial, vertical and azimuthal dimensions respectively[1], at rotation rates of $\omega = 1.0$ rad s^{-1} (the simulation used for comparison in section 2.3) and $\omega = 2.764$ rad s^{-1}. As for the Fluidity-ICOM simulations, the MORALS configurations were initialised with zero velocity and with a linear stratification.

The final mid-height temperature fields from these Fluidity-ICOM simulations are shown in figures 4.1 and 4.2. A series of increasing wavenumber solutions are observed. At $\omega = 0.101$ rad s^{-1} (and also in the non-rotating $\omega = 0$ case, not shown) the solution is axisymmetric. At $\omega = 0.608$ rad s^{-1} to $\omega = 2.503$ rad s^{-1} progressively increasing mode number solutions, of modes 2, 3 and 4, are observed. The mode transition points are consistent with experimental observations (see, for example, figure 4.10). At $\omega = 2.764$ rad s^{-1} a mode 3 solution is once again observed. This is characteristic of the existence of multiple-equilibria for this system, and is again consistent with the experimental observations.

4.1.1. Comparison with MORALS

We initially perform a comparison of the Fluidity-ICOM and MORALS simulations. Since a rigid lid boundary condition is applied, it follows from the continuity equation (2.3b) that the horizontal component of the vertically averaged velocity field is non-divergent. Hence one can define a depth averaged (barotropic) stream function Ψ_D where, in cylindrical coordinates (r, θ, z):

[1]MORALS data courtesy of Dr. R. M. B. Young, Department of Physics, University of Oxford.

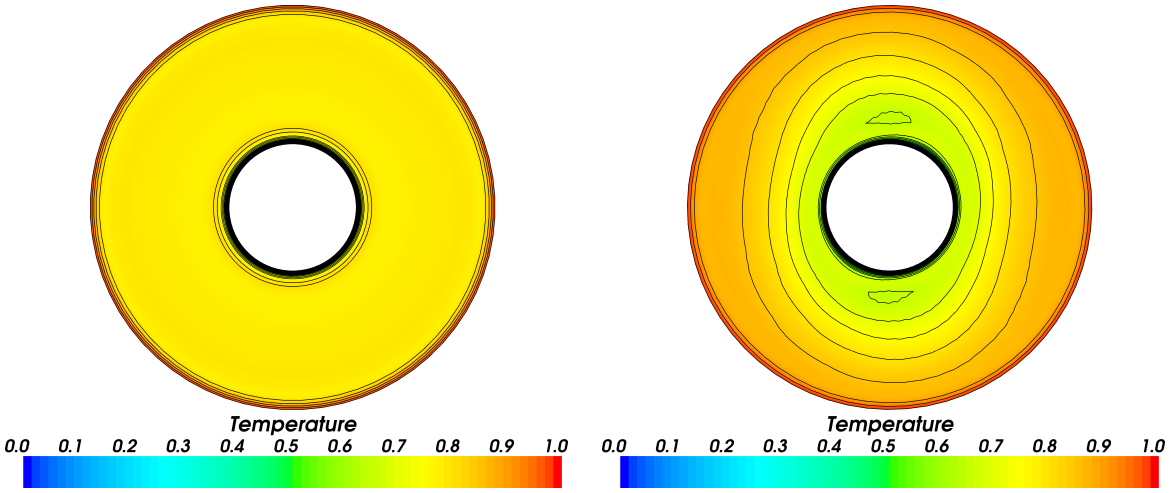

$\omega = 0.101$ rad s^{-1}, $\Theta = 59$, $Ta = 3.3 \times 10^4$ $\omega = 0.608$ rad s^{-1}, $\Theta = 1.6$, $Ta = 1.2 \times 10^6$

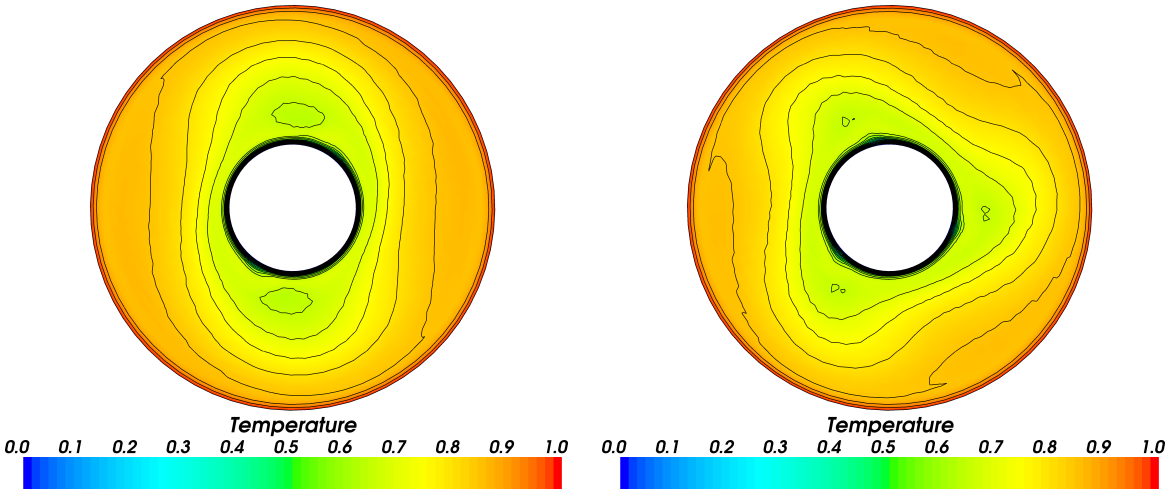

$\omega = 0.745$ rad s^{-1}, $\Theta = 1.1$, $Ta = 1.8 \times 10^6$ $\omega = 0.838$ rad s^{-1}, $\Theta = 0.85$, $Ta = 2.3 \times 10^6$

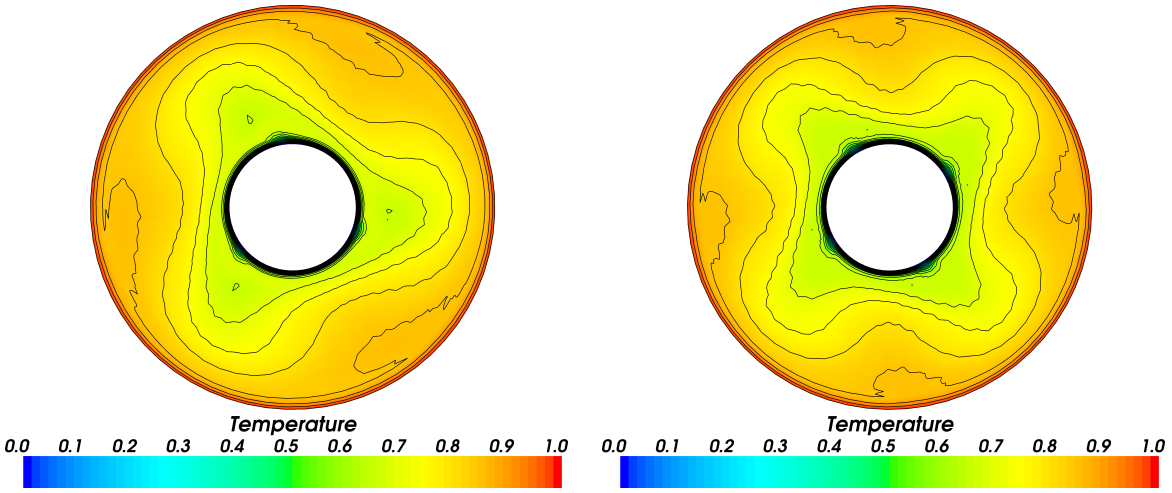

$\omega = 1.0$ rad s^{-1}, $\Theta = 0.60$, $Ta = 3.3 \times 10^6$ $\omega = 1.247$ rad s^{-1}, $\Theta = 0.39$, $Ta = 5.1 \times 10^6$

Figure 4.1.: Mid-height temperature field for fixed structured mesh Fluidity-ICOM simulations of resolution 24 × 24×64 in the radial, vertical, and azimuthal dimensions respectively at increasing rotation rates, and using parameters as given in table 2.1. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$. Θ and Ta are the thermal Rossby and Taylor numbers respectively, as defined by equations (1.5) and (1.6).

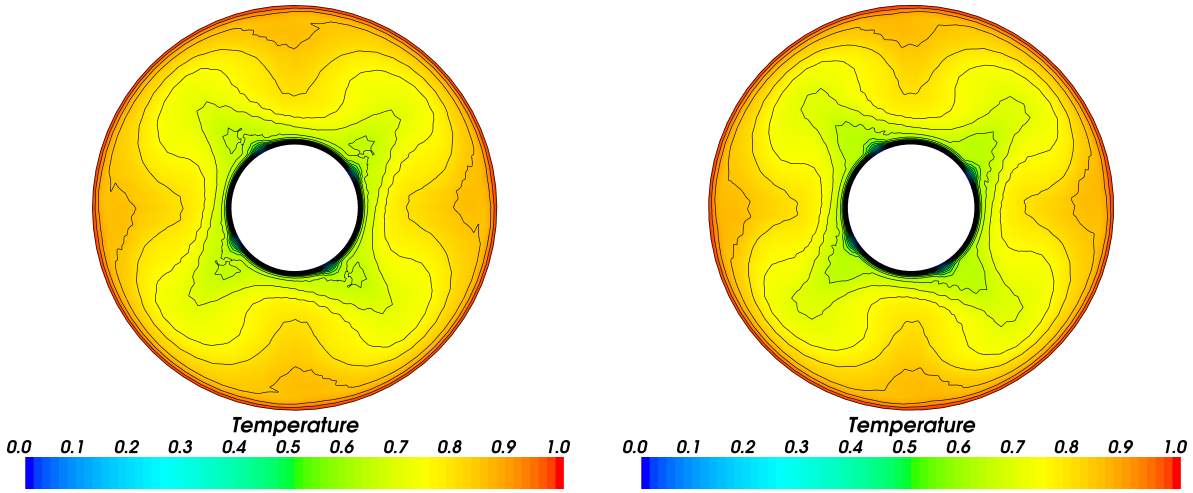

$\omega = 1.522$ rad s^{-1}, $\Theta = 0.26$, $Ta = 7.6 \times 10^6$ $\omega = 1.756$ rad s^{-1}, $\Theta = 0.19$, $Ta = 1.0 \times 10^7$

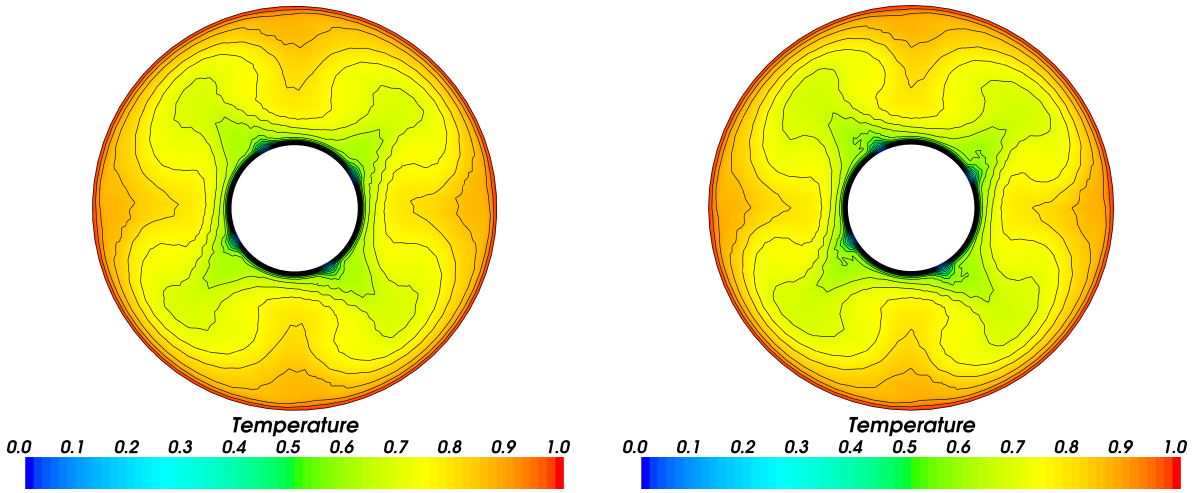

$\omega = 2.014$ rad s^{-1}, $\Theta = 0.15$, $Ta = 1.3 \times 10^7$ $\omega = 2.256$ rad s^{-1}, $\Theta = 0.12$, $Ta = 1.7 \times 10^7$

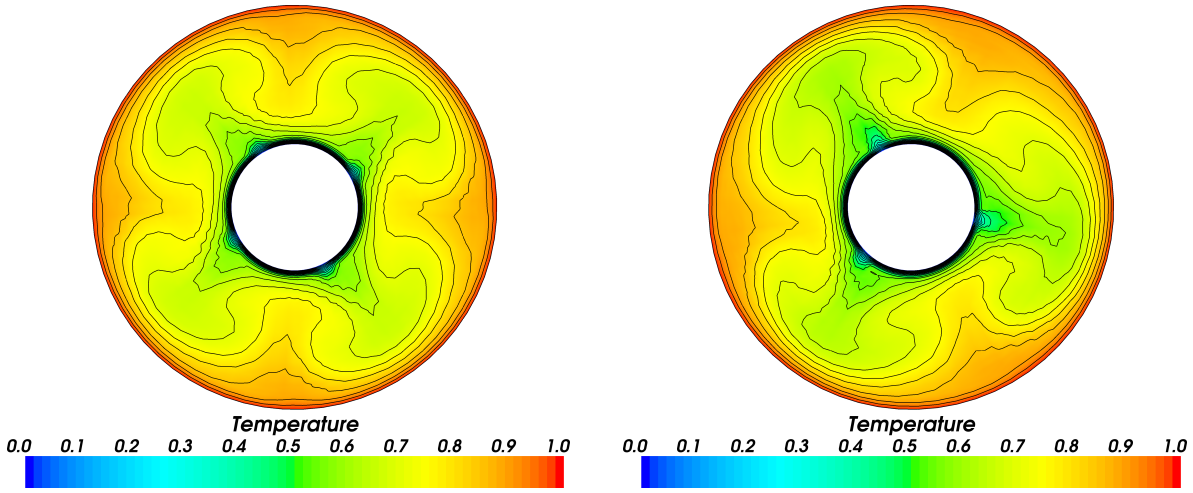

$\omega = 2.503$ rad s^{-1}, $\Theta = 0.096$, $Ta = 2.0 \times 10^7$ $\omega = 2.764$ rad s^{-1}, $\Theta = 0.078$, $Ta = 2.5 \times 10^7$

Figure 4.2.: Mid-height temperature field for fixed structured mesh Fluidity-ICOM simulations of resolution 24 × 24×64 in the radial, vertical, and azimuthal dimensions respectively at increasing rotation rates, and using parameters as given in table 2.1. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$. Θ and Ta are the thermal Rossby and Taylor numbers respectively, as defined by equations (1.5) and (1.6).

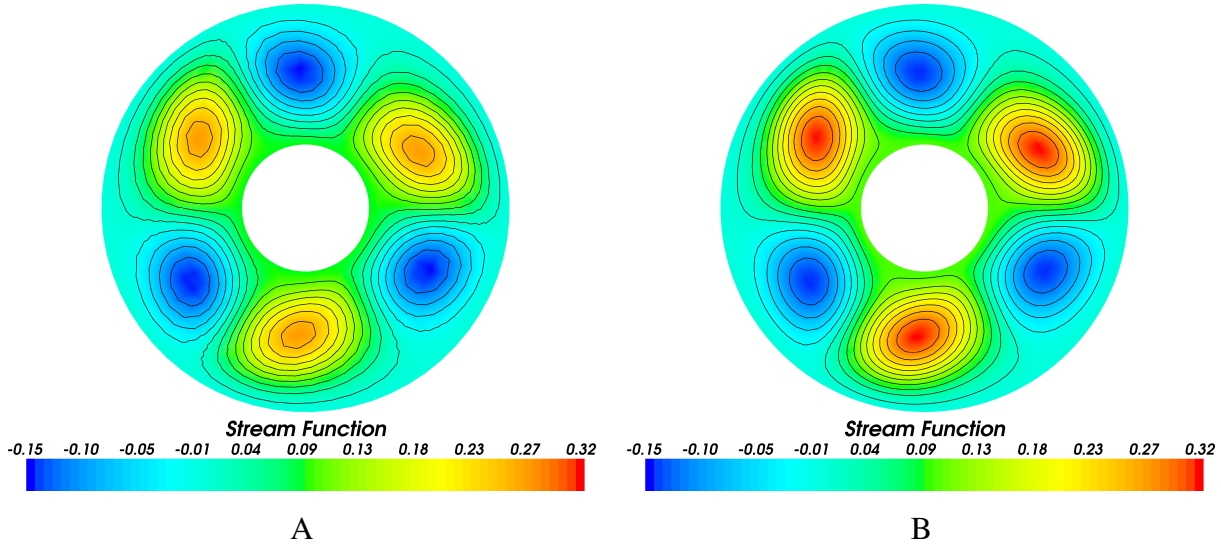

Figure 4.3.: Depth averaged stream function as defined by equation (4.1) at rotation rate $\omega = 1$ rad s^{-1}. A: Fluidity-ICOM simulation with $24 \times 24 \times 64$ nodes in the radial, vertical and azimuthal dimensions respectively. B: Corresponding MORALS simulation with $32 \times 32 \times 128$ nodes in the radial, vertical and azimuthal dimensions respectively. Stream function units are cm^2 s^{-1}. Figures shown on a common colour scale.

$$u_{D,r} = \frac{1}{r}\frac{\partial \Psi_D}{\partial \theta}, \qquad u_{D,\theta} = -\frac{\partial \Psi_D}{\partial r}, \tag{4.1}$$

and where $\boldsymbol{u}_D$ is the horizontal component of the depth averaged velocity:

$$(\boldsymbol{u}_D)^T = \left(\frac{\int_{z=0}^{z=d} u_r dz}{d}, \frac{\int_{z=0}^{z=d} u_\theta dz}{d}, 0 \right). \tag{4.2}$$

Here $\boldsymbol{u}_D$ is defined on the domain Ω_L forming the upper lid surface of the annular chamber. The computation of Ψ_D using a continuous Galerkin method, together with a discussion of the treatment of the boundary conditions, is presented in appendix A.5.

The depth averaged stream function Ψ_D, computed using a P_1 discretisation, is shown for the $\omega = 1.0$ rad s^{-1} and $\omega = 2.764$ rad s^{-1} Fluidity-ICOM and MORALS simulations in figures 4.3 and 4.4. At $\omega = 1.0$ rad s^{-1} there is a good qualitative agreement between the two models. This agreement is less good for the $\omega = 2.764$ rad s^{-1} case, with the eddies being considerably more anisotropic in the Fluidity-ICOM simulation. In addition, at $\omega = 2.764$ rad s^{-1} the depth averaged stream function is observed to have a greater magnitude in the MORALS simulation. This indicates that the baroclinic eddies are significantly weaker in the Fluidity-ICOM simulation relative to those observed in MORALS.

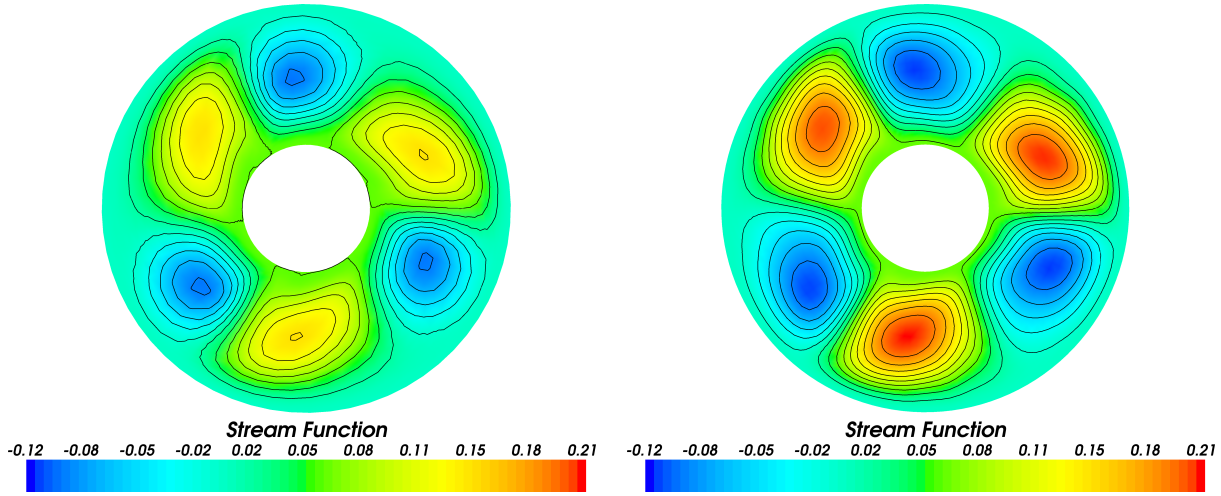

Figure 4.4.: Depth averaged stream function as defined by equation (4.1) at rotation rate $\omega = 2.764$ rad s^{-1}. A: Fluidity-ICOM simulation with $24 \times 24 \times 64$ nodes in the radial, vertical and azimuthal dimensions respectively. B: Corresponding MORALS simulation with $32 \times 32 \times 128$ nodes in the radial, vertical and azimuthal dimensions respectively. Stream function units are cm^2 s^{-1}. Figures shown on a common colour scale.

Due to the multiply connected topology of the annular domain, it again follows from the continuity equation (2.3b) that the meridional component of the azimuthally averaged velocity is divergence free. Hence one can define an azimuthally averaged (meridional) stream function Ψ_A where, in cylindrical coordinates:

$$u_{A,r} = \frac{1}{r}\frac{\partial \Psi_A}{\partial z}, \qquad u_{A,z} = -\frac{1}{r}\frac{\partial \Psi_A}{\partial r}, \tag{4.3}$$

and where $\boldsymbol{u}_A$ is the meridional component of the azimuthally averaged velocity:

$$(\boldsymbol{u}_A)^T = \left(\frac{\int_{\theta=0}^{\theta=2\pi} u_r d\theta}{2\pi}, 0, \frac{\int_{\theta=0}^{\theta=2\pi} u_z d\theta}{2\pi} \right). \tag{4.4}$$

$(\boldsymbol{u}_A)^T$ is defined on the domain Ω_S consisting of a vertical cross-section of the annular chamber, with $a \leq r \leq b$ and $0 \leq z \leq d$. The computation of Ψ_A using a continuous Galerkin method is presented in appendix A.6. Note that the azimuthally averaged stream function, as defined by equation (4.3), is a Stokes stream function [Batchelor, 1967]. In particular, note that the dimensions of the azimuthally averaged stream function and depth averaged stream function differ, due to the factor $1/r$ in the former.

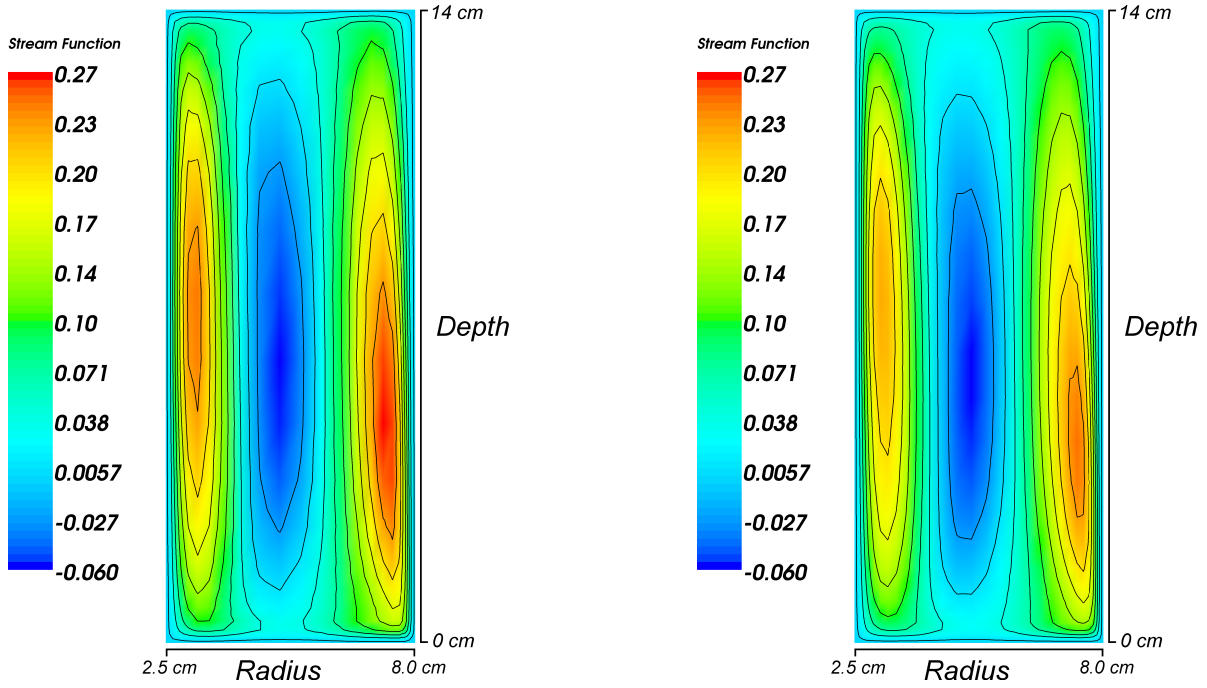

Figure 4.5.: Azimuthally averaged stream function as defined by equation (4.3) at rotation rate $\omega = 1.0$ rad s^{-1}. A: Fluidity-ICOM simulation with $24 \times 24 \times 64$ nodes in the radial, vertical and azimuthal dimensions respectively. B: Corresponding MORALS simulation with $32 \times 32 \times 128$ nodes in the radial, vertical and azimuthal dimensions respectively. Stream function units are cm^3 s^{-1}. Figures shown on a common colour scale.

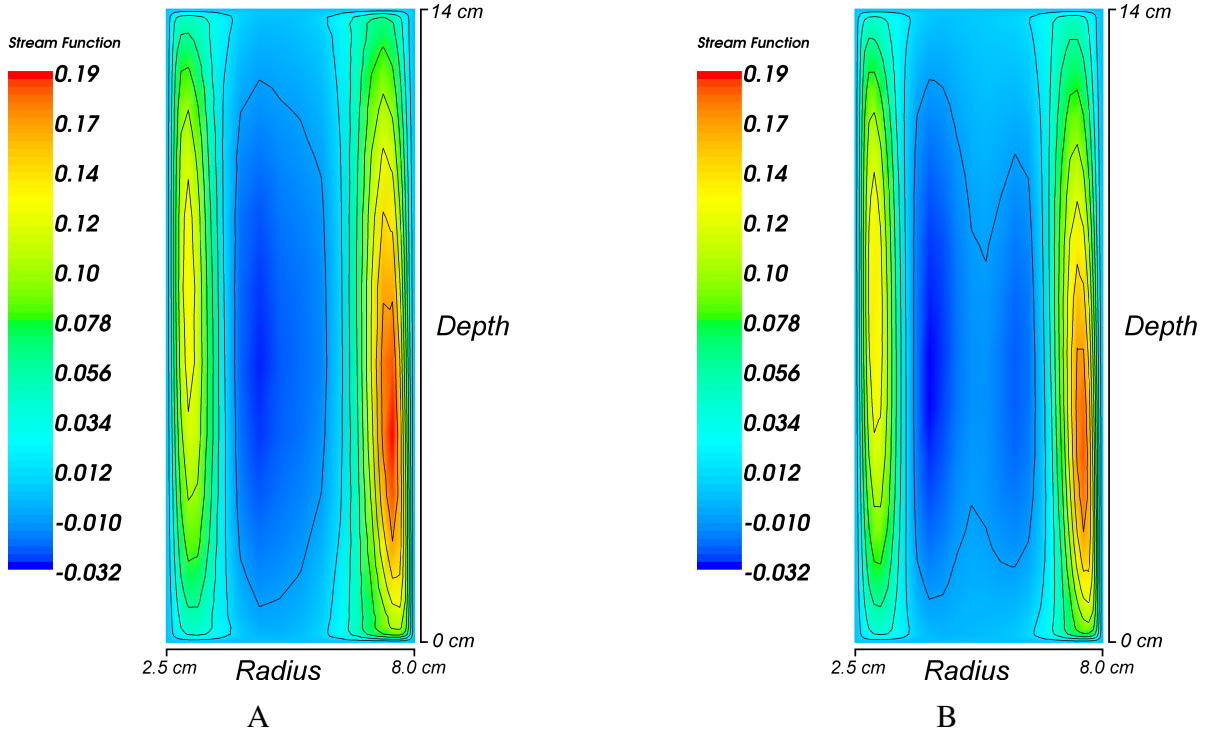

Figure 4.6.: Azimuthally averaged stream function as defined by equation (4.3) at rotation rate $\omega = 2.764$ rad s^{-1}. A: Fluidity-ICOM simulation with $24 \times 24 \times 64$ nodes in the radial, vertical and azimuthal dimensions respectively. B: Corresponding MORALS simulation with $32 \times 32 \times 128$ nodes in the radial, vertical and azimuthal dimensions respectively. Stream function units are cm^3 s^{-1}. Figures shown on a common colour scale.

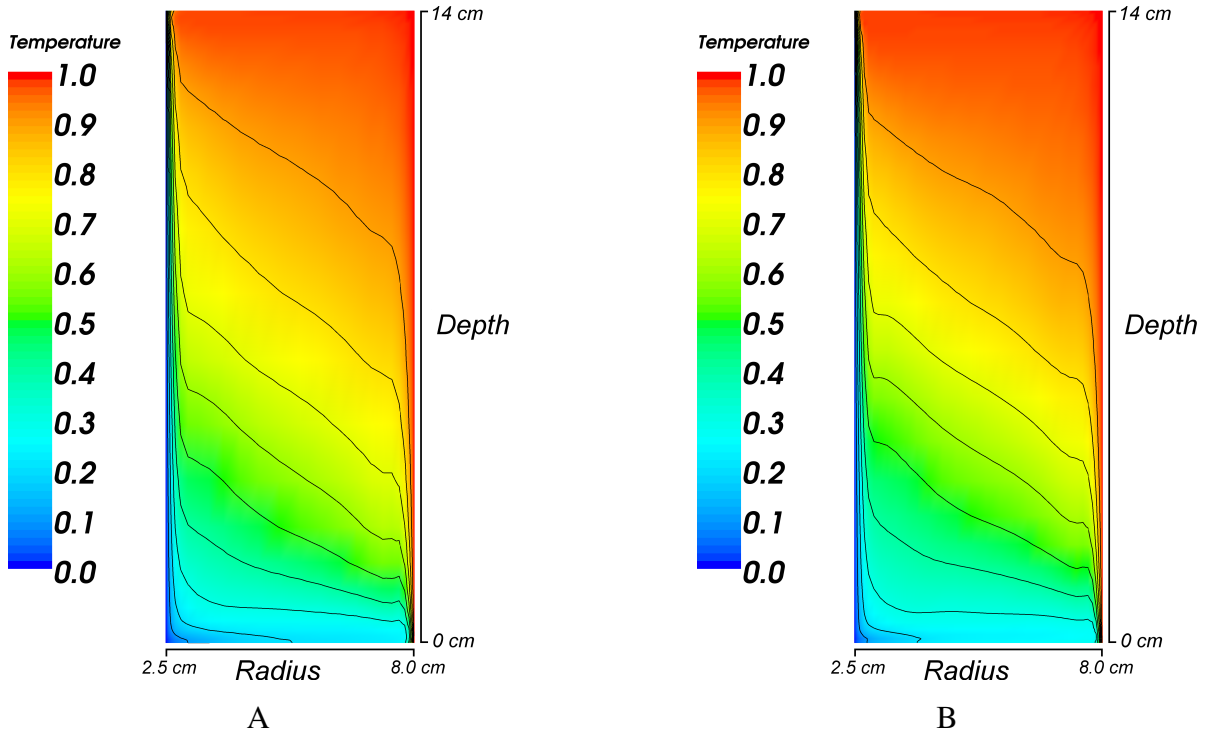

Figure 4.7.: Azimuthally averaged temperature field at rotation rate $\omega = 1.0$ rad s^{-1}. A: Fluidity-ICOM simulation with $24 \times 24 \times 64$ nodes in the radial, vertical and azimuthal dimensions respectively. B: Corresponding MORALS simulation with $32 \times 32 \times 128$ nodes in the radial, vertical and azimuthal dimensions respectively. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$.

The azimuthally averaged stream function Ψ_A, computed using a P_1 discretisation, is shown for the $\omega = 1.0$ rad s^{-1} and $\omega = 2.764$ rad s^{-1} Fluidity-ICOM and MORALS simulations in figures 4.5 and 4.6. At both rotation rates two separate circulations, located towards the inner and outer wall boundary regions, are observed. These correspond to direct thermally driven circulations [Jacoby et al., 2011]. A much weaker indirect thermally driven Ferrel-type cell is observed towards the tank centre at both rotation rates. There is an excellent qualitative agreement between MORALS and Fluidity-ICOM at $\omega = 1.0$ rad s^{-1}. At $\omega = 2.764$ rad s^{-1} there are some deviations. In particular, for the MORALS simulation, two distinct minima in the azimuthally averaged stream function are observed in the interior region. This additional interior structure is not seen for the Fluidity-ICOM simulation.

The azimuthally averaged temperature fields are shown for the $\omega = 1.0$ rad s^{-1} and $\omega = 2.764$ rad s^{-1} Fluidity-ICOM and MORALS simulations in figures 4.7 and 4.8. At $\omega = 1.0$ rad s^{-1} a clear thermal boundary layer is observed at the inner and outer side walls, with sloping isopycnals in the interior. At this rotation rate there is excellent agreement between Fluidity-ICOM and MORALS, with deviations in the azimuthally averaged temperature limited primarily to the near side-wall boundary layer region and the lower Ekman layer. These deviations are attributed to the increased mesh resolution used for the MORALS simulation. At $\omega = 2.764$ rad s^{-1} the azimuthally

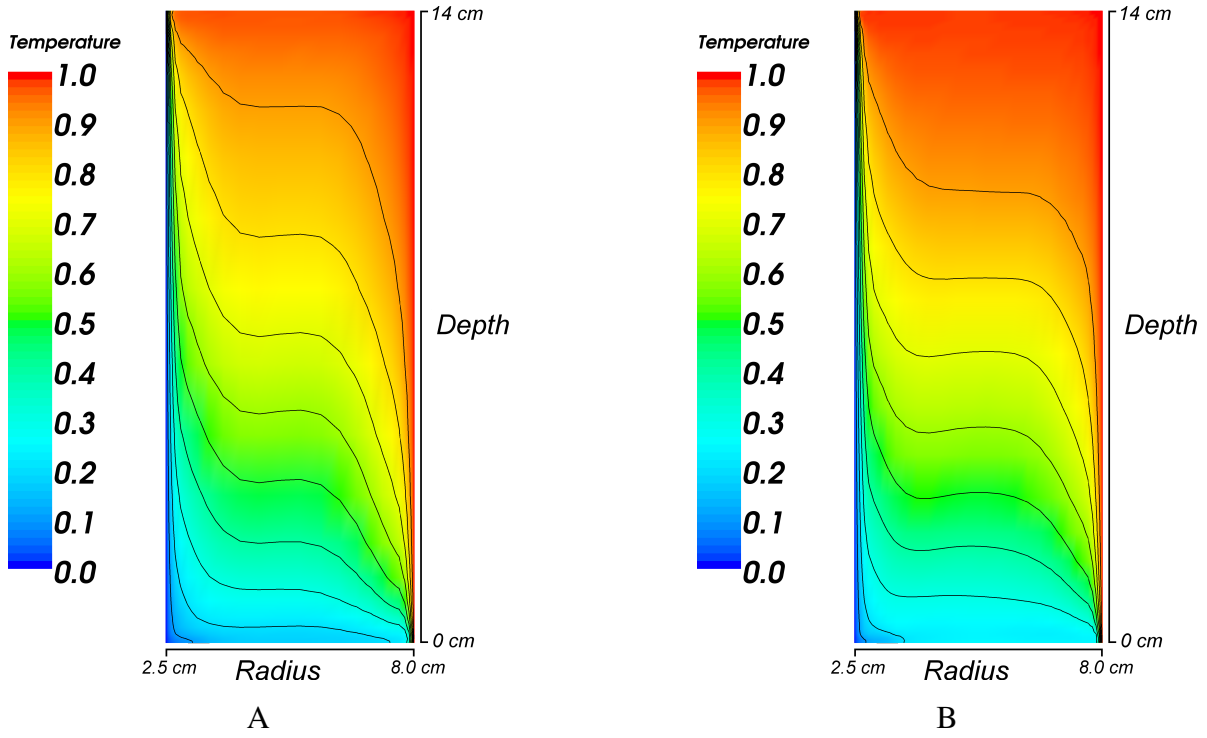

Figure 4.8.: Azimuthally averaged temperature field at rotation rate (4.3) at rotation rate $\omega = 2.764$ rad s^{-1}. A: Fluidity-ICOM simulation with $24 \times 24 \times 64$ nodes in the radial, vertical and azimuthal dimensions respectively. B: Corresponding MORALS simulation with $32 \times 32 \times 128$ nodes in the radial, vertical and azimuthal dimensions respectively. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$.

averaged temperature contours (and hence the isopycnals) are near horizontal in the interior. The agreement between the MORALS and Fluidity-ICOM simulations is less good. In particular, the agreement is poor in the region of the lower Ekman layer.

