

Coupling the Planetary Boundary Layer to the Large Scale Dynamics of the Atmosphere: The Impact of Vertical Discretisation

Submitted by

Dan Holdaway

to the University of Exeter as a thesis for the degree of Doctor of Philosophy in
Mathematics, May 2010.

This thesis is available for Library use on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

I certify that all material in this thesis which is not my own work has been identified and that no material is included for which a degree has previously been conferred upon me.

.....
Dan Holdaway

Abstract

Accurate coupling between the resolved scale dynamics and sub-grid scale physics is essential for accurate modelling of the atmosphere. Previous emphasis has been towards the temporal aspects of this so called physics-dynamics coupling problem, with little attention towards the spatial aspects. When designing a model for numerical weather prediction there is a choice for how to vertically arrange the required variables, namely the Lorenz and Charney-Phillips grids, and there is ongoing debate as to which is the optimal. The Charney-Phillips grid is considered good for capturing the large scale dynamics and wave propagation whereas the Lorenz grid is more suitable for conservation. However the Lorenz grid supports a computational mode. In the first half of this thesis it is argued that the Lorenz grid is preferred for modelling the stably stratified boundary layer. This presents the question: which grid will produce most accurate results when coupling the large scale dynamics to the stably stratified planetary boundary layer? The second half of this thesis addresses this question.

The normal mode analysis approach, as used in previous work of a similar nature, is employed. This is an attractive methodology since it allows one to pin down exactly why a particular configuration performs well. In order to apply this method a one dimensional column model is set up, where horizontally wavelike solutions with a given wavenumber are assumed. Applying this method encounters issues when the problem is non normal, as it will be when including boundary layer terms. It is shown that when addressing the coupled problem the lack of orthogonality between eigenvectors can cause mode analysis to break down. Dynamical modes could still be interpreted and compared using the eigenvectors but boundary layer modes could not. It is argued that one can recover some of the usefulness of the methodology by examining singular vectors and singular values; these retain the appropriate physical interpretation and allow for valid comparison due to orthogonality between singular vectors.

Despite the problems in using the desirable methodology some interesting results have been gained. It is shown that the Lorenz grid is favoured when the

boundary layer is considered on its own; it captures the structures of the steady states and transient singular vectors more accurately than the Charney-Phillips grid. For the coupled boundary layer and dynamics the Charney-Phillips grid is found to be most accurate in terms of capturing the steady state. Dispersion properties of dynamical modes in the coupled problem depend on the choice of horizontal wavenumber. For smaller horizontal wavenumber there is little to distinguish between Lorenz and Charney-Phillips grids, both the frequency and structure of dynamical modes is captured accurately. Dynamical mode structures are found to be harder to interpret when using larger horizontal wavenumbers; for those that are examined the Charney-Phillips grid produces the most sensible and accurate results. It is found that boundary layer modes in the coupled problem cannot be concisely compared between the Lorenz and Charney-Phillips grids due to the issues that arise with the methodology. The Lorenz grid computational mode is found to be suppressed by the boundary layer, but only in the boundary layer region.

Acknowledgements

First and foremost I would like to thank my supervisors John Thuburn and Nigel Wood. They have been a tremendous help and support throughout the course of my PhD and I'm forever grateful for everything I have learnt from them. Thanks also to the Met Office and to the EPSRC for jointly funding the work.

I would like to thank my parents Sue and Graham and brother Dom for always being there for me and for their continued and generous support. Thanks also to my extended family for their support and always showing such a keen interest.

Thanks to all the people I have known and lived with while in Exeter. Special thanks to my house mate Andy and to K4, Joe, Mercedes, Jamie and Tessa.

Finally a loud thank you to all of my fellow PhD students for making PhD life so entertaining!

Contents

Contents	5
Acknowledgements	5
List of Figures	8
List of Tables	14
1 Introduction and Motivation	16
1.1 Introduction	16
1.2 Dynamics	17
1.3 Physics Parametrisation	18
1.3.1 The Atmospheric Boundary Layer	19
1.4 Physics-Dynamics Coupling	21
1.4.1 Stably Stratified Planetary Boundary Layer	23
1.4.2 Grid Staggering	25
1.4.3 Grid Spacing	27
1.5 Governing Equations	28
1.6 Methodolgy	30
1.6.1 Linearisation	31
1.6.2 Dispersion Relation	32
1.6.3 Eigendecomposition	35
1.6.4 Singular Value Decomposition (SVD)	36
1.6.5 Discrete and Continuous Spectra	38
1.7 Thesis Outline	40

I	Vertical Discretisations for the Stably Stratified Planetary Boundary Layer	42
2	The Boundary Layer Model	44
2.1	Governing Equations	45
2.1.1	Reynolds Averaging	46
2.1.2	Boussinesq Approximation	52
2.2	Model Closure and Parametrisation	54
2.2.1	Parametrisation	55
2.2.2	The Surface Boundary Conditions	57
2.3	Subsidence to Balance the Boundary Layer Diffusion	59
2.4	Test Cases	62
2.5	Discretisation and Averaging	63
2.5.1	Choice of Vertical Staggering	64
2.5.2	Grid Spacing	70
3	Steady State Solution	77
3.1	Newton Iteration	80
3.2	High Resolution Convergence	83
3.3	Lorenz Configuration	88
3.4	Charney-Phillips Option I (Ri computed at z_w levels)	92
3.5	Charney-Phillips Option II (Ri computed at z_ρ levels)	98
3.6	Charney-Phillips Option III (Ri computed at all levels)	101
3.7	Comparison of various types of grid spacing	105
3.8	Concluding Remarks on Steady State Solutions	112
4	Transient Solution	115
4.1	Transient Equations and Discretisation	117
4.2	Matrix Normality	119
4.2.1	The form of Matrix A	121
4.3	Singular Value Decomposition (SVD)	122
4.3.1	Energetics and The Energy-Norm	125

4.3.2	Form of the Spectrum	129
4.4	Testing Methodology	131
4.5	Physical Interpretation of Eigenvectors and Singular Vectors	133
4.6	Comparison of High Resolution Singular Vectors	142
4.7	Comparison of Lorenz and Charney-Phillips Low Resolution	147
4.7.1	Comparison with Charney-Phillips Option I-iii	158
4.7.2	Shallower Boundary Layers and Alternative Grid Spacing	162
5	Summary of Part I	168
II	Comparison of Vertical Discretisations Whilst Coupling Boundary Layer and Dynamics	176
6	The Fully Compressible Dynamics-Boundary Layer Model	179
6.1	Form of the Vertical Pressure Gradient	181
6.2	Linearisation	182
6.2.1	Steady State	182
6.2.2	Transients	186
6.2.3	Energy Norm	190
6.3	Case studies	193
7	Examination of Mode Types and Interaction	195
7.1	Isothermal - Dynamics only Case	196
7.1.1	Capturing Modes with Singular Vectors	205
7.1.2	Lorenz Grid Computational Mode	210
7.1.3	Shorter Horizontal Wavelengths	210
7.2	Isothermal Case with Stretched Grid	213
7.2.1	Lorenz Grid Computational Mode	215
7.2.2	Shorter Horizontal Wavelength	216
7.2.3	Mode Structure	218
7.3	Coupled Reference State, Dynamics only in the Linearisation	223
7.3.1	Dispersion Relation	224