At $\omega = 1.0$ rad s^{-1} the final tank averaged temperature is $(0.7316\Delta T)+T_A$ for the Fluidity-ICOM simulation and $(0.7327\Delta T) + T_A$ for the MORALS simulation. For $\omega = 1.0$ rad s^{-1} Fluidity-ICOM and MORALS agree on the tank averaged temperature to within $10^{-3}\Delta T$. At $\omega = 2.764$ rad s^{-1} the final tank averaged temperature is $(0.6919\Delta T) + T_A$ for the Fluidity-ICOM simulation and $(0.7141\Delta T) + T_A$ for the MORALS simulation. Hence, at $\omega = 2.764$ rad s^{-1}, there is a much more significant disagreement between Fluidity-ICOM and MORALS regarding the heat content of the system. This contributes to the relatively poor qualitative agreement between in the Fluidity-ICOM and MORALS azimuthally averaged temperature profiles.

The potential energy distributions are shown for the $\omega = 1.0$ rad s^{-1} and $\omega = 2.764$ rad s^{-1} Fluidity-ICOM and MORALS simulations in figure 4.9. At $\omega = 1.0$ rad s^{-1} the potential energy distributions are almost coincident, demonstrating the excellent agreement between the two models at this rotation rate. At $\omega = 2.764$ rad s^{-1} there are significant differences between the two models.

Specifically, the potential energy distribution of the Fluidity-ICOM simulation is, relative to the MORALS simulation, biased towards higher potential energies.

The baroclinic eddies act to release potential energy from the system, and act to tilt the isopycnals towards the horizontal (more precisely, towards the geopotential contours). At $\omega = 2.764$ rad s^{-1} the depth averaged stream function indicates that the eddies observed in the Fluidity-ICOM simulation are weaker than those observed in the corresponding MORALS simulation. This leads to an increased system potential energy for the Fluidity-ICOM simulation, as observed in the potential energy distribution. Due to the annular geometry, the heat content of the rotating annulus system is lower for a state with linearly stratified horizontally aligned temperature contours than for a diffusion dominated state with vertically aligned temperature contours. Hence, at $\omega = 2.764$ rad s^{-1}, the lower tank average temperature observed in the Fluidity-ICOM simulation, relative to that observed for the corresponding MORALS simulation, is again consistent with weaker baroclinic eddies.

Hence there is good agreement between Fluidity-ICOM and MORALS at $\omega = 1.0$ rad s^{-1}. At $\omega = 2.764$ rad s^{-1} the agreement is less good, with the Fluidity-ICOM simulation exhibiting weaker baroclinic eddies.

4.1.2. Comparison with laboratory data

The standard annulus geometry used for these simulations, with $a = 2.5$cm, $b = 8.0$cm, $d = 14.0$cm and with a rigid lid, has been the subject of a number of experimental studies. In particular, for these fluid parameters (notwithstanding the difference in viscosity noted in section 2.3.1), laboratory data for a number of thermal forcings and rotation rates are available – see Hignett et al. [1985], Hignett [1985] and Read [2003] for further details. These data include measurements of the system heat transport. Heat transport measurements were obtained using calorimetric measurements of the thermal forcing, and have a quoted error of ±2.5% [Read, 2003]. Mid-radius mid-height temperature measurements are also available, obtained from an in-situ thermocouple array.

Data for a total of 239 separate experiments in this annular geometry and with these fluid parameters were retrieved from the laboratory data archive. One of these experiments was conducted with a thermal forcing of $\Delta T = 2$ K, 98 with a thermal forcing of $\Delta T = 4$ K, and 140 with a thermal forcing of $\Delta T = 10$ K. For the $\Delta T = 4$ K experiments, 73 contained measurements of the heat transport, and 96 contained mid-radius mid-height temperature measurements from an in-situ ther-

A

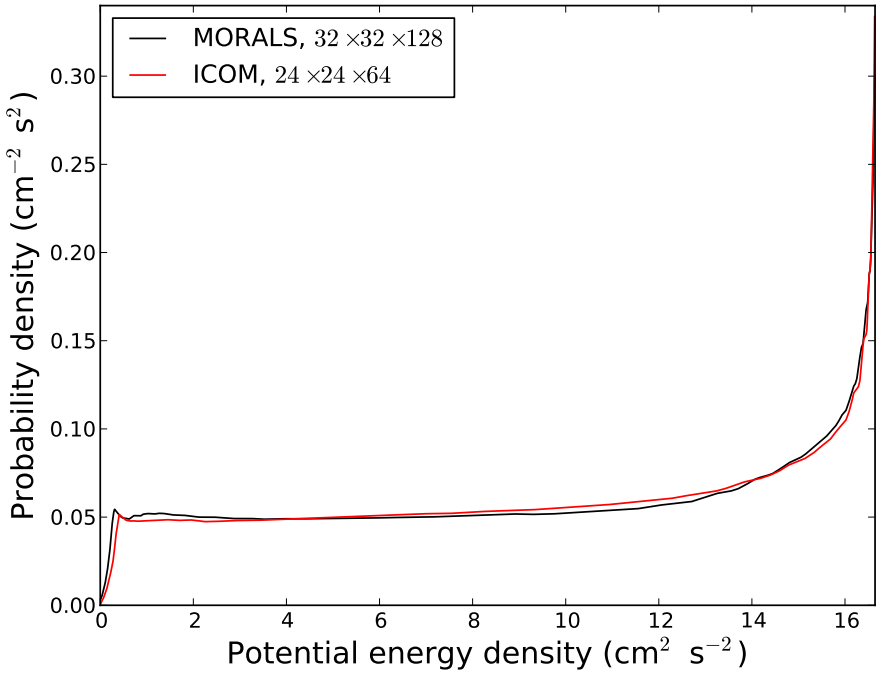

B

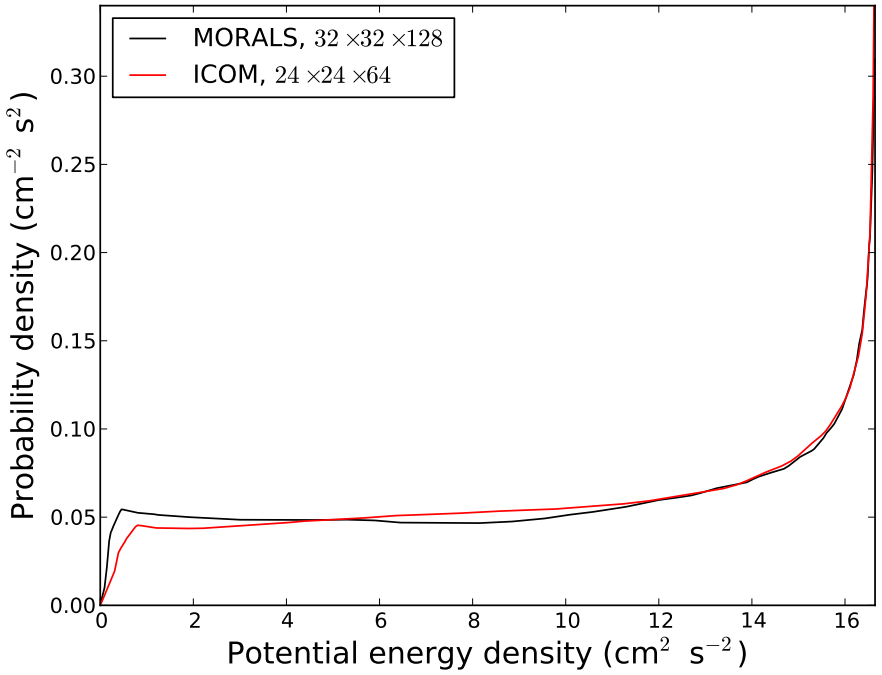

Figure 4.9.: Probability density function for the potential energy density $U = -g\alpha(T - T_A)z$ in the rotating annulus, simulated using Fluidity-ICOM with a fixed mesh at resolution 24×24×64 nodes in the radial, vertical and azimuthal dimensions respectively, and using MORALS at a resolution $32 \times 32 \times 128$ nodes in the radial, vertical and azimuthal dimensions respectively. The background potential energy density gz is excluded, and the potential energy density zero point chosen so that $\min(U) = 0$. For a randomly selected fluid parcel, the probability of the parcel having a potential energy density between U and $U + dU$ is $\mathcal{P}(U)\,dU$, where $\mathcal{P}(U)$ is the probability density function. A is at rotation rate $\omega = 1.0$ rad s^{-1}and B at rotation rate $\omega = 2.764$ rad s^{-1}.

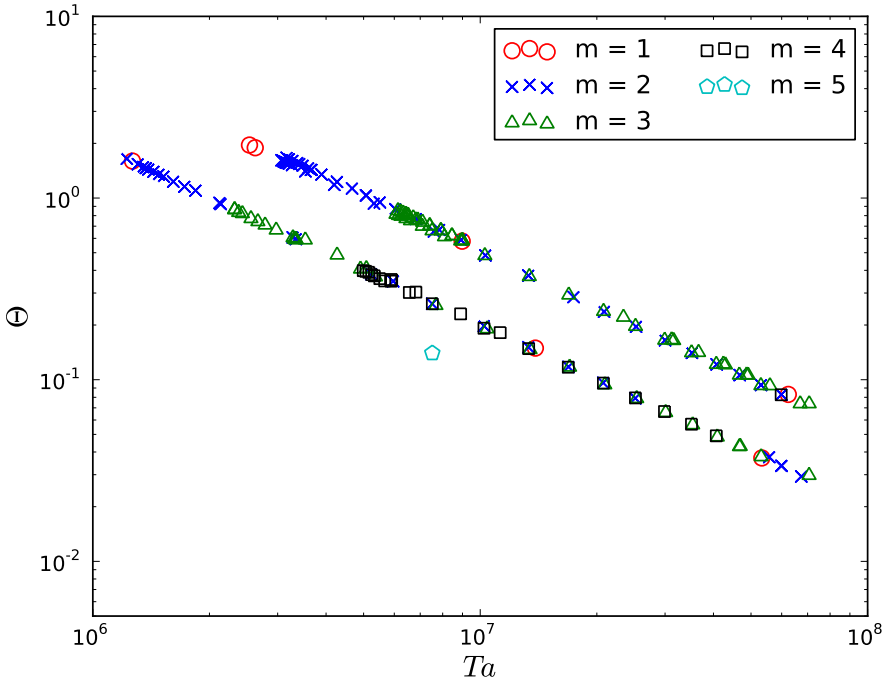

Figure 4.10.: Sketch regime diagram showing the mid-radius mid-height dominant Fourier modes, derived from the experimental thermocouple array data. Here m is the number of the dominant Fourier mode. Θ and Ta are the thermal Rossby and Taylor numbers respectively, as defined by equations (1.5) and (1.6). The single $m = 5$ experiment has a thermal forcing of $\Delta T = 2$K, the lower line of experiments $\Delta T = 4$K, and the upper line of experiments $\Delta T = 10$K.

mocouple array. For the $\Delta T = 10$ K experiments, 71 contained measurements of the heat transport, and 132 contained mid-radius mid-height temperature measurements from an in-situ thermocouple array. The $\Delta T = 2$ K experiment contained both heat transport and thermocouple array measurements. The dominant Fourier modes for these experiments are shown in figure 4.10.

From these data, four key diagnostics were computed. The non-dimensional Nusselt number, expressing the heat transport relative to the heat transport of a solid of equal conductivity to the fluid, was calculated from the measurements of the system heat flow. Expressed in terms of the fluid temperature field this takes the form as per equation (2.46). Expressed in terms of a boundary heat flow this takes the form [Read, 2003]:

$$Nu = \frac{\ln b/a}{2\pi d\Delta T}\frac{H_I}{\kappa\rho c}, \tag{4.5}$$

where κ is the fluid diffusivity, ρ the fluid density, c the specific heat capacity, and H_I the heat flow (a power) through the tank inner wall. For consistency with the definition in equation (2.46), a negative Nusselt number corresponds to flow out of the domain through the inner wall. The

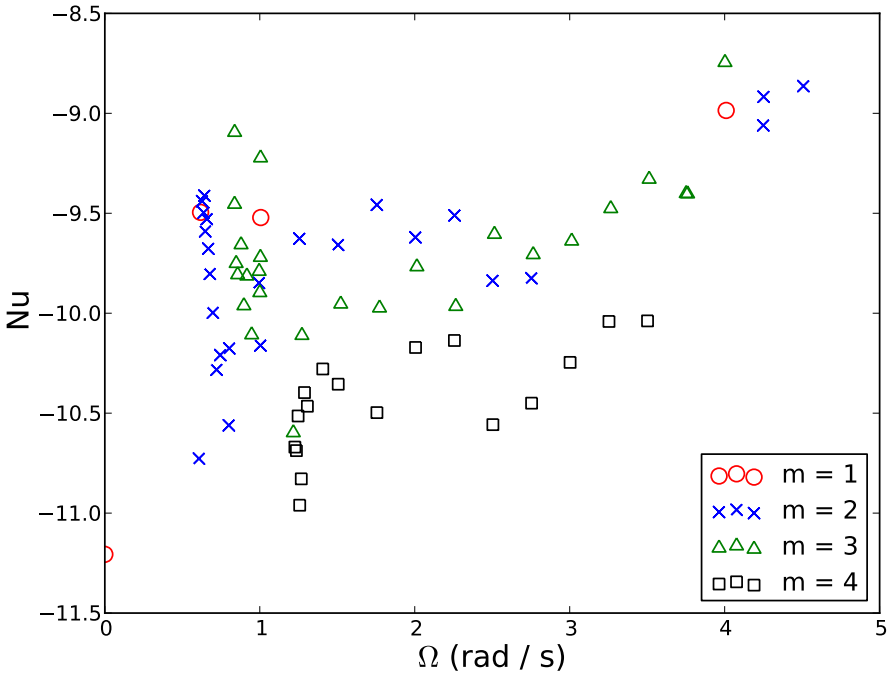

Figure 4.11.: Experimental heat transport measurements for heat flow data extracted from the laboratory data archive (without error bars), for experiments with a thermal forcing of $\Delta T = 4$ K. Here m is the number of the dominant mid-radius mid-height temperature Fourier mode.

measured heat transports for the $\Delta T = 4$ K experiments are shown in figure 4.11. Note that, for a given rotation rate, the maximum possible heat transport of the system is remarkably constant [Hide and Mason, 1975, Read, 2003]. Also note that, for a given mode number, as the rotation rate is increased the heat transport initially increases, reaches a maximum, and then decreases. At higher rotation rates, higher mode numbers are associated with increased heat transport. Hence baroclinic instability can be viewed as a ventilation mechanism, whereby the heat transport of the system is maximised, and maintained at close to the non-rotating value.

A Fourier decomposition of the mid-radius mid-height temperature measurements was used to compute the mid-radius mid-height temperature mode amplitudes. An example of this Fourier analysis is shown in figure 4.12. For systems exhibiting variation of the dominant mode amplitude with time (for example, due to amplitude vacillation), the period and amplitude of vacillation were computed via a least squares harmonic fit, via calculation of a Lomb-Scargle periodogram [Scargle, 1982, Press and Rybicki, 1989]. Note that previous studies (see, for example, Young and Read [2008a]) have calculated the magnitude of amplitude vacillation via measurements of the peak-to-peak variation of the dominant mode amplitude, rather than the amplitude of a least squares fit. From this calculation of the magnitude of amplitude vacillation, one can define a vacillation index [Young and Read, 2008a]:

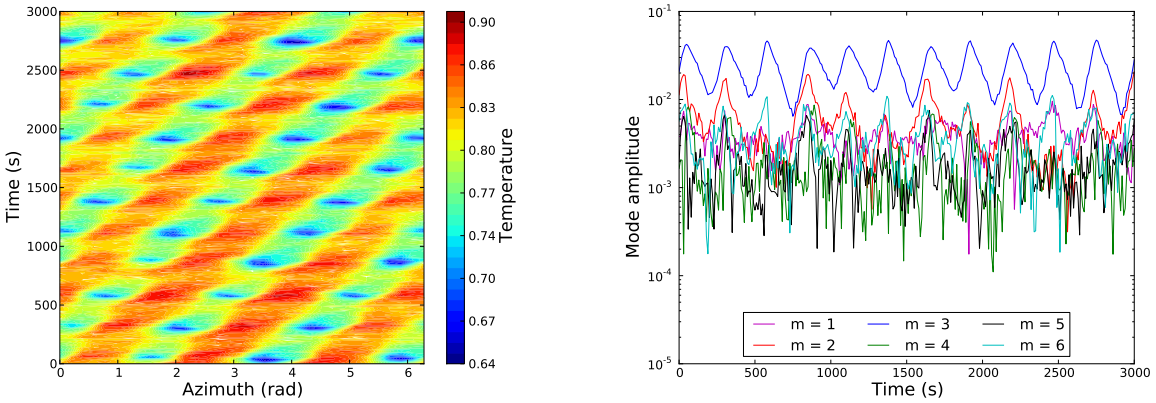

Figure 4.12.: Left: Mid-radius mid-height temperature measurements for an experiment at $\omega = 0.838$ rad s^{-1} and $\Delta T = 4$ K. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$. Right: Amplitude of Fourier modes derived from these mid-radius mid-height temperature measurements. The Fourier amplitudes show a dominant mode 3 with amplitude vacillation, and with intermittent bursts of a weaker mode 2. The laboratory noise level is at a Fourier amplitude of $\sim 10^{-3}/\Delta T$.

$$I_v = \frac{2a_{I,m}}{a_m} \tag{4.6}$$

where a_m is the mean amplitude of the dominant Fourier mode, and $a_{I,m}$ is the amplitude of the least squares harmonic fit to the time varying amplitude of the dominant Fourier mode. It follows from this definition that $0 \leq I_v \leq 2$. The measured vacillation indices, derived from the mid-radius mid-height temperature measurements for the $\Delta T = 4$ K experiments, are shown in figure 4.13. Note that increases in the vacillation index are associated with mode transition points. The vacillation index also increases at higher rotation rates, indicating the approach towards chaotic and turbulent dynamics.

The measured heat transport for fixed mesh Fluidity-ICOM simulations of the thermally driven rotating annulus are shown in figure 4.14, with laboratory observations and values derived from the MORALS simulations for comparison. Note that the laboratory error bars are primarily an indication of the experimental error [Read, 2003]. At lower rotation rates, the heat transports for the fixed mesh Fluidity-ICOM simulations are close to the experimental observations – indeed, for several of the lower rotation rate simulations the simulated heat transport is within the experimental error. At higher rotation rates the agreement is progressively less good. By contrast, the two MORALS simulations have system heat transports that are extremely close to the experimental values, for both $\omega = 1$ rad s^{-1} and $\omega = 2.764$ rad s^{-1}. The decrease in heat transport observed for the fixed mesh Fluidity-ICOM simulations, as the rotation rate is increased, could be due to a

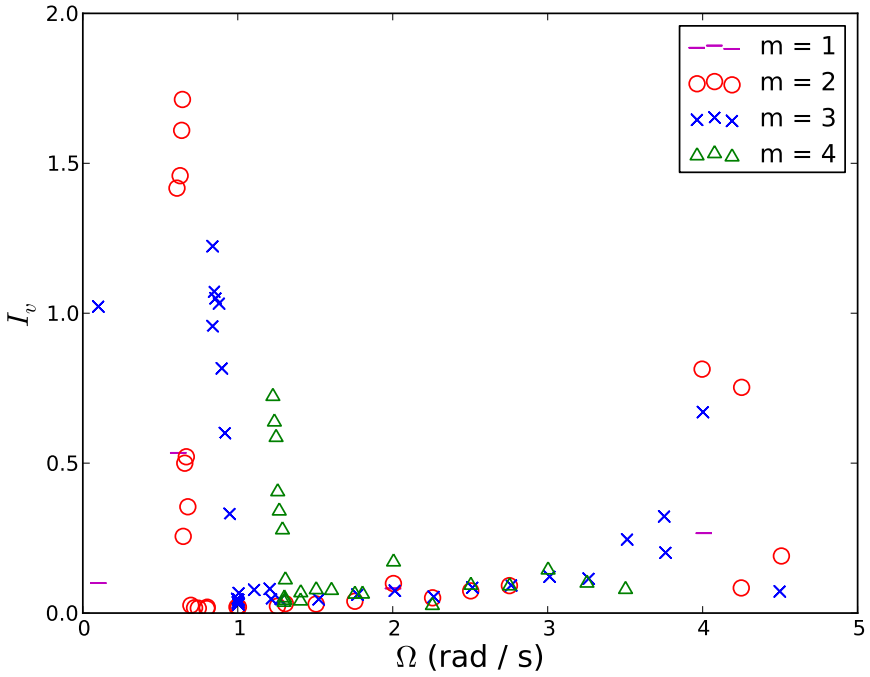

Figure 4.13.: Mid-radius mid-height temperature vacillation indices, as defined by equation (4.6), for experiments with a thermal forcing of $\Delta T = 4$ K. Here m is the number of the dominant temperature mid-radius mid-height temperature Fourier mode.

systematic weakening of the baroclinic eddies, as indicated in the comparison with MORALS in section 4.1.1.

The dominant mode amplitudes for the mid-height mid-radius temperature for fixed mesh Fluidity-ICOM simulations of the thermally driven rotating annulus are shown in figure 4.15, with laboratory observations and values derived from the MORALS simulations used in section 4.1.1 for comparison. For all except the highest rotation rate $\omega = 2.764$ rad s^{-1} case the fixed mesh Fluidity-ICOM simulations show a good agreement with the experimental observations. As for the heat transport measurements above, the dominant mode amplitudes of the two MORALS simulations show an excellent agreement with the experimental observations.

The corresponding dominant mode drift periods are shown in figure 4.16, with laboratory observations and values derived from the MORALS simulations used in section 4.1.1 for comparison. The drift periods for the fixed mesh Fluidity-ICOM simulations are generally lower than the corresponding observations, particularly at higher rotation rates, as previously observed in section 2.3.1. The drift periods for the two MORALS simulations are also consistently lower than the experimental observations.

The vacillation index for the mid-height mid-radius temperature for fixed mesh Fluidity-ICOM simulations of the thermally driven rotating annulus are shown in figure 4.17, with laboratory ob-

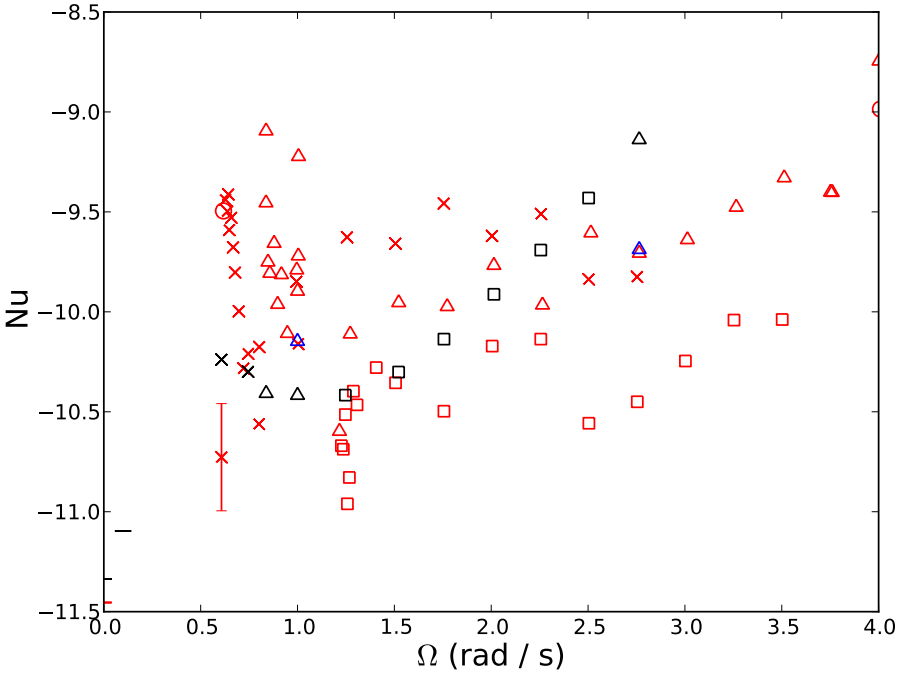

Figure 4.14.: Thermally driven rotating annulus heat transport measurements. Laboratory observations are in red (with one typical error bar), fixed mesh Fluidity-ICOM simulations in black, and MORALS simulations in blue. The model values are averaged over 1000 s. A line denotes an axisymmetric system, a circle denotes mode one, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four.

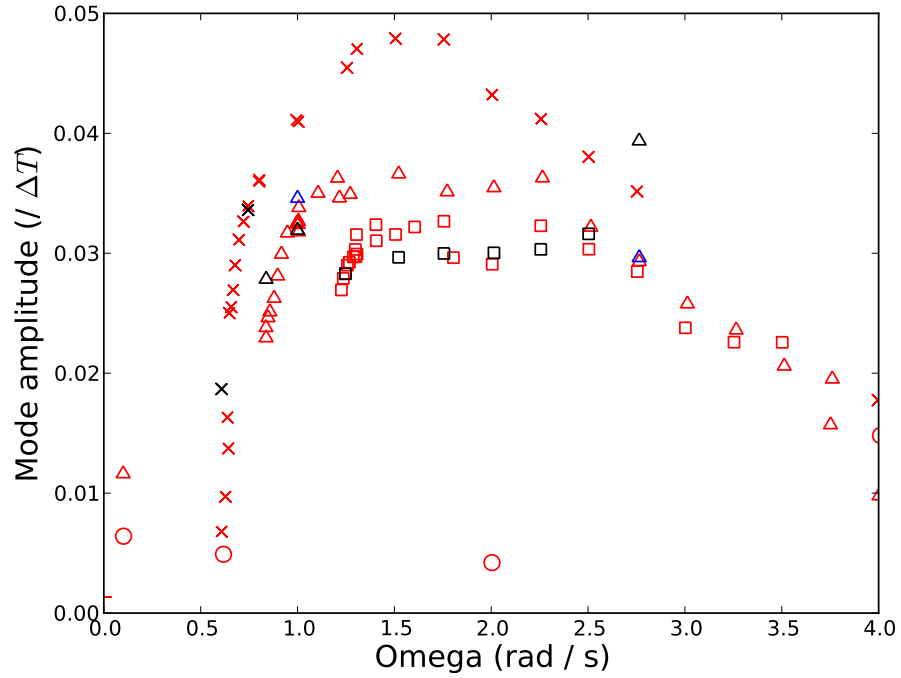

Figure 4.15.: Amplitude of the dominant Fourier mode for the mid-radius mid-height temperature. Laboratory observations are in red, fixed mesh Fluidity-ICOM simulations in black, and MORALS simulations in blue. The laboratory values are averaged over each thermocouple array data set. The model values are averaged over 1000 s. A line denotes an axisymmetric system, a circle denotes mode one, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four.

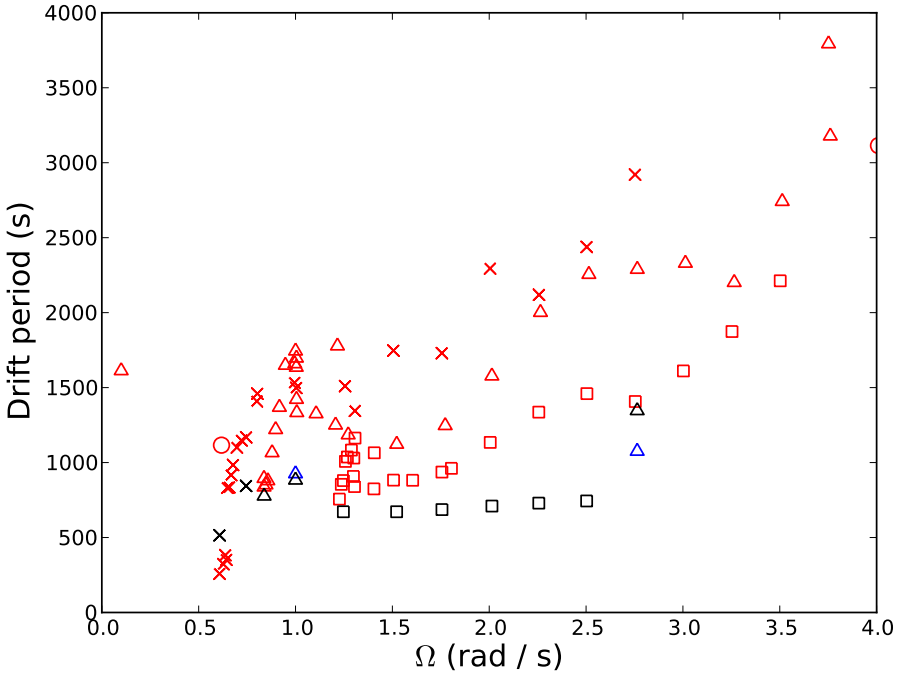

Figure 4.16.: Drift period of the dominant Fourier mode mode for the mid-radius mid-height temperature. Laboratory observations are in red, fixed mesh Fluidity-ICOM simulations in black, and MORALS simulations in blue. The laboratory values are averaged over each thermocouple array data set. An experiment with a vary large drift period (at $\omega = 4.0$ rad s^{-1}), and experiments with very small mode amplitude, are not shown. The model values are averaged over 1000 s. A line denotes an axisymmetric system, a circle denotes mode one, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four.

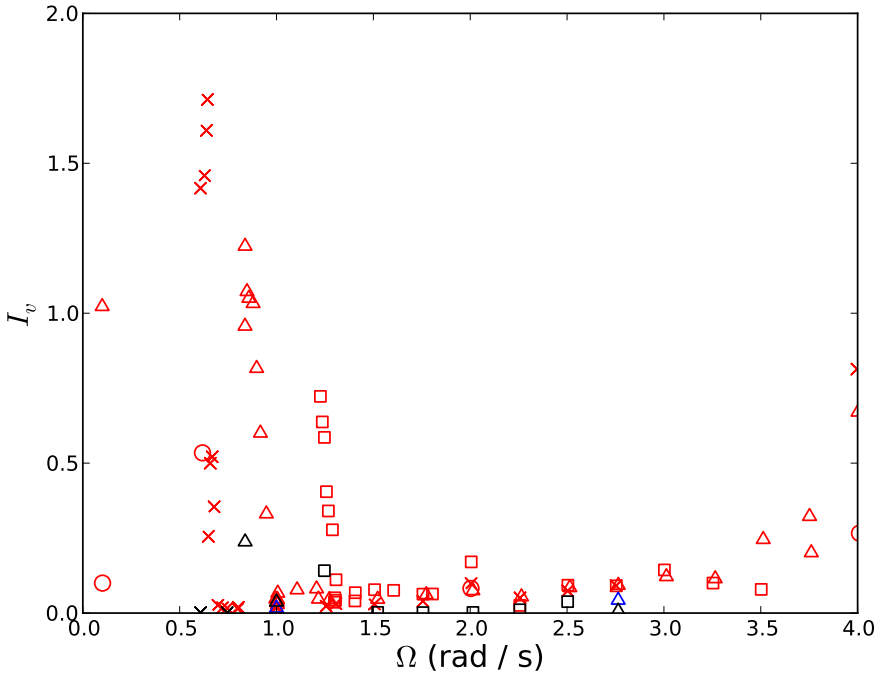

Figure 4.17.: Mid-radius mid-height temperature vacillation indices, as defined by equation (4.6). Laboratory observations are in red, fixed mesh Fluidity-ICOM simulations in black, and MORALS simulations in blue. A circle denotes mode one system, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four.

servations and values derived from the MORALS simulations used in section 4.1.1 for comparison. Two of the 13 fixed mesh Fluidity-ICOM simulations, at rotation rates of $\omega = 0.838$ rad s^{-1} and $\omega = 1.247$ rad s^{-1}, exhibit clear amplitude vacillation. The 2 simulations exhibiting amplitude vacillation have a rotation rate that corresponds well with the transition points observed in experiment. However, the fixed mesh Fluidity-ICOM simulations show significantly weaker amplitude vacillation than observed in the experiments, with correspondingly lower vacillation indices. The two MORALS simulations are not in an amplitude vacillating regime, and hence no comparison can be made for this model using these data. Note that highly complex vacillating regimes have been observed in previous MORALS simulations with differing parameters (see, for example, Young and Read [2008a]).

Hence fixed mesh Fluidity-ICOM simulations of the thermally driven rotating annulus with a rigid lid, with fluid parameters as in table 2.1, and for rotation rates $\omega = 0$ case, to $\omega = 2.764$ rad s^{-1}, have a reasonable qualitative agreement with experiment. The model is able to represent baroclinic instability, with the dominant modes observed corresponding well to those observed in experiment. The system heat transport corresponds well to experimental values at lower rotation rates, and progressively less well at higher rotation rates. The dominant mode amplitude for the mid-radius mid-height temperature corresponds well to experimental values. The regimes

observed in the Fluidity-ICOM simulations are significantly more steady than those observed in experiment, with only relatively weak amplitude vacillation observed.

4.2. Adaptive mesh Fluidity-ICOM configuration

The previous section compared fixed mesh Fluidity-ICOM simulations of the thermally driven rotating annulus with laboratory observations and MORALS simulations. While Fluidity-ICOM agrees reasonably well with experiment at lower rotation rates, it agrees less well at higher rotation rates. Hence there is, at higher rotation rates, scope for improvement of the model by the application of dynamic mesh adaptivity.

The Fluidity-ICOM simulations in the previous section were conducted on a hyperbolic tangent stretch mesh derived from the mesh used by MORALS. Since MORALS compares well against laboratory data at low rotation rates, this is a challenging test for dynamic mesh adaptivity – the method must seek an improvement over a mesh that is already known, with MORALS, to lead to a good representation of the system.

There are a large number of free parameters involved in configuring the use of mesh adaptivity in Fluidity-ICOM simulations. Broadly speaking, there are four main areas to configure: the frequency of mesh optimisation, the mesh quality metric, the selection of an optimisation scheme, and the choice of a mesh-to-mesh interpolant for the system prognostic fields.

In section 4.2.1 the period of mesh optimisation is considered. In section 4.2.2 the metric formulation is described, and in section 4.2.3 interpolants are tested and compared. Finally, in section 4.2.4 the chosen configuration for thermally driven rotating annulus simulations is outlined.

4.2.1. Frequency of mesh optimisation

The frequency of mesh optimisation in oceanographic applications is discussed in Munday et al. [2010]. It is noted that the frequency of mesh optimisation is unlikely to affect results when applied to systems with steady state solutions. By contrast, systems with non-steady solutions are likely to be sensitive to this parameter. In Alauzet et al. [2003] an iterative adaption procedure for simulations of transient phenomena is described. In this procedure, an optimised mesh is generated at a time $t + \Delta t$, based upon the fields over a time t to $t + \Delta t$, and is used to restart the simulation at time t. This procedure is iterated to convergence, to within a user specified tolerance. Due to the in-

creased computational cost of repeated time integrations, this method is not applied here. Instead, the mesh is optimised with a fixed frequency, using a metric based upon the fields at the time of mesh optimisation. Note, therefore, that the the mesh optimisation lags behind the evolving flow structure.

In choosing a frequency of mesh optimisation, a compromise must be made between the need to maintain a mesh optimised to the current model dynamics, and the additional cost and interpolation artefacts associated with the optimisation procedure itself. This parameter is not tested in this chapter. For all mesh adaptive simulations discussed in this chapter, the mesh was adapted after every 20 model timesteps.

4.2.2. Metric formulation

As described in section 1.1.3, the mesh optimisation library is controlled via a mesh quality metric M, defining the size of elements in the optimised space. A displacement $\boldsymbol{v}$ has, in this metric, a length $\sqrt{(\boldsymbol{v}^T M \boldsymbol{v})}$. Here we consider a metric tensor derived from estimates of the current field interpolation error. More advanced goal-driven metrics, constraining the error in a particular goal via the use of an adjoint model (for example, as in Power et al. [2006]), were not available using Fluidity-ICOM at the time of this project.

In one dimension, the interpolation of a (sufficiently smooth) continuous field q by collocation, to yield a discrete P_1 field q^δ, has an associated ∞-norm interpolation error bounded by [Davis, 1975, Axelsson and Barker, 1984]:

$$\epsilon_e \leq \hat{\gamma} h_e^2 \max_{x' \in e} \left| \left. \frac{\partial^2 q}{\partial x^2} \right|_{x'} \right|, \tag{4.7}$$

where ϵ_e is the ∞-norm interpolation error for element e, $\hat{\gamma}$ is a mesh independent constant, and h_e is the element size. In multiple dimensions this becomes [Axelsson and Barker, 1984, Alauzet et al., 2006]:

$$\epsilon_e \leq \hat{\gamma} \max_{\boldsymbol{x} \in e} \left[\max_{\boldsymbol{v} \in \{\boldsymbol{x}_1 - \boldsymbol{x}_2 | \boldsymbol{x}_1, \boldsymbol{x}_2 \in e\}} \left(\boldsymbol{v}^T \left| H(\boldsymbol{x}) \right| \boldsymbol{v} \right) \right], \tag{4.8}$$

where ϵ_e is the ∞-norm interpolation error for element e, $\boldsymbol{x}$ is a point in e, $\boldsymbol{v}$ is a displacement between two points in e, and H is the Hessian (matrix of second derivatives) of q. $|H|$ is a matrix with the same eigenvectors as H, and with the corresponding eigenvalues equal to the absolute value of the eigenvalues of H. From this it can shown that the mesh chosen to interpolate q to yield a P_1 field q^δ, such that the number of nodes in the mesh is minimised, has edges $\boldsymbol{h}_i$ such that [Piggott et al., 2005, Alauzet et al., 2006]:

$$|\boldsymbol{h}_i| = \left(\frac{\boldsymbol{h}_i^T \, |H| \, \boldsymbol{h}_i}{\hat{\epsilon}} \right)^{-\frac{1}{2}}, \tag{4.9}$$

and where the $\boldsymbol{h}_i$ are aligned in the direction of the eigenvectors of H. $\hat{\epsilon}$ is a parameter determining the number of nodes in the mesh, with $N \propto \hat{\epsilon}^{-\frac{2}{D}}$ in D dimensions [Alauzet et al., 2006]. From this it follows that the associated metric tensor is:

$$M = \frac{|H|}{\hat{\epsilon}}. \tag{4.10}$$

Note that ∞-norm interpolation error driven mesh adaptivity is therefore equivalent to field curvature driven adaptivity.

More generally, if we instead consider a metric that acts to constrain the L_p norm interpolation error of q, the metric tensor becomes, in D dimensions [Alauzet et al., 2006]:

$$M = \det(H)^{\frac{-1}{2p+D}} \, \frac{|H|}{\hat{\epsilon}}. \tag{4.11}$$

Equation (4.11) reduces to equation (4.10) as $p \to \infty$. In applying equation (4.11) to a numerical model, H can be approximated from the discrete field q^δ. This approximation, together with a user specified $\hat{\epsilon}$, yields a mesh quality metric tensor.

Given a number of fields q_i^δ, each with metric tensors M_i, a single metric tensor suitable for controlling mesh adaptivity can be formed via metric superposition: formation of a single metric M such that:

$$\boldsymbol{v}^T M \boldsymbol{v} \geq \max_i \left(\boldsymbol{v}^T M_i \boldsymbol{v} \right). \tag{4.12}$$

For further details of metric superposition, see Pain et al. [2001]

From the metric tensor, an objective functional defining the quality of an individual element can be defined. For example, the Fluidity-ICOM mesh adaptivity library uses the Pain et al. [2001] functional:

$$\mathcal{F}_e = \sum_{l \in \mathcal{L}_e} (r_l - 1)^2 + \left(\frac{\alpha}{\rho_e} - 1 \right)^2, \tag{4.13}$$

where $\mathcal{L}_e$ is the set of edges in element e, r_l is the length of edge l measured using M, ρ_e is the minimum in-sphere radius for element e measured using M, and α is the minimum in-sphere radius for an ideal simplex ($1/\left(2\sqrt{6}\right)$ for an ideal tetrahedron). An ideal element therefore has an objective functional value of $\mathcal{F}_e = 0$. This objective functional can be used to determine the relative improvement of local topological operations (h-adaptivity) and nodal perturbations (a form of r-adaptivity).

Figure 4.18 demonstrates meshes generated using an L_2 interpolation error driven metric, and an L_∞ interpolation error driven metric, for a field with both high curvature and low curvature features. As noted in Alauzet et al. [2006], when using the L_∞ metric mesh resolution is preferentially concentrated on the feature of highest field curvature. Little mesh resolution is concentrated in the features of lower curvature. When using the L_2 metric, mesh resolution is concentrated in both the dominant high curvature features and the weaker features. Hence Alauzet et al. [2006] conclude that L_p interpolation error driven metrics of lower p are more (quote) *"more appropriate to capture weak phenomena in simulations involving large amplitude phenomena"*.

Figure 4.19 demonstrates meshes generated using an L_2 interpolation error driven metric, and an L_∞ interpolation error driven metric, for a rotating annulus simulation (using output from the $\omega = 2.764$ rad s^{-1} fixed simulation from section 4.1 – a mode 3 simulation). The metric is derived from all three components of the velocity field, with equal tolerances $\hat{\epsilon}_{u_x} = \hat{\epsilon}_{u_y} = \hat{\epsilon}_{u_z}$, and with the three resulting metrics combined via superposition as described in Pain et al. [2001]. The mesh generated using an L_2 metric has 50821 nodes, a minimum edge length of 0.023 cm, and a maximum element anisotropy (ratio of maximum to minimum element edge length) of 200. The mesh

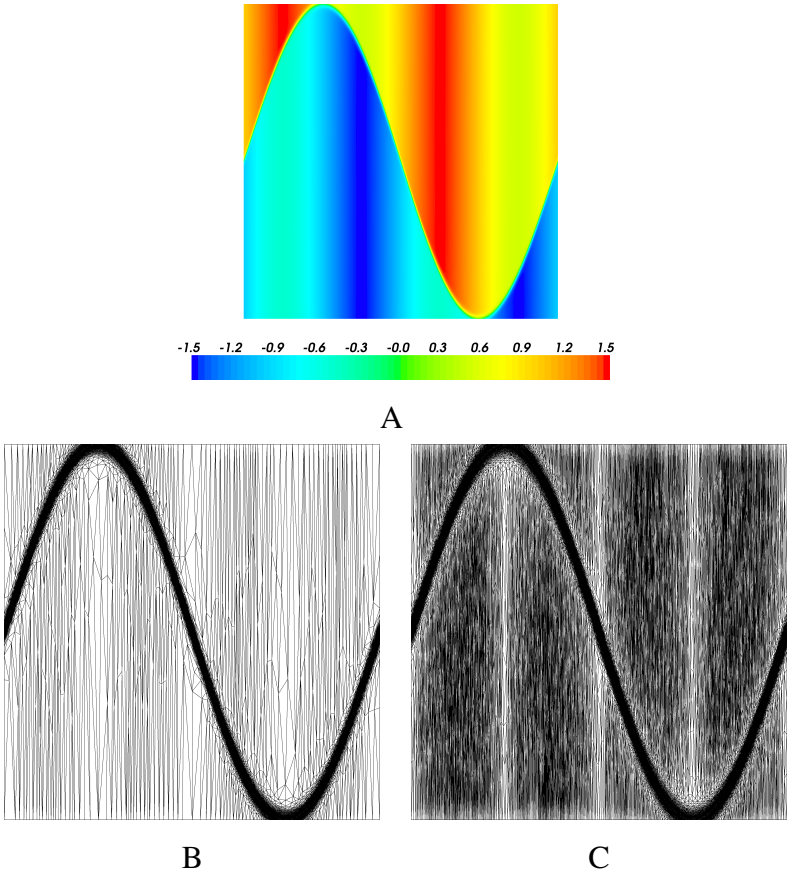

Figure 4.18.: Optimisation of a mesh to represent a field $\phi = 0.5 \sin(4\pi x) + \tanh[100(y - 0.5 \sin(2\pi x) - 0.5)]$, in a domain $0 \leq x \leq 1$, $0 \leq y \leq 1$, based on Alauzet et al. [2006] figure 4. Optimisations are performed using the ani2d mesh adaptivity library [Vasilevskii and Lipnikov, 1999]. A: The field for which the mesh is optimised. B: Mesh with 20137 nodes, generated using an L_∞ interpolation error driven metric. C: Mesh with 20656 nodes, generated using an L_2 interpolation error driven metric.

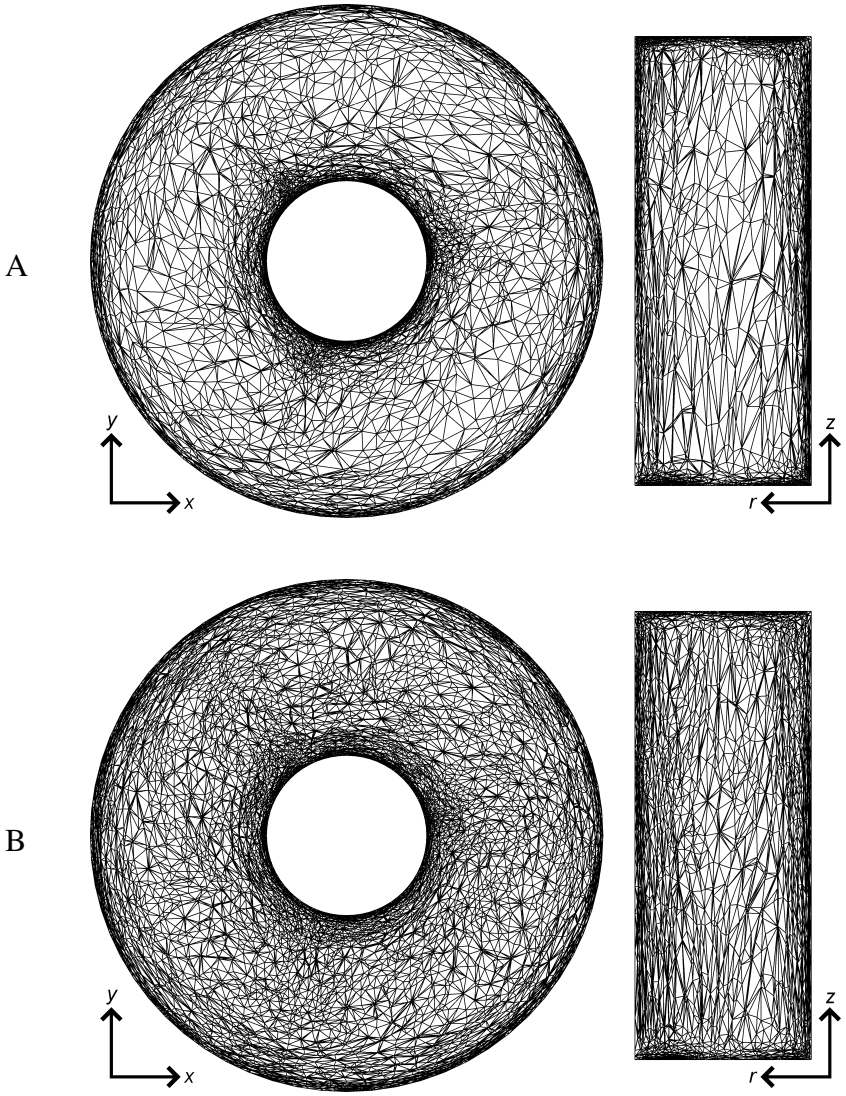

Figure 4.19.: Vertical and horizontal sections through meshes optimised to constrain the interpolation error in the velocity field for a thermally driven rotating annulus simulation. The fields used for mesh optimisation are taken from the $\omega = 2.764$ rad s^{-1} fixed mesh simulation in section 4.1. Optimisations are performed using the Fluidity-ICOM mesh adaptivity library [Pain et al., 2001]. A: Optimisation constraining the L_∞ interpolation error. The mesh has 52294 nodes. B: Optimisation constraining the L_2 interpolation error. This mesh has 50821 nodes

generated using an L_∞ metric has 52294 nodes, a minimum edge length of 0.020 cm, and a maximum anisotropy of 210. Note that the mesh generated using the L_∞ metric primarily concentrates resolution in the boundary layers, with a relatively low resolution in the interior region. Moreover, this mesh primarily concentrates resolution in the narrowest, and hence highest curvature, regions of the boundary layers. It can be seen, for example, that there is relatively low resolution in the lower outer wall boundary region. By contrast, the mesh generated using the L_2 metric extends resolution away from the boundary region, with more uniform resolution in the boundary layers and increased resolution in the interior. Since the thermally driven rotating annulus contains both high curvature features (the boundary layers) and low curvature features (the interior), the use of an L_p adaptivity metric with lower p is therefore more appropriate than the use of an L_∞ metric. All simulations described in this chapter use an L_2 interpolation error driven adaptivity metric with a tolerance of $\hat{\epsilon} = 0.1$ (cgs units) for all three velocity field components.

4.2.3. Interpolant

The final stage in applying dynamic mesh adaptivity is the selection of a mesh-to-mesh interpolant. Interpolants are discussed in detail in chapter 3. An additional issue is encountered when the model is run on parallel resources. Fluidity-ICOM is parallelised using MPI, with a number of processes designated approximately equal sized partitions. Partitions share "halo" regions, with data in these regions communicated between processes as required. This is described in further detail in appendix A.1. Applying dynamic mesh adaptivity to a simulation decomposed into a number of partitions raises two issues. Firstly, a shared halo region must be preserved by the mesh optimisation procedure, in order for the validity of the domain decomposition to be maintained – sufficient information regarding the decomposition must be preserved by the mesh optimisation. This can be achieved by locking the shared halo region, and excluding it from the mesh optimisation procedure. Secondly, the simulation must be re-load-balanced, to distribute the newly optimised mesh evenly between the processes. This is achieved using a parallel re-load-balancing step [Karypis and Kumar, 1998]. The re-load-balancing step can further be used to perturb the halo region, so as to expose this region (locked in the previous mesh optimisation) in a further mesh optimisation step. Iteration of this procedure allows the entire mesh to be optimised. In three dimensions three such adapt / re-load-balance iterations are required in order to process the entire mesh [Piggott et al., 2008b]. Hence, when dynamic mesh adaptivity is applied in parallel (and in three dimensions), the solution fields are interpolated *three times*. This makes an appropriate choice of interpolants essential. In addition, it means that the results of a parallel dynamic mesh adaptive Fluidity-ICOM

simulation are not necessarily equivalent to those of an otherwise identically configured serial dynamic mesh adaptive Fluidity-ICOM simulation.

We now proceed to test the influence of mesh-to-mesh interpolation in simulations of the thermally driven rotating annulus. In order to test only the influence of interpolation, and allow the influence from the choice of the mesh adaptivity metric to be excluded, a series of unstructured annulus meshes of locally equal resolution were generated. A first unstructured mesh was generated of resolution locally equal to that of a structured mesh of resolution $24 \times 24 \times 64$ in the radial, vertical, and azimuthal dimensions respectively, with a hyperbolic tangent stretch mesh to resolve the boundary layers as used in section 4.1 . This unstructured mesh was generated using the Fluidity-ICOM mesh adaptivity library, with a mesh adaptivity metric derived from the polar decomposition of the Jacobian mapping elements in the structured mesh to the ideal tetrahedron [Formaggia and Perotto, 2001, Micheletti and Perotto, 2006] (as described in section 2.3, and demonstrated in figure 2.8). A second mesh was generated in an analogous manner, using a metric derived from the first unstructured mesh. This process was iterated to generate 10 such unstructured meshes, with each successive pair of meshes having approximately equal local mesh resolution. These 10 meshes had (in order of creation): 37129, 35913, 34515, 33141, 31870, 30639, 29443, 28371, 27395, and 26523 nodes respectively. There is a systematic decrease in the number of mesh nodes in each iteration of this procedure, as a result of the sub-optimality of meshes generated using the mesh adaptivity library.

The simulation was configured as in section 4.1, and with a rotation rate of $\omega = 1.0$ rad s^{-1}. After every 20 model timesteps, the computational mesh was switched to the next unstructured mesh in the sequence of unstructured meshes generated by the above procedure (looping to the first unstructured mesh after every 200 timesteps), with the velocity, pressure and temperature fields interpolated to the new unstructured mesh. This procedure therefore tests the influence of mesh-to-mesh interpolation in Fluidity-ICOM simulations of the thermally driven rotating annulus, when applied between computational meshes of approximately equal local mesh resolution. Three interpolants were tested: collocation for all fields, Galerkin projection for all fields, and the geostrophic balance preserving interpolant described in chapter 3 for velocity, combined with Galerkin projection for temperature. In applying Galerkin projection, the Dirichlet boundary conditions for velocity and temperature were applied in the interpolation, via a strongly applied boundary condition. In applying the geostrophic balance preserving interpolant, the Helmholtz decomposition of the Coriolis acceleration was discretised with a balanced pressure decomposition method, using lumping of the velocity mass matrix, and with K matrix stabilisation, for reasons as previously noted in sections 3.4 and 3.5. In addition, two further fixed mesh "control" simulations were conducted on

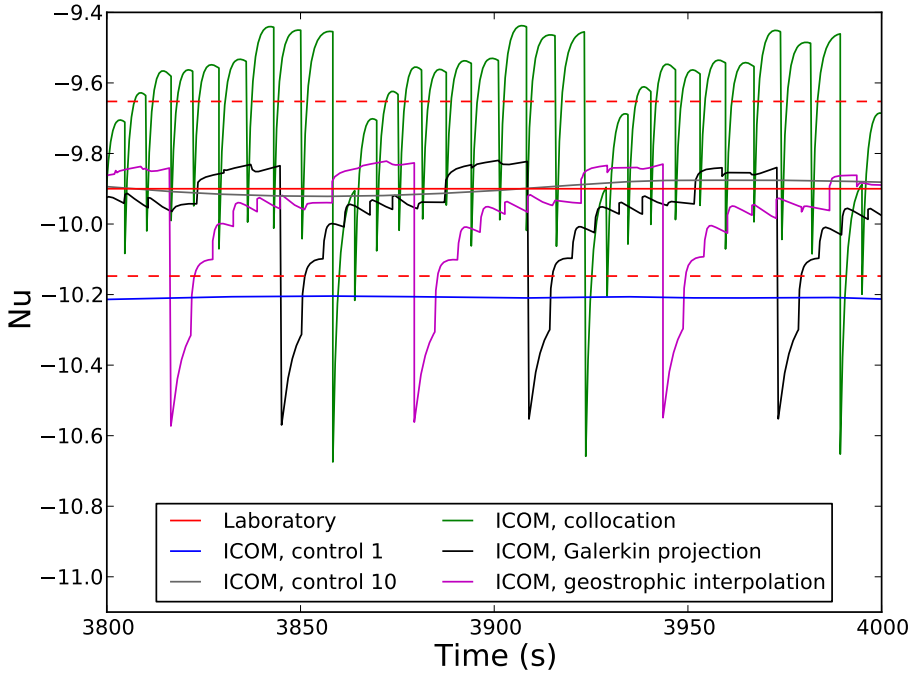

Figure 4.20.: Thermally driven rotating annulus heat transport measurements at rotation rate $\omega = 1.0$ rad s^{-1}, with repeated interpolation between approximately equal resolution unstructured meshes. The interpolants tested are: collocation, Galerkin projection with strong boundary conditions applied, and geostrophic balance preserving interpolation combined with Galerkin projection (with strong boundary conditions) for temperature. Laboratory observations are shown in red, with error bars. "Control 1" and "Control 10" refer to the fixed mesh simulations conducted on the highest and lowest resolution meshes respectively.

the highest and lowest resolution unstructured meshes, with no mesh-to-mesh interpolation. These provide an indication of the range of variation expected purely as a result of the variation in resolution between the unstructured meshes used in the test. The simulations were integrated in serial for 4000 s of simulated laboratory time, with a varying timestep targeting a maximum Courant number of unity.

The heat transport for these simulations is shown in figure 4.20. When using collocation, clear and severe interpolation induced shocks are observed after every interpolation. These shocks lead to the heat transport exceeding the values observed in the two fixed mesh control simulations. When applying Galerkin projection the shocks are significantly reduced, and the heat transport of the system is largely confined between the values in the two control simulations. The remaining sudden change in heat transport occurs when interpolating from the lowest resolution mesh, with 26523 nodes, to the original mesh, with 37129 nodes. The heat transport observed when using the geostrophic balance preserving interpolant is extremely similar to that observed when applying Galerkin projection.

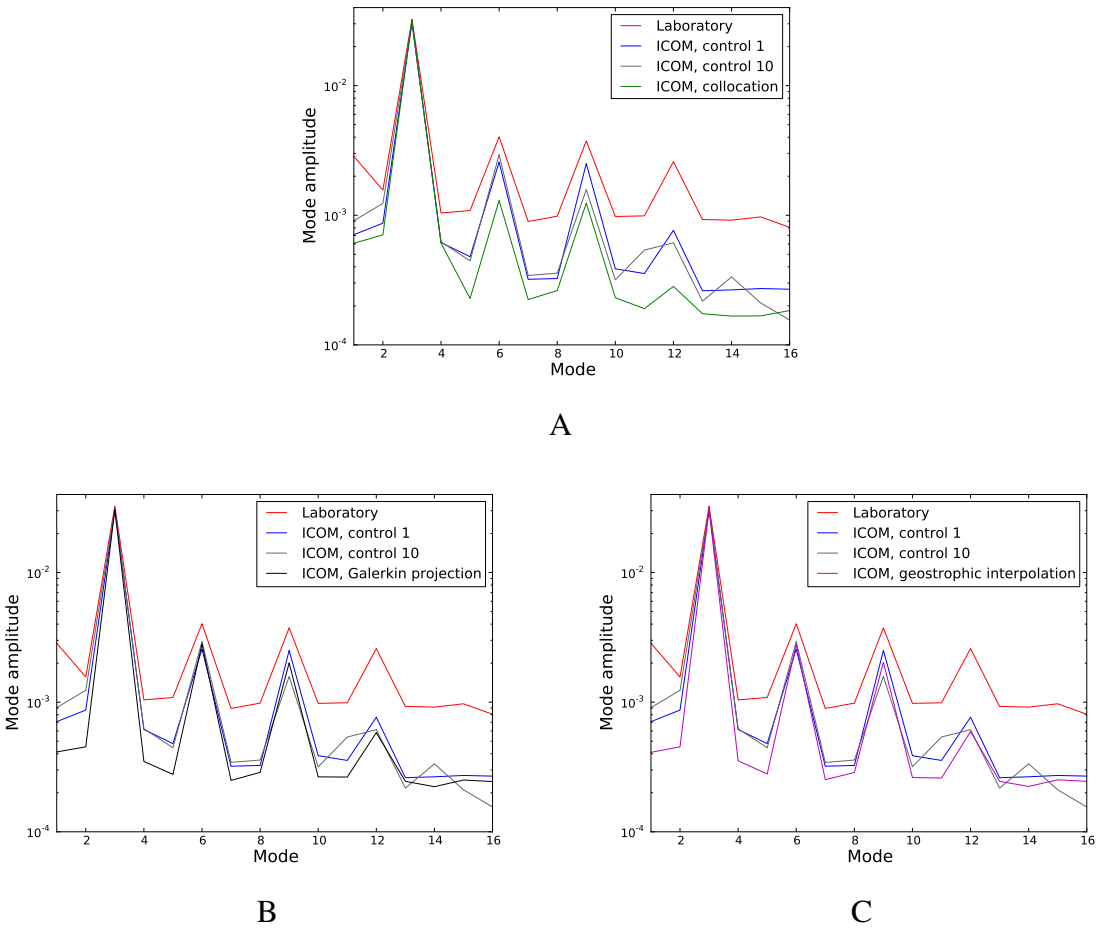

Figure 4.21.: Time-average mid-height temperature Fourier mode amplitudes at rotation rate $\omega = 1.0$ rad s^{-1}, with repeated interpolation between approximately equal resolution unstructured meshes. Laboratory observations are shown in red. "Control 1" and "Control 10" refer to the fixed mesh simulations conducted on the highest and lowest resolution meshes respectively. A: Collocation. B: Galerkin projection, with strong boundary conditions. C: Geostrophic balance preserving interpolant, combined with Galerkin projection (with strong boundary conditions) for temperature.

The time-averaged mode amplitudes for these simulations, derived from the mid-radius mid-heigh temperature, are shown in figure 4.21. The mode amplitudes when applying Galerkin projection generally correspond well to the control simulations – in particular, the amplitude of the dominant mode and its harmonics lie close to or between those observed in the control simulations. The mode amplitudes observed when applying collocation are generally lower than those observed in the control simulations, particularly at higher mode numbers. This may be caused by the increased numerical diffusion associated with collocation, as noted in sections 3.1 and 3.4. The mode amplitudes when using the geostrophic balance preserving interpolant are almost identical to those observed when applying Galerkin projection.

For the fixed mesh control simulations the minimum and maximum temperatures are observed to be bounded by the range of the thermal forcing. The interpolation test using collocation also maintains a bounded temperature field. The simulations applying Galerkin projection and the

geostrophic balance preserving interpolant do not maintain the temperature bounds – the maximum temperature observed was $T = 1.002\Delta T + T_A$ in each case. However, these overshoots were corrected after a single model timestep, and were therefore deemed to be tolerable. The conservative bounding scheme of Farrell et al. [2009] was not tested.

Hence it is concluded that, for dynamic mesh adaptive simulations of the thermally driven rotating annulus using a stabilised P_1P_1 discretisation, mesh-to-mesh Galerkin projection of velocity, temperature, and pressure (with the Dirichlet boundary conditions for velocity and temperature applied in the Galerkin projection in the strong sense), is preferable to the use of collocation. The use of a geostrophic balance preserving interpolant offers no significant benefits, at increased computational cost.

4.2.4. Configuration

The configuration used for mesh adaptive simulations is therefore as follows. The final state of the fixed mesh simulations of section 4.1 were used to initialise dynamic mesh adaptive simulations. Mesh optimisation was applied every 20 model timesteps using the Fluidity-ICOM mesh adaptivity library. The adaptivity metric was derived from all three components of the velocity field, using a metric constraining the L_2 interpolation error. The velocity, pressure, and temperature fields were interpolated via Galerkin projection, with strong boundary conditions (no-slip and thermal forcing) applied in the interpolation. The simulations were conducted in parallel with 8 MPI processes (using the Oxford Supercomputing Centre "redqueen" machine). Hence three iterations of mesh optimisation combined with parallel re-load-balancing were applied.

For additional computational efficiency, only the element-wise topological optimisation component of the Fluidity-ICOM mesh adaptivity (*h*-adaptivity) was enabled, and nodal perturbations (*r*-adaptivity) was disabled. This was found to have only a relatively small influence on the mesh quality, at greatly reduced computational cost. For a typical rotating annulus configuration with mesh optimisation applied in serial, applying only *h*-adaptivity, rather than full *hr*-adaptivity, increased the mean element functional (as defined by equation (4.13)) by 40%, and the maximum element functional by only 25%, while being 46% faster.

4.3. Adaptive mesh Fluidity-ICOM simulations

The fixed mesh simulations described in section 4.1 were resumed with mesh adaptivity enabled, using the configuration described in the previous section. The simulations were integrated for a further 2000 s of simulated laboratory time. The model timestep size was adapted after every timestep to target a maximum Courant number of unity.

Examples of the meshes generated by these simulations are shown in figure 4.22. A mid-radius mid-height temperature Hovmöller diagram for the $\omega = 0.838$ rad s^{-1}simulation is shown in figure 4.23, and the corresponding time varying number of mesh nodes in figure 4.24. As previously noted in section 4.1, the fixed mesh simulation at this rotation rate exhibits amplitude vacillation. The dynamic mesh adaptive simulation also exhibits amplitude vacillation. The number of mesh nodes is observed to vary in time, with increases and decreases in the mesh resolution corresponding to variations in the baroclinic wave amplitude.

The average number of mesh nodes for all mesh adaptive simulations are shown in figure 4.25. Note that the simulations at higher rotation rate generally have an increased number of mesh nodes – an exception is a decrease in the number of mesh nodes for a higher rotation rate mode 3 simulation, relative to lower rotation rate mode 4 simulations. Hence dynamic mesh adaptivity acts to increase the mesh resolution for the higher Reynolds number simulations, in order to resolve the more active dynamics.

Figure 4.26 shows the depth averaged stream functions for the $\omega = 1.0$ rad s^{-1} and $\omega = 2.764$ rad s^{-1} dynamic mesh adaptive simulations. These were computed via a Galerkin projection of the depth integral of the velocity field, as described in appendix A.4, onto the surface of the structured MORALS meshes used in section 4.1. The depth averaged stream function was computed as described in section 4.1.1 and in appendix A.5, using a P_1 continuous Galerkin discretisation. Comparison with figure 4.4 shows that the $\omega = 2.764$ rad s^{-1} dynamic mesh adaptive simulation has, relative to the fixed mesh simulation, a more favourable comparison with the result from MORALS. In particular, the amplitude of the depth averaged stream function is increased in the dynamic mesh adaptive simulation, relative to the corresponding fixed mesh simulation. Figure 4.27 shows the potential energy distribution for the $\omega = 1.0$ rad s^{-1} and $\omega = 2.764$ rad s^{-1} dynamic mesh adaptive simulations, with the results from the corresponding MORALS and fixed mesh Fluidity-ICOM simulations for comparison. As previously noted in section 4.1.1, the fixed mesh simulation at $\omega = 2.764$ rad s^{-1} has a bias towards higher potential energies, corresponding to a decrease in the amplitude of the baroclinic eddies. The dynamic mesh adaptive simulation

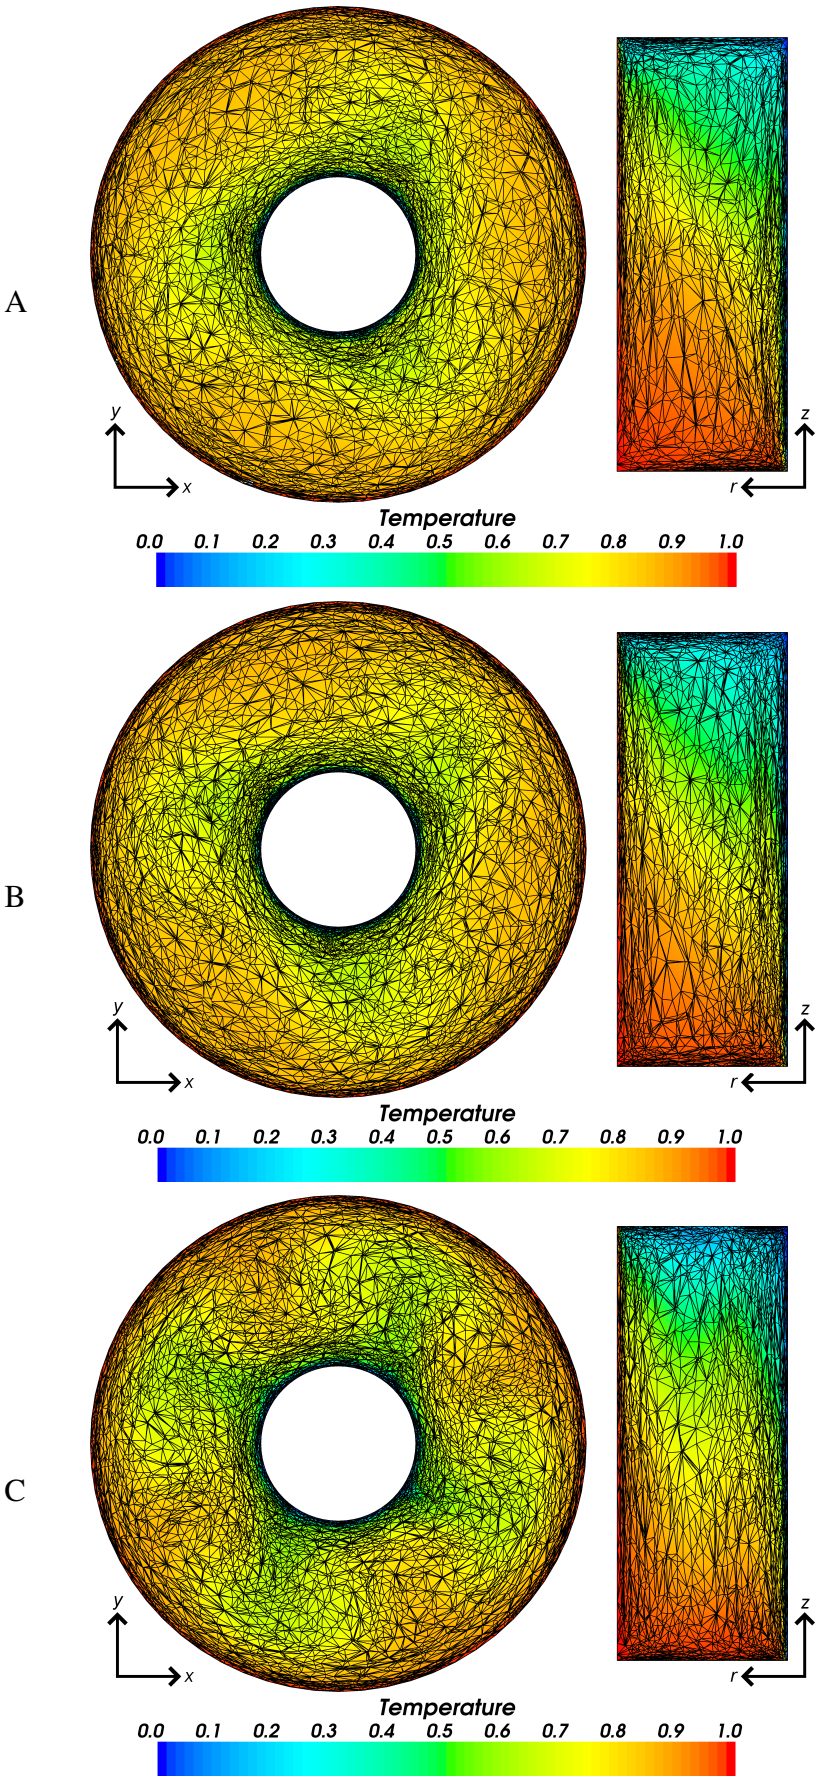

Figure 4.22.: Vertical and horizontal sections through the final meshes for three dynamic mesh adaptive annulus simulations of the thermally driven rotating annulus. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$. A: $\omega = 0.838$ rad s^{-1}, amplitude vacillating mode 3. This mesh has 41808 nodes. B: $\omega = 1.0$ rad s^{-1}, steady mode 3. This mesh has 43196 nodes. C: $\omega = 2.289$ rad s^{-1}, steady mode 4. This mesh has 53771 nodes.

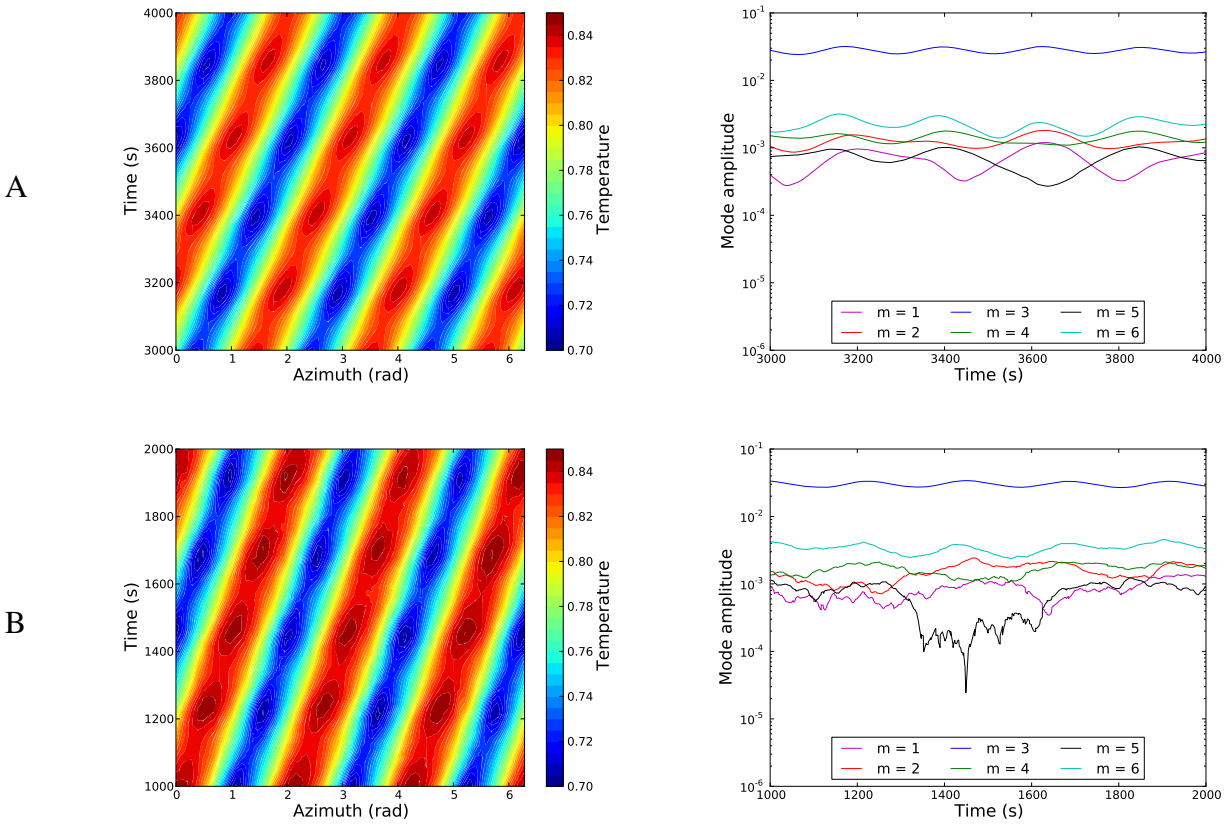

Figure 4.23.: Mid-radius mid-height temperature measurements for the $\omega = 0.838$ rad s^{-1} Fluidity-ICOM simulations of the thermally driven rotating annulus. A: Fixed mesh simulation at resolution $24 \times 24 \times 64$ nodes in the radial, vertical and azimuthal dimensions respectively. B: Dynamic mesh adaptive simulation. Hovmöller diagrams are shown on the left (with a common colour scale), and the amplitude of the first six Fourier modes on the right, where m is the mode number. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$.