7.3.2	Unstable and Decaying Modes	228
7.4	Fully Coupled Linearisation	233
7.4.1	Mode Structures	238
7.4.2	Comments on the Methodology	248
7.4.3	Mode Tracking: Motivation	251
7.4.4	Mode Tracking: Methodology	255
7.5	Lorenz Grid Computational Mode	258
7.6	Chapter Summary	261
8	Comparison of Lorenz and Charney-Phillips Grids	266
8.1	Steady State Comparison	268
8.1.1	High Resolution Steady States	268
8.1.2	Operational Resolution Steady States	270
8.2	Coupled Transient Comparison	274
8.2.1	Examination of Dynamical Modes	276
8.2.2	Examination of Damped Modes	288
8.3	Comparing Transients Using the SVD	295
8.4	Chapter Summary	299
9	Overall Summary	302
9.1	Methodology	304
9.2	Vertical Configurations	308
9.3	Concluding Discussion	311
	Bibliography	313

List of Figures

1.1	Arakawa A- to D-grids for the shallow water equations	33
1.2	Lorenz v Charney-Phillips for isothermal Euler equations	34
2.1	Lorenz and Charney-Phillips grid schematic	65
2.2	Distribution of points for log and log-linear grids	76
3.1	High resolution Lorenz and Charney-Phillips with different stretching	84
3.2	High resolution errors for different grid stretching	85
3.3	θ comparison of log and log-linear high res grids	86
3.4	High resolution error convergence rates, Lorenz grid	87
3.5	High resolution error convergence rates, Charney-Phillips grid . . .	88
3.6	Lorenz grid steady states against high resolution - velocity	90
3.7	Lorenz grid steady states against high resolution - potential temper- ature	91
3.8	Charney-Phillips option I steps	92
3.9	Lorenz versus Charney-Phillips steady states against high resolution - velocity	93
3.10	Lorenz versus Charney-Phillips steady states against high resolution - potential temperature	94
3.11	Comparison of different versions of Charney-Phillips option I steady states - velocity	95
3.12	Comparison of different version of Charney-Phillips option I steady states - potential temperature	96
3.13	Test case comparing option I configurations	97

3.14	Charney-Phillips option II steps	98
3.15	Lorenz and Charney-Phillips option II steady states against high resolution - velocity	99
3.16	Lorenz and Charney-Phillips option II steady states against high resolution - potential temperature	100
3.17	Charney-Phillips option III steps	102
3.18	Flow diagram for Charney-Phillips Option II(d)-i	104
3.19	Flow diagram for Charney-Phillips Option III(d)	105
3.20	Lorenz and Charney-Phillips option III steady states against high resolution - velocity	106
3.21	Lorenz and Charney-Phillips option III steady states against high resolution - potential temperature	107
3.22	Comparison of steady states for different grid spacing schemes - Lorenz	108
3.23	Comparison of steady states for different grid spacing schemes - potential temperature	109
3.24	Comparison of steady states for different grid spacing schemes - Charney-Phillips	110
3.25	Comparison of steady states for different grid spacing schemes - potential temperature	111
3.26	Comparison of steady states for different grid spacing schemes - near surface	112
4.1	Successive eigenvectors, showing orthogonality issue	120
4.2	A leading singular vector for θ'	129
4.3	Boundary layer 5 eigenvalues in the complex plane	134
4.4	Boundary layer 5 eigenvalues by real part	135
4.5	Selection of eigenvectors for boundary layer 5	136
4.6	Boundary layer 5 high resolution singular values	137
4.7	Boundary layer 5 inertio type singular vector	138
4.8	Boundary layer 5 θ' singular vector	139
4.9	Boundary layer 5 leading boundary layer singular vector	140

4.10	Boundary layer 5 highly damped boundary layer singular vector . . .	141
4.11	Boundary layer 1 high resolution singular values	142
4.12	Comparison of singular values for boundary layer 5	150
4.13	Comparison of eigenvalues for boundary layer 5	150
4.14	Comparison of singular vectors for leading boundary layer singular vector, boundary layer 5	152
4.15	Comparison of singular vectors for leading boundary layer singular vector, boundary layer 5	154
4.16	Comparison of singular vectors for leading boundary layer singular vector, boundary layer 5	155
4.17	Comparison of singular vectors for leading boundary layer singular vector, boundary layer 5	156
4.18	Comparison of energy norm singular vectors for leading boundary layer singular vector, boundary layer 5	157
4.19	Comparison of Lorenz and Charney-Phillips I-iii by eigenvalues . . .	159
4.20	Leading boundary layer singular vector from Lorenz and Charney- Phillips I-iii	160
4.21	Leading boundary layer singular vector from Lorenz and Charney- Phillips I-iii	161
4.22	Leading boundary layer singular vector comparison for boundary layer 1	166
4.23	Leading boundary layer singular vector comparison for boundary layer 4	167
7.1	Isothermal Dispersion Relation Charney-Phillips	197
7.2	Isothermal Dispersion Relation Lorenz	198
7.3	Isothermal Acoustic Mode	201
7.4	Isothermal Inertio-Gravity Mode (Gravity End)	203
7.5	Isothermal Inertio-Gravity Mode (Inertial End)	204
7.6	Isothermal Rossby Mode	205
7.7	Isothermal Acoustic Singular Vectors	206

7.8	Isothermal Rossby Singular Vectors	207
7.9	Isothermal Computational Mode	211
7.10	Isothermal Dispersion, Larger Horizontal Wavenumber, Charney-Phillips	212
7.11	Isothermal Dispersion, Larger Horizontal Wavenumber, Lorenz	212
7.12	Isothermal Dispersion Relation, Stretched Grid, Charney-Phillips	215
7.13	Isothermal Dispersion Relation, Stretched Grid, Lorenz	215
7.14	Isothermal Stretched Grid Computational Mode	216
7.15	Isothermal Dispersion Relation, Stretched Grid and Larger Horizontal Wavenumber, Charney-Phillips	217
7.16	Isothermal Dispersion Relation, Stretched Grid and Larger Horizontal Wavenumber, Lorenz	217
7.17	Isothermal Inertio-Gravity Mode, Stretched Grid, Larger Horizontal Wavenumber	219
7.18	Isothermal Rossby Mode, Stretched Grid, Larger Horizontal Wavenumber	220
7.19	Isothermal Rossby Singular Vector, Stretched Grid, Larger Horizontal Wavenumber	222
7.20	Frequencies For Coupled But No Transient Boundary Layer, Lorenz Grid	225
7.21	Frequencies For Coupled But No Transient Boundary Layer, Larger Horizontal Wavenumber, Lorenz Grid	226
7.22	Eigenvalues For Coupled But No Transient Boundary Layer, Lorenz Grid	229
7.23	Growth/Decay Rates For Coupled But No Transient Boundary Layer, Lorenz Grid	230
7.24	Frequency For Coupled But No Transient Boundary Layer, Larger Horizontal Wavenumber, Lorenz Grid	230
7.25	Growth/Decay Rates For Coupled But No Transient Boundary Layer, Larger Horizontal Wavenumber, Lorenz Grid	231