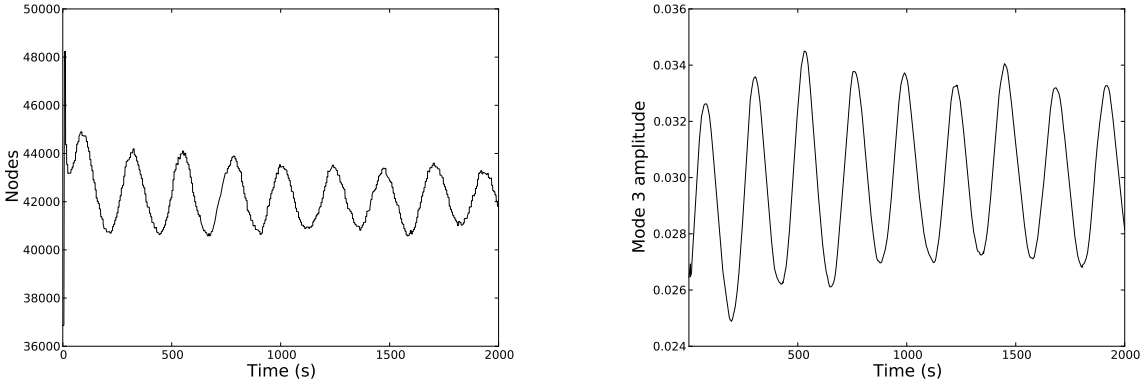

Figure 4.24.: Left: Number of mesh nodes as a function of time for the $\omega = 0.838$ rad s^{-1} dynamic mesh adaptive Fluidity-ICOM simulation of the thermally driven rotating annulus. Right: Amplitude of the third Fourier mode, derived from the mid-radius mid height temperature. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$.

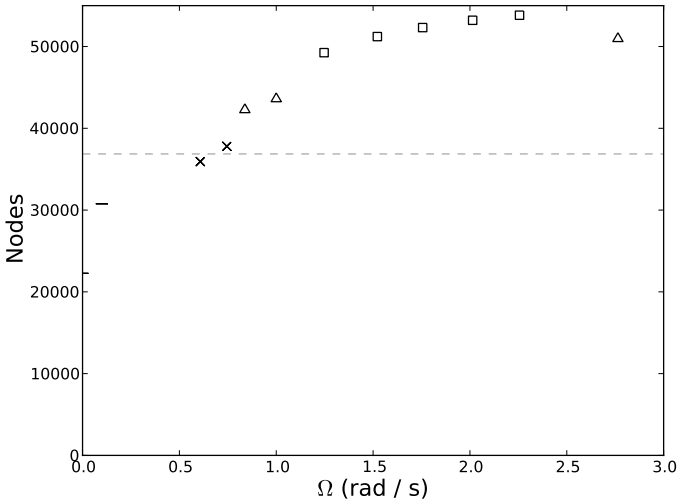

Figure 4.25.: Average number of mesh nodes for dynamic mesh adaptive Fluidity-ICOM simulations of the thermally driven rotating annulus. A line denotes an axisymmetric system, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four. The horizontal lines shows the number of nodes in the fixed mesh simulations described in section 4.1

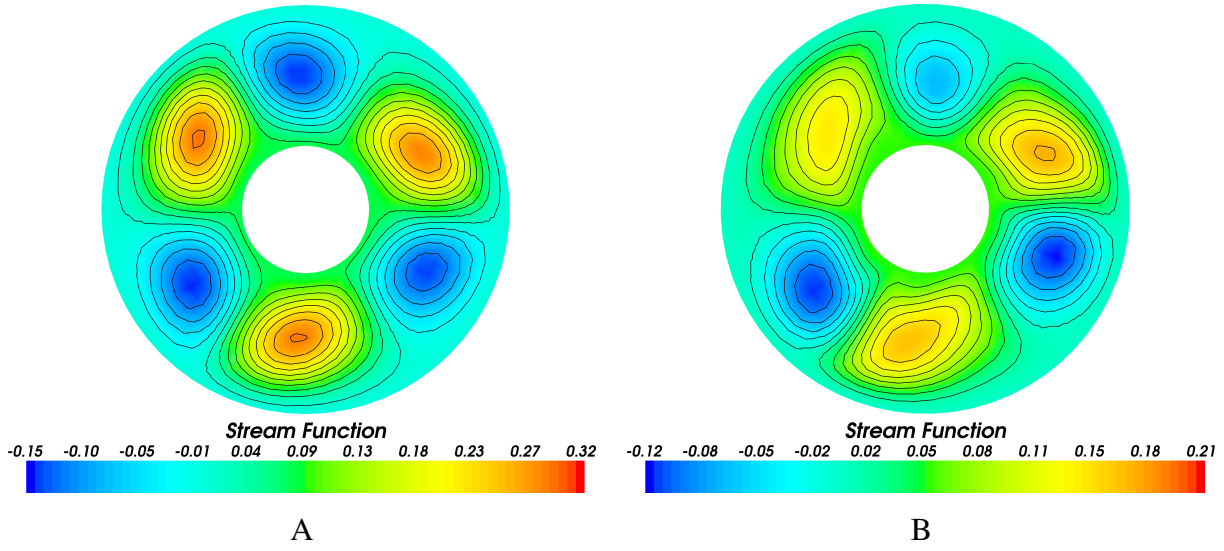

Figure 4.26.: Depth averaged stream function as defined by equation (4.1) for dynamic mesh adaptive Fluidity-ICOM simulations of the thermally driven rotating annulus. A: Rotation rate $\omega = 1.0$ rad s^{-1}. B: Rotation rate $\omega = 2.764$ rad s^{-1}. Colour scales are chosen to match figures 4.3 and 4.4 respectively.

shows a reduction of this bias, indicating that the dynamic mesh adaptive simulation has a more accurate representation of baroclinic instability.

The Fourier amplitudes derived from the mid-radius mid-height temperature were computed for the dynamic mesh adaptive simulations. The amplitude of the dominant Fourier modes are shown in figure 4.28, with laboratory observations and the results from the fixed mesh simulations for comparison. It can be seen that the dynamic mesh adaptive simulations are generally associated with increased mode amplitudes relative to the fixed mesh equivalents. These amplitudes are, at higher rotation rates, significantly larger than those observed in the laboratory.

A

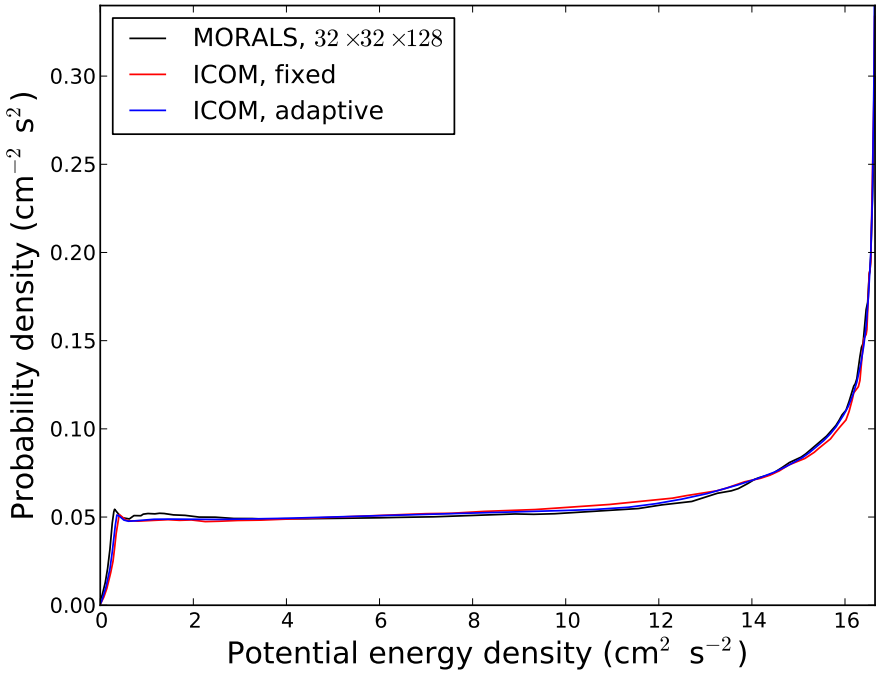

B

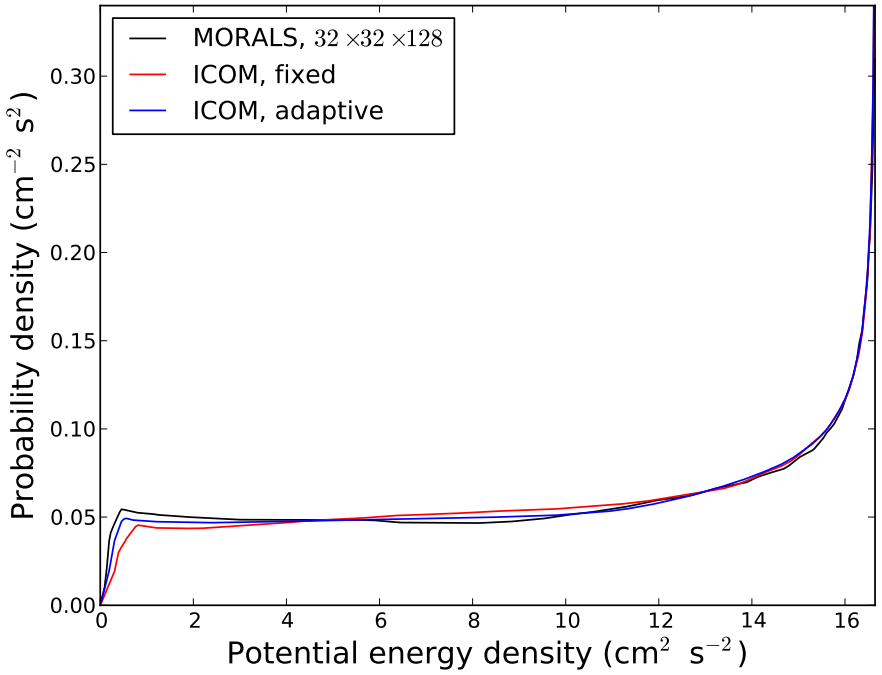

Figure 4.27.: Probability density function for the potential energy density $U = -g\alpha(T - T_A)z$ in the rotating in the rotating annulus for dynamic mesh adaptive Fluidity-ICOM simulations, fixed mesh Fluidity-ICOM simulations, and MORALS simulations. The background potential energy density gz is excluded, and the potential energy density zero point chosen so that $\min(U) = 0$. For a randomly selected fluid parcel, the probability of the parcel having a potential energy density between U and $U + dU$ is $\mathcal{P}(U)\,dU$, where $\mathcal{P}(U)$ is the probability density function. A is at rotation rate $\omega = 1.0$ rad s^{-1}and B at rotation rate $\omega = 2.764$ rad s^{-1}.

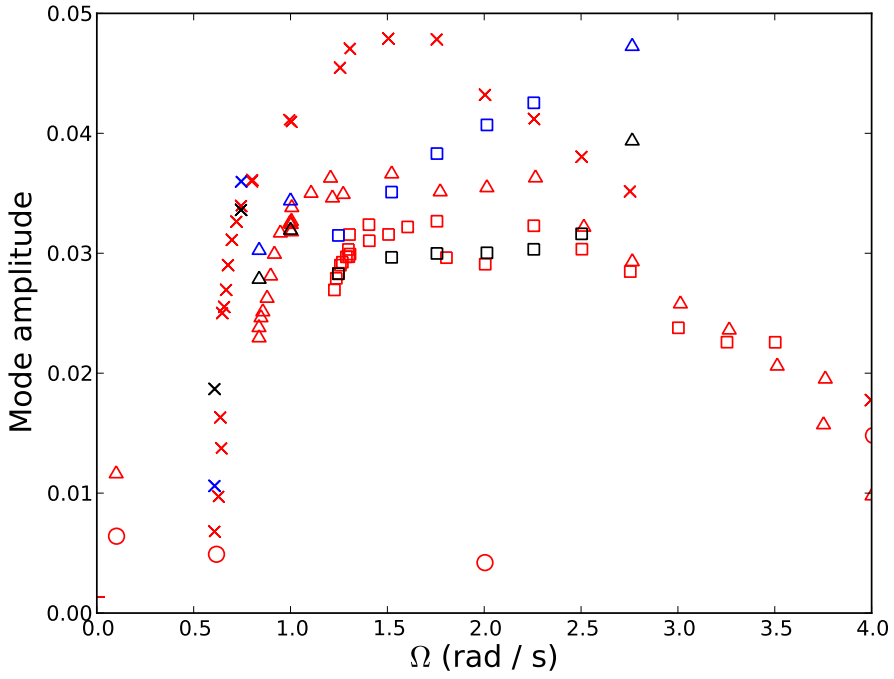

Figure 4.28.: Amplitude of the dominant Fourier mode for the mid-radius mid-height temperature. The amplitudes are computed for the normalised temperature, $(T - T_A)/\Delta T$. Laboratory observations are in red, fixed mesh Fluidity-ICOM simulations in black, and dynamic mesh adaptive Fluidity-ICOM simulations in blue. The laboratory values are averaged over each thermocouple array data set. The model values are averaged over 1000 s. A line denotes an axisymmetric system, a circle denotes mode one, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four.

The corresponding dominant mode drift periods are shown in figure 4.29, with laboratory observations and the results from the fixed mesh simulations for comparison. The dynamic mesh adaptive simulations are generally associated with increased drift periods relative to the fixed mesh equivalents, although the drift periods are still generally somewhat below the experimental observations, particularly at higher rotation rates.

The vacillation indices, as computed from equation (4.6), are shown in figure 4.30, with laboratory observations and the results from the fixed mesh simulations for comparison. It was previously noted in section 4.1.2 that the fixed mesh Fluidity-ICOM simulations show only weak amplitude vacillation. Two of the 13 fixed mesh simulations exhibited amplitude vacillation. With the application of dynamic mesh adaptivity, one additional amplitude vacillating system is observed, at $\omega = 0.608$ rad s^{-1}. The remaining two simulating exhibiting amplitude vacillation, with $\omega = 0.838$ rad s^{-1} and $\omega = 1.247$ rad s^{-1}, have, in the dynamic mesh adaptive simulations, weaker amplitude vacillation than observed in the corresponding fixed mesh simulations. In particular, the $\omega = 1.247$ rad s^{-1} dynamic mesh adaptive simulation has a very low vacillation index, and may be approaching a steady solution. Hence the introduction of dynamic mesh adaptivity is not observed to lead to more a physically realistic magnitude of amplitude vacillation.

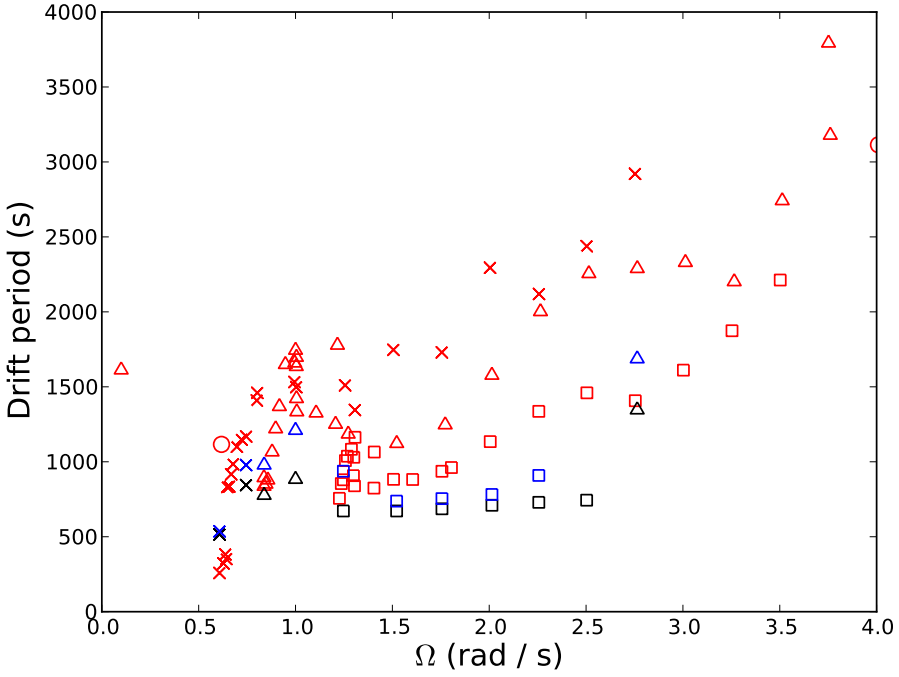

Figure 4.29.: Drift period of the dominant Fourier mode mode for the mid-radius mid-height temperature. Laboratory observations are in red, fixed mesh Fluidity-ICOM simulations in black, and dynamic mesh adaptive Fluidity-ICOM simulations in blue. The laboratory values are averaged over each thermocouple array data set. An experiment with a vary large drift period (at $\omega = 4.0$ rad s^{-1}), and experiments with very small mode amplitude, are not shown. The model values are averaged over 1000 s. A line denotes an axisymmetric system, a circle denotes mode one, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four.

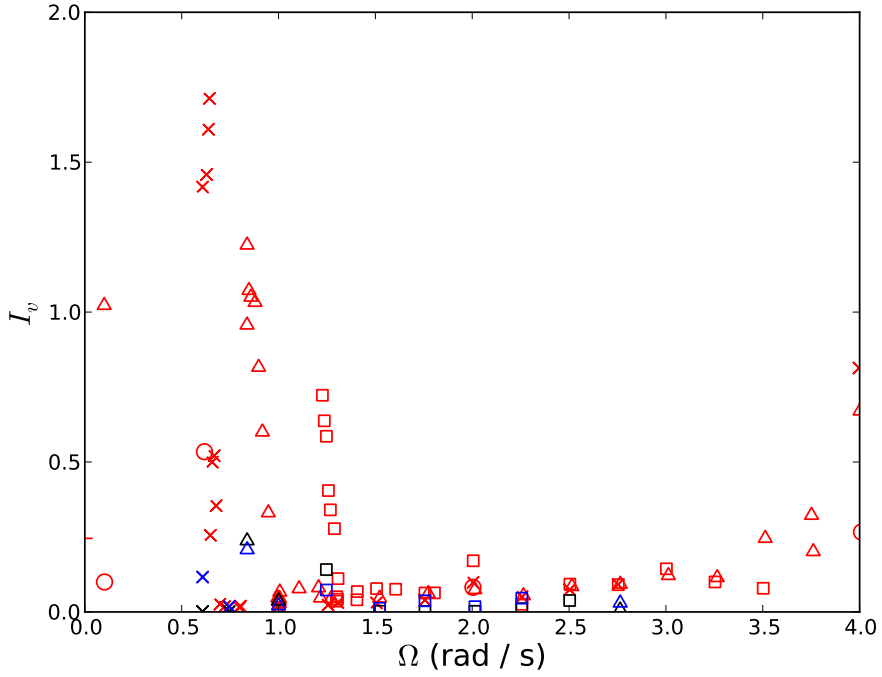

Figure 4.30.: Mid-radius mid-height temperature vacillation indices, as defined by equation (4.6). Laboratory observations are in red, fixed mesh Fluidity-ICOM simulations in black, and dynamic mesh adaptive Fluidity-ICOM simulations in blue. A circle denotes mode one system, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four.

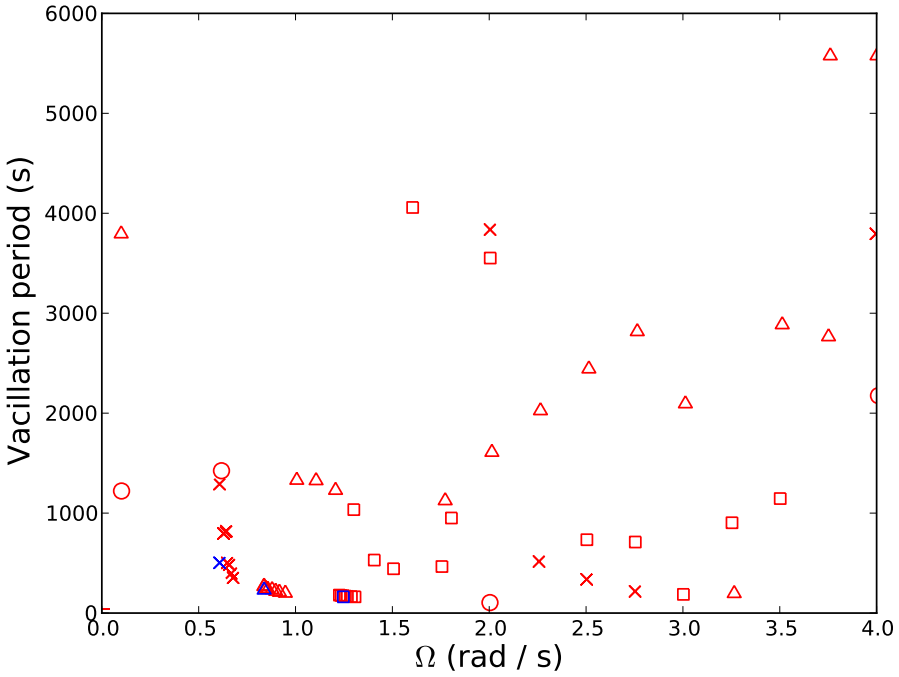

Figure 4.31.: Mid-radius mid-height temperature vacillation periods. Laboratory observations are in red and dynamic mesh adaptive Fluidity-ICOM simulations in blue. A circle denotes mode one system, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four. Only experiments and simulations with a vacillation index (as defined by equation (4.6)) above 0.05 are shown.

The corresponding vacillation periods are shown in figure 4.31, with laboratory observations for comparison. The vacillation periods for the three dynamic mesh adaptive simulations exhibiting amplitude vacillation show an excellent agreement with experiment. The vacillation periods for the fixed mesh Fluidity-ICOM simulations are not shown, but are virtually coincident with the adaptive mesh results.

The system heat transport for the dynamic mesh adaptive simulations are shown in figure 4.32. The dynamic mesh adaptive simulations are observed to have a significantly decreased system heat transport, relative to the fixed mesh simulations. At higher rotation rates, this leads to the heat transport of the dynamic mesh adaptive simulations being significantly lower than laboratory observations.

A cause of this reduction in heat transport is shown in figure 4.33, for the $\omega = 1.0$ rad s^{-1} case. This shows the spatial distribution of the heat transport through the tank inner wall for the fixed mesh and adaptive mesh simulations. Note that, in the fixed mesh simulation, two primary contributions to the total heat transport can be identified: a large heat flux confined close to the upper boundary, and a more extensive but lower amplitude flux with a distinctive mode three signature. The former is associated with the Ekman transport, and the latter with the baroclinic eddy heat flux. In the adaptive mesh simulation the Ekman transport contribution is suppressed.

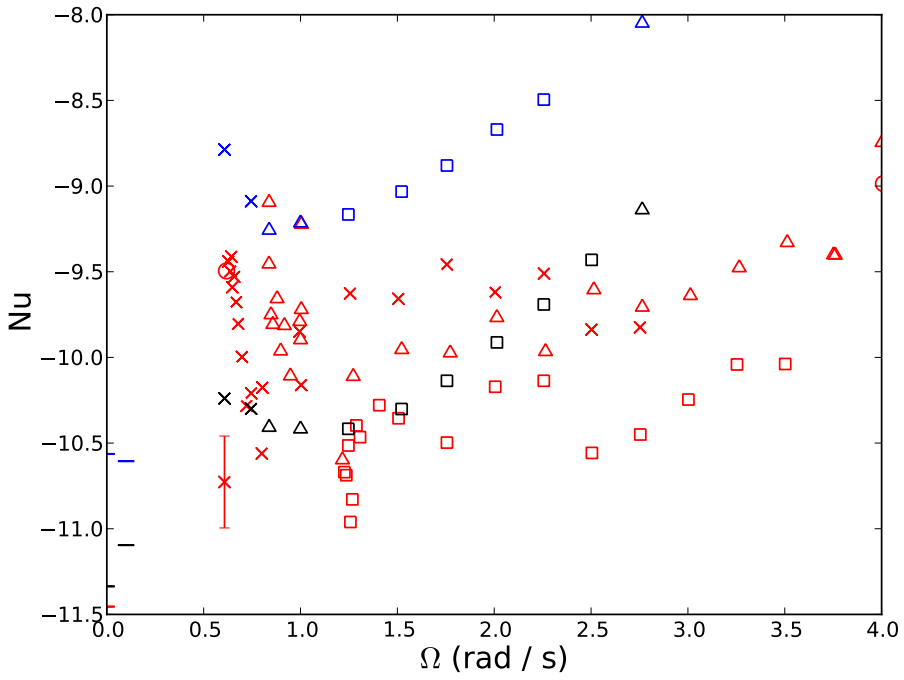

Figure 4.32.: Thermally driven rotating annulus heat transport measurements. Laboratory observations are in red (with one typical error bar), fixed mesh Fluidity-ICOM simulations in black, and dynamic mesh adaptive Fluidity-ICOM simulations in blue. The laboratory values are averaged over each thermocouple array data set. The model values are averaged over 1000 s. A line denotes an axisymmetric system, a circle denotes mode one, a cross denotes mode two, a triangle denotes mode three, and a square denotes mode four.

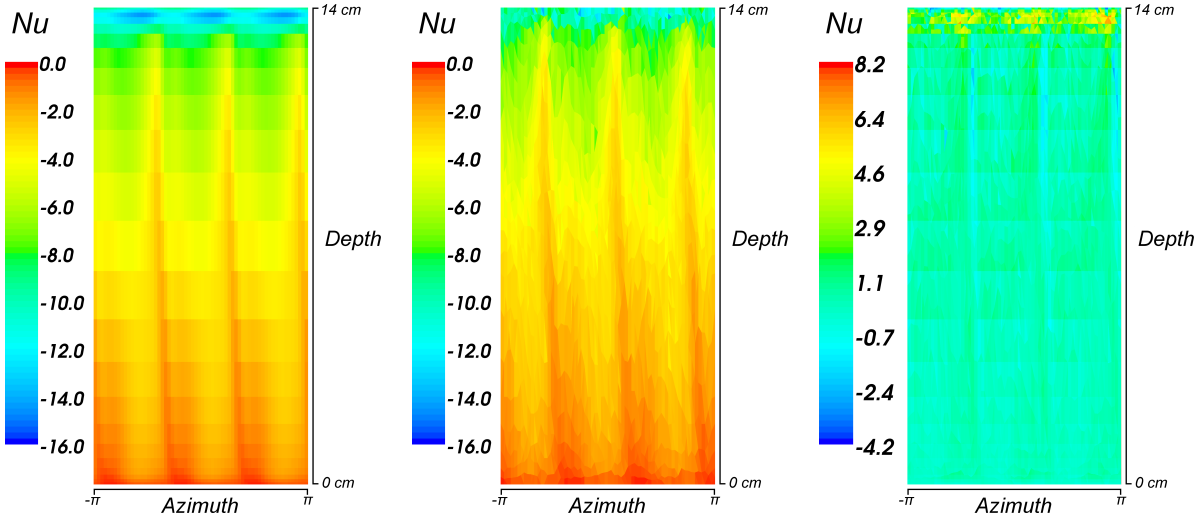

Figure 4.33.: Heat flux through the tank inner wall, expressed as a non-dimensional Nusselt number $Nu = [(\ln b/a) / (2\pi d \Delta T)] [2\pi a d (\nabla T \cdot \hat{\boldsymbol{n}})]$, for Fluidity-ICOM simulations of the thermally driven rotating annulus at rotation rate $\omega = 1.0$ rad s^{-1}. A: Fixed mesh Fluidity-ICOM. B: Dynamic mesh adaptive Fluidity-ICOM. C: Difference between A and B. This difference is displayed on the intersection mesh (the "supermesh") of the meshes used in A and B (see section 3.1).

One possible cause for this suppression is identified in figure 4.34. This shows the azimuthal Fourier mode amplitudes for the temperature field, as a function of radius and height, for the dominant mode 3 and for mode 4. These are shown for the $\omega = 1.0$ rad s^{-1} fixed mesh and adaptive mesh Fluidity-ICOM simulations, and for the corresponding MORALS simulation. For all three simulations, two primary maxima can be identified in the mode 3 amplitude: one near mid-radius, in the lower half of the tank, and one in the upper region of the inner wall thermal boundary layer. A smaller signal can also be seen in the lower Ekman layer. The mode 3 amplitude results are qualitatively very similar for all three simulations. The MORALS simulation has an extremely low mode 4 amplitude. For the fixed mesh Fluidity-ICOM simulation, two mesh scale maxima can be observed in the mode 4 amplitude: one at the top of the inner wall, and one at the bottom of the outer wall. The dynamic mesh adaptive Fluidity-ICOM simulation also exhibits these peaks in mode 4 amplitude, but with approximately 5 times the magnitude. In these two regions the Ekman boundary layer flux impinges upon the lateral boundaries. These are strong advection regions, and this strong advection appears to lead to a numerical instability in the fixed mesh Fluidity-ICOM simulation. For the chosen adaptivity metrics dynamic mesh adaptivity actively reduces resolution in these regions, leading to an increase in the magnitude of the instability. Since this instability is located in the region where the Ekman layer flux impinges upon the inner boundary, the associated Ekman layer heat transport is disrupted, leading to a poor representation of the overall system heat transport.

Hence, the introduction of dynamic mesh adaptivity is, at $\omega = 2.764$ rad s^{-1}, observed to lead to a modest qualitative improvement in the depth averaged stream function, with an increase in the strength of the baroclinic eddies and a more favourable comparison with a corresponding MORALS simulation. The introduction of dynamic mesh adaptivity is also, at this rotation rate, observed to lead to a potential energy distribution with a more favourable comparison against MORALS results. This indicates an improved representation of baroclinic instability. The amplitude of the dominant mid-radius mid-height temperature modes are observed to be increased, relative to fixed mesh simulations, leading to a slightly less favourable comparison with laboratory observations. The representation of amplitude vacillation is not improved – the dynamic mesh adaptive simulations generally show weaker amplitude vacillation than the fixed mesh simulations. In addition, the application of dynamic mesh adaptivity is observed to lead to a substantial decrease of the system heat transport, and a very poor comparison with experimental observations. This is attributed to a numerical instability acting to disrupt the Ekman transport. Dynamic mesh adaptivity under-resolves the flow in these regions, and hence degrades the quality of the solution as measured using this diagnostic. It is possible that the numerical instability is the result of the

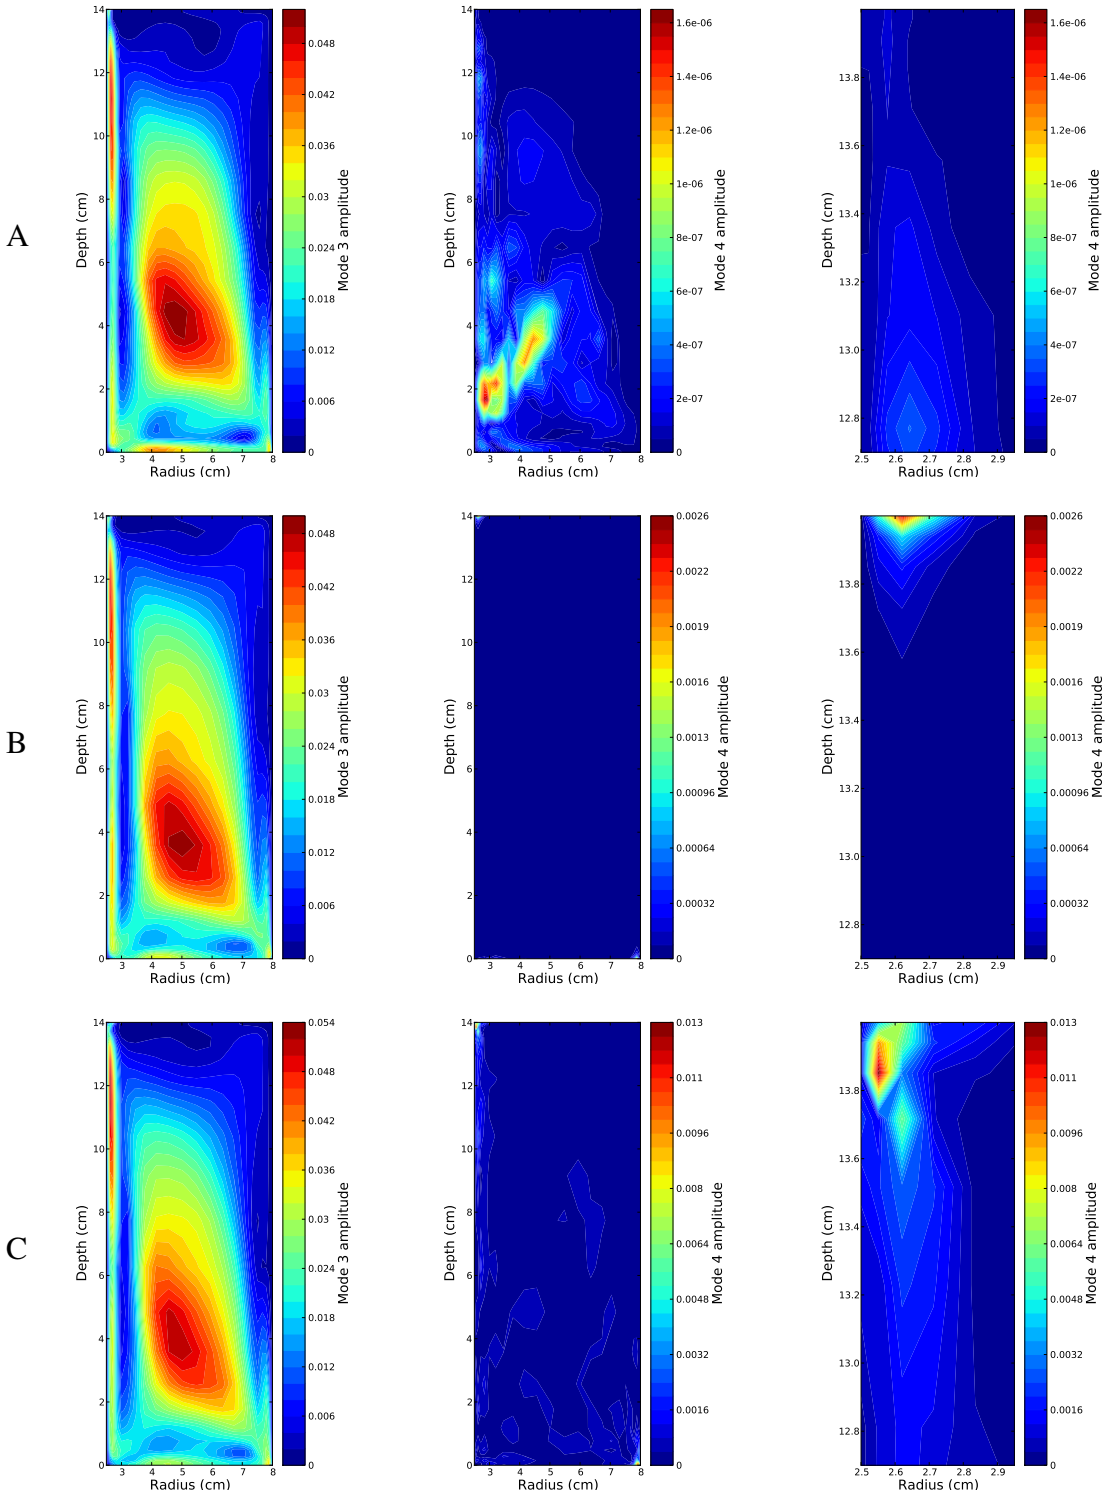

Figure 4.34.: Azimuthal Fourier modes for the temperature field in simulations of the thermally driven rotating annulus at rotation rate $\omega = 1.0$ rad s^{-1}. Mode 3 amplitudes are shown on the left and mode 4 amplitudes in the centre. On the right the mode 4 amplitudes for just the upper left region of the domain are shown. The mode amplitudes are computed for the normalised temperature, $(T - T_A)/\Delta T$. A: MORALS. B: Fixed mesh Fluidity-ICOM. C: Dynamic mesh adaptive Fluidity-ICOM.

use of a centred advection scheme for momentum. More advanced upwinding schemes, such as the streamline upwind Petrov Galerkin (SUPG) method [Hughes, 1987a], were not available for use with the momentum equation in Fluidity-ICOM at the time of this project. Such schemes may help to reduce or remove the observed instability.

4.4. Computational cost

The primary purpose of the validation exercise described in this chapter is to identify potential shortcomings in the application of dynamic mesh adaptive numerical methods to geophysical systems. For completeness, the computational cost of these simulations is briefly discussed.

The Fluidity-ICOM simulations described in sections 4.1 and 4.3 were executed in parallel on a 512 core cluster (the Oxford Supercomputing Centre "redqueen" machine) with 4 cores per cluster node. Each node had two 2.67GHz Intel Xeon CPU and 2GB of RAM. Fluidity-ICOM was compiled with version 10.1 of the Intel Fortran compiler. Each simulation was conducted using 8 MPI processes. The run times described below are the total CPU run times for the simulations – i.e. a simulation executed in parallel for 1000 s has a total CPU run time of 8000 s. Due to the additional overheads associated with performing simulations in parallel (in particular, the cost of inter-process communication and the increased cost of input/output), serial simulations are expected to have a slightly lower total CPU time as compared against an otherwise identically configured parallel simulation.

The fixed mesh simulations described in section 4.1 have a computational cost of between 62 and 520 CPU seconds per simulated (lab time) second. The dynamic mesh adaptive simulations described in section 4.3 have a computational cost of between 86 and 600 CPU seconds per simulated (lab time) second. The exact value of the computational cost is dependent upon the system rotation rate, the configuration of the linear solvers, and (in the dynamic mesh adaptive case) the number of mesh nodes. For comparison, a MORALS mesh simulation at resolution $24 \times 24 \times 64$, at a rotation rate of $\omega = 1.0$ rad s^{-1}, and integrated in serial, has a computational cost of approximately 1 CPU second per simulated second[2].

It is likely that the cost of the Fluidity-ICOM simulations can be reduced somewhat by the use of a P_1 balanced pressure decomposition, rather than a P_2 balanced pressure decomposition as used in this chapter. As noted in section 2.3.3, this leads to a significant reduction in the cost of the Fluidity-ICOM simulations with no significant loss of solution quality. Despite this, for the

[2]MORALS data courtesy of Dr. R. M. B. Young, Department of Physics, University of Oxford.

thermally driven rotating annulus at lower rotation rates, it is clear that Fluidity-ICOM has a very significantly increased cost relative to MORALS, with no observed increase in accuracy.

4.5. Summary

In this chapter the application of dynamic mesh adaptivity to numerical simulations of the thermally driven rotating annulus have been described. Fluidity-ICOM was tested using the stabilised P_1P_1 element pair with a P_2 balanced pressure decomposition. In section 4.1 simulations were conducted on a fixed hyperbolic tangent stretch mesh, in the style of the mesh used by MORALS. In section 4.3, dynamic mesh adaptivity was enabled. Since it is known that MORALS has a favourable comparison with laboratory observations at lower rotation rates, comparison of the fixed mesh simulations against the adaptive mesh simulations represents a particularly challenging test of the method, and enables shortcomings to be identified.

In the fixed mesh simulations it was found that Fluidity-ICOM is able to represent baroclinic instability, with observed wave modes corresponding well to those observed in experiment. At $\omega = 1.0$ rad s^{-1} Fluidity-ICOM shows a favourable qualitative comparison with MORALS. At $\omega = 2.764$ rad s^{-1} Fluidity-ICOM shows a less favourable comparison with MORALS, with a bias towards higher potential energies and a weaker depth averaged stream function, indicating weaker baroclinic eddies in the Fluidity-ICOM simulation. Quantitative comparison of the heat transport observed in the Fluidity-ICOM simulations against available laboratory data shows a favourable comparison at lower rotation rates, and a progressively less favourable comparison at higher rotation rates. In addition, the fixed mesh Fluidity-ICOM simulations exhibit only very weak amplitude vacillation, of a magnitude significantly lower than observed in experiment.

The introduction of dynamic mesh adaptivity leads to a moderate improvement in the potential energy distribution at $\omega = 2.764$ rad s^{-1}, as compared against MORALS. However, quantitative comparison against laboratory data is much less favourable. In particular, a systematically reduced heat transport is observed. A numerical instability is observed in the two regions where the Ekman layer flux impinges upon the lateral boundaries. Dynamic mesh adaptivity under-resolves the flow in these regions, increasing the impact of the instability. The instability disrupts the Ekman layer flux, and degrades the simulated heat transport.

The dynamic mesh adaptive simulations described in this chapter use a metric derived only from the velocity field, and not from the temperature field. It might be expected that deriving a mesh adaptivity metric from the temperature field, and super-imposing this with a metric derived

from the velocity field, may lead to increased accuracy in the otherwise under-resolved side-wall regions of the Ekman layers. Given the high computational cost of these simulations, noted in section 4.4, it is difficult to tune the adaptivity metrics for this problem. It is possible that goal-based adaptivity metrics, directly acting to constrain the error in the system heat transport, may provide additional insight into more appropriate metric choices for this system.

In each of the dynamic mesh adaptive simulations described in section 4.3, mesh optimisation was applied ~ 400 times. Hence each dynamic mesh adaptive simulation was conducted on ~ 400 unrelated meshes. The model is therefore required to be numerically stable, and to run without error, on ~ 400 unrelated unstructured meshes. As observed in this chapter, the presence of any numerical instability, possibly of small magnitude on one mesh, may become significant on later optimised meshes. The need for particular attention in developing robust numerical methods for these applications is therefore clear.

Finally, the computational cost of these simulations was discussed in section 4.4. Fixed and adaptive mesh Fluidity-ICOM simulations of the thermally driven rotating annulus, with these system parameters, are typically of order one hundred times slower than real time when conducted in parallel on 8 MPI processes. MORALS require approximately 1 CPU second per simulated second at an equivalent resolution. Hence the computational cost of Fluidity-ICOM simulations of the thermally driven rotating annulus is, for these system parameters, prohibitive. Hence, while these tests are useful for identifying shortcomings in the application of these method to geophysical systems, it is clear that unstructured dynamic mesh adaptive methods should in general be applied for more challenging systems, such as systems with complex geometry or at higher Reynolds number.

For the remainder of this thesis Fluidity-ICOM will be applied to systems which are inaccessible or challenging to simulate using more traditional structured mesh methods. The extensibility of the model enables systems with complex boundaries and at higher Reynolds numbers to be accessed, albeit at significant computational cost. In the following chapter, Fluidity-ICOM simulations of the thermally driven rotating annulus at much higher Taylor number, and in a domain with sloping bounding topography, will be described.

Chapter 5.

High Taylor number simulations

In the previous chapter potential problems in applying dynamic mesh adaptivity to simulations of the thermally driven rotating annulus were identified and discussed. In particular, it was identified that dynamic mesh adaptive simulations of the thermally driven rotating annulus are sensitive to numerical instability, and have a high computational cost. For the parameters investigated, no significant improvements in the solution quality were observed when moving from hyperbolic tangent stretch meshes, tuned to resolve the boundary layers, to full dynamic mesh adaptive simulations.

This chapter discusses simulations of an annulus system that is known to be challenging to simulate numerically: a system with high Taylor number (i.e. high rotational Reynolds number) and with sloping upper and lower boundaries. This system requires the ability to conform to boundaries that are not aligned with the coordinate system, and the ability to simulate complex three-dimensional flows. Hence, while there are problems with applying dynamic mesh adaptivity to geophysical systems, and while Fluidity-ICOM has a high computational cost relative to structured mesh numerical methods, the extensibility of the model can be exploited to simulate previously inaccessible phenomena.

The simulations in this chapter are based on the experiments described in Wordsworth et al. [2008]. In Wordsworth et al. [2008] thermally driven rotating annulus experiments are conducted at high Taylor number, with and without sloping upper and lower boundaries, in chaotic and turbulent flow regimes. To leading order, linearly sloping boundaries are, for barotropic eddies, equivalent to a β-plane. Hence these experiments are used to model geostrophic turbulence on a β-plane. In Wordsworth et al. [2008] it is noted that the introduction of sloping boundaries promotes the formation of internal jet-like structures. More detailed analysis of the experiments reveals that the eddies and mean zonal flow exchange energy directly, with a net non-local energy transfer from

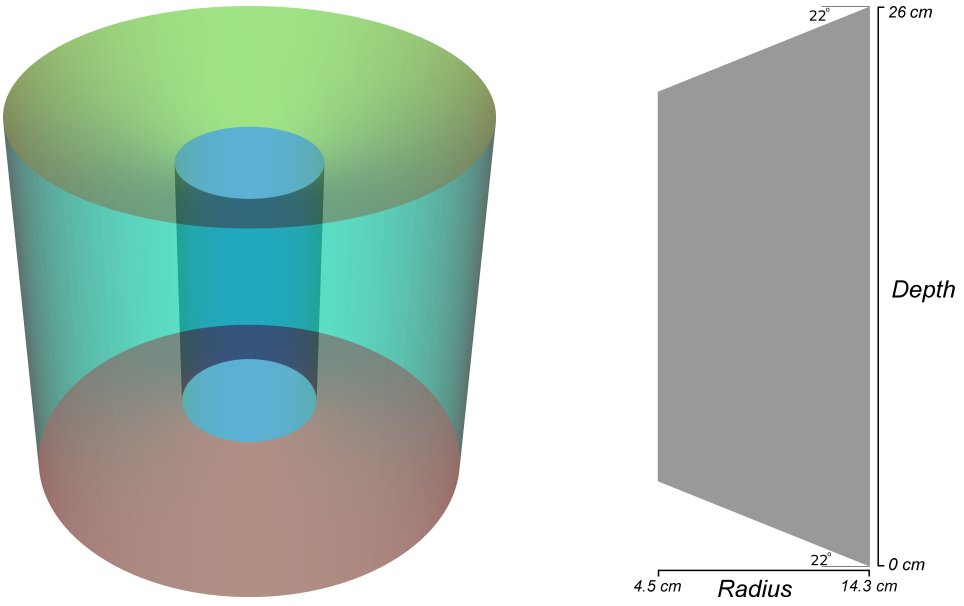

Figure 5.1.: Left: Perspective view of the annular domain with sloping upper and lower boundaries. Right: Vertical slice through the domain, with domain measurements.

high azimuthal wavenumbers to the mean flow. This indicates that the barotropic eddies transfer energy to the mean flow directly, without a more conventional local (in Fourier space) "cascade" type energy transfer mechanism (see, for example, Salmon [1978] and Salmon [1988]).

In this chapter Fluidity-ICOM is used to simulate the thermally driven rotating annulus with sloping upper and lower boundaries and at high Taylor number. These simulations are then used to examine jet formation in a high Taylor number continuously stratified fluid on a β-plane, and to test the conclusions of Wordsworth et al. [2008] concerning the locality of energy transfer. In section 5.1 the model configuration is outlined and the system parameters supplied. In section 5.2 a qualitative overview of Fluidity-ICOM simulations of this system is described. Finally, section 5.3 contains a more detailed analysis of the simulations, with emphasis on the eddy-mean-flow interaction.

5.1. Configuration

The simulations in this chapter are based on the experiments described in Wordsworth et al. [2008] (and also the DPhil work of Wordsworth [2008]). The simulations are conducted in an annular domain with sloping upper and lower boundaries, as shown in figure 5.1. The upper and lower boundaries slope at 22° in the positive-β sense – that is, the effect of the sloping boundaries on barotropic eddies is equivalent to a positive β effect. The fluid parameters are given in table 5.1.

Parameter	Symbol	Value
Minor radius	a	4.5 cm
Major radius	b	14.3 cm
Maximum depth	d	26.0 cm
Kinematic viscosity	ν	2.04×10^{-2} cm^2 s^{-1}
Thermal diffusivity	κ	1.3×10^{-3} cm^2 s^{-1}
Volumetric expansion coefficient	α	3.16×10^{-4} K^{-1}
Gravitational acceleration	g	981 cm s $^{-2}$
Thermal forcing	$\Delta T = T_B - T_A$	2 °C
Rate of rotation	ω	1.3 rad s^{-1}, 2.6 rad s^{-1}
Topographic beta	β	0.095 rad cm^{-1} s^{-1} at $\omega = 1.3$ rad s^{-1}, 0.19 rad cm^{-1} s^{-1} at $\omega = 2.6$ rad s^{-1},
Thermal Rossby number	Θ	0.099 at $\omega = 1.3$ rad s^{-1}, 0.025 at $\omega = 2.6$ rad s^{-1}
Taylor number	Ta	5.7×10^7 at $\omega = 1.3$ rad s^{-1}, 2.3×10^8 at $\omega = 2.6$ rad s^{-1}

Table 5.1.: System parameters used for simulations of the thermally driven rotating annulus with sloping upper lower boundaries and at high Taylor number. The parameters are as used in Wordsworth et al. [2008] (fluid A), with the thermal diffusivity derived from data in Fowlis and Rossby [1964]. The topographic β parameter is computed as $\beta = 4\omega \tan 22° / \overline{d}$, where $\overline{d}$ is the mean tank depth [Wordsworth et al., 2008].

The model configuration is as described in sections 4.1 and 4.2. In particular, these simulations were conducted using the stabilised P_1P_1 finite element pair with a P_2 balanced pressure decomposition, and with a control volume finite element discretisation for temperature. Dynamic mesh adaptivity was applied after every 20 timesteps, using an L_2 interpolation error driven adaptivity metric with a tolerance of $\hat{\epsilon} = 0.1$ (cgs units) for all three velocity field components. Fields were interpolated using Galerkin projection, with strong boundary conditions applied in the interpolation procedure. The simulations were conducted with a varying timestep size, targeting a maximum Courant number of unity.

Two simulations were conducted, at rotation rates of $\omega = 1.3$ rad s^{-1} ($\Theta = 0.099$, $Ta = 5.7 \times 10^7$) and $\omega = 2.6$ rad s^{-1} ($\Theta = 0.025$, $Ta = 2.3 \times 10^8$). Note that the thermal Rossby number of these systems is of a similar order to those described in chapter 4, but that the Taylor number is significantly higher (by a factor of $\gtrsim 10$ in the $\omega = 2.6$ rad s^{-1} case). The simulation at $\omega = 1.3$ rad s^{-1} was integrated in parallel using 24 MPI processes for a total of 8992 s of simulated laboratory time, equivalent to ~ 80 times the Ekman spin-up time $T_{Ek} = d/\sqrt{2\omega\nu}$ [Wordsworth, 2008]. The rotation rate was then increased to $\omega = 2.6$ rad s^{-1}, and the model integrated for a further 4000 s of simulated laboratory time, equivalent to ~ 50 times the Ekman spin-up time. The simulations were conducted on the Oxford Supercomputing Centre "redqueen" machine. The $\omega =$

2.6 rad s^{-1} simulation exhibited a numerical instability, similar to the CFL instability observed in some of the fixed mesh simulations in section 4.1. This instability disappeared after approximately 700 s of simulated laboratory time.

In the following sections, time-averages are taken over the final 1000 s of each simulation. Model fields were output on every 20th model timestep. Hence time-averages were computed using field data from every 20th model timestep.

Combined, these two simulations required approximately 7 processor months of computing time. It was therefore infeasible to perform a more complete suite of simulations, as conducted experimentally in Wordsworth et al. [2008]. The simulations described in this chapter required approximated 1500 s CPU seconds per simulated (lab time) second. This is $\sim 3-24$ times the cost of the simulations described in the previous chapter.

5.2. Overview

5.2.1. Computational meshes

The final meshes for the two simulations are shown in figures 5.2 and 5.3. Node counts for these simulations are shown in figure 5.4.

For the $\omega = 1.3$ rad s^{-1} simulation, the final mesh has 122225 nodes, 672509 elements, a minimum edge length of 0.037 cm, and a maximum edge length of 4.0 cm. For the $\omega = 2.6$ rad s^{-1} simulation, the final mesh has 128737 nodes, 708494 elements, a minimum edge length of 0.042 cm, and a maximum edge length of 4.8 cm. Structured mesh Fluidity-ICOM simulations on hyperbolic tangent stretch MORALS meshes were not numerically stable with these system parameters.

The use of dynamically adaptive meshes by Fluidity-ICOM and (fixed) non-uniform meshes by MORALS makes it difficult to estimate the computational cost of these simulations relative to MORALS. A uniform mesh of global resolution equal to the minimum resolution utilised in the unstructured Fluidity-ICOM meshes would contain $\sim 10^8$ nodes – of order ~ 1000 times as many as used in the dynamic mesh adaptive Fluidity-ICOM simulations. It is likely, however, that the use of an appropriately stretched mesh may enable simulations of these flow regimes with significantly fewer degrees of freedom. An unmodified version of MORALS has a hard-coded upper resolution limit of $32 \times 32 \times 128$ nodes in the radial, vertical, and azimuthal dimensions respectively, corresponding to 131072 degrees of freedom. A slightly modified version of MORALS has been known

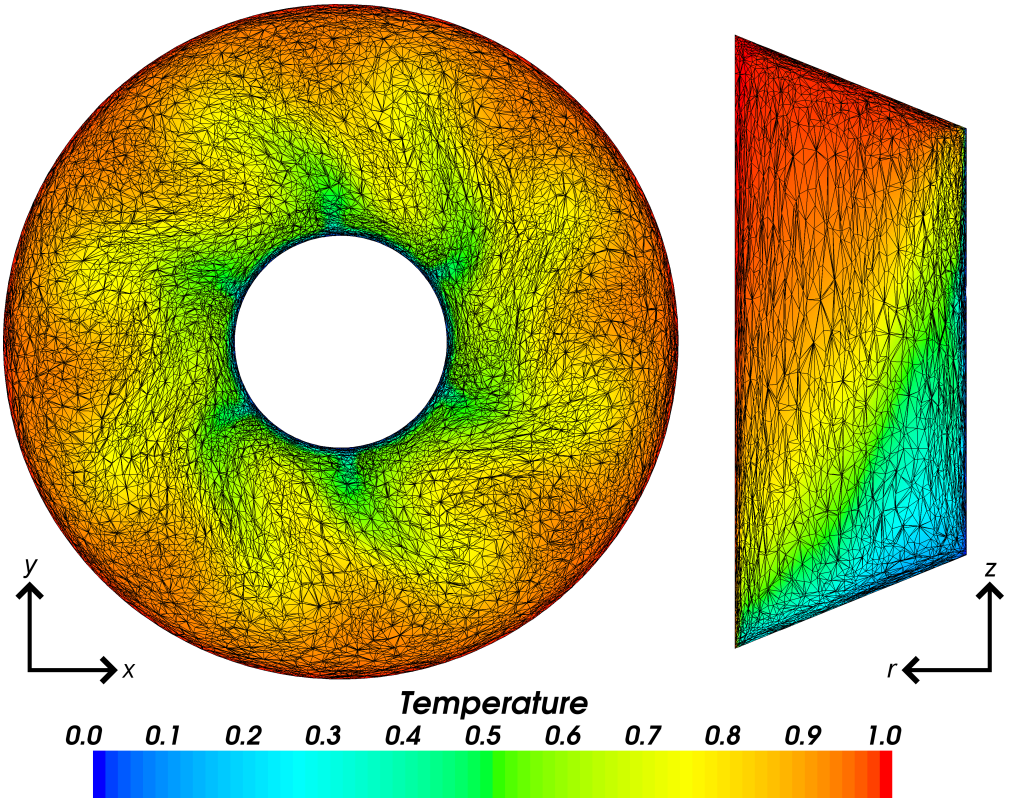

Figure 5.2.: Vertical and horizontal sections through the final mesh of a high Taylor number Fluidity-ICOM simulation at $\omega = 1.3$ rad s^{-1}. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$. This mesh has 122225 nodes.

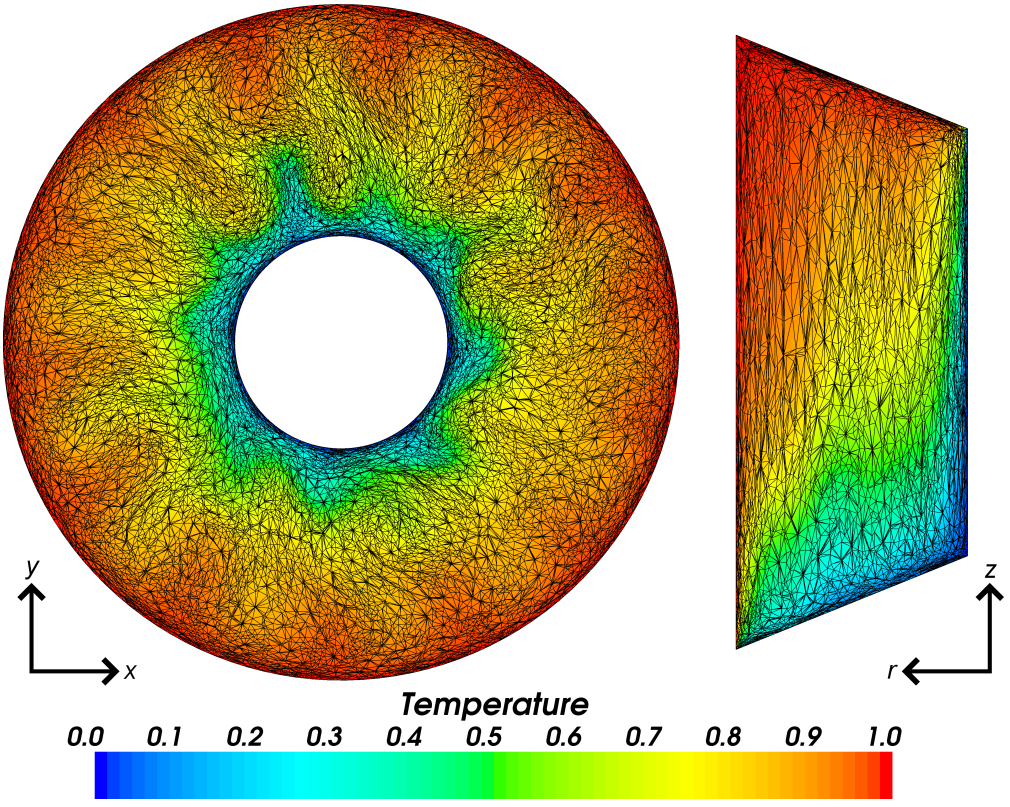

Figure 5.3.: Vertical and horizontal sections through the final mesh of a high Taylor number Fluidity-ICOM simulation at $\omega = 2.6$ rad s^{-1}. The temperature shown here is the normalised temperature, $(T - T_A)/\Delta T$. This mesh has 128737 nodes.

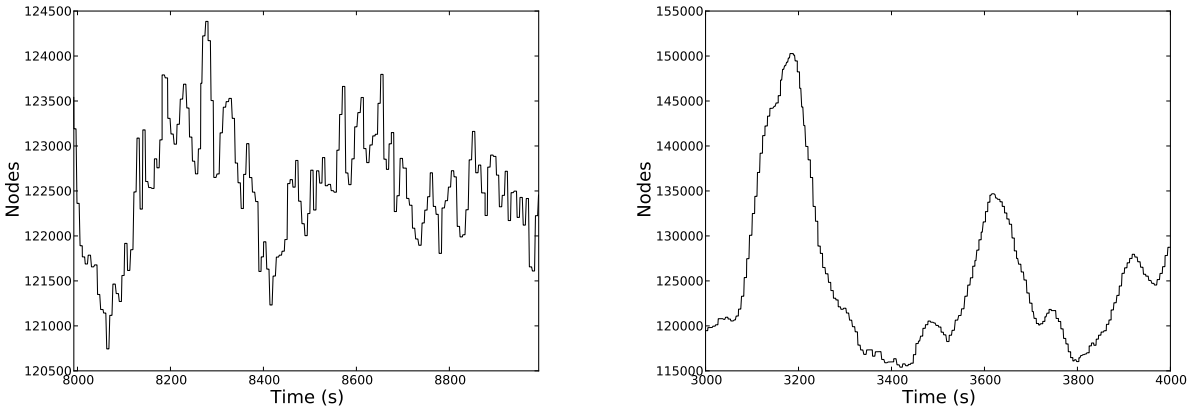

Figure 5.4.: Time varying node counts for the high Taylor number Fluidity-ICOM simulations, shown for the final 1000 s of each simulation. Left: $\omega = 1.3$ rad s^{-1}. Right: $\omega = 2.6$ rad s^{-1}.

to function at resolution $96 \times 96 \times 256$ nodes in the radial, vertical, and azimuthal dimensions respectively, corresponding to 2359296 degrees of freedom (R. Young, personal communication). However, MORALS simulations for the flow regimes described in this chapter are not available, and hence the computational cost of Fluidity-ICOM relative to a comparable structured finite difference model is not known. The presence of sloping topography in these simulations makes any comparison with MORALS particularly difficult.

As noted in section 3.1, Galerkin projection does not necessarily preserve field bounds. The temperature field range observed in these simulations was indeed not bounded by the thermal forcing. The maximum overshoot observed was $0.19\Delta T$. The overshoots were located in the upper inner wall region – the same region in which a numerical instability was observed in section 4.3. While the magnitude of these overshoots is significant, the overshoot typically occurred in only one part in 10^6 of the domain volume. It was noted in the previous chapter (in section 4.3) that the system heat transport for the lower Taylor number configuration may have been affected by numerical issues in this region of the domain. Despite this, interior diagnostics such as the baroclinic wave amplitude and drift rates were not found to be adversely affected. Hence the influence of this numerical artefact on the interior wave dynamics of the system is expected to be very small.

5.2.2. Flow features

The final mid-height velocity fields and the vertical component of the mid-height vorticity are shown for these simulations in figures 5.5 and 5.6. Isopycnals of the final system state are shown in figures 5.7 and 5.8.

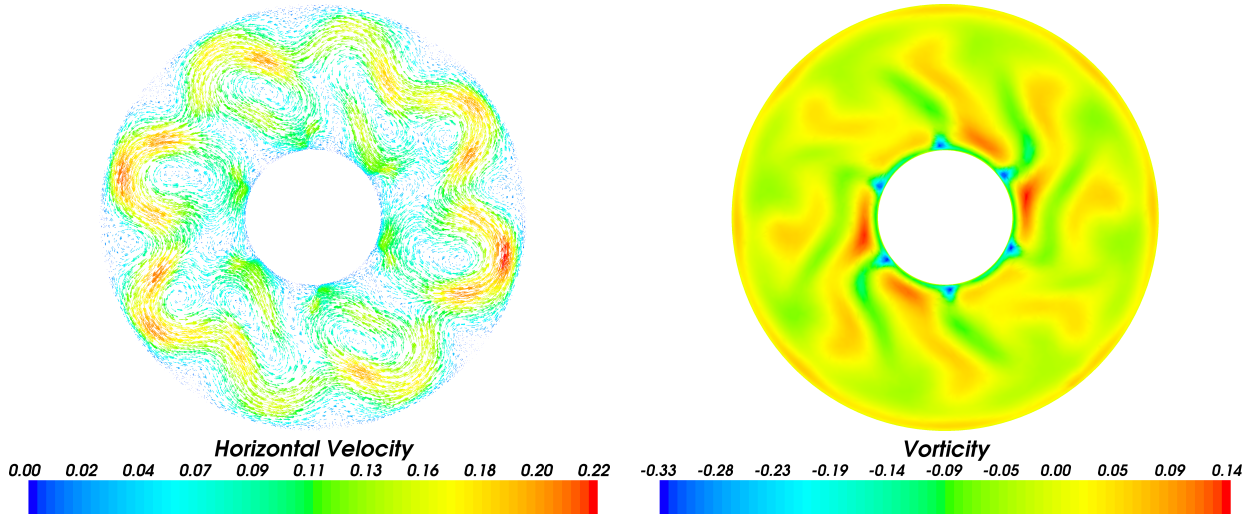

Figure 5.5.: Left: Final mid-height horizontal velocity for a high Taylor number Fluidity-ICOM simulation at $\omega = 1.3$ rad s^{-1}. Velocity units are cm s^{-1}. Right: Vertical component of the mid-height vorticity. Vorticity units are s^{-1}.

Considering the lower rotation rate of $\omega = 1.3$ rad s^{-1}, it can be seen that the system contains two distinct wavenumbers. Towards the outer wall a wave-like mode 8 is observed. This drifts in the retrograde ("westward") direction, consistent with westward phase propagation of Rossby waves on a (topographic) β-plane, where the "westward" direction is opposite in sense to the direction of rotation. Note that this retrograde propogation persists over regions of "eastward" azimuthal mean flow. Towards the inner wall, a more active and unstable mode 6 is observed. Features in this region also drift in the retrograde direction, although with a higher drift speed. These flow features are broadly similar to those observed in Wordsworth [2008]. At the higher rotation rate of $\omega = 2.6$ rad s^{-1} the system is considerably more unstable. Towards the outer wall a wave-like mode 14 is observed. This wave-like region is now unstable, and shedding of eddies is observed. The region towards the inner wall is very active at this rotation rate. Features again propagate in the retrograde direction, with the outer wave-like disturbance drifting at a lower speed than the disturbances in the inner wall region.

The velocity field was interpolated, using collocation, onto a structured MORALS mesh of resolution $32 \times 32 \times 128$ in the radial, vertical, and azimuthal dimensions respectively (with a σ-coordinate stretching in the vertical). The interpolated field was then averaged in the azimuthal direction and in time. Collocation was used in the calculation of diagnostics described in this chapter as it has a greatly reduced computational cost relative to, for example, Galerkin projection. In sections 3.1 and 4.2.3 it was concluded that collocation is associated with a relatively high numerical diffusion as compared with Galerkin projection. However the interpolation error from the single application of collocation, for the purpose of computing diagnostics, is not expected

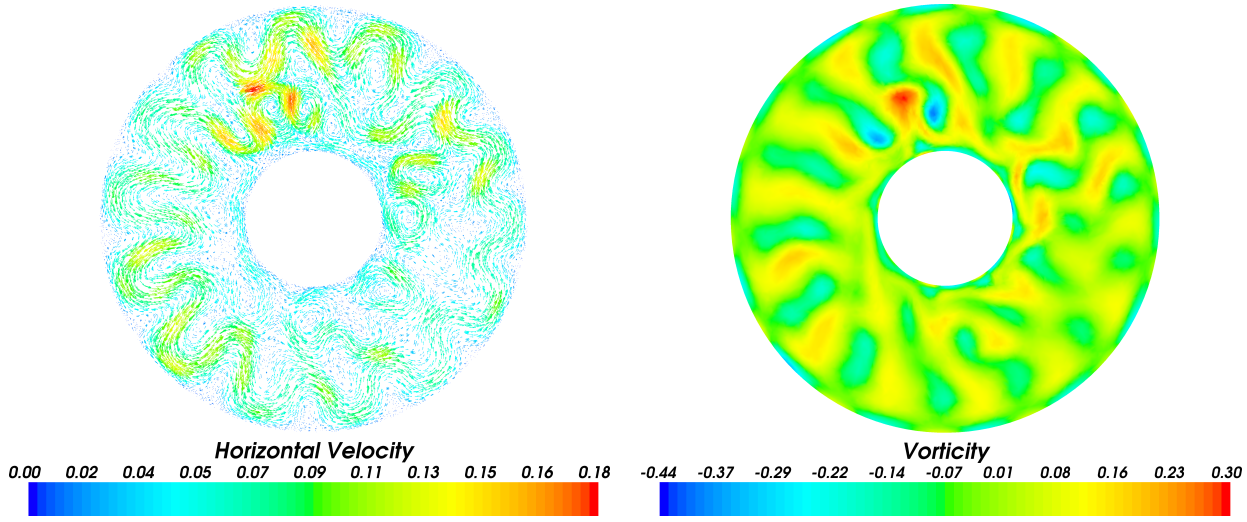

Figure 5.6.: Left: Final mid-height horizontal velocity for a high Taylor number Fluidity-ICOM simulation at ω = 2.6 rad s^{-1}. Velocity units are cm s^{-1}. Right: Vertical component of the mid-height vorticity. Vorticity units are s^{-1}.

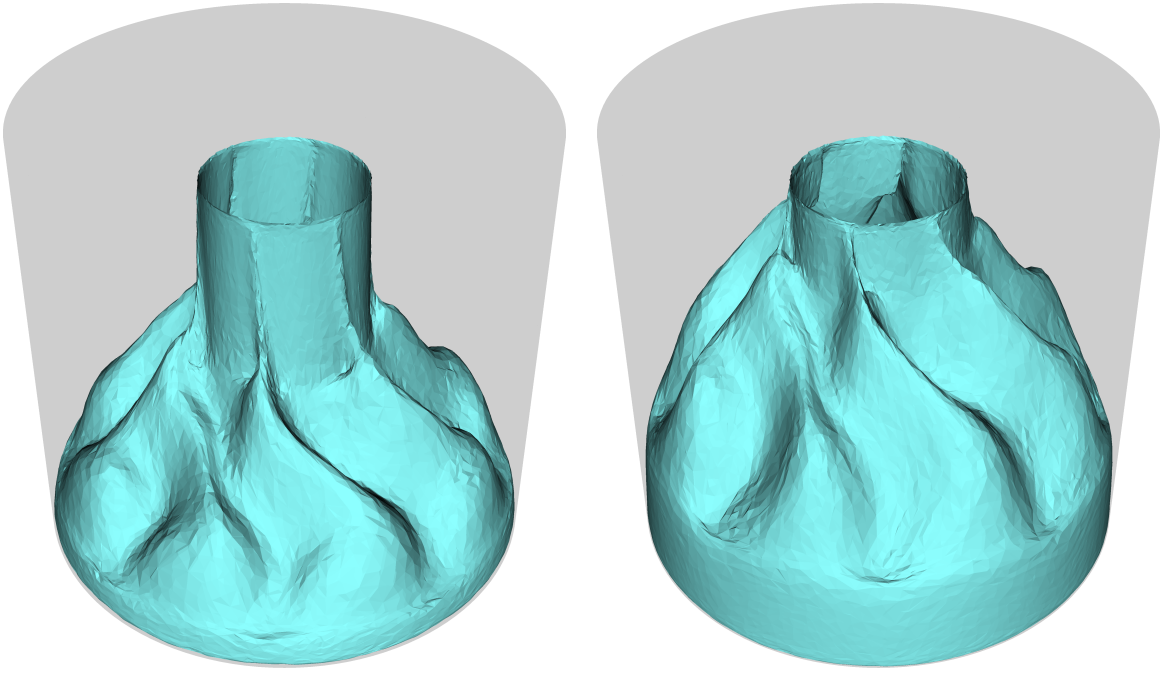

Figure 5.7.: Two isopycnals for the final state of a high Taylor number ICOM simulation at ω = 1.3 rad s^{-1}. Left: $(T - T_A) = \frac{1}{2}\Delta T$ isopycnal. Right: $(T - T_A) = \frac{3}{4}\Delta T$ isopycnal.

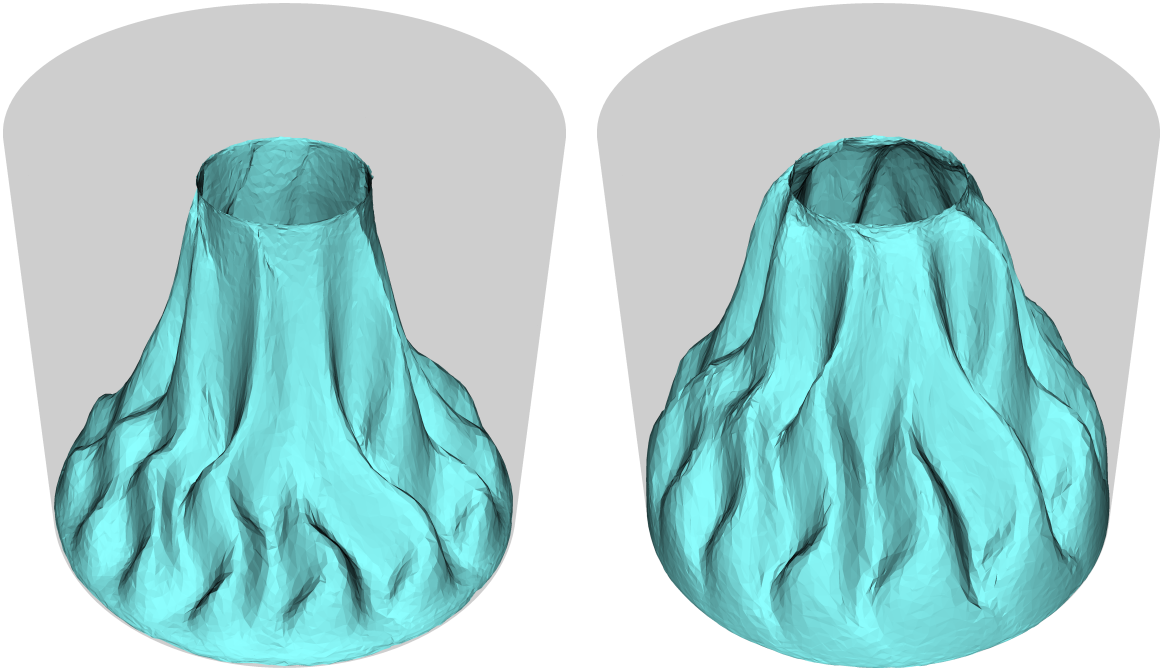

Figure 5.8.: Two isopycnals for the final state of a high Taylor number ICOM simulation at $\omega = 2.6$ rad s^{-1}. Left: $(T - T_A) = \frac{1}{2}\Delta T$ isopycnal. Right: $(T - T_A) = \frac{3}{4}\Delta T$ isopycnal.

to be significantly higher than applying Galerkin projection (see, for example, the collocation / Galerkin projection comparison in Farrell [2009]).