7.26	Frequency For Coupled But No Transient Boundary Layer, Larger Horizontal Wavenumber, Lorenz Grid	231
7.27	Eigenvalues For Fully Coupled, Lorenz Grid	234
7.28	Growth/Deacy Rates For Fully Coupled, Lorenz Grid	235
7.29	Frequencies For Fully Coupled, Lorenz Grid	235
7.30	Growth/Deacy Rates For Fully Coupled, larger Horizontal Wavenum- ber, Lorenz Grid	237
7.31	Frequencies For Fully Coupled, larger Horizontal Wavenumber, Lorenz Grid	237
7.32	Fully Coupled Inertio-Gravity Mode	240
7.33	Fully Coupled Rossby Mode	242
7.34	Fully Coupled Damped Mode	243
7.35	Growth/Deacy Rates For Fully Coupled, Boundary Layer 5, Lorenz Grid	247
7.36	Frequencies For Fully Coupled, Boundary Layer 5, Lorenz Grid . . .	247
7.37	Fully Coupled Inertio-Gravity Singular Vector	250
7.38	Fully Coupled Rossby Singular Vector	251
7.39	Mode Tracking From Boundary Layer Off to Boundary On in the Transients	257
7.40	Boundary Layer Off Modes That Become Damped	258
7.41	Computational Mode in Fully Coupled, Boundary Layer Off in Tran- sients	259
7.42	Two Grid Wave Schematic	260
7.43	Computational Mode in Fully Coupled	261
8.1	High Resolution Steady State Density	269
8.2	High Resolution Steady State Dynamics Versus Boundary Layer . .	270
8.3	Lorenz Versus Charney-Phillips Coupled Reference State U and V .	271
8.4	Lorenz Versus Charney-Phillips Coupled Reference State $\theta^{(r)}$ and $\rho^{(r)}$	272
8.5	Lorenz vs Charney-Phillips Boundary Layer 1 Off in Transients . .	277
8.6	Lorenz vs Charney-Phillips Boundary Layer 1 On in Transients . .	278

8.7	Frequency Lorenz vs Charney-Phillips Boundary Layer 5 On in Transients	279
8.8	Lorenz vs Charney-Phillips Boundary Layer 1 On in Transients	281
8.9	Lorenz vs Charney-Phillips Coupled Inertio-Gravity Mode	283
8.10	Lorenz vs Charney-Phillips Coupled Rossby Mode	285
8.11	Lorenz vs Charney-Phillips Coupled $k = \frac{2\pi}{10^3}$ Mode	287
8.12	Lorenz vs Charney-Phillips Coupled BL1 damped modes	290
8.13	Growth/Decay Lorenz vs Charney-Phillips Boundary Layer 5 On in Transients	291
8.14	Lorenz vs Charney-Phillips Coupled Boundary Layer Mode	293
8.15	Lorenz versus Charney-Phillips Boundary Layer Mode in Singular Vectors	296
8.16	Lorenz versus Charney-Phillips Acoustic Mode in Singular Vectors	298

List of Tables

2.1	Test cases by boundary depths	63
2.2	Charney-Phillips averaging choices	68
2.3	Values of α in geometric grid.	71
4.1	Boundary layer 5, comparison of high resolution singular vectors by orthogonality test	146
4.2	Boundary layer 5, comparison of low and high resolution singular vectors by orthogonality test	149
4.3	Comparison of Lorenz and Charney-Phillips I-iii by orthogonality .	158
4.4	Comparison of Lorenz and Charney-Phillips by orthogonality, bound- ary layer 1	163
4.5	Comparison of Lorenz and Charney-Phillips by orthogonality, bound- ary layer 2	164
4.6	Comparison of Lorenz and Charney-Phillips by orthogonality, bound- ary layer 3	164
4.7	Comparison of Lorenz and Charney-Phillips by orthogonality, bound- ary layer 4	165
7.1	Dominant Energy Variables	206
7.2	Continuous Spectrum Search, Fully Coupled	254
8.1	μ For Most Damped Boundary Layer 1	290
8.2	μ For Least Damped Boundary Layer 5 BL modes	295

Chapter 1

Introduction and Motivation

1.1 Introduction

The atmosphere has a highly multi-scale structure. Spatially, processes such as micro-scale turbulence can occur on scales of less than 1 metre while simultaneously large-scale planetary motion can influence regions of up to ten thousand kilometres. These processes develop over a considerable range of temporal scales, the smallest scale turbulence will occur over a period of the order 1 second while the largest scale processes may be evolving over a period of the order of weeks or even months.

There are a set of governing equations for the atmosphere which describe the multitude of physical processes that can occur. The complex nature of the fluid flows in the atmosphere means analytical solutions are only available for highly idealised cases, no solution to the full governing equations is currently known. Instead the equations must be solved numerically. The number of grid points that can be used in a model is dictated by the amount of computing power that is available, the length of the forecast and frequency at which forecast information is required. An institute such as the UK Met Office produce operational numerical weather forecasts using a global model. They are currently able to employ a global model having a horizontal grid spacing of 25km in the mid latitudes and with 70 vertical levels reaching up to 80km high. The time step is 15 minutes. Any model can only be said to resolve the scales which are larger than the grid spacing. A

model with this kind of resolution will be unable to resolve many of the important processes, such as small-scale turbulence and convection, that occur either below its spatial resolution or evolve quicker than the time step. Processes that have scale larger than the grid-scale are colloquially referred to as the ‘dynamics’ while sub grid-scale processes that are unresolved are referred to as the ‘physics’. Even though their scale is relatively small the physics include some important, influential processes that occur in the atmosphere. Capturing boundary layer turbulence, for example, is crucial for accurate representation of the daily weather and near surface pollution; representing convective clouds is important as they can produce heavy rain and lead to destructive thunderstorms. Both of these processes, and many more, can be sub grid-scale.

1.2 Dynamics

The dynamics, or dynamical core, consists of all the resolved-scale processes that occur in the atmosphere. This includes all the types of large-scale wave oscillations that are possible, such as acoustic waves, gravity waves, inertial waves and Rossby waves as well as all the types of nonlinear interactions. Any numerical model that is used should be capable of representing the most important of these structures accurately as they have the potential to drive the climate and weather systems of the planet. Rossby waves are considered the most important large-scale wave in the atmosphere [34], as they propagate and become distorted the large-scale of their motion will generate the cyclones and anticyclones that can have such governance on the weather. If the model cannot accurately represent the features of the Rossby wave, for example if the speed of propagation is too slow, or the structure mis-represented, then the model may place large-scale weather systems incorrectly, thereby reducing the value of the forecast.

It should be noted that sometimes it is actually desirable to ignore certain dynamics for improved accuracy and reduced computational effort. For example, it may be argued that acoustic waves have little meteorological significance and so a model that did not capture them could be beneficial. Fast propagation of

acoustic waves might otherwise demand a very small time step for stability [23]. These approximations are known as filtered equation sets, probably the most well known of which is the Boussinesq approximation.

The large-scale waves that occur in the atmosphere arise due to perturbations acting against a restoring mechanism. Acoustic waves occur when a perturbation to pressure occurs, this perturbation is then restored by the tendency of fluid in regions of high pressure to move towards regions of low pressure. As one part of the fluid moves to restore the perturbation so another perturbation occurs with the same resulting restoration, subsequently the perturbation propagates away from the source in the form of a wave. Different waves are possible due to different restoring mechanisms, gravity waves have gravity, or stratification, as the restoring mechanism, inertial waves are restored by rotational forces. Rossby waves occur when a perturbation occurs in potential vorticity, restored by the law of conservation of potential vorticity. Each type of wave motion can be recognised in the measured atmospheric fields and will dominate in the field most associated to the perturbation and restoration force. Acoustic waves would be observed as dominant perturbations in the pressure field, gravity waves as perturbations in buoyancy, inertial waves in the wind fields and Rossby waves in the potential vorticity.

1.3 Physics Parametrisation

Although the physics cannot be resolved by the grid, much is known about the physical processes that occur at this scale. Using this knowledge mathematical models can be constructed that represent the interaction that occurs between the resolved and sub grid processes. Including these representations in the model will ensure that their influence on the overall flow is considered. Since the only thing known about the atmosphere numerically at any given time comes from the quantities stored on the grid, i.e. the resolved processes, the physics needs to be represented by writing the sub grid process in terms of the large-scale flow. This process of determining unknown quantities in terms of the known parameters is called parametrisation. All kinds of processes need to be parametrised in order to perform

an accurate forecast, such as boundary layer turbulence, convection, clouds, land interaction, gravity wave drag and radiation. Parametrisation in this sense refers to the representation of processes which are otherwise not readily accounted for.

Once a process has been identified as requiring a parametrisation a mathematical theory needs to be developed based on the physical properties of that type of phenomenon. Numerical testing is then performed and can be checked by examining observations. The overall numerical model will, in general, consist of the dynamical core plus the physics parametrisations.

1.3.1 The Atmospheric Boundary Layer

The atmospheric boundary layer is the turbulent region in the lowest part of the troposphere. Its structure is highly distinguishable from the free atmosphere and is defined as the region in which flow characteristics are directly influenced by interactions with the surface of the Earth. Accurate numerical representation of the structure and evolution of the boundary layer is important not only because it helps describe everyday weather conditions but also because it describes the transfer of heat, moisture and momentum from the surface to the free atmosphere. Further it can trap pollutants, affecting air quality, it largely defines how much of the energy in solar radiation is transmitted to the rest of the atmosphere and it has a large influence on ocean currents through wind stress. The study of the boundary layer has the added motivation that it inevitably requires numerical modelling of turbulence, something notoriously difficult to achieve accurately due to range of scales and complex motions involved.

The structure and depth of the atmospheric boundary layer is often governed by shear driven turbulence, this is either suppressed or enhanced by the stratification and hence evolves with a diurnal cycle. During the day the majority of the high energy short-wave radiation from the sun passes through the atmosphere and acts to heat the surface of the earth, in turn heating the air closest to the surface. This warmer less dense air then travels convectively upwards to regions of higher density, cooler air, enhancing the turbulence. This in conjunction with the turbulence

generated by shear, existing due to surface friction acting against the wind, can produce turbulence high above the surface, up to the order of a kilometre. As the sun begins to set the surface cools and the air in the boundary layer becomes stably stratified, i.e. with cool air below warm air. The stable stratification acts to suppress the shear generated turbulence. Meanwhile however shear generated turbulence may be increased, for example by the nocturnal jet [62] so a turbulent region still exists. Sometimes, if the shear is small, the turbulence in the boundary layer can be governed mainly by the surface heat flux. The overall height of the boundary layer is generally lower at night, of the order of a few hundred metres. The daytime boundary layer is referred to as the convective or unstably stratified boundary layer while the nocturnal boundary layer is known as the stably stratified boundary layer. A further neutral case can occur where the potential temperature is constant throughout the depth of the boundary layer, this may be useful for simplifying the model and assisting in analysis.

Frictional Damping in the Atmosphere

Turbulent exchange of momentum with the surface in the atmospheric boundary layer has a tendency to slow the flow of air and so it is often referred to as frictional damping. This frictional damping can produce some interesting and important phenomena in the atmosphere. One example of an effect of frictional damping is when a component of the near surface wind becomes directed across isobar towards low pressure, this will occur in the majority of the boundary layer region. Flow towards regions of low pressure results in mass flux towards the centre of cyclones and away from the centre of anticyclones. By the principle of mass conservation this results in ascent at the top of the boundary layer in a cyclone and descent at the top of the boundary layer in an anticyclone. The process of ascent and descent is known as Ekman pumping and leads to the often overcast or clear skies that are associated with low and high pressure respectively. The rate of this ascent or descent may be of the order of a few millimetres per second for synoptic scale flow [34].

A further process that can result from the boundary layer friction is cyclone spin-down. The process of Ekman transport in the boundary layer, as described above, is balanced by an ascent above the boundary layer. Similarly, for cyclones, there is an outwardly radial flow above the boundary layer balancing the inwardly radial flow in the boundary layer. As this outward flow occurs fluid with high angular momentum is replaced by fluid with low angular momentum, causing a spin-down of the cyclone. Typical spin-down time scales are of the order of a few days [34].

The time scales as given by [34] are computed for barotropic cyclones, of more interest is the effect the boundary layer can have on baroclinic cyclones. This question is addressed in the studies of [1] and [5]. In these papers the spin down of baroclinic cyclones is studied and attributed to either potential vorticity generation in the boundary layer or surface stress. Spin down time scales are found to be of a similar order to that found for barotropic cyclones.

In this thesis the influence of the boundary layer on cyclone development will not be explicitly examined, however it will aid in the study to have an understanding of these mechanisms and the associated time scales.

1.4 Physics-Dynamics Coupling

An important current challenge in numerical weather prediction is to obtain accurate coupling between models formulated for sub grid-scale physics and the models formulated for resolved large-scale dynamics. This is a difficult problem to overcome not only due to the differences in spatial scale but also in the temporal-scale. A number of authors have provided discussion for the appropriate time stepping scheme for the coupled problem, [15, 20, 21, 22, 60, 61, 76, 79].

So far little emphasis has been placed on the spatial aspects of the physics-dynamics coupling problem and it is here that this thesis aims to assist in addressing this aspect of the challenge. When constructing a numerical model for weather and climate prediction there is considerable choice for how to construct the model in space, namely the three connected issues of vertical coordinate system, variable

selection and variable arrangement.