The resulting azimuthally averaged stream function and the azimuthal component of velocity are shown in figures 5.9 and 5.10. Examining the azimuthal component of velocity, a strong vertical shear is observed at both rotation rates. Multiple overturning cells are observed in the azimuthally averaged stream function. The cells have a greater magnitude, and the interior cells are somewhat more well-defined, for the lower rotation rate $\omega = 1.3$ rad s^{-1} case.

Hence in these two simulations, planetary wave jet-like features are observed in the outer wall region. The region towards the inner wall is considerably more active, particularly in the higher rotation rate $\omega = 2.6$ rad s^{-1} case. The $\omega = 2.6$ rad s^{-1} case is observed to be chaotic, but not fully turbulent.

5.3. Analysis

5.3.1. Eddy fluxes

In this and the following sections, we define an eddy mean-flow decomposition of a field q:

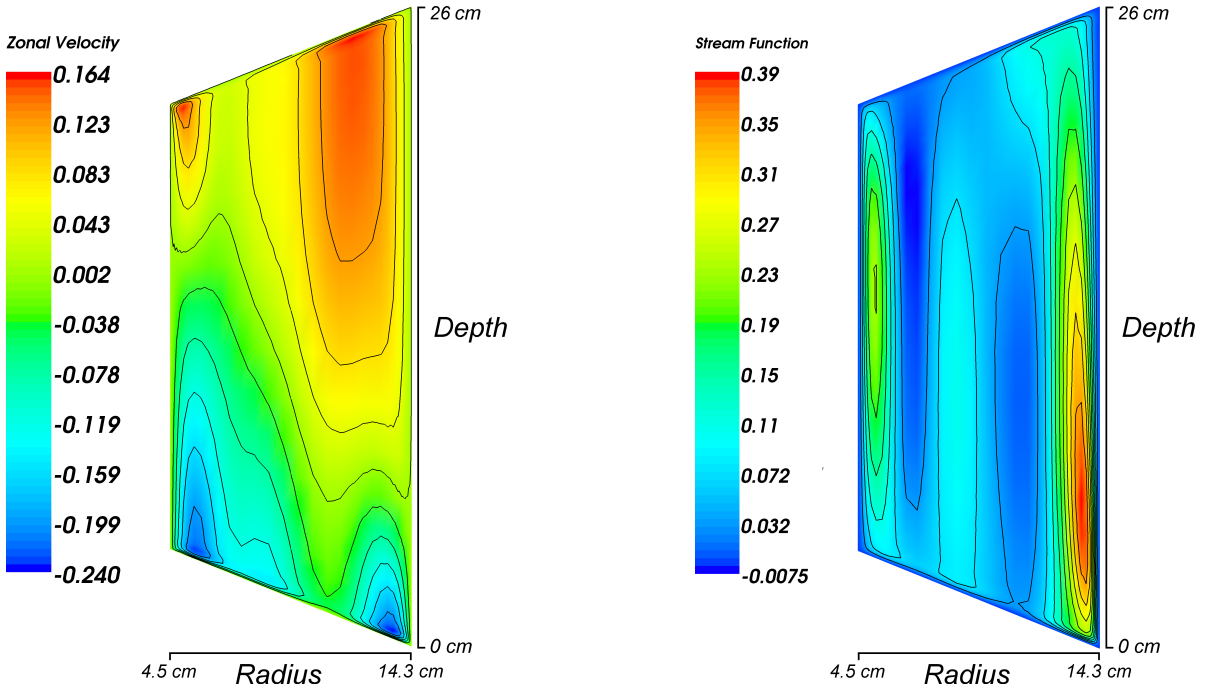

Figure 5.9.: Time and azimuthally averaged velocity field for a high Taylor number Fluidity-ICOM simulation at $\omega = 1.3$ rad s^{-1}. Left: Azimuthal velocity component. Units are cm s^{-1}. Right: Azimuthally averaged stream function as defined by equation (4.3). Units are cm^3 s^{-1}.

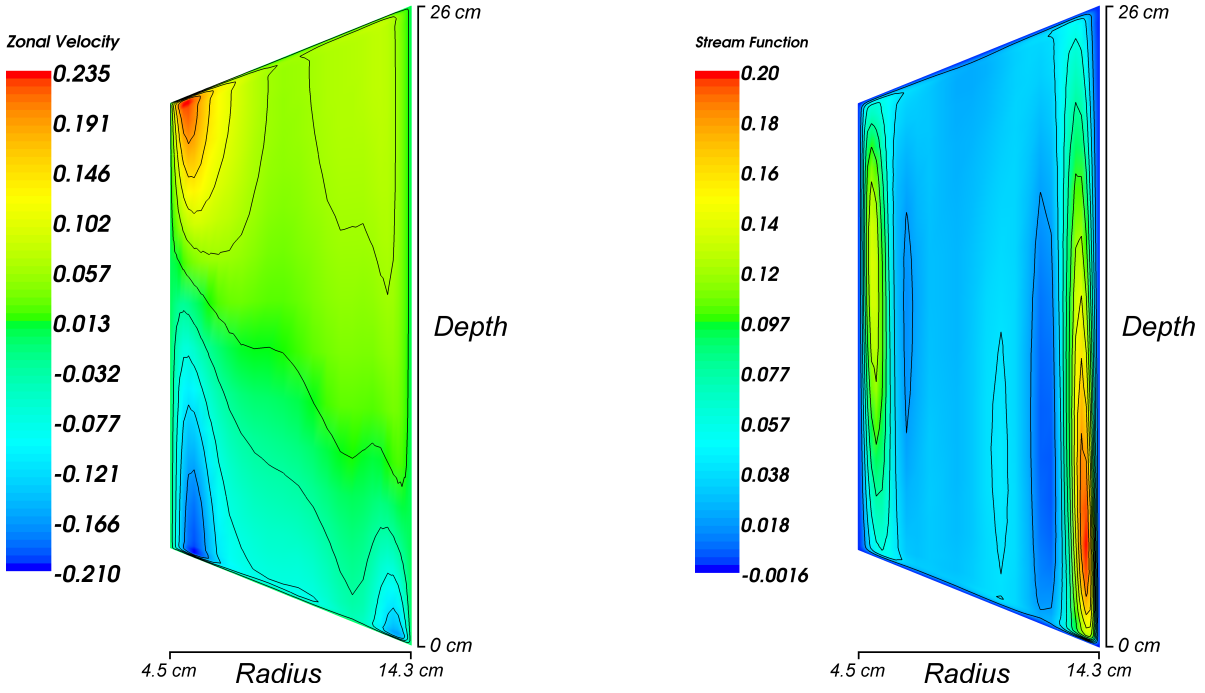

Figure 5.10.: Time and azimuthally averaged velocity field for a high Taylor number Fluidity-ICOM simulation at $\omega = 2.6$ rad s^{-1}. Left: Azimuthal velocity component. Units are cm s^{-1}. Right: Azimuthally averaged stream function as defined by equation (4.3). Units are cm^3 s^{-1}.

$$q = \overline{q} + q' \tag{5.1}$$

where $\overline{q}$ is the azimuthal mean and q' is the deviation from this mean. We further define $\langle \ldots \rangle$ to denote "time-averaged". For simplicity only eddy fluxes derived from the mid-height fields are considered. A two dimensional MORALS mesh of resolution 32×128 in the radial and horizontal dimensions respectively was generated, and the fields interpolated onto this mesh via collocation.

The time and azimuthally averaged mean azimuthal velocity $\langle \overline{u_\theta} \rangle$, together with the time and azimuthally averaged eddy momentum flux $\left\langle \overline{u_r' u_\theta'} \right\rangle$, are shown for the mid-height velocity fields of the two Fluidity-ICOM simulations in figures 5.11 and 5.12. Also shown is the time and azimuthally averaged eddy heat flux $\left\langle \overline{u_r' T'} \right\rangle$.

In an eddy mean-flow decomposition of the form (5.1) the (negative) divergence of the eddy momentum correlations appears as a forcing term in the equation for the evolution of the mean flow [Vallis, 2006]. $\left\langle \overline{u_r' u_\theta'} \right\rangle$ represents radial transport of azimuthal momentum by the eddies. Hence a correlation between convergence of $\left\langle \overline{u_r' u_\theta'} \right\rangle$ and the mean azimuthal velocity $\langle \overline{u_\theta} \rangle$ suggests a forcing of the mean flow by the eddies.

For the $\omega = 1.3$ rad s^{-1} case, convergence of the eddy momentum flux is observed at $r = 12.0$ cm. This corresponds well with a maximum positive azimuthal velocity in this region. At $r = 6.2$ cm divergence of the eddy momentum flux is observed, occurring in a region of negative azimuthal velocity. These observations suggest that the two jet-like features observed in the azimuthal velocity may be forced by the eddies. A further convergence of the eddy momentum flux is observed at $r = 5.3$ cm, albeit with small magnitude, indicating that the negative azimuthal velocity jet loses momentum to the eddies in this region.

For the $\omega = 2.6$ rad s^{-1} case, convergence of the eddy momentum flux is observed at $r = 12.4$ cm. This corresponds well with a maximum in the positive azimuthal velocity in this region. At $r = 8.7$ cm divergence of the eddy momentum flux is observed, and corresponds with a minimum in the positive azimuthal velocity. At $r = 7.2$ cm convergence in the eddy momentum flux is again observed, and corresponds with a second maximum in the positive azimuthal velocity. These observations suggest a more complicated interaction between the eddies and the mean flow, with the jets being forced by the eddies at some radii, and losing energy to the eddies at others. The alternating sign of the eddy momentum flux may be associated with the shedding of eddies by the outer wave-like feature, noted in section 5.2.2. At lower radii, a divergence of the eddy momentum

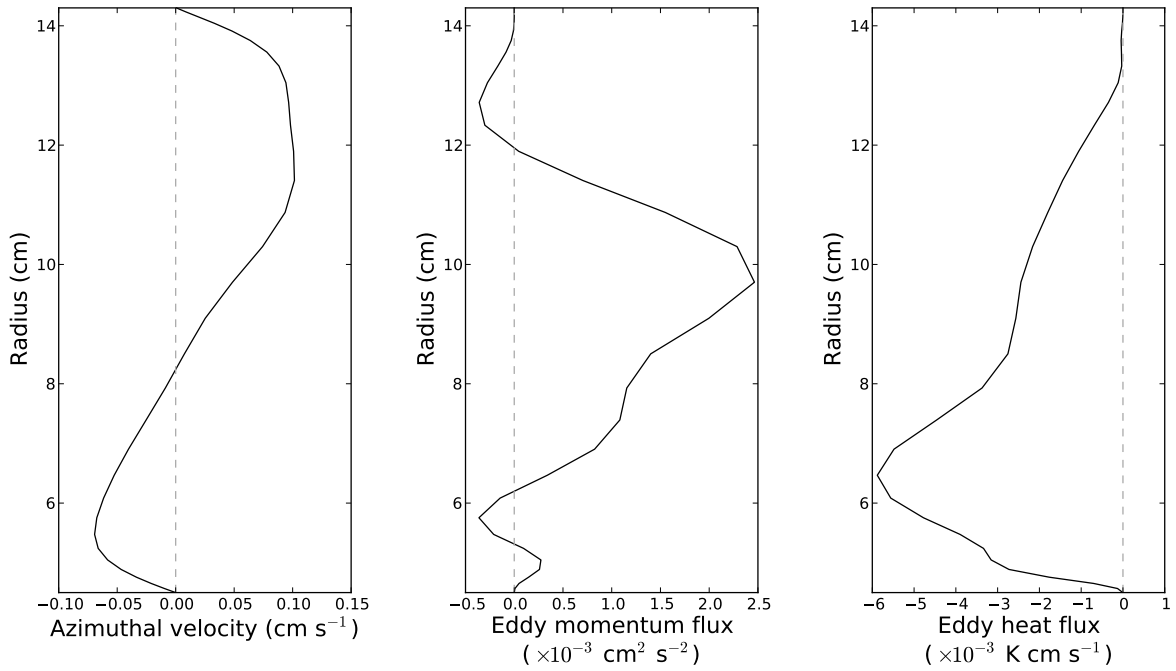

Figure 5.11.: Time and azimuthally averaged mid-height azimuthal velocity, eddy momentum flux, and eddy heat flux, for a high Taylor number Fluidity-ICOM simulation at $\omega = 1.3$ rad s^{-1}.

flux at $r = 5.8$ cm corresponds with a negative azimuthal velocity, indicating that the inner wall jet is forced by the eddies.

Note that these results are not entirely consistent with experimental observations. In Wordsworth [2008], experiments conducted with sloping topography always showed a correlation between the sign of eddy momentum flux convergence and the sign of the azimuthal velocity – the experimental observations always suggested that the jets were forced by the eddies. The functional form of the eddy momentum fluxes also does not appear as in experiment.

The eddy heat flux is strictly negative at both $\omega = 1.3$ rad s^{-1} and $\omega = 2.6$ rad s^{-1}. Hence the eddies everywhere act to ventilate the system, and to transport heat from the warm outer wall to the cooled inner wall – i.e. the eddy heat flux is strictly down-gradient. This is consistent with the experimental observations described in Wordsworth [2008].

5.3.2. Spectral analysis

For systems with multiple scales in the radial dimension it is necessary to consider the system spectra in both the radial and azimuthal dimensions (rather than simply the azimuthal dimension, as considered in computing Fourier amplitudes in chapter 4). In a two dimensional annulus geometry

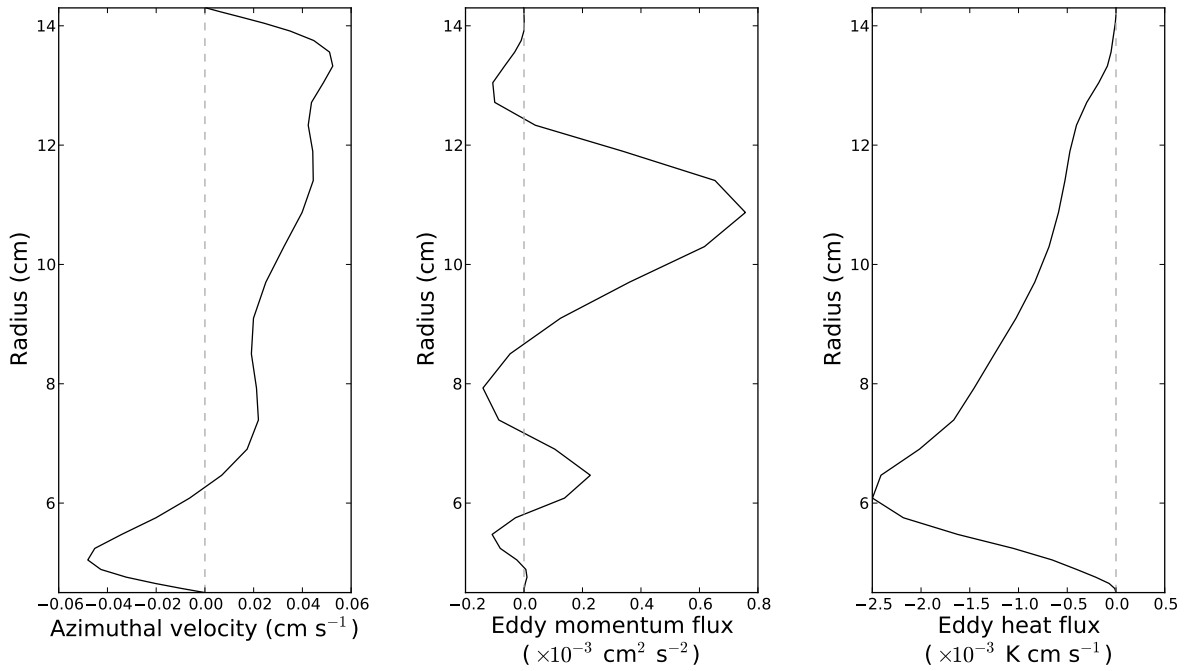

Figure 5.12.: Time and azimuthally averaged mid-height azimuthal velocity, eddy momentum flux, and eddy heat flux, for a high Taylor number Fluidity-ICOM simulation at $\omega = 2.6$ rad s^{-1}.

the equivalent orthogonal decomposition of a scalar field q takes the form [Arfken and Weber, 2001, Wordsworth, 2008]:

$$q(r,\theta) = \sum_{m=-\infty}^{m=\infty} \sum_{n=1}^{n=\infty} \left[a^J_{mn} J_m(\gamma_{mn} r) + a^Y_{mn} Y_m(\gamma_{mn} r) \right] e^{im\theta}, \tag{5.2}$$

where J_m and Y_m are Bessel functions of the first and second kind respectively, γ_{mn} are constants dependent upon the boundary conditions for q, and a^J_{mn} and a^Y_{mn} are the coefficients to be determined. a_{mn} is the amplitude of mode m in the azimuthal dimension and n in the radial dimension. Note that the a_{mn} include the azimuthal phase information, and hence are complex. If q satisfies the boundary conditions:

$$q(r = a, \theta) = q(r = b, \theta) = 0, \tag{5.3}$$

then this leads to:

$$q(r,\theta) = \sum_{m=-\infty}^{m=\infty} \sum_{n=1}^{n=\infty} a_{mn} R(\gamma_{mn} r)\, e^{im\theta}, \tag{5.4}$$

where:

$$R(\gamma_{mn} r) = J_m(\gamma_{mn} r) + \frac{a^Y_{mn}}{a^J_{mn}} Y_m(\gamma_{mn} r). \tag{5.5}$$

From equation (5.2) and the boundary conditions (5.3) it follows that γ_{mn} is, for a given m, the nth root of the transcendental equation [Arfken and Weber, 2001]:

$$J_m(\gamma_{mn} a)\, Y_m(\gamma_{mn} b) - J_m(\gamma_{mn} b)\, Y_m(\gamma_{mn} a) = 0 \tag{5.6}$$

and a^Y_{mn}/a^J_{mn} is given by:

$$\frac{a^Y_{mn}}{a^J_{mn}} = -\frac{J_m(\gamma_{mn} a)}{Y_m(\gamma_{mn} a)} = -\frac{J_m(\gamma_{mn} b)}{Y_m(\gamma_{mn} b)}. \tag{5.7}$$

These basis functions were computed for the annular geometry in figure 5.1. The coefficients γ_{mn} were found by solving equation (5.6) using the SciPy [1] Python library, via the scipy.linalg.fsolve routine. The radial basis functions R_m were then normalised, using Gaussian quadrature via the scipy.integrate.quadrature routine, to yield an orthonormal set.

The components of the velocity field satisfy the boundary conditions (5.3), and hence can be decomposed according to equation (5.4). From this one can define a spectral kinetic energy:

$$E = \tfrac{1}{2}\left(u_r u_r^* + u_\theta u_\theta^*\right), \tag{5.8}$$

[1] Available from http://www.scipy.org/.

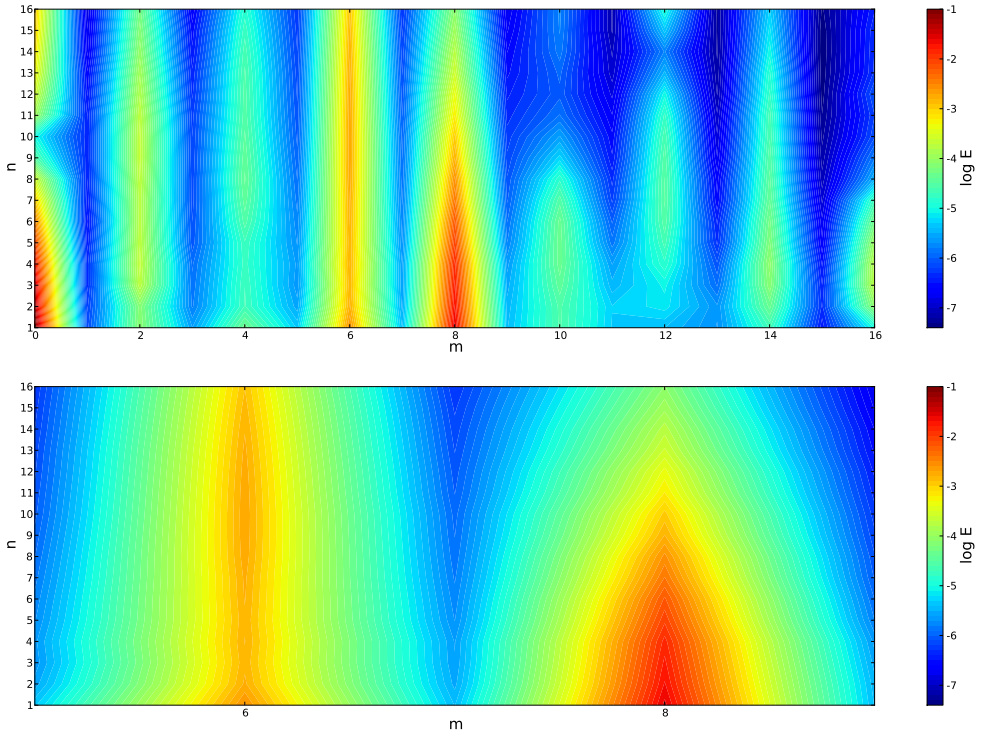

Figure 5.13.: Time-averaged mid-height kinetic energy spectrum for a high Taylor number Fluidity-ICOM simulation at $\omega = 1.3$ rad s^{-1}. m and n are the azimuthal and radial mode numbers respectively. Above: The positive azimuthal modes. Below: Section focusing on the dominant azimuthal modes. Note in particular the secondary peak with azimuthal mode $m = 6$ and radial mode $n \sim 10$.

where u_r and u_θ are the velocity components in the radial and azimuthal dimensions respectively, and $(\ldots)^*$ denotes complex conjugate. This allows one to compute a two dimensional energy spectrum E_{mn}.

The energy spectra were computed for the two Fluidity-ICOM simulations at rotation rates $\omega = 1.3$ rad s^{-1} and $\omega = 2.6$ rad s^{-1}. The spectra were computed for the mid-height fields, with the fields interpolated via collocation onto a structured two dimensional MORALS mesh of resolution 32×128 nodes in the radial and azimuthal dimensions respectively. These spectra are shown in figures 5.13 and 5.14.

In the $\omega = 1.3$ rad s^{-1} simulation three primary peaks in the kinetic energy spectrum are observed. There is a large maximum in the mean, mode zero, azimuthal mode, spread over a range of lower radial modes. A second maximum is observed at mode 8 in the azimuthal dimension, peaking at mode 1 in the radial dimension. This is identified as the outer wall wave-like feature in this system. An additional maximum at mode 6 in the azimuthal dimension is observed, with a peak in the radial dimension at mode ~ 10. This is identified as the inner wall train of eddies in this system. A series of secondary maxima are also observed in all even numbered azimuthal modes, in each case over a broad range of radial modes.

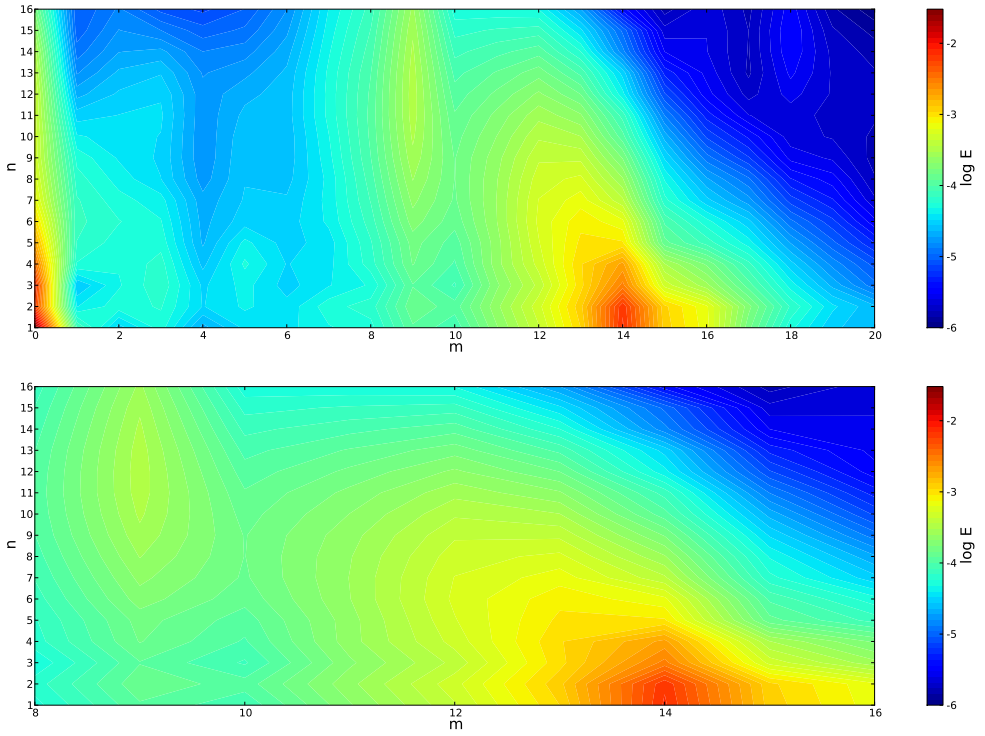

Figure 5.14.: Time-averaged mid-height kinetic energy spectrum for a high Taylor number Fluidity-ICOM simulation at $\omega = 2.6$ rad s^{-1}. m and n are the azimuthal and radial mode numbers respectively. Above: The positive azimuthal modes. Below: Section focusing on the dominant azimuthal modes. Note in particular the secondary peak at azimuthal mode $m = 9$ and radial mode $n \sim 12$.

In the $\omega = 2.6$ rad s^{-1} simulation three primary peaks in the kinetic energy spectrum are again observed. There is again a large maximum in the mean azimuthal mode, over a range of lower radial modes. A large maximum is observed in the mode 14 azimuthal mode and mode 1 radial mode. A much weaker additional maximum is observed in the mode 9 azimuthal mode, peaking at higher $\sim$ 12 modes in the radial dimension. These can again be identified as the outer and inner wall features in this flow respectively. No significant peaks in harmonics are observed at this rotation rate.

These spectra can be compared against the experimental measurements in Wordsworth [2008] and Wordsworth et al. [2008]. The simulated results are broadly consistent with the experimental observations. However, the spectra observed in the simulations appear to match well with the spectra of experiments conducted at lower rotation rates (see, for example, figure 6.15 of Wordsworth [2008]). Hence it is likely that the model is adding artificial diffusion, possibly due to a lack of model resolution or as a result of the application of dynamic mesh adaptivity.

Following Wordsworth [2008], we now consider resonant triad energy exchange in this system. The analysis is simplified by considering only barotropic azimuthal modes. If we consider only the advection term in the momentum equation (2.3a), then it follows from a Fourier series expansion

of u_r and u_θ, together with a Fourier transform of the momentum equation that [James et al., 1981, Wordsworth et al., 2008]:

$$\mathcal{P}_m = \frac{\partial \langle E_m \rangle}{\partial t} = -\sum_{m=p+q} \left\langle T_{mpq} \right\rangle, \tag{5.9}$$

where:

$$T_{mpq} = \tfrac{1}{2}\left(u^*_{r,m} a_{pq} + u^*_{\theta,m} b_{pq} + u_{r,m} a^*_{pq} + u_{\theta,m} b^*_{pq}\right), \tag{5.10}$$

$$a_{pq} = u_{r,p} \frac{\partial u_{r,q}}{\partial r} + \frac{iq u_{\theta,p} u_{r,q}}{r} - \frac{u_{\theta,p} u_{\theta,q}}{r}, \tag{5.11}$$

$$b_{pq} = u_{r,p} \frac{\partial u_{\theta,q}}{\partial r} + \frac{iq u_{\theta,p} u_{\theta,q}}{r} + \frac{u_{r,p} u_{\theta,q}}{r}. \tag{5.12}$$

$\mathcal{P}_m$ is the time-averaged energy flux into the barotropic azimuthal mode m resulting from resonant triad interactions from all barotropic azimuthal modes. T_{mpq} is the energy flux into the barotropic azimuthal mode m resulting from a resonant triad interaction between the barotropic azimuthal modes p and q.

The velocity field was interpolated, using collocation, onto a structured MORALS mesh of resolution $32 \times 32 \times 128$ in the radial, vertical, and azimuthal dimensions respectively (with a σ-coordinate stretching in the vertical). From this, the time-averaged resonant triad coefficients $\left\langle T_{mpq} \right\rangle$ were computed from the depth averaged velocity fields. Examples of the resulting energy fluxes $\mathcal{P}_m$, computed from the triad coefficients T_{mpq} for restricted p and q, are shown in figures 5.15 and 5.16. In these figures dominant local interactions would lead to large amplitudes in and adjacent to the shaded regions – that is, the modes would gain energy due to resonant triad interactions with modes of similar wavenumber. By contrast, dominant non-local interactions would lead to large amplitudes far from the shaded regions. As an example, in figure 5.15 the mean flow (mode zero) gains energy due to triad interactions with modes -8 to -5 and $+5$ to $+8$, while the energy gain due to triad interactions with modes -4 to $+4$ is more than an order of magnitude smaller. This indicates predominantly non-local energy transfer. The local energy fluxes, resulting from the triad coefficients T_{mpq} with one of $p = m \pm 1$ or $q = m \pm 1$, are shown in figure 5.17.

For the $\omega = 1.3$ rad s^{-1} simulation the dominant exchange is a net transfer of energy from the dominant azimuthal modes (modes 6 and 8, contained within the upper right sub-figure) to the mean azimuthal flow. Other exchanges are at least an order of magnitude lower in amplitude. A small local transfer (left sub-figure in figure 5.17) is observed from mode 8, but no significant exchange to the lowest wavenumbers or mean flow is observed – that is, the lowest wavenumbers do not gain significant energy from these interactions.

For the $\omega = 2.6$ rad s^{-1} simulation the dominant exchanges are again a net transfer of energy from the dominant azimuthal modes (modes 9 and 14, contained within the lower left and lower right sub-figures in figure 5.16) to the mean flow. There is also in general a significant transfer from the dominant azimuthal modes to small wavenumbers. The local energy exchanges (right figure in figure 5.17) are now somewhat more significant than in the $\omega = 1.3$ rad s^{-1} case. There is a suggestion of local energy transfer from mode 14, but again no local transfer to the lowest wavenumbers is observed.

These results can be compared against the experimental measurements in Wordsworth [2008] and Wordsworth et al. [2008]. The simulated results are again broadly consistent with the experimental observations (see, for example, figure 6.16 of Wordsworth [2008]). However, once again, the simulation results compare well against experiments conducted at lower rotation rate.

The Rossby radius of deformation $Ro = \sqrt{g\alpha\Delta T d}/\omega$ [Wordsworth, 2008] for these simulations are 3.1 cm and 1.5 cm for the $\omega = 1.3$ rad s^{-1} and $\omega = 2.6$ rad s^{-1} configurations respectively. These correspond to azimuthal modes $\gtrsim 9$ at $\omega = 1.3$ rad s^{-1}and $\gtrsim 18$ at $\omega = 2.6$ rad s^{-1}. Significant local energy transfer between barotropic-baroclinic energy exchange is not expected on scales significantly larger than the deformation scale – mixed barotropic-baroclinic triads (between two baroclinic modes and one barotropic mode) exchange energy only between the two baroclinic modes on these scales [Salmon, 1988]. Hence significant barotropic-baroclinic energy exchange is not expected, in these simulations, for azimuthal wavenumbers lower than $\lesssim 9$. It is concluded from this analysis that the primary energy transfer is from the higher wavenumbers to the mean flow, via a direct non-local transfer of energy.

5.4. Summary

In this chapter the first numerical simulations of the thermally driven rotating annulus at high Taylor number ($\gtrsim 10^8$), combined with sloping upper and lower boundaries, have been described. These simulations were conducted using fully unstructured dynamic mesh adaptivity. The additional

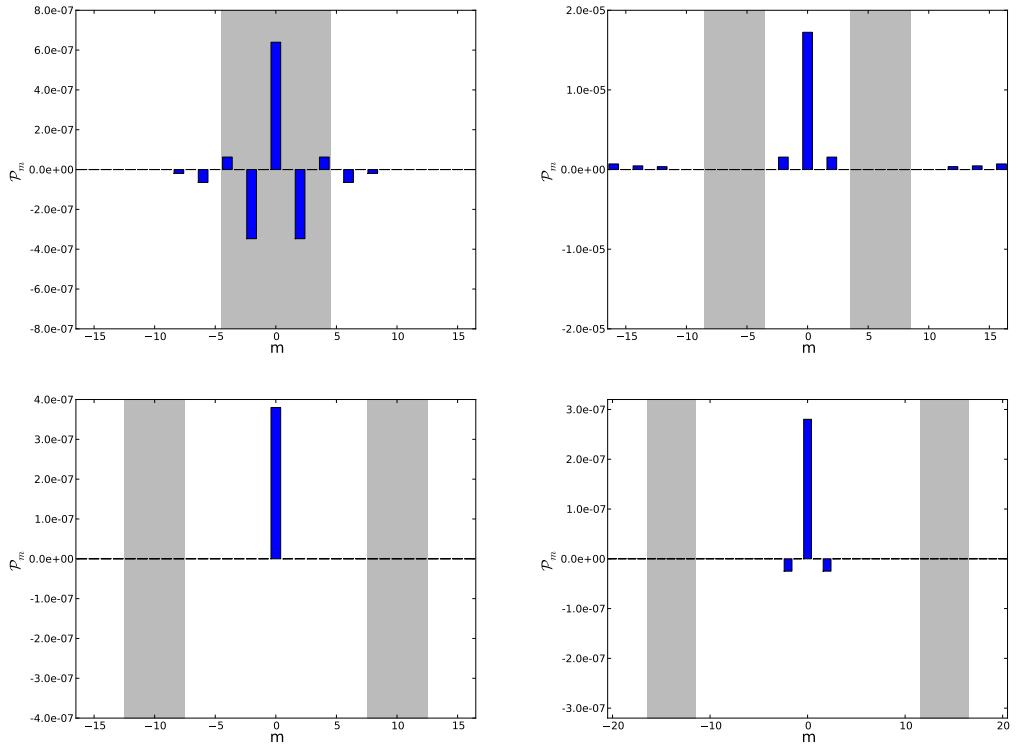

Figure 5.15.: Time and radially averaged resonant triad kinetic energy flux $\mathcal{P}_m$ for a high Taylor number Fluidity-ICOM simulation at $\omega = 1.3$ rad s^{-1}. In each case, only the triad coefficients T_{mpq} with p and q in the shaded region are considered. Upper left: $p, q \in [-4, 4]$. Upper right: $|p|, |q| \in [5, 8]$. Lower left: $|p|, |q| \in [9, 12]$. Lower right: $|p|, |q| \in [13, 16]$.

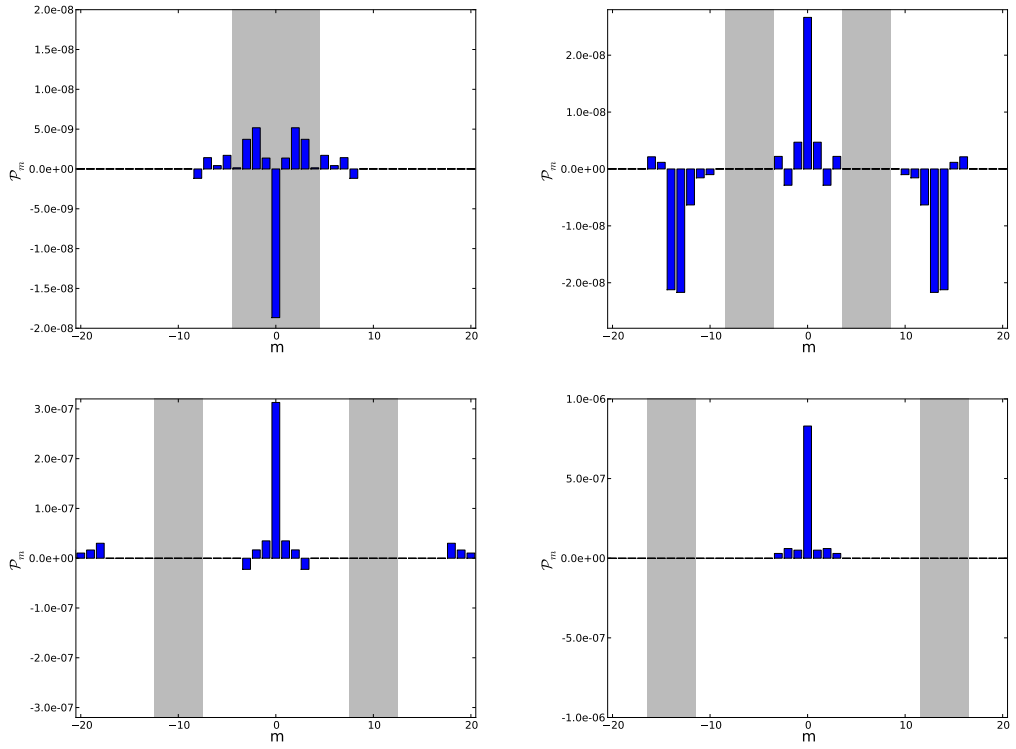

Figure 5.16.: Time and radially averaged resonant triad kinetic energy flux $\mathcal{P}_m$ for a high Taylor number Fluidity-ICOM simulation at $\omega = 2.6$ rad s^{-1}. In each case, only the triad coefficients T_{mpq} with p and q in the shaded region are considered. Upper left: $p, q \in [-4, 4]$. Upper right: $|p|, |q| \in [5, 8]$. Lower left: $|p|, |q| \in [9, 12]$. Lower right: $|p|, |q| \in [13, 16]$.