Examples for the choice of vertical coordinate may include: height-based, height-based terrain-following [13]; pressure or mass-based terrain-following [44]; isentropic-based [35]; or Lagrangian coordinate [46, 40]. A particular coordinate system may be beneficial due to a simplification that is afforded in the governing equations. The variables required in order to obtain a closed calculation of the governing equations include the three components of velocity and two thermodynamics variables. The two thermodynamic variables can be chosen from, for example, potential temperature, temperature, pressure, density or entropy. The arrangement of the variables divides into choice of horizontal staggering and choice of vertical staggering. For a number of reasons it may be beneficial to store model variables at different places in space. For example, if an equation required the product of pressure gradient and a velocity it would be beneficial to the accuracy of that equation if the pressure were staggered relative to (placed halfway between) the velocity, then a second order finite difference numerical gradient would automatically be in the right place without the need for averaging. Of course this needs to be considered in the context of a whole system of equations where a particular staggering might be beneficial to one part of an equation whilst damaging to another part. Examples of the horizontal staggering include the classic Arakawa A- to E-grids [3]. Popular in atmospheric modelling, due to good representation of the dispersion relation, and in operational use at the Met Office [59], is the C-grid. In the vertical there are two common choices of staggering the Lorenz [48] and Charney-Phillips [12] grids. There is ongoing debate as to which staggering is the best option, for example the Met Office now employ the Charney-Phillips configuration [59] whereas the ECMWF [73, 7] use a variation of the Lorenz configuration. These issues of spatial arrangement are inextricably connected with each other. For a particular vertical coordinate it may be preferable to use certain thermodynamic variables and as a result a certain grid staggering.

Considering the breadth of choice available, it is likely that some heuristic argument would be applied, based on the features that a model should be capable of

capturing in order to make a decision. The issue with this however is that the full equation set is highly complex and so differences between two rival configurations may not be immediately apparent or clear cut. In order to address this question [68] constructed 168 test cases covering three types of vertical coordinate, every combination of two from a choice of five thermodynamic variables and a number of different vertical staggers, including Lorenz and Charney-Phillips. They concentrated on the vertical configuration since the horizontal staggering has been well studied, e.g. [27]. The methodology that was used, which will be discussed presently, allowed them to systematically check every combination and grade any configuration from optimal to problematic; giving clear and concise conclusions.

The work of [68] was for the inviscid case, i.e. one where only the dynamics are captured in the model but without small-scale physics. The overall aim of this thesis will be to extend the analysis to cover a particular physics-dynamics coupled problem. This is a question of interest due to potential contradictions between the optimal configuration found for the dynamics on its own and the optimal configuration for the physics on its own. For example [68] found Charney-Phillips to give the optimal configuration for the dynamics. However it is likely that Lorenz grid is preferable for the boundary layer parametrisation e.g. [13].

This thesis will focus on the choice of the stably stratified planetary boundary layer for the physics since this is where there is the greatest apparent conflict in the choice of vertical staggering. The dynamical mechanisms can be represented by modelling the inviscid Navier-Stokes equations (Euler equations). The aim is to investigate the spatial aspects of coupling these and in particular the best choice of vertical staggering, the Lorenz or Charney-Phillips grid.

1.4.1 Stably Stratified Planetary Boundary Layer

The choice of vertical staggering is most relevant in a problem where local variability in the vertical is strong. For a stably stratified boundary layer buoyancy and wind speed can vary dramatically with height; for a convective boundary layer quantities will generally be well mixed and so will not vary as dramatically in the

vertical. The parametrisation of the stably stratified boundary layer relies on a local quantity, known as the gradient Richardson number; it will be shown that it is due to this quantity that a conflict in the choice of vertical staggering arises between the stably stratified boundary layer and the dynamics.

The main aim of this thesis will be to explore the conflicts in the choice of staggering. However, in addition to the question of staggering, there are a number of other ongoing challenges related to the study of the stably stratified boundary layer. One of these challenges is in setting up a suitable stable numerical scheme. Due to the way that sub grid processes are parametrised, modelling of the stably stratified boundary layer represents a nonlinear diffusion problem. The relative time scales of this nonlinear diffusion, in comparison with the typical time step of a global model, means that care is required in order to set up a properly stable and accurate numerical scheme. A typical time step in a global model may be around 15 minutes. The time scale of the diffusion depends on the square of the vertical grid spacing divided by the eddy viscosity, both of which vary with height. Typically diffusion may occur over periods ranging from a few seconds up to hours. Explicit time stepping would only be conditionally stable, and since the time scale of the diffusion may be quite small, the condition would be unlikely to hold for the kinds of time steps required for global models. An implicit scheme would be required in order to gain unconditional stability, however implicit schemes would be quite computationally demanding. Further to this, care is required when choosing the order for the time stepping. A second order accurate scheme will provide high levels of accuracy if the diffusion time scale is long compared to the time step in the model. If the time scale is short then the diffusion is not represented well and a first order accurate scheme may be beneficial for damping noise. As outlined by [80], the problem when it comes to designing the numerical scheme, is that diffusion occurring at time scales much longer and much shorter than the model's time step can occur simultaneously in the domain. The challenge is to design a numerical scheme that is stable, accurate and that does not require excessive computational effort.

In addition to the work of [80] the problem of setting up a stable and accurate time stepping scheme for the stably stratified boundary layer was also addressed by [38] and [29]. In both papers linear stability analysis is presented for nonlinear diffusion equations. [38] examine several different implicit and explicit schemes and show them to have a number of limitations, they also present novel examples of some alternative schemes with good stability properties. In [29] two novel schemes are presented for giving stability in the numerics. The approach employed is one which examines the rate of change of the diffusion operator with respect to the model variables through the time step.

This thesis does not intend to specifically address the question of numerical stability for modelling the stably stratified boundary layer. However, in the process of addressing the question of staggering, the numerics of the boundary layer will be examined, including some the time scales of linear variability. Doing so will help to build an overall understanding of the stably stratified planetary boundary layer and provide insight into the important scales in the problem.

1.4.2 Grid Staggering

The principal focus for this thesis is the effect of different vertical staggering and different vertical spacing for the coupled problem. It was previously noted that the C-grid is the horizontal staggering of choice in many operational forecast models and is as a result of its beneficial properties [3, 65]. The choice of vertical staggering is also a topic of attention [2, 14, 27, 45, 64, 68, 69] but less so within the context of physics-dynamics coupling. A previous investigation of differences between the Lorenz and Charney-Phillips grid in a coupled environment can be seen in [81]. In this paper the differences are examined for a three level minimal hurricane model and for a particular rapid vortex intensification event. In their study they pay particular attention to the Lorenz grid computational mode, discussed presently, and the errors that it can introduce into the vortex structure as the intensification matures.

Previous work on the dynamics, namely [64, 68], has shown that Charney-

Phillips is the preferred grid for the dynamics. This is due to the optimal representation of propagation of all kinds of waves and also due to no spurious modes existing. The Lorenz grid is preferred for the stably stratified boundary layer parametrisation due to the dependency on the Richardson number which can be found without averaging on the Lorenz grid. The investigation now has to be made into which grid is preferred when these two large and small-scale processes are coupled.

This question exists within a wider context of choice of staggering, the “density-moisture-potential temperature conundrum”. Optimal wave propagation requires density staggered relative to potential temperature, optimal moisture conservation requires moisture to be stored with density and physical parametrisations often require moisture to be stored with potential temperature.

Computational Mode

For the dynamics only situation the Lorenz grid gives fairly good wave propagation properties, although not as good as the Charney-Phillips grid. However the main downside is that the Lorenz grid supports an additional computational mode [33, 68, 69]. This is effectively a spurious solution to the equation, due to a property of the numerics, and is non-physical. In the Euler equations this is due to an averaging that is required on a thermodynamic variable in the vertical momentum conservation equation. A wave, whose wavelength is exactly twice that of the grid spacing, in this thermodynamic variable will be averaged to zero and become invisible to the model, this results in that solution having non-physical zero propagation. The presence of a computational mode can lead to a number of problems, for example manifesting itself as unwanted baroclinic instability [4] or giving non-physical inconsistency between discretised and continuous equations [58].

The computational mode is known to exist in the dynamics only case but how it behaves in a situation where the atmospheric boundary layer is coupled to the dynamics needs to be determined. It could be hypothesised that the strong diffusion

mechanism in the boundary layer will act to dampen the computational mode throughout the domain. If this were the case then it would be highly beneficial for the Lorenz grid. Further to the examination of the computational mode for the coupled boundary layer problem it will also be interesting to attempt to quantify how a resting state with shear and non-constant temperature gradient influences the computational mode.

1.4.3 Grid Spacing

The question of optimal grid staggering is connected also to the question of grid spacing. In the previous work investigating the influence of the spatial arrangement on the dynamics a uniform vertical grid is often used to simplify the mathematics, i.e. it can help analysis of the discretisation [68]. Using a uniform grid also helps in finding normal mode structures and dispersion relations since each normal mode will maintain its own scale throughout the domain. In the atmosphere using a uniform grid may be good for capturing the large-scale homogeneous dynamics but overall would not yield accurate simulations since computational limitations would demand a large grid spacing. If the Met Office arranged the vertical levels of their global model uniformly the spacing would be over a kilometre. Many of the small-scale features in the atmosphere occur due to the interaction with the ground, i.e. in the boundary layer. In order to capture these accurately it is beneficial to have a finer resolution here.

In order to achieve a grid with varying resolution some analytical function of spacing or height is derived, this can then be adapted depending on where the most model levels are required. An example of such is a geometric stretching. It should be noted that the most accurate results are likely to be found when variation of spacing is smooth and when all levels, staggered and un-staggered, vary together, i.e. not so that one level is positioned half-way between another [39].

The thesis will include discussion of the grid spacing, the effect of a smooth variation of the spacing against placing staggered levels halfway between unstaggered levels and how damaging the use of a uniform grid is for boundary layer

structures.

1.5 Governing Equations

The governing equations that are used to model the coupled boundary layer-dynamics are the Reynolds averaged Navier-Stokes equations. Equations for conservation of each of the three components of momentum, along with an equation for conservation of potential temperature and density are required. Through the nonlinear advection and viscosity terms the full Navier-Stokes equations include all the mechanisms for boundary layer flow. However it would not be sensible to model numerically in this way due to the small-scales involved, it would effectively require a direct numerical simulation.

Instead of trying to employ direct numerical simulation the Reynolds averaging is used, which reduces the model variables into a mean part plus fluctuations around that mean part. By considering these two parts of the system separately allows for the model to be somewhat simplified. Both parts of the system evolve temporally but the mean state will evolve considerably slower than the fluctuations. Over a sufficient temporal period the fluctuation part can be neglected since the mean of any fluctuation is approximately zero. Of course the approximation is only valid while the mean part of the system remains approximately constant, if too long a time period is considered the mean part would change. By assuming this correct temporal period has occurred allows one to reduce the model to its mean state. Due to the nonlinear interactions in the equations there will be terms which involve products of fluctuating terms, which will not be zero over the averaging period. Indeed without these nonlinear interactions it would not be possible to properly model turbulence, an inherently nonlinear process.

The terms that are obtained as a result of the Reynolds averaging are turbulent fluxes of that parameter, for each equation there will be three fluxes, representing turbulent flux in three dimensions. The turbulent fluxes describe how the quantity measured by that equation will be transported by turbulent eddies. This transport will be sub grid-scale and thus unknown in terms of the mean part variables so

requires a parametrisation. The advantage of using Reynolds averaging over modelling the Navier-Stokes equations in their natural form is that terms have been gained which can be written in terms of the mean flow through a closure, rather than needing to capture all relevant scales.

Generally assumptions can be made for the situation in order to minimise the number of sub grid processes that need to be parametrised. In the stably stratified boundary layer for example, it can be assumed that the vertical gradient of vertical turbulent flux will dominate the horizontal gradient of turbulent flux. This is due to the relative vertical and horizontal scales, in the stable boundary layer the depth will be shallow and so eddies will have larger horizontal scale than vertical scale, resulting in larger vertical gradients. It can also be assumed that transport by the turbulent eddies will dominate transport by molecular viscosity. A further assumption often made in the boundary layer is that the turbulent vertical transport of momentum and heat will dominate the transport by the mean vertical wind so that no equation for conservation of vertical momentum is required.

Two distinct equation sets will be considered: one for modelling the boundary layer on its own and one for coupling the dynamics to the boundary layer. Both equation sets can be derived using the Reynolds averaging technique. When studying the boundary layer on its own a number of filtering techniques will be used so that dynamical mechanisms are not captured, allowing for the study of the boundary layer mechanism on its own. This enables not just an examination of optimal grid configuration for the boundary layer but also to establish the most appropriate methodology for analysis. In addition to the complete nonlinear behaviour, the fully coupled system of equations can be thought of as supporting four distinct mechanisms: Rossby waves, inertio-gravity waves, acoustic waves and boundary layer diffusion, the former three being dynamical processes. It may also be hypothesised that further behaviour resulting from the coupling may be present. For the boundary layer only case the dynamical waves are filtered out by ignoring latitude dependence in Coriolis, having no horizontal dependence, no vertical wind and being in the Boussinesq framework. The equation set for the boundary

layer then represents a 1D column of the model. When examining the coupled problem a 1D column model can still be used whilst allowing for all the dynamical mechanisms by assuming disturbances with certain horizontal wavenumber.

The equations are discretised and solved numerically on a grid. A second order accurate finite differencing is used for numerical derivatives. It has been previously noted [45] that inaccuracies that occur due to choosing a particular vertical staggering may be overcome by choosing higher order finite difference schemes. However, considering their additional expense, moving to higher order schemes may not always improve model performance due to the relatively low order of continuity of the atmosphere [32]. Indeed many operational models employ second order schemes so the question of how the staggering options perform for second order is still relevant. In addition a higher order scheme will not alleviate computational mode problems.