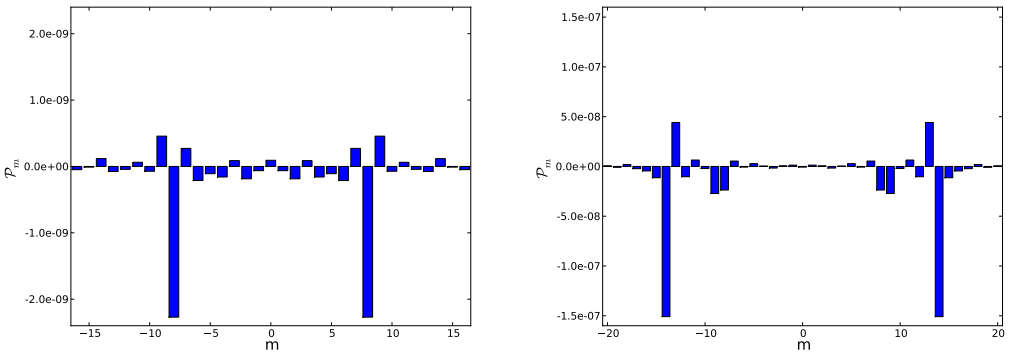

Figure 5.17.: Time and radially averaged local resonant triad kinetic energy flux $\mathcal{P}_m$ for high Taylor number Fluidity-ICOM simulations. The triad coefficients T_{mpq} are restricted to the local modes with $p = m \pm 1$ or $q = m \pm 1$. Left: $\omega = 1.3$ rad s^{-1}. Right: $\omega = 2.6$ rad s^{-1}.

geometric flexibility offered by the unstructured mesh methods enables simulations in complex geometries, such as in domains with upper and lower boundary topography, without any special considerations. Furthermore, dynamic mesh adaptivity was utilised to enable simulations at vastly reduced computational cost, relative to simulations conducted on uniform meshes. Simulations on non-uniform boundary layer resolving MORALS meshes, with comparable numbers of degrees of freedom to the dynamic mesh adaptive simulations, were not numerically stable with these system parameters. While the Fluidity-ICOM simulations described in this chapter undoubtedly have a high computational cost, the numerical methods implemented in Fluidity-ICOM allow simulations of this particularly challenging system to be attempted. The computational cost of these simulations relative to an existing finite difference model, and in particular a finite difference model with a non-uniform grid (such as MORALS), was not investigated. This would enable a more precise assessment of the competitiveness of Fluidity-ICOM in these higher Reynolds number flow regimes.

The simulations offer a very favourable qualitative comparison with respect to previous experimental observations. In particular, the flow is observed to be divided into two distinct regions, one towards the inner wall and one towards the outer wall, with differing azimuthal wavenumbers in each region. Towards the outer wall a planetary wave jet-like structure is observed. The inner wall region is considerably more active. Analysis of the energy spectra and the barotropic resonant triads indicate that the numerical simulations are, when compared against experimental observations and for the external parameters used, biased towards lower Taylor numbers. Hence it is likely that implicit numerical diffusion, possibly due to a lack of model resolution or resulting from the mesh-to-mesh interpolation step in the application of dynamic mesh adaptivity, is significant for these simulations.

Further analysis of the barotropic resonant triads indicates a direct non-local energy transfer, from the dominant azimuthal modes to the mean flow and low azimuthal wavenumbers. This confirms the experimental observations of Wordsworth et al. [2008]. There is some evidence for a small local energy transfer from the dominant azimuthal wavenumber towards lower wavenumbers, although this is of a relatively low magnitude and is not observed to transfer energy to the lowest wavenumbers. A local energy cascade mechanism is therefore not observed in these simulations.

Owing to the high computational cost of the numerical simulations described in this chapter, it was infeasible to attempt still higher Taylor numbers. Simulations conducted in a more fully turbulent regime would clearly be of interest. Furthermore, the sloping upper and lower boundaries imposed in these simulations act as a topographic β-plane, but such a topographic β effect applies, strictly speaking, only to barotropic flows. The numerical simulations were observed to be highly baroclinic. Further simulations conducted with flat boundaries and with a true β-plane would therefore be of interest.

Chapter 6.

Conclusions

This thesis details the application of unstructured mesh finite element methods with dynamic mesh adaptivity, using the Imperial College Ocean Model (Fluidity-ICOM), to numerical simulations of the thermally driven rotating annulus. The thermally driven annulus system was used to develop and test novel numerical schemes, and unstructured dynamic mesh adaptive methods were used to enable simulations of a particularly challenging rotating annulus system. This chapter provides a summary of the work described in this thesis, together with a number of topics for possible future research in this area.

6.1. Summary

The methods utilised throughout this thesis, of arbitrary anisotropic unstructured meshes, and dynamic optimisation of the meshes as the flows develop, undoubtedly have significant potential in the field of geophysical fluid dynamics modelling. The ability to conform to arbitrary bounding topography enables simulations in highly complex geometries, without any special considerations – the extensibility of the finite element method in this regard is an inherent feature of the approach. Furthermore, computational meshes of wildly varying resolution can be used, possibly with highly anisotropic elements, to enable more accurate simulations per mesh degree of freedom. Varying the mesh resolution dynamically allows the application of these methods to transient systems. Hence unstructured dynamic mesh adaptivity has the potential for efficient simulations of transient phenomena, with multiple flow scales, in complex geometries. There is interest in testing the applicability of these methods to geophysical systems, and in particular to ocean modelling, as part of the Imperial College Ocean Model.

However, ocean modelling using unstructured mesh finite element methods is significantly less well-established, and in particular less well tested, than more traditional structured mesh finite difference methods. Even the choice of discretisation used by the dynamical core is a topic of active research: P_1P_1 elements are discussed in Piggott et al. [2008a], hexahedral elements with piecewise linear velocity basis functions and piecewise constant pressure basis functions (the Q_1P_0 element pair) are discussed in Ford et al. [2004a], and more recently the mixed discontinuous-continuous $P_{1\mathrm{DG}}P_2$ element is introduced in Cotter et al. [2009b]. A method for the efficient solution of the elliptic pressure projection equation on meshes that are unstructured in the vertical in tight aspect ratio domains has only very recently been described in Kramer et al. [2010]. With regard to dynamic mesh adaptivity, the efficient projection of solution fields, while preserving key dynamical properties such as field integrals, has again only recently been described in Farrell et al. [2009] and Farrell and Maddison [2011].

Chapters 2 and 3 focus on tackling some of these limitations. Firstly, in chapter 2 it was noted that any numerical scheme that fails to yield an accurate solution for the dynamical pressure is susceptible to a very poor representation of geostrophic and hydrostatic balance. In σ-coordinate modelling it is known that this can lead to issues with the accurate representation of hydrostatic balance in the presence of steep topography. It was shown that large physical balance errors may be inherent in certain finite element discretisation of the Navier-Stokes equations. In particular, pressure stabilisation methods, required for some LBB unstable element pairs (including the P_1P_1 element pair), lead to intolerable pressure gradient errors. Balanced pressure decomposition methods have previously been suggested, for example in Piggott et al. [2008b], as methods by which these issues can be addressed. In chapter 2, the first complete physical and mathematical description of balanced pressure decomposition methods is supplied. A novel benchmark test is used to quantify the increase in accuracy resulting from the application of these methods, exploiting the independence of rotation of barotropic simply connected flows on an f-plane. Finally, the thermally driven rotating annulus is used to demonstrate the utility of the method for a real geophysical system.

In chapter 3 this idea is extended to dynamic mesh adaptive simulations. It is noted that, while a system on one computational mesh may be in exact geostrophic balance, there is no guarantee that the system remains in balance after interpolation onto a new computational mesh. A new interpolant is presented, whereby the Coriolis acceleration is Helmholtz decomposed on the donor mesh, interpolated subject to appropriate constraints, and recomposed on the target mesh. This method can guarantee exactly steady geostrophic balance preservation with the $P_{1\mathrm{DG}}P_2$ discretisation of the linearised shallow-water equations. Several numerical examples are used to demonstrate

the applicability of this property. It was determined that this approach leads to no significant benefits over an L_2 optimal (Galerkin) projection for linearised shallow-water ocean modelling using a P_1P_1 finite element pair. Further testing using the thermally driven rotation annulus, in section 4.2.3, also indicates that a Galerkin projection of solution fields is sufficient for P_1P_1 Navier-Stokes simulations.

Fluidity-ICOM simulations of the thermally driven rotating annulus using dynamic mesh adaptivity are described in detail in chapter 4. Simulations on fixed, boundary layer resolving, structured meshes were conducted using the stabilised P_1P_1 finite element pair (with a balanced pressure decomposition method applied). These were compared against two higher resolution simulations conducted using the C-grid finite difference model MORALS. At a lower rotation rate the two models showed excellent agreement, while at a higher rotation rate the simulations showed significant disagreement, with indications of weaker baroclinic eddies in the Fluidity-ICOM simulation. When applying dynamic mesh adaptivity, a modest improvement in the representation of the baroclinic eddies was observed, as compared against MORALS. However, the solution quality was decreased as measured by other diagnostics: in particular, the system heat transport was significantly and non-physically reduced. This was attributed to numerical instabilities, of small magnitude in the fixed structured mesh simulations, becoming more significant in the dynamic mesh adaptive simulations. In addition to this, these simulation had a very high computational cost, with no significant increase in quality over the structured C-grid finite difference model MORALS. A more appropriate choice of metric may lead to more favourable results, but the high computational cost of the simulations makes it difficult to choose and tune adaptivity metrics.

Noting the limitations of unstructured dynamic mesh adaptive methods exposed in chapter 4, the simulations were extended in chapter 5 to a previously challenging system: a high Taylor number rotating annulus system with a topographic β-plane. The computational cost of these simulations was extremely high, requiring many months of processor time, but the unstructured dynamic mesh adaptive methods did enable simulations of this particularly challenging system to be attempted. The numerical simulations showed a favourable comparison against experiment, albeit with a bias towards lower Taylor numbers, indicating implicit numerical diffusion in the Fluidity-ICOM simulations. The simulations also allowed jet formation mechanisms to be tested. In particular, a key conclusion of Wordsworth et al. [2008] was confirmed: the transfer of eddy energy, in Fourier space, was non-local, with energy transferred directly to the mean flow. Higher Taylor number and more fully turbulent regimes were not attempted due to the high computational cost of the simulations.

6.2. Future work

The work presented in this thesis only touches upon the potential applications of these methods. However, significant challenges remain. Dynamic mesh adaptivity, in particular, requires the model to run without error and without significant numerical instability on hundreds or possibly thousands of unrelated computational meshes. Even small instabilities, of small magnitude on one mesh, may become significant on later optimised meshes. An additional concern is the very high computational cost of the approach. For the simulations described in chapter 4, a full regime diagram reconstruction could be performed using MORALS, with several hundred permutations of the system external parameters, for each simulation conducted using Fluidity-ICOM. Restated, the Fluidity-ICOM simulations were found to have a computational cost several orders of magnitude higher than MORALS, for no significant increase in accuracy, at least for the lower rotation rates tested. There are no alternative model simulations against which to make a comparison for the simulations described in chapter 5. However, the range of simulations which could be attempted was severely limited by the high cost of each individual simulation.

It should be noted that the simulations in chapter 4 are at a relatively modest Reynolds number ($Re \sim 600$ for the highest rotation $\omega = 2.764$ rad s^{-1} simulations). It might be expected that the benefits of dynamic mesh adaptivity will be more significant at higher Reynolds number. One might, for example, expect that the computational cost per unit accuracy of dynamic mesh adaptive simulations would become more competitive, with respect to fixed mesh simulations, as the Reynolds number is increased. An analysis of the scaling of the computational cost with Reynolds number, specifically when these methods are applied to baroclinic systems, should therefore be considered in future work. From this, the systems for which dynamic mesh adaptivity is an efficient approach could be determined – the Reynolds number and computational cost of competitive dynamic mesh adaptive unstructured mesh simulations could be clearly identified.

In addition, there are still a number of open questions regarding the configuration of dynamic mesh adaptive simulations for geophysical problems. The choice of dynamical core is currently a topic of active research. The choice of adaptivity metrics for geophysical simulations is an open question – should one constrain the error in the prognostic tracers (temperature or salinity), or the velocity and pressure fields? Should one constrain the errors in the velocity field, or the (potentially noisy) vorticity field? Goal-based error metrics, using an adjoint model in order to constrain the error in a particular quantity, may help to answer these questions, either by being applied directly (albeit at the cost of running an adjoint model) or indirectly by being used to assess the quality of other more heuristic metrics. Some progress is made in this thesis with respect to the choice

of mesh-to-mesh interpolants, although for very long time integrations the preservation of other qualities, such as energy, may become necessary. A further question, not addressed in this thesis, is the application of physical parameterisations to unstructured meshes. Physical schemes are often tuned for particular mesh resolutions. The development of robust physical models that can be applied with multiple mesh resolutions, and with mesh resolutions that change in time, is a particularly challenging area for future research.

Finally, unstructured dynamic mesh adaptive methods are not competitive in all circumstances. While the simulations described in chapter 4 are primarily used to test the numerical methods, and the possible issues with dynamic mesh adaptivity in particular, it is clear that dynamic mesh adaptive methods are not required in order to simulate low rotation rate thermally driven rotating annulus regimes. Fixed structured mesh finite difference methods, such as implemented in MORALS, are far more efficient, and show a very favourable comparison with experiment. By contrast, there are some systems which it may be difficult or impossible to simulate using fixed structured mesh methods. The two-layer rotating annulus experiments described in Williams et al. [2003], for example, may be a more appropriate system for the application of dynamic mesh adaptive methods. In these experiments, a two layer rotating annulus system is formed using two immiscible fluids of differing densities. The system is rotated about its axis of symmetry, and is mechanically forced by differential rotation of the upper annulus lid. This causes the layer interface to slope, leading to a tilt between isobars and isopycnals and ultimately (for appropriate system parameters) to baroclinic instability. In a numerical model of this system, resolution is required in the region of the interface. Since the interface evolves in time, this resolution cannot be placed a-priori. Hence, for this system, dynamic mesh adaptive methods have the potential to enable simulations of otherwise inaccessible phenomena.

With this in mind, future simulations of the thermally driven rotating annulus using unstructured mesh methods, possibly with dynamic mesh adaptivity, should focus on systems which cannot easily be attempted with more conventional structured mesh methods. In particular, there is scope for still higher Taylor number simulations, enabling energy transfer and jet formation to be tested in more fully turbulent flow regimes. Such simulations would require significant computational resources. Alternatively, the geometric flexibility of the unstructured mesh methods could be exploited to enable simulations with complex topologies and topography – for example, in simulations of the thermally driven rotating annulus with periodic bottom topography, and with full and partial radial barriers. In the latter case, the unstructured mesh approach should be tested against σ-coordinate and curvilinear coordinate finite difference modelling.

It is therefore concluded that, while dynamic mesh adaptivity is an extremely powerful numerical tool, clear issues remain in the application of the method to geophysical systems. Key unanswered numerical questions remain regarding both the details of the approach, and where the approach should be attempted. Dynamic mesh adaptivity may be a competitive option (or even the only reasonable option) in some cases, but currently is not competitive in some applications. For simulations of the thermally driven rotating annulus, the minimum problem size for which dynamic mesh adaptivity is competitive appears to be large.

Appendix A.

Unstructured mesh diagnostics

This appendix describes methods for computing several unstructured mesh diagnostics that were utilised or developed over the course of this DPhil project. These diagnostics include: heat flow computation, point field evaluation (or "detectors"), depth integration, depth averaged (barotropic) stream functions, and azimuthally averaged (meridional) stream functions.

A.1. Fluidity-ICOM parallelism

Many of the diagnostics described in this appendix are parallelised. For clarity, the MPI parallelism method used by Fluidity-ICOM is briefly described.

In a parallel Fluidity-ICOM simulation the domain is divided into a series of partitions, with each partition associated with a single MPI process. Nodes in this partitioning are designated either "owned", meaning that the owning process is responsible for computing the value of fields on these nodes, or "non-owned", meaning that the process associated with an adjacent partition is responsible for computing the value of fields on these nodes. The owned nodes visible to adjacent processes, together with the non-owned nodes, form a shared halo region. The size of this shared halo region is dependent upon the stencil of the finite element discretisation used for a particular field. Field data is communicated in this halo region as required in order to maintain the validity of computations on all processes. From this node ownership, one can further define a unique element ownership, whereby the process owning a given element is defined to be the lowest ranked process owning any nodes in the element. An example of node ownership for a P_1 field is shown in figure A.1.

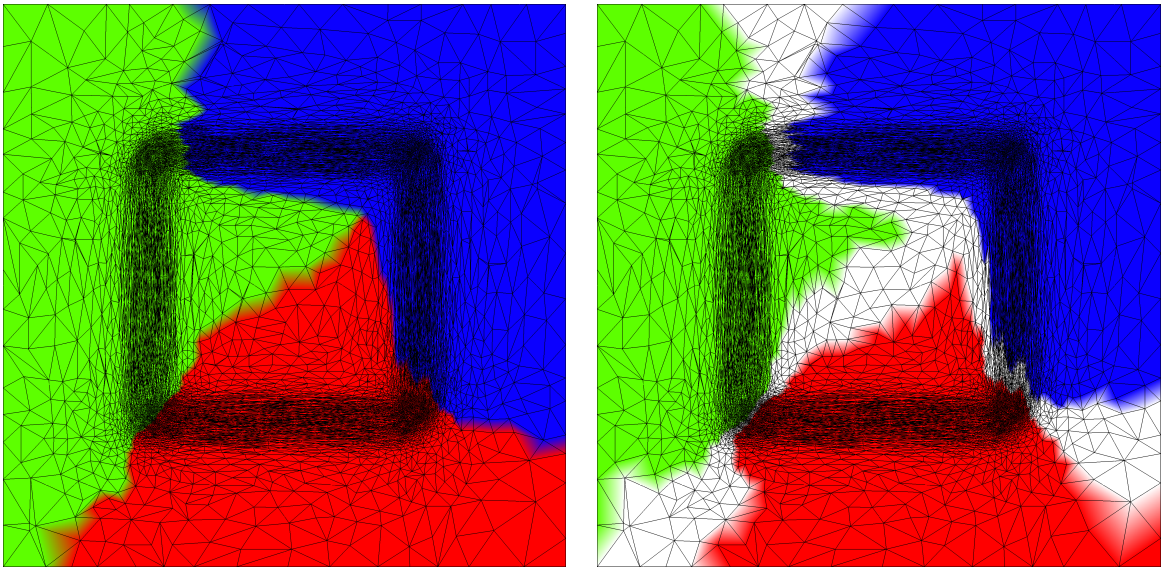

Figure A.1.: Left: P_1 node ownership for a three partition domain decomposition of an unstructured mesh using the METIS graph partitioning library [Karypis and Kumar, 1998]. Right: As left, with the shared halo region coloured in white.

A.2. Heat flow

The heat flow for the rotating annulus can be defined in terms of a non-dimensional Nusselt number, expressing the heat transport of the system relative to a solid of equal conductivity (see section 2.3). Expressed in terms of the diffusive flow through the tank inner wall this takes the form:

$$Nu = \frac{\ln b/a}{2\pi d \Delta T} \int_{\partial\Omega_I} \nabla T \cdot \hat{\boldsymbol{n}}. \tag{A.1}$$

where the terms are defined as in table 2.1, $\partial\Omega_I$ is the annulus tank inner wall, and $\hat{\boldsymbol{n}}$ is a unit normal outward on $\partial\Omega_I$. For a discrete temperature field T^δ with basis functions ϕ_i this becomes:

$$Nu = \frac{\ln b/a}{2\pi d \Delta T} \sum_i \int_{\partial\Omega_I} (\nabla\phi_i \cdot \hat{\boldsymbol{n}})\, \tilde{T}. \tag{A.2}$$

where $\tilde{T}$ denotes the nodal values for T^δ. The integrand on the right-hand-side is well-defined for any Lagrange interpolating polynomial basis functions of degree greater than zero (discontinuous or continuous). Hence the Nusselt number can be computed directly from equation (A.2).

For parallel simulations, each process computes the integral in equation (A.2) over only those inner wall surface elements that are attached to owned elements. The Nusselt number is then computed using a collective sum reduction operation.

A.3. Detectors

Detectors are used in Fluidity-ICOM simulations to evaluate field data at specified locations within the simulation domain, with timestep level precision. For example, arrays of detectors can be used to create virtual thermocouple arrays, with the temperature field output at each location in the array at each model timestep, as used in section 2.3 and in chapter 4.

In order to evaluate the values of fields for a given detector, it is necessary to identify the element in the mesh containing the detector location. This procedure is challenging for general unstructured meshes. Testing all elements in the computational mesh for containment of the detector point (a brute force search) has prohibitive computational cost. In addition, due to the use of finite precision arithmetic, it is challenging to formulate a point containment "predicate", determining whether a given point is contained within a given element, that is both unique and covering. For example, for a point close to element interfaces, it is possible for a finite precision point containment predicate to find zero containing elements (i.e. the predicate is not covering), one containing element (as desired), or multiple containing elements (i.e. the predicate is not unique), as illustrated in figure A.2. This issue is not resolved by using an "epsilon-ball" type tolerance, which reduces the number of zero containing element cases at the cost of increasing the number of multiple containing element cases. An epsilon-ball tolerance approach can also increase precision errors when evaluating fields within an element determined to "contain" the detector point. See Kettner et al. [2008] for a discussion of surprising errors that can arise when using finite precision arithmetic in computational geometry.

An algorithm determining a unique element containing a given detector point is formulated as follows. We seek the element in a finite element mesh containing a given detector point. We assume that this detector point is known to lie somewhere within the computational domain. We further assume the existence of a finite precision containment predicate, determining whether a given element contains the detector point. If the containment predicate determines that the detector point is not contained within the given element, the predicate returns a measure of the distance of the point from the element. For simplex meshes, this predicate was derived from the detector point barycentric coordinates (or area coordinates) within each test simplex. If, for a given test simplex,

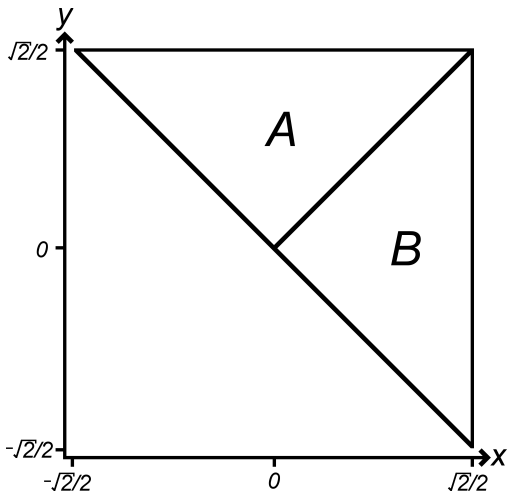

Cartesian coordinate	Barycentric coordinates for A	Barycentric coordinates for B
$\begin{pmatrix} 0.1 \\ 0.1 \end{pmatrix}$	$\begin{pmatrix} 0.0 \\ 0.858579 \\ 0.141421 \end{pmatrix}$	$\begin{pmatrix} 0.0 \\ 0.141421 \\ 0.858577 \end{pmatrix}$
$\begin{pmatrix} 0.1001 \\ 0.1001 \end{pmatrix}$	$\begin{pmatrix} -5.551115 \times 10^{-17} \\ 0.858437 \\ 0.141563 \end{pmatrix}$	$\begin{pmatrix} -2.775558 \times 10^{-17} \\ 0.141563 \\ 0.858437 \end{pmatrix}$

Figure A.2.: Illustration of the point containment problem using a finite precision containment predicate. The test is performed for a four node, two element triangular mesh as shown. The containment predicate is derived from the barycentric coordinates of the test point within each test element, with a point defined to be contained within the triangle if all barycentric coordinate components are non-negative. The barycentric coordinates, computed using double precision arithmetic, are shown for two different test points. For the coordinate $(0.1, 0.1)^T$ all barycentric coordinate components are non-negative, for both element A and element B. Hence the test point is determined to be contained by both elements. For the coordinate $(0.1001, 0.1001)^T$ the barycentric coordinates within element A and element B each contain a negative component. Hence this test point is determined to be contained by no elements.

all barycentric coordinate components are non-negative, the point is defined to be contained within the simplex. Otherwise, the distance of the point from the simplex is returned as the absolute value of the minimum barycentric coordinate component.

First, the elements are filtered with an R-tree spatial indexing algorithm as described in Guttman [1984], using the libspatialindex library [1]. Only elements whose bounding box (with an extension of 10% in each dimension to account for finite precision errors) contains the detector point are considered. Each of these elements is then tested for detector point containment in turn. The first element determined, using the finite precision containment predicate, to contain the detector point is returned as the containing element. If no such element is found, the element found to have the smallest measure (or the first element found to have the smallest measure) of the distance of the detector point from the element is returned as the element containing the detector point.

For parallel simulations it is necessary also to identify a unique process owning the detector. The process with the lowest rank for which the containment predicate returns an owner is designated the owner of the detector. If no such process exists, the process with an element giving the smallest measure (or the lowest ranked process with an element giving the smallest measure) of the distance of the detector point from the element is designated the owner of the detector.

A.4. Vertical integration

For a computational mesh with nodes aligned in the vertical dimension, vertical integrals can be computed using simple integration along each vertically aligned node "chain". This simple approach cannot be used for meshes that are unstructured in the vertical dimension. Instead, we consider the Galerkin projection of the vertical integral. This can be computed using the method described in Farrell [2009] (section 5.4.1.3).

The method is briefly outlined as follows. For a field q^δ represented using basis functions ϕ_i on a D dimensional mesh $\mathcal{M}(\Omega)$ on a domain Ω, the Galerkin projection of the vertical integral q_I^δ represented using basis functions $\phi_{S,i}$ on a $D-1$ dimensional mesh $\mathcal{M}_S(\partial\Omega_S)$ on a domain $\partial\Omega_S$ is given by:

$$\int_{\partial\Omega_S} \phi_{S,i} q_I^\delta = \int_{\partial\Omega_S} \phi_{S,i} \left[\int q^\delta dz \right]. \tag{A.3}$$

[1] Available from http://trac.gispython.org/spatialindex.

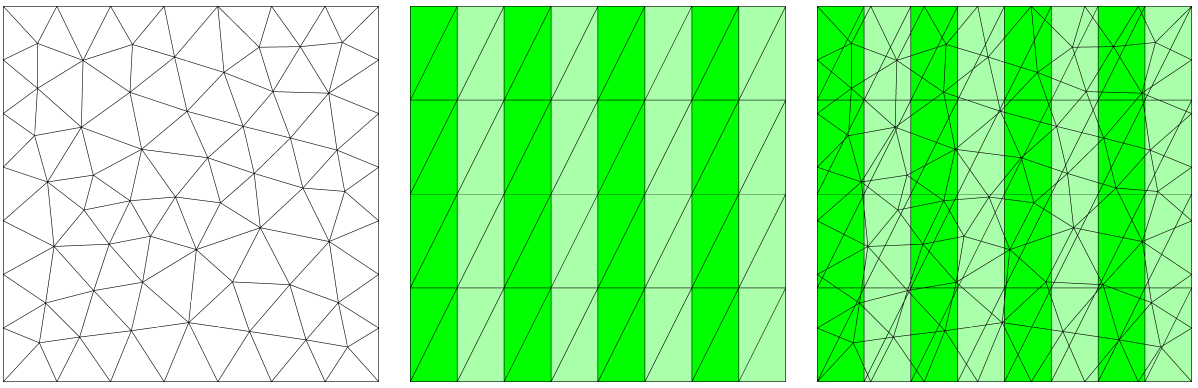

Figure A.3.: Construction of an intersection mesh suitable for assembling the vertical integration Galerkin projection equation (A.5). A: The mesh supporting the field to be integrated. B: An extrusion mesh, supporting the vertical extrusion of the surface mesh basis functions onto which the vertical integral is projected. C: Supermesh of A and B, highlighting the columns of extruded elements.

Re-expressed, this becomes:

$$\int_{\partial\Omega_S} \phi_{S,i} q_I^\delta = \int_\Omega \Pi\left(\phi_{S,i}\right) q^\delta, \tag{A.4}$$

where $\Pi\left(\phi_{S,i}\right)$ is the vertical extrusion of the basis function $\phi_{S,i}$. This leads to:

$$M_S \tilde{q}_I = V\tilde{q}, \tag{A.5}$$

where:

$$\left(M_S\right)_{ij} = \int_{\Omega_S} \phi_{S,i}\phi_{S,j} \qquad V_{ij} = \int_\Omega \Pi\left(\phi_{S,i}\right)\phi_j, \tag{A.6}$$

and where $\tilde{q}_I$ and $\tilde{q}$ are the nodal values for q_I^δ and q^δ respectively. M_S is the surface mesh mass matrix and V is matrix of the inner products of the vertical extrusion of the surface mesh basis functions $\phi_{S,i}$ for q_I^δ, with the volume mesh basis functions ϕ_i for q^δ. As for the Galerkin projection of fields between meshes (see section 3.1), it is challenging to assemble the matrix V, as $\Pi\left(\phi_{S,i}\right)$ and ϕ_i are not supported on the same mesh.

It is possible to assemble the right-hand-side of equation (A.5) using the local supermeshing approach of Farrell and Maddison [2011]. For each element e_S in the surface mesh $\mathcal{M}_S$ ($\partial\Omega_S$), a vertical extrusion mesh $\mathcal{M}_{S,e}(\Omega)$ supporting $\Pi(\phi_{S,i})$ is generated. The extruded mesh $\mathcal{M}_{S,e}(\Omega)$ is then intersected with the volume mesh $\mathcal{M}(\Omega)$ to form a supermesh. The contribution to the right-hand-side of equation (A.5) for element e_S can be computed via integration over this supermesh. This process is illustrated in figure A.3. As for Galerkin projection between meshes, pairs of intersecting elements from $\mathcal{M}_{S,e}(\Omega)$ and $\mathcal{M}(\Omega)$ can be considered separately, and their contribution to the right-hand-side of equation (A.5) computed individually, without generating the supermesh in its entirety.

A.5. Depth averaged stream function

For the thermally driven rotating annulus with a rigid lid, it follows that the from the continuity equation (2.3b) that the horizontal component of the vertically averaged velocity field is non-divergent. The conservative component of the Helmholtz decomposition of the depth averaged stream function is therefore zero (see section 2.1.1), and the depth averaged velocity can be expressed as the curl of a vector potential. Hence the depth averaged velocity can be defined in terms of a depth averaged (barotropic) stream function Ψ_D where, in cylindrical coordinates (r, θ, z):

$$u_{D,r} = \frac{1}{r}\frac{\partial \Psi_D}{\partial \theta}, \qquad u_{D,\theta} = -\frac{\partial \Psi_D}{\partial r}, \tag{A.7}$$

or equivalently:

$$\hat{\boldsymbol{z}} \times \boldsymbol{u_D} = \nabla_H \Psi_D, \tag{A.8}$$

where $\nabla_H = (\partial/\partial r, (1/r)\,\partial/\partial\theta, 0)^T$ is the horizontal gradient operator and where $\boldsymbol{u}_D$ is the horizontal component of the depth averaged velocity:

$$(\boldsymbol{u}_D)^T = \left(\frac{\int_{z=0}^{z=d} u_r dz}{d}, \frac{\int_{z=0}^{z=d} u_\theta dz}{d}, 0 \right). \tag{A.9}$$

Here $\boldsymbol{u}_D$ is defined on the domain Ω_L forming the upper surface of the domain. For meshes with nodes aligned in the vertical dimension, the integrals in equation (A.9) can be computed via integration along each vertically aligned node "chain". For more general vertically unstructured meshes, $\boldsymbol{u}_D$ can be computed via the unstructured mesh depth integration approach presented in section A.4.

We proceed to formulate a continuous Galerkin discretisation of equation (A.8). Multiplying equation (A.8) by a test function ϕ, integrating over Ω_L, and applying Green's theorem yields the weak form:

$$\int_{\Omega_L} \nabla_H \phi \cdot \nabla_H \Psi_D - \int_{\partial\Omega_L} \phi \left[\nabla_H \Psi_D - \hat{\boldsymbol{z}} \times \boldsymbol{u}_D\right] \cdot \hat{\boldsymbol{n}} = \int_{\Omega_L} \nabla_H \phi \cdot (\hat{\boldsymbol{z}} \times \boldsymbol{u}_D) \quad \forall \phi, \tag{A.10}$$

where $\partial\Omega_L$ bounds Ω_L and $\hat{\boldsymbol{n}}$ is a unit normal outward on $\partial\Omega_L$. Applying strong Dirichlet boundary conditions for Ψ_D (see below) the boundary integral vanishes. Restricting ϕ and Ψ_D to finite-dimensional spaces to yield the finite-dimensional ϕ^δ and Ψ_D^δ leads to a discrete diagnostic equation for the depth averaged stream function:

$$\int_{\Omega_L} \nabla_H \phi^\delta \cdot \nabla_H \Psi_D^\delta = \int_{\Omega_L} \nabla_H \phi^\delta \cdot \left(\hat{\boldsymbol{z}} \times \boldsymbol{u}_D^\delta\right) \quad \forall \phi^\delta, \tag{A.11}$$

Choosing ϕ^δ and Ψ_D^δ to be from the same (C^0 continuous) space complete a continuous Galerkin discretisation. P_1 basis functions are used for computations of Ψ_D^δ in this thesis.

Since $\boldsymbol{u}_D$ satisfies the no-slip boundary condition on $\partial\Omega_L$, Ψ_D is constant over each connected sub-domain of $\partial\Omega_L$. For the rotating annulus the domain Ω_L, forming the upper lid surface of the annular chamber, is multiply connected, and hence Ψ_D takes separate values on the outer and inner annulus walls. Choosing $\Psi_D = 0$ on the outer wall it follows that at the inner wall Ψ_D is equal to the net depth averaged azimuthal flux. Hence the boundary conditions are:

$$\Psi_D = 0 \quad \text{on } \partial\Omega_{L,O}, \tag{A.12a}$$

$$\Psi_D = \int_{x=a}^{x=b} u_{D,y} dx \Big|_{y=0} \qquad \text{on } \partial\Omega_{L,I}, \tag{A.12b}$$

where $\partial\Omega_{L,O}$ is the outer wall boundary and $\partial\Omega_{L,I}$ is the inner wall boundary. These boundary conditions were applied in the strong sense, via removal of the basis functions associated with boundary condition nodes.

A.6. Azimuthally averaged stream function

Due to the multiply connected topology of the annular domain, it again follows from the continuity equation (2.3b) that the meridional component of the azimuthally averaged velocity is divergence free. Hence, as in the depth averaged case, the azimuthally averaged velocity can be defined in terms of an azimuthally averaged (meridional) stream function Ψ_A where, in cylindrical coordinates:

$$u_{A,r} = \frac{1}{r}\frac{\partial \Psi_A}{\partial z}, \qquad\qquad u_{A,z} = -\frac{1}{r}\frac{\partial \Psi_A}{\partial r}, \tag{A.13}$$

or equivalently:

$$-\hat{\boldsymbol{\theta}} \times r\boldsymbol{u_A} = \nabla_M \Psi_A, \tag{A.14}$$

where $\nabla_M = (\partial/\partial r, 0, \partial/\partial z)^T$ and $\boldsymbol{u_A}$ is the meridional component of the azimuthally averaged velocity:

$$(\boldsymbol{u}_A)^T = \left(\frac{\int_{\theta=0}^{\theta=2\pi} u_r d\theta}{2\pi}, 0, \frac{\int_{\theta=0}^{\theta=2\pi} u_z d\theta}{2\pi} \right). \tag{A.15}$$

$(\boldsymbol{u}_A)^T$ is defined on the domain Ω_S consisting of a vertical cross-section of the annular chamber, with $a \leq r \leq b$ and $0 \leq z \leq d$. For meshes with nodes aligned in the azimuthal dimension, the integrals in equation (A.15) can be computed via integration along each azimuthally aligned node

"chain". For more general unstructured meshes $\boldsymbol{u}_A$ can, in principle, be calculated via an analogous method to the vertical integration approach presented in section A.4.

Following an analogous derivation to that for the depth averaged stream function described in section A.5, this leads to the finite element discretisation:

$$\int_{\Omega_S} \nabla_M \phi^\delta \cdot \nabla_M \Psi_A^\delta = -\int_{\Omega_S} \nabla_M \phi^\delta \cdot \left(\hat{\boldsymbol{\theta}} \times r\boldsymbol{u}_A^\delta\right) \quad \forall \phi^\delta, \tag{A.16}$$

Choosing ϕ^δ and Ψ_A^δ to be from the same (C^0 continuous) space yields a continuous Galerkin discretisation. P_1 basis functions are for computations of Ψ_A^δ in this thesis.