1.6 Methodolgy

In order to examine the choice of vertical configuration for the coupled boundary layer dynamics problem an appropriate methodology will be required.

A popular method for comparing competing numerical configurations, and seen in a number of studies e.g. [2, 81], relies on examination of model output after equivalent simulations. For example in this case the full Reynolds averaged equations would be coded up and run with a given set of initial conditions for two competing vertical configurations. Runs would be performed at operational resolution and then after a given duration be compared with the output from a high resolution version of the run. Alternatively a realistic situation could be considered and compared to observations. Although these provide useful testing methodologies they can be somewhat ambiguous and would not necessarily provide much insight into exactly why a particular vertical configuration performed well. A more systematic methodology involves expanding the system into its steady and linear time dependent (transient) parts and considering the two separately. This allows for techniques such as modal analysis to be employed to examine the different waves and scales supported by the system. Examining the behaviour of different linear

structures can provide clear and concise insight into how different configurations can be expected to perform for a given problem.

1.6.1 Linearisation

In linearisation all of the model variables are written as a sum of the steady state part plus a transient part. The transient part represents a small amplitude perturbation away from the steady state. All terms in the equations are expanded in this way and the resulting formulation is split to create equations containing only steady parts and equations containing only first order transients. Transient terms are relatively small in comparison to the steady terms and so expanded terms which include products of two or more transient variables can be neglected. The steady state equations and the transient equations are then examined individually with the appropriate technique. Reducing the system in this way should allow one to pinpoint exactly where a certain vertical configuration performs well and derive from those results the implications for the overall model.

The steady state part of the system is independent of time and will have overall structure similar to that seen when examining the structure of the atmosphere. Generating the steady state can be done with a method such as false time-stepping or the Newton method.

The time dependent transient part of the system represents the linear evolution of the system and has wavelike solutions. When a system of equations is solved discretely there will be a set of wavelike structures that can be supported; any transient linear evolution can be projected onto the set. This set of waves are referred to as the normal modes of the solution [66, 67, 68, 64]. With the appropriate technique for the situation each mode in the set can be examined independently. Comparing the structure, frequency and growth or decay rate of each mode allows one to derive the implications for how the transient behaviour will be represented discretely. Studying normal modes allows for a highly systematic approach for the comparison of different vertical configurations. For this kind of approach to be most successful requires a suitably posed set of equations and this will need to be

taken into consideration in this thesis. It should be noted that this methodology does not allow for examination of the model’s full nonlinear response, nevertheless the linear wavelike response is important. Large-scale atmospheric flow is in hydrostatic and geostrophic balance; accurate representation of the structure of the acoustic and inertio-gravity waves, that propagate as the flow adjusts towards balance after a perturbation, is considered essential. Further, accurate representation of the frequency and structure of the energetically dominant Rossby waves is also considered essential.

In idealised cases the steady and transient parts of the system may be found analytically, then solutions from different configurations can be compared with the analytical solution. For more complicated systems this is not currently possible and the techniques for finding the steady and transient parts of the system need to be performed for a high resolution grid as well. When candidate configurations are compared at high resolution there should be indistinguishable difference between the configurations, otherwise differences at low resolution cannot be properly quantified.

1.6.2 Dispersion Relation

The dispersion relation for a set of linearised equations describes frequency against wavenumber for all the resolved waves. It is a very useful relationship to obtain since when it is compared with an analytical or a high resolution truth solution it will describe which waves will propagate most accurately. Comparing scales in this way can lead to some very clear insight when comparing different configurations. Figure 1.1 and Figure 1.2 demonstrate how studying the dispersion relation can be highly informative as to why a particular staggering works well. Figure 1.1 describes the error between the dispersion relation for the A- to D-grids and the analytical dispersion relation for the f -plane linearised shallow water equations. The plots are for Rossby radius of deformation chosen to be 10,954m and with horizontal grid spacing of 10,000m. The C-grid can be seen to give a good representation of the dispersion relation in comparison to the other grids, across the majority of

normalised horizontal wave numbers. This demonstrates why it is a popular choice of staggering in atmospheric modelling. Similarly the B-grid, which is also quite popular, can be seen to give a good representation of the propagation of the largest scale waves. In situations where the Rossby radius of deformation is reduced in comparison to the grid spacing the B-grid is considered more accurate e.g. [37]. For a complete study for this kind of problem see e.g. [26].




Figure 1.1: Dispersion error for the Arakawa A- to D-grids for the f -plane linearised shallow water equations.

Repeating exactly the work of [68] Figure 1.2 shows the dispersion relation for the discrete normal modes of the dynamics only Euler equations with isothermal reference state. The Charney-Phillips grid can be seen to capture the frequency of small-scale Rossby modes more accurately than the Lorenz grid. Note that the computational mode is not shown on the plot since it has zero frequency. The methodology used by [68] is that which will be employed here.




Figure 1.2: Dispersion relation for the discrete normal modes for the dynamics only isothermal equations, from top to bottom the branches represent acoustic, inertio-gravity and Rossby modes.

Frequency rates are important for neutral modes, i.e. modes that do not grow or decay in amplitude. For non neutral modes the rate of growth or decay is also important. For a simple enough reference state the large-scale Rossby, inertio-gravity and acoustic dynamical modes are neutral modes; important information is the speed at which the waves will propagate, given by the dispersion relation. For the modes which represent the influence of the boundary layer diffusion there will be decay and the usefulness of a vertical configuration needs to be determined also from the rate of that decay. If a configuration fails to damp particular modes fast enough then it could cause problems for the model as that mode will continue to exist incorrectly in the solution; the problem could then be exacerbated if the structure of that mode was also captured inaccurately.

An ideal vertical configuration for the coupled boundary layer and dynamics model would capture the decay rate of boundary layer modes most accurately, represent the frequencies of the neutral dynamical modes most accurately, capture all mode structure most accurately, have no computational mode and place all model variables at levels most suited to meeting any conservation properties. Unfortunately no such configuration is currently known. The question now is how the best

configurations for some of these things, Lorenz and Charney-Phillips, can perform when all the features are combined.

For sufficiently simple situations, such as the f -plane shallow water equations and the isothermal resting state Euler equations, properties such as dispersion relation can be obtained analytically. Further to this it is also possible to derive analytical descriptions of the discrete versions. For the more complicated situation of the coupled dynamics-boundary layer, information about the mode frequency, damping and structure needs to be obtained numerically and for this a method such as eigendecomposition is required.

1.6.3 Eigendecomposition

Firstly the entire linearised transient equation set can be written in matrix form,

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}. \quad (1.1)$$

The vector \mathbf{x} contains all the transient model parameters, the matrix \mathbf{A} contains the coefficients of the transient parameters, including steady state variables. Seeking horizontally wavelike temporal solutions of the form $\mathbf{x} \propto \exp(ikx + ily + \lambda t)$ allows (1.1) to be written in eigenvalue form,

$$\lambda\mathbf{x} = \mathbf{A}\mathbf{x}. \quad (1.2)$$

Eigenvectors \mathbf{x} represent the complex mode structures and complex eigenvalues $\lambda = \mu - i\omega$ represent the frequency ω and/or growth or decay rate μ of the corresponding mode. The decomposition $\mathbf{A} = \mathbf{X}\mathbf{\Lambda}\mathbf{X}^{-1}$ is used to find the solutions. Eigenvectors are the columns of \mathbf{X} and eigenvalues are the diagonal elements of $\mathbf{\Lambda}$. The imaginary part of the eigenvalue represents the frequency of the corresponding mode; if k is positive then positive imaginary part denotes eastward propagation and negative imaginary part denotes westward propagation. If the real part is positive the corresponding mode will grow, if negative it will decay. Neutral dynamical modes will have purely imaginary corresponding eigenvalue. Boundary layer modes have

real part dominating imaginary part, representing the fact that the diffusion of those structures dominates their frequency. Note that in the presence of a complex reference state they will not necessarily be purely real due to the propagation induced by shear and stratification.

Normal Matrices

When the operator, or system of equations, is normal the matrix \mathbf{A} in which it is expressed will also be normal. A normal matrix has the property $\mathbf{A}\mathbf{A}^* = \mathbf{A}^*\mathbf{A}$, where \mathbf{A}^* is the conjugate transpose of \mathbf{A} . When this is the case the eigenvectors form an orthogonal basis and represent the physical attributes of the system through normal modes. In many fluid dynamical applications, particularly those that give rise to turbulence, the operator is non normal. When the matrix is far from normal its eigenvectors may be far from orthogonal, meaning it becomes difficult to obtain any physical interpretation of the system using eigenvalues and eigenvectors. If this is the case it may not necessarily be wise to examine eigenvalues and eigenvectors [71]. Since this implies solutions will be non normal when the matrix is far from normal it makes no sense to even describe the transient solutions as “normal modes”, instead they are referred to as eigenmodes.

In order to recover the systematic methodology, that makes the eigendecomposition so popular, other decompositions that return a normal basis may be useful.

1.6.4 Singular Value Decomposition (SVD)

Perhaps the most promising decomposition for recovering some of the systematic methodology afforded by examining normal modes is the Singular Value Decomposition (SVD). Rather than decompose the matrix into eigenvectors and eigenvalues it decomposes into input (right) singular vectors, output (left) singular vectors and singular values, the decomposition is given by $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$. Despite the structure of \mathbf{A} the decomposition always returns singular vectors forming an orthonormal basis, with corresponding real and positive singular values. It is effectively a generalisation of the eigendecomposition that is more suitable for non-normal or non-square

operators. Although a generalisation of the eigendecomposition, the two are nevertheless related and there are two important links to note. Firstly when the matrix \mathbf{A} is normal, Hermitian and sign-definite the eigendecomposition is equivalent to the SVD, that is to say that input singular vectors are identical to the output singular vectors and are equal to the eigenvectors and that singular values are equal to eigenvalues. For a Hermitian matrix $\mathbf{A} = \mathbf{A}^*$; if \mathbf{A} is positive definite then $\mathbf{z}^* \mathbf{A} \mathbf{z} > 0$ for any non-zero complex vectors \mathbf{z} , if \mathbf{A} is negative definite then $\mathbf{z}^* \mathbf{A} \mathbf{z} < 0$. A sign-definite matrix is Hermitian and a Hermitian matrix is normal. Secondly the SVD of any matrix \mathbf{A} is equal to the eigendecomposition of $\mathbf{A} \mathbf{A}^*$ and $\mathbf{A}^* \mathbf{A}$.

Singular vectors were first used in meteorology by [49], where they were used as a tool for calculating how the error in a model initialisation grows throughout the forecast. For the last two decades or so, as computing resources have become sufficient for calculating the decomposition, the use of singular vectors and singular values has been seen in a large number of weather and climate model analysis applications from ensemble initialisation e.g. [52] and error growth estimation e.g. [53] to non-normal stability analysis e.g. [57].

From the various applications it can be seen that the singular vectors and singular values can play an important role in describing the behaviour of \mathbf{A} . For example the largest singular value has a corresponding singular vector that would be perturbed the most when multiplied by the matrix of coefficients \mathbf{A} [43]. This provides us with some mathematical basis upon which the method can be used to compare different configurations when a non-normal matrix is obtained.

An important consideration when applying the SVD is the norm in which to work, also discussed in [43]. Unlike the eigenvalues the singular values will be sensitive to the norm. It can be shown that a system which conserves energy, such as the isothermal resting state dynamics only case, will be normal and Hermitian when written in the energy norm. Since it is understood that the SVD is equivalent to the eigendecomposition when the system is normal, Hermitian and sign-definite it would make most sense to work in the energy norm when calculating singular

values and singular vectors. Although the system will dissipate energy through the boundary layer diffusion, working in the energy norm will still likely give the closest match between the two decompositions. If the total energy in the system can be given by the matrix multiplication $E = \mathbf{x}^T \mathbf{E} \mathbf{x}$ then working in the energy norm transforms \mathbf{x} so that $\mathbf{E} \equiv \mathbf{I}$. If working in the energy norm, singular vectors should provide the structures which when multiplied by \mathbf{A} give the greatest perturbation in terms of energy and therefore those that are likely to be of most physical relevance.

Despite singular values being positive and real, the decomposition would still be valid for a situation of purely neutral modes. Rather than the largest singular values corresponding to the singular vector that would undergo the largest growth or decay it would correspond to the singular vector that would undergo the fastest oscillation. From another perspective one could just as well construct the problem so that eigenvalues were all real so that the eigenvalue with largest real part corresponded to the eigenvector undergoing the fastest oscillation. In the case where some neutral and some decaying modes are present the ability of the singular values to describe the physical behaviour needs to be determined. They will likely pick the most dominant from either case since they are incapable of distinguishing between real and imaginary part. In most situations the behaviour should be recoverable from the singular vectors.

It will require careful numerical testing to be sure of the extent to which singular values and vectors can be used to understand the properties of \mathbf{A} in the coupled boundary layer dynamics case. It is clear that under the right conditions the eigendecomposition will produce results with well understood physical interpretation. It is less clear however, under what conditions the SVD will be able to produce results from which a clear physical interpretation, and thus comparison of grid configurations, is possible.

1.6.5 Discrete and Continuous Spectra

A further issue to consider when seeking wavelike solutions for fluid flow problems which exhibit non-trivial reference state and complex viscous terms is the form of

the spectrum. The spectrum here represents the range of scales supported by the continuous equations, or the possible wavelengths of the wavelike solutions. Certain problems, such as that considered by [68], clearly have discrete spectra; solutions are shown to be sinusoidal so only waves with certain wavelengths can exist whilst satisfying