Since Ω_S is simply connected, the boundary condition $\Psi_A = 0$ on $\partial\Omega_S$, where $\partial\Omega_S$ bounds Ω_S, can be applied. As for the depth averaged stream function, this boundary condition was applied in the strong sense, via removal of the basis functions associated with boundary condition nodes.

Appendix B.

Fluidity-ICOM development

Fluidity-ICOM is an in-development ocean model. Very substantial development work occurred during the course of this DPhil project, with contributions from a number of developers. This development work included: a new model interface, re-implementation of the Navier-Stokes solver, re-implementation and extension of the balanced pressure decomposition methods (see chapter 2), new tracer diffusion schemes, re-implementation of detector based diagnostics (see section A.3), re-implementation and extension of the Fluidity-ICOM parallelisation scheme (see section A.1), new mesh-to-mesh interpolants (see chapter 3), and updates and extensions of the interfaces to external libraries (including the mesh adaptivity libraries).

Fluidity-ICOM uses a distributed development model, with coding contributions maintained using the subversion version control system[1]. In addition, Fluidity-ICOM implements automated build testing and verification using the Buildbot system [2]. Development of Fluidity-ICOM is rapid – it is common for more than 50 commits (modifications) to be made within a single week.

B.1. Contributions to Fluidity-ICOM

A total of 2408 model contributions were made to the Fluidity-ICOM subversion repository as part of this DPhil project[3]. The following is a list of the more significant contributions, together with the associated source files modified. In all cases further contributions were made by other Fluidity-ICOM developers. Source file names refer to Fluidity-ICOM subversion version r14166.

[1] Available from http://subversion.tigris.org.
[2] Available from http://buildbot.net/trac.
[3] Number of commits to the main Fluidity-ICOM subversion repository, the "trunk", up to revision r14166 on 29/08/10.

Balanced pressure decomposition methods:

Main source files: `assemble/Geostrophic_Pressure.F90`
`assemble/Hydrostatic_Pressure.F90`

Re-implementation and extension of the balanced pressure decomposition methods described in Pain et al. [2005] and Piggott et al. [2008a] and discussed in detail in chapter 2.

Geostrophic balance preserving interpolation:

Main source files: `assemble/Geostrophic_Pressure.F90`

Helmholtz decomposed geostrophic balance preserving interpolation discussed in detail in chapter 3

Spud-Diamond model interface:

Main source files: `libspud/diamond/bin/diamond`
`libspud/diamond/diamond/debug.py`
`libspud/diamond/diamond/dialogs.py`
`libspud/diamond/diamond/interface.py`
`libspud/include/cspud.h`
`libspud/include/spud_enums.h`
`libspud/include/spud_config.h`
`libspud/include/spud.h`
`libspud/python/libspud.py.in`
`libspud/src/fspud.f90`
`libspud/src/spud.cpp`
`libspud/src/spud_interfaces.cpp`

XML based options file system with an automatically generated graphical user interface. See Ham et al. [2009] for a complete description of this system.

Adaptivity interfaces:

Main source files:
```
assemble/Adapt_Integration.F90
assemble/Adapt_State.F90
assemble/Mba2d_Integration.F90
assemble/Interpolation_manager.F90
assemble/Node_Locking.F90
assemble/Sam_integration.F90
assemble/Surface_Id_Interleaving.F90
```

Re-implementation of the mesh adaptivity library interfaces. Primarily to enable the use of custom mesh-to-mesh interpolants.

Parallelisation:

Main source files:
```
femtools/Halo_Data_Types.F90
femtools/Halos.F90
femtools/Halos_Allocates.F90
femtools/Halos_Base.F90
femtools/Halos_Communications.F90
femtools/Halos_Debug.F90
femtools/Halos_Derivation.F90
femtools/Halos_IO.cpp
femtools/Halos_Numbering.F90
femtools/Halos_Ownership.F90
femtools/Halos_Registration.F90
femtools/Halos_Repair.F90
femtools/Surface_Labels.F90
fldecomp/fldtriangle.cpp
include/fldecomp.h
include/Halos_IO.h
```

Updates to enable parallel simulations using the Spud-Diamond interface (above). Re-implementation and extension of the low level parallel data structures for use in the adaptivity interface re-implementation (above), and to enable more general finite element discretisations to be used in parallel.

Supermesh construction:

Main source files:
```
include/Element_Intersection.cpp
femtools/Intersection_finder.F90
femtools/Supermesh.F90
femtools/Supermesh_Assembly.F90
include/Element_Intersection.h
include/MeshDataStream.h
tools/Vertical_Integration.F90
```

Supermesh construction via element intersection, for use in Galerkin projection of fields between meshes, as described in section 3.1, and for vertical integration, as described in section A.4. Also includes subroutines for integration of arbitrary bilinear forms between non-matching meshes ("supermesh assembly").

Surface integral diagnostics:

Main source files:
```
femtools/Diagnostic_variables.F90
femtools/Surface_Integrals.F90
```

Surface integral computations, such as for the heat transport diagnostic described in section A.2.

Detector diagnostics:

Main source files:
```
include/Element_Intersection.cpp
femtools/Diagnostic_variables.F90
femtools/Node_Owner_Finder.cpp
femtools/Node_Owner_Finder_Fortran.F90
femtools/Node_Ownership.F90
femtools/Picker_Data_Types.F90
femtools/Pickers_Allocates.F90
femtools/Pickers_Base.F90
femtools/Pickers_Deallocates.F90
femtools/Pickers.F90
femtools/Pickers_Inquire.F90
include/Element_Intersection.h
include/Node_Owner_Finder.h
include/MeshDataStream.h
```

Re-implementation of point detector diagnostics to use the precision robust algorithm described in section A.3.

Diagnostic fields:

Main source files:
```
assemble/Diagnostics_fields_wrapper.F90
assemble/Vorticity_Diagnostics.F90
diagnostics/Binary_Operators.F90
diagnostics/Diagnostic_Fields_New.F90.in
diagnostics/Diagnostic_Fields_Interfaces.F90
diagnostics/Diagnostic_Source_Fields.F90
diagnostics/Differential_Operators.F90
diagnostics/Field_Copies.F90
diagnostics/Metric_Diagnostics.F90
diagnostics/Momentum_Diagnostics.F90
diagnostics/Parallel_Diagnostics.F90
diagnostics/Python_Diagnostics.F90
diagnostics/Surface.F90
femtools/Diagnostics_fields.F90
scripts/make_diagnostic_fields
```

Addition of new on-line diagnostic fields. Also creation of an automated diagnostic field system, whereby generic diagnostic fields (e.g. gradient, divergence) can be added to the model via the addition of a single subroutine, and the specification of options for the new diagnostic field. This automated system implements diagnostic field dependency resolution using a topological sort algorithm.

Shallow-water model:

Main source files:
```
main/Simple_Shallow_Water.F90
```

Linearised shallow-water solver supporting multiple finite element pairs. Used for the P_1P_1 geostrophic balance preserving interpolation test in section 3.4.

Advection-diffusion solver:

Main source files:
```
assemble/Advection_Diffusion_CG.F90
assemble/Upwind_Stabilisation.F90
```

Continuous Galerkin finite element solver for the advection-diffusion equations, with streamline upwind Petrov Galerkin stabilisation. While this advection-diffusion solver was developed during this DPhil project, it was not used for any simulations discussed in this thesis.

Bibliography

A. Adcroft and D. Marshall. How slippery are piecewise-constant coastlines in numerical ocean models? *Tellus A*, 50:95–108, 1998.

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

F. Alauzet, P. L. George, B. Mohammadi, P. Frey, and H. Borouchaki. Transient fixed point-based unstructured mesh adaptation. *International Journal for Numerical Methods in Fluids*, 43:729–745, 2003.

F. Alauzet, A. Loseille, A. Dervieux, and P. Frey. Multi-dimensional continuous metric for mesh adaptation. In *Proceedings of the 15th International Meshing Roundtable*, pages 191–214, 2006.

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

A. Arakawa and M. J. Suarez. Vertical differencing of the primitive equations in sigma coordinates. *Monthly Weather Review*, 111:34–45, 1983.

G. B. Arfken and H. J. Weber. *Mathematical methods for physicists*. Harcourt Academic Press, 2001.

O. Axelsson and V. A. Barker. *Finite element solution of boundary value problems: theory and computation*. Academic Press, 1984.

I. Babuška. Error-bounds for finite element method. *Numerische Mathematik*, 16(4):322–333, 1971.

D. Bacon, N. N. Ahmad, T. Dunn, S. G. Gopalakrishnan, M. S. Hall, and A. Sarma. Hurricane track forecasting with OMEGA. *Natural Hazards*, 41:457–470, 2007.

D. P. Bacon, N. N. Ahmed, Z. Boybeyi, T. J. Dunn, M. S. Hall, P. C. S. Lee, R. A. Sarma, M. D. Turner, K. T. Waight III, S. H. Young, and J. W. Zack. A dynamically adapting weather and dispersion model: The operational multiscale environment model with grid adaptivity (OMEGA). *Monthly Weather Review*, 128:2044–2076, 2000.

D. P. Bacon, N. N. Ahmad, Z. Boybeyi, T. J. Dunn, S. G. Gopalakrishnan, M. S. Hall, Y. Jin, P. C. S. Lee, D. E. Mays, A. Sarma, M. D. Turner, T. R. Wait, K. T. Waight III, S. H. Young, and J. W. Zack. Dynamically adapting unstructured triangular grids: A new paradigm for geophysical fluid dynamics modelling. *Proceedings of the Indian National Science Academy A*, 69(5):457–471, 2003.

S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.

S. Balay, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang. *PETSc Users Manual*. Argonne National Laboratory, anl-95/11 – revision 3.0.0 edition, 2008.

S. Balay, K. Buschelman, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang. PETSc Web page. http://www.mcs.anl.gov/petsc, 2009.

V. Barcilon. Role of the Ekman layers in the stability of the symmetric regime obtained in a rotating annulus. *Journal of the Atmospheric Sciences*, 21:291–299, 1964.

F. Bassi and S. Rebay. A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier-Stokes equations. *Journal of Computational Physics*, 131: 267–279, 1997.

G. K. Batchelor. *An introduction to fluid mechanics*. Cambridge University Press Cambridge, 1967.

A. Beckmann and D. B. Haidvogel. Numerical simulation of flow around a tall isolated seamount. Part I: Problem formulation and model accuracy. *Journal of Physical Oceanography*, 23:1736–1753, 1993.

M. J. Bell and A. A. White. The Stability of internal baroclinic jets: Some analytical results. *Journal of the Atmospheric Sciences*, 45(18):2571–2590, 1988.

E. D. Berger, K. S. McKinley, R. D. Blumofe, and P. R. Wilson. Hoard: A scalable memory allocator for multithreaded applications. *ACM SIGPLAN Notices*, 35(11):117–128, 2000.

M. J. Berger and J. Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics*, 53:484–512, 1984.

P. E. Bernard, N. Chevaugeon, V. Legat, E. Deleersnijder, and J. F. Remacle. High-order h-adaptive discontinuous Galerkin methods for ocean modelling. *Ocean Dynamics*, 57(2):109–121, 2007.

J. Berntsen and Ø. Thiem. Estimating the internal pressure gradient errors in a σ-coordinate ocean model for the Nordic Seas. *Ocean Dynamics*, 57(4):417–429, 2007.

E. Blayo and L. Debreu. Adaptive mesh refinement for finite-difference ocean models: First experiments. *Journal of Physical Oceanography*, 29:1239–1250, 1999.

O. Botella and R. Peyret. Benchmark spectral results on the lid-driven cavity flow. *Computers & Fluids*, 27(4):421–433, 1998.

F. P. Bretherton. Critical layer instability in baroclinic flows. *Quarterly Journal of the Royal Meteorological Society*, 92(393):325–334, 1966.

F. Brezzi. On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers. *RAIRO Analyse numrique*, 8(2):129–151, 1974.

A. N. Brooks and T. J. R. Hughes. Streamline upwind / Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations. *Computer Methods in Applied Mechanics and Engineering*, 32:199–259, 1982.

C. H. Bruneau and M. Saad. The 2D lid-driven cavity problem revisited. *Computers & Fluids*, 35: 326–348, 2006.

K. Bryan. A numerical method for the study of the circulation of the world ocean. *Journal of Computational Physics*, 4:347–376, 1969.

A. A. Castrejón-Pita and P. L. Read. Baroclinic waves in an air-filled thermally driven rotating annulus. *Physical Review E*, 75(2):026301, 2007.

J. G. Charney. The dynamics of long waves in a baroclinic westerly current. *Journal of Meteorology*, 4(5):135–162, 1947.

C. Chen, H. Liu, and R. C. Beardsley. An unstructured grid, finite-volume, three-dimensional, primitive equations ocean model: Application to coastal ocean and estuaries. *Journal of Atmospheric and Oceanic Technology*, 20:159–186, 2003.

C. Chen, H. Huang, R. C. Beardsley, H. Liu, Q. Xu, and G. Cowles. A finite volume numerical approach for coastal ocean circulation studies: Comparisons with finite difference models. *Journal of Geophysical Research*, 112, 2007.

A. J. Chorin. A numerical method for solving incompressible viscous flow problems. *Journal of Computational Physics*, 135:118–125, 1967.

P. C. Chu and C. Fan. Sixth-order difference scheme for sigma coordinate ocean models. *Journal of Physical Oceanography*, 27(9):2064–2071, 1997.

P. C. Chu and C. Fan. Hydrostatic correction for sigma coordinate ocean models. *Journal of Geophysical Research*, 108(C6), 2003.

R. Comblen, S. Legrand, E. Deleersnijder, and V. Legat. A finite element method for solving the shallow water equations on the sphere. *Ocean Modelling*, 28(1-3):12–23, 2009.

R. Cools and A. Haegemans. Algorithm 824: CUBPACK: A package for automatic cubature; framework description. *ACM Transactions on Mathematical Software*, 29(3):287–296, 2003.

C. J. Cotter and D. A. Ham. Numerical wave propagation for the triangular $P_{1\mathrm{DG}}P_2$ finite element pair. *Journal of Computational Physics*, 230(8):2806–2820, 2011.

C. J. Cotter, D. A. Ham, and C. C. Pain. A mixed discontinuous / continuous finite element pair for shallow-water ocean modelling. *Ocean Modelling*, 26(1-2):86–90, 2009a.

C. J. Cotter, D. A. Ham, C. C. Pain, and S. Reich. LBB stability of a mixed Galerkin finite element pair for fluid flow simulations. *Journal of Computational Physics*, 228(2):336–348, 2009b.

R. Courant, K. Friedrichs, and H. Lewy. Ueber die partiellen Differenzengleichungen der mathematischen Physik. *Mathematische Annalen*, 100(1):32–74, 1928.

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

S. Danilov, G. Kivman, and J. Schröter. A finite-element ocean model: Principles and evaluation. *Ocean Modelling*, 6(2):125–150, 2004.

S. Danilov, G. Kivman, and J. Schröter. Evaluation of an eddy-permitting finite-element ocean model in the North Atlantic. *Ocean modelling*, 10(1-2):35–49, 2005.

S. Danilov, Q. Wang, M. Losch, D. Sidorenko, and J Schröter. Modeling ocean circulation on unstructured meshes: comparison of two horizontal discretizations. *Ocean Dynamics*, 58(5-6): 365–374, 2008.

P. J. Davis. *Interpolation and approximation*. Dover, 1975.

F. M. Denaro. On the application of the Helmholtz–Hodge decomposition in projection methods for incompressible flows with general boundary conditions. *International Journal for Numerical Methods in Fluids*, 43:43–69, 2003.

E. T. Eady. Long waves and cyclone waves. *Tellus*, 1(3):33–52, 1949.

D. H. Eberly. *3D game engine design*. Morgan Kaufmann, 2001.

A. El Hraiech, H. Borouchaki, P. Villon, and L. Moreau. Mechanical field interpolation. In *AIP Conference Proceedings*, volume 778, pages 201–208. AIP, 2005.

S. S. Elliott. *Numerical studies of baroclinic instability*. PhD thesis, University of Oxford, 1995.

R. D. Falgout and U. M. Yang. hypre: A library of high performance preconditioners. *Computational Science*, 2331/2002:632–641, 2009.

L. Farnell and R. A. Plumb. Numerical integration of flow in a rotating annulus I: Axisymmetric model. Occasional Note Met O 21 75/3, Geophysical Fluid Dynamics Laboratory, Meteorological Office, Bracknell, Berkshire., December 1975.

L. Farnell and R. A. Plumb. Numerical integration of flow in a rotating annulus II: Three dimensional model. Occasional Note Met O 21 76/1, Geophysical Fluid Dynamics Laboratory, Meteorological Office, Bracknell, Berkshire., February 1976.

P. E. Farrell. *Galerkin projection of discrete fields via supermesh construction*. PhD thesis, Imperial College London, 2009.

P. E. Farrell and J. R. Maddison. Conservative interpolation between volume meshes by local Galerkin projection. *Computer Methods in Applied Mechanics and Engineering*, 200(1-4):89–100, 2011.

P. E. Farrell, M. D. Piggott, C. C. Pain, G. J. Gorman, and C. R. Wilson. Conservative interpolation between unstructured meshes via supermesh construction. *Computer Methods in Applied*

Mechanics and Engineering, 198:2632–2642, 2009.

R. Ford. Realistic-bathymetry, small-dissipation solutions for the North Atlantic / Carribean. In P. Müller and D. Henderson, editors, *Topographic Effects in the Ocean: Proceedings, Hawaiian Winter Workshop, University of Hawaii at Manoa*, pages 163–167, January 1995.

R. Ford, C. C. Pain, M. D. Piggott, A. J. H. Goddard, C. R. E. de Oliveira, and A. P. Umpleby. A nonhydrostatic finite-element model for three-dimensional stratified oceanic flows. Part I: Model formulation. *Monthly Weather Review*, 132(12):2816–2831, 2004a.

R. Ford, C. C Pain, M. D Piggott, A. J. H. Goddard, C. R. E. de Oliveira, and A. P. Umpleby. A nonhydrostatic finite-element model for three-dimensional stratified oceanic flows. Part II: Model validation. *Monthly Weather Review*, 132(12):2832–2844, 2004b.

L. Formaggia and S. Perotto. New anisotropic a priori error estimates. *Numerische Mathematik*, 89(4):641–667, 2001.

W. W. Fowlis and R. Hide. Thermal convection in a rotating annulus of liquid: Effect of viscosity on the transition between axisymmetric and non-axisymmetric flow regimes. *Journal of the Atmospheric Sciences*, 22(5):541–558, 1965.

W. W. Fowlis and H. T. Rossby. A compilation of selected data on the temperature dependence of some physical properties of water, glyerol-water solutions and mercury. Laboratory note, Department of Geology and Geophysics, Massachusetts Institute of Technology, Cambridge 39, Massachusetts, December 1964. Hydrodynamics of rotating fluids project, R. Hide principle investigator.

W. G. Früh and P. L. Read. Wave interactions and the transition to chaos of baroclinic waves in a thermally driven rotating annulus. *Philosophical Transactions: Mathematical Physical and Engineering Sciences*, 355(1722):101–153, 1997.

D. Fultz. Experimental analogies to atmospheric motions. *Compendium of meteorology*, pages 1235–1248, 1951.

D. Fultz, R. R. Long, G. V. Owens, W. Bohan, R. Kaylor, and J. Weil. Studies of thermal convection in a rotating cylinder with some implications for large-scale atmospheric motions. *Meteorological monographs*, 4(21), 1959.

J. M. Gary. Estimate of truncation error in transformed coordinate, primitive equation atmospheric models. *Journal of the Atmospheric Sciences*, 30(2):223–233, 1973.

P. L. George and H. Borouchaki. *Delaunay triangulation and meshing: Application to finite elements*. Hermes, 1998.

C. Geuzaine and J. F. Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11): 1309–1331, 2009.

A. E. Gill. The boundary-layer regime for convection in a rectangular cavity. *Journal of Fluid Mechanics*, 26:515–536, 1966.

J. Grandy. Conservative remapping and region overlays by intersecting arbitrary polyhedra. *Journal of Computational Physics*, 148(2):433–466, 1999.

P. M. Gresho. On the theory of semi-implicit projection methods for viscous incompressible flow and its implementation via a finite element method that also introduces a nearly consistent mass matrix. Part 1: Theory. *International Journal for Numerical Methods in Fluids*, 11(5):587–620, 1990.

P. M. Gresho and R. L. Sani. *Incompressible flow and the finite element method*, volume 1: advection-diffusion. John Wiley & Sons, 2000.

S. M. Griffies, C. Böning, F. O. Bryan, E. P. Chassignet, R. Gerdes, H. Hasumi, A. Hirst, A. M. Treguier, and D. Webb. Developments in ocean climate modelling. *Ocean Modelling*, 2(3-4): 123–192, 2000.

A. Guttman. R-trees: A dynamic index structure for spatial searching. In *Proceedings of the 1984 ACM SIGMOD International Conference on Management of Data*, pages 47–57. ACM, 1984.

D. A. Ham and C. J. Cotter. A mixed discontinuous / continuous Galerkin formulation of the Imperial College Ocean Model: Incompressible Flow. *In preparation*, 2011.

D. A. Ham, J. Pietrzak, and G. S. Stelling. A scalable unstructured grid 3-dimensional finite volume model for the shallow water equations. *Ocean Modelling*, 10(1-2):153–169, 2005.

D. A. Ham, S. C. Kramer, G. S. Stelling, and J. Pietrzak. The symmetry and stability of unstructured mesh C-grid shallow water models under the influence of Coriolis. *Ocean Modelling*, 16 (1-2):47–60, 2007.

D. A. Ham, P. E. Farrell, G. J. Gorman, J. R. Maddison, C. R. Wilson, S. C. Kramer, J. Shipton, G. S. Collins, C. J. Cotter, and M. D. Piggott. Spud 1.0: Generalising and automating the user interfaces of scientific computer models. *Geoscientific Model Development*, 2:33–42, 2009.

M. W. Heinstein and T. A. Laursen. A three dimensional surface-to-surface projection algorithm for non-coincident domains. *Communications in Numerical Methods in Engineering*, 19(6): 421–432, 2003.

V. E. Henson and U. M. Yang. BoomerAMG: A parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics*, 41:155–177, 2002.

R. Hide. Some experiments on thermal convection in a rotating liquid. *Quarterly Journal of the Royal Meteorological Society*, 79(339):161, 1953.

R. Hide. *Some laboratory experiments on free thermal convection in a rotating fluid subject to a horizontal temperature gradient and their relation to the theory of the global atmospheric circulation*, pages 196–221. London: Royal Meteorological Society, 1969.

R. Hide. On the effects of rotation on fluid motions in cylindrical containers of various shapes and topological characteristics. *Dynamics of Atmospheres and Oceans*, 27(1-4):243–256, 1998.

R. Hide. Geomagnetism, "Vacillation", Atmospheric Predicability and "Deterministic Chaos". *Paths of Discovery*, 18:257–254, 2006.

R. Hide and P. J. Mason. Baroclinic waves in a rotating fluid subject to internal heating. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 268(12):201–232, 1970.

R. Hide and P. J. Mason. Sloping convection in a rotating fluid. *Advances in Physics*, 24(1): 47–100, 1975.

P. Hignett. Characteristics of amplitude vacillation in a differentially heated rotating fluid annulus. *Geophysical & Astrophysical Fluid Dynamics*, 31(3-4):247–281, 1985.

P. Hignett, A. A. White, R. D. Carter, W. D. N. Jackson, and R. M. Small. A comparison of laboratory measurements and numerical simulations of baroclinic wave flows in a rotating cylindrical annulus. *Quarterly Journal of the Royal Meteorological Society*, 111(467):131–154, 1985.

T. J. R. Hughes. Recent progress in the development and understanding of SUPG methods with special reference to the compressible Euler and Navier-Stokes equations. *International Journal for Numerical Methods in Fluids*, 7(11):1261–1275, 1987a.

T. J. R. Hughes. *The finite element method: Linear static and dynamic finite element analysis*. Prentice-Hall International, 1987b.

T. J. R. Hughes, M. Mallet, and A. Mizukami. A new finite element formulation for computational fluid dynamics: II. Beyond SUPG. *Computer Methods in Applied Mechanics and Engineering*, 54(3):341–355, 1986.

T. N. L Jacoby, P. L. Read, P. D. Williams, and R. M. B Young. Generation of inertia-gravity waves in the rotating thermal annulus by a localised boundary layer instability. *Geophysical and Astrophysical Fluid Dynamics*, 105(2-3):161–181, 2011.

I. N. James, P. R. Jonas, and L. Farnell. A combined laboratory and numerical study of fully developed steady baroclinic waves in a cylindrical annulus. *Quarterly Journal of the Royal Meteorological Society*, 107(451):51–78, 1981.

X. Jiao and M. T. Heath. Common-refinement-based data transfer between non-matching meshes in multiphysics simulations. *International Journal for Numerical Methods in Engineering*, 61 (14):2402–2427, 2004a.

X. Jiao and M. T. Heath. Overlaying surface meshes, Part I: Algorithms. *International journal of computational geometry & applications*, 14(6):403–419, 2004b.

G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, 20(1):359–392, 1998.

L. Kettner, K. Mehlhorn, S. Pion, S. Schirra, and C. Yap. Classroom examples of robustness problems in geometric computations. *Computational Geometry*, 40(1):61–78, 2008.

S. C. Kramer, C. J. Cotter, and C. C. Pain. Solving the Poisson equation on small aspect ratio domains using unstructured meshes. *Ocean Modelling*, 35(3):253–263, 2010.

O. A. Ladyzhenskaya. *The mathematical theory of viscous incompressible flow*. Gordon and Breach Science Publishers Inc., 1969.

D. Y. Le Roux, A. Staniforth, and C. A. Lin. Finite elements for shallow-water equation ocean models. *Monthly Weather Review*, 126(7):1931–1951, 1998.

D. Y. Le Roux, V. Rostand, and B. Pouliot. Analysis of Numerically Induced Oscillations in 2D Finite-Element Shallow-Water Models Part I: Inertia-Gravity Waves. *SIAM Journal on Scientific Computing*, 29(1):331–360, 2007.

B. Leimkuhler and S. Reich. *Simulating Hamiltonian dynamics*. Cambridge University Press, 2004.

Edward N. Lorenz. *The nature and theory of the general circulation of the atmosphere*. World Meteorological Organisation, 1967.

R. A. Luettich Jr, J. J. Westerink, and N. W. Scheffner. ADCIRC: An Advanced Three-Dimensional Circulation Model for shelves, coasts, and estuaries. Report 1. Theory and Methodology of ADCIRC-2DDI and ADCIRC-3DL. Technical Report DRP-92-6, Department of the Army, US Army Corps of Engineers, Washington DC, November 1992.

J. R. Maddison, C. J. Cotter, and P. E. Farrell. Geostrophic balance preserving interpolation in mesh adaptive linearised shallow-water ocean modelling. *Ocean Modelling*, 37(1-2):35–48, 2011a.

J. R. Maddison, D. P. Marshall, C. C. Pain, and M. D. Piggott. Accurate representation of geostrophic and hydrostatic balance in unstructured mesh finite element ocean modelling. *Ocean Modelling, in press*, 2011b. doi: doi:10.1016/j.ocemod.2011.04.009.

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

J. Marshall, C. Hill, L. Perelman, and A. Adcroft. Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling. *Journal of Geophysical Research*, 102(C3):5733–5752, 1997b.

P. J. Mason. Baroclinic waves in a container with sloping end walls. *Philosophical Transactions for the Royal Society of London. Series A, Mathematical, Physical and Engineering Sciences*, 278(1284):397–445, 1975.

J. D. McCalpin. A comparison of second-order and fourth-order pressure gradient algorithms in a σ-co-ordinate ocean model. *International Journal for Numerical Methods in Fluids*, 18(4): 361–383, 2005.

M. E. McIntyre and W. A. Norton. Potential vorticity inversion on a hemisphere. *Journal of the Atmospheric Sciences*, 57(9):1214–1235, 2000.

G. L. Mellor, T. Ezer, and L. Y. Oey. The pressure gradient conundrum of sigma coordinate ocean models. *Journal of Atmospheric and Oceanic Technology*, 11(4):1126–1134, 1994.

F. Mesinger. On the convergence and error problems of the calculation of the pressure gradient force in sigma coordinate models. *Geophysical and Astrophysical Fluid Dynamics*, 19:105–117, 1982.

S. Micheletti and S. Perotto. Reliability and efficiency of an anisotropic Zienkiewicz–Zhu error estimator. *Computer Methods in Applied Mechanics and Engineering*, 195(9-12):799–835, 2006.

D. R. Munday, D. P. Marshall, and M. D. Piggott. Idealised flow past an island in a dynamically adaptive finite element model. *Ocean Dynamics*, 60(4):835–850, 2010.

C. C. Pain, A. P. Umpleby, C. R. E. De Oliveira, and A. J. H. Goddard. Tetrahedral mesh optimisation and adaptivity for steady-state and transient finite element calculations. *Computer Methods in Applied Mechanics and Engineering*, 190(29-30):3771–3796, 2001.

C. C. Pain, M. D. Piggott, A. J. H. Goddard, F. Fang, G. J. Gorman, D. P. Marshall, M. D. Eaton, P. W. Power, and C. R. E. de Oliveira. Three-dimensional unstructured mesh ocean modelling. *Ocean Modelling*, 10(1-2):5–33, 2005.

M. D. Piggott, C. C. Pain, G. J. Gorman, P. W. Power, and A. J. H. Goddard. h, r, and hr adaptivity with applications in numerical ocean modelling. *Ocean Modelling*, 10(1-2):95–113, 2005.

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

M. D. Piggott, C. C. Pain, G. J. Gorman, D. P. Marshall, and P. D. Killworth. *Unstructured adaptive meshes for ocean modeling*, pages 383–408. American Geophysical Union, 2008b.

P. W. Power, M. D. Piggott, F. Fang, G. J. Gorman, C. C. Pain, D. P. Marshall, A. J. H. Goddard, and I. M. Navon. Adjoint goal-based error norms for adaptive mesh ocean modelling. *Ocean Modelling*, 15(1-2):3–38, 2006.

W. H. Press and G. B. Rybicki. Fast algorithm for spectral analysis of unevenly sampled data. *The Astrophysical Journal*, 338:277–280, 1989.

Q. G. Rayer. *An experimental investigation of heat transfer by large scale motions in rotating fluids*. PhD thesis, University of Oxford, 1992.

Q. G. Rayer. A numerical investigation of the flow in a fully blocked differentially heated rotating fluid annulus. *International Journal of Modern Physics C – Computational Physics and Physical Computation*, 5(2):203–206, 1994.

Q. G. Rayer, D. W. Johnson, and R. Hide. Thermal convection in a rotating fluid annulus blocked by a radial barrier. *Geophysical & Astrophysical Fluid Dynamics*, 87(3-4):215–252, 1998.

P. L. Read. *Rotating annulus flows and baroclinic waves*, chapter IV.3, pages 185–214. Springer, 1992.

P. L. Read. A combined laboratory and numerical study of heat transport by baroclinic eddies and axisymmetric flows. *Journal of Fluid Mechanics*, 489:301–323, 2003.

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

S. H. Risch. *Large-scale interactions in baroclinic flow with topography*. PhD thesis, University of Oxford, 1999.

W. S. Russell and P. R. Eiseman. A boundary conforming structured grid for global ocean circulation studies. *International Journal for Numerical Methods in Fluids*, 28(5):761–788, 1998.

Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving non-symmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869,

1986.

R. Salmon. Two-layer quasi-geostrophic turbulence in a simple special case. *Geophysical & Astrophysical Fluid Dynamics*, 10(1):25–52, 1978.

R. Salmon. *Lectures on geophysical fluid dynamics*. Oxford University Press, 1988.

J. D. Scargle. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal*, 263:835–853, 1982.

J. Shen. Hopf bifurcation of the unsteady regularized driven cavity flow. *Journal of Computational Physics*, 95(1):228–245, 1991.

B. Sitte and C. Egbers. *Higher order dynamics of baroclinic waves*, volume 549, pages 355–375. Springer-Verlag, 2000.

A. St-Cyr, C. Jablonowski, J. M. Dennis, H. M. Tufo, and S. J. Thomas. A comparison of two shallow water models with non-conforming adaptive grids. *Monthly Weather Review*, 136(6): 1898–1922, 2008.

C. Staquet and J. J. Riley. On the velocity field associated with potential vorticity. *Dynamics of Atmospheres and Oceans*, 14:93–123, 1989.

P. K. Sweby. High resolution schemes using flux limiters for hyperbolic conservation laws. *SIAM Journal on Numerical Analysis*, 21(5):995–1011, 1984.

Y. Tanizume, H. Yamashita, and E. Nakamae. Tetrahedral elements generation using topological mapping and space dividing for 3-D magnetic field FEM. *IEEE Transactions on Magnetics*, 26 (2):775–778, 1990.

R. Temam. Une méthode d'approximation de la solution des équations de Navier-Stokes. *Bulletin de la Société Mathématique de la France*, 96:115–152, 1968.

J. F. Thompson, F. C. Thames, and C. W. Mastin. Automatic numerical generation of body-fitted curvilinear coordinate system for field containing any number of arbitrary two-dimensional bodies. *Journal of Computational Physics*, 15(3):299–319, 1974.

G. Umgiesser, D. M. Canu, A. Cucco, and C. Solidoro. A finite element model for the Venice Lagoon. Development, set up, calibration and validation. *Journal of Marine Systems*, 51(1-4): 123–145, 2004.

G. K. Vallis. *Atmospheric and oceanic fluid dynamics: Fundamentals and large-scale circulation*. Cambridge University Press, 2006.

P. Vaněk, J. Mandel, and M. Brezina. Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing*, 56(3):179–196, 1996.

Y. V. Vasilevskii and K. N. Lipnikov. An adaptive algorithm for quasioptimal mesh generation. *Computational Mathematics and Mathematical Physics*, 39(9):1468–1486, 1999.

F. Vettin. Ueber den aufsteigenden Luftstrom, die Entstehung des Hagels und der Wirbel-Stürme. *Annalen der Physik*, 178(10):246–255, 1857.

M. R. Wells, P. A. Allison, G. J. Hampson, M. D. Piggott, and C. C. Pain. Modelling ancient tides: the Upper Carboniferous epi-continental seaway of Northwest Europe. *Sedimentology*, 52(4): 715–735, 2005a.

M. R. Wells, P. A. Allison, M. D. Piggott, C. C. Pain, G. J. Hampson, and C. R. E. De Oliveira. Large sea, small tides: the Late Carboniferous seaway of NW Europe. *Journal of the Geological Society*, 162(3):417–420, 2005b.

H. Weyl. The method of orthogonal projection in potential theory. *Duke Mathematical Journal*, 7 (1):411–444, 1940.

G. P. Williams. Numerical integration of the three-dimensional Navier-Stokes equations for incompressible flow. *Journal of Fluid Mechanics*, 37:727–750, 1969.

P. D. Williams. *Nonlinear interactions of fast and slow modes in rotating, stratified fluid flows*. PhD thesis, University of Oxford, 2003.

P. D. Williams, P. L. Read, and T. W. N. Haine. Spontaneous generation and impact of inertia-gravity waves in a stratified, two-layer shear flow. *Geophysical Research Letters*, 30(24):2255, 2003.

P. D. Williams, T. W. N. Haine, and P. L. Read. Stochastic resonance in a nonlinear model of a rotating, stratified shear flow, with a simple stochastic inertia-gravity wave parameterization. *Nonlinear Processes in Geophysics*, 11(1):127–135, 2004.

P. D. Williams, T. W. N. Haine, P. L. Read, S. R. Lewis, and Y. H. Yamazaki. QUAGMIRE v1.3: A quasi-geostrophic model for investigating rotating fluids experiments. *Geoscientific Model Development*, 2:13–32, 2009.

R. T. Williams and O. C. Zienkiewicz. Improved finite element forms for the shallow-water wave equations. *International Journal for Numerical Methods in Fluids*, 1(1):81–97, 1981.

C. Wilson. *Modelling multiple-material flows on adaptive unstructured meshes*. PhD thesis, Imperial College London, 2009.

R. D. Wordsworth. *Theoretical and experimental investigations of turbulent jet formation in planetary fluid dynamics*. PhD thesis, University of Oxford, 2008.

R. D. Wordsworth, P. L. Read, and Y. H. Yamazaki. Turbulence, waves, and jets in a differentially heated rotating annulus experiment. *Physics of Fluids*, 20(12):126602, 2008.

R. M. B. Young and P. L. Read. Flow transitions resembling bifurcations of the logistic map in simulations of the baroclinic rotating annulus. *Physica D: Nonlinear Phenomena*, 237(18): 2251–2262, 2008a.

R. M. B. Young and P. L. Read. Breeding and predictability in the baroclinic rotating annulus using a perfect model. *Nonlinear Processes in Geophysics*, 15(3):469–487, 2008b.

O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu. *The finite element method: Its basis and fundamentals*. Butterworth-Heinemann, 2005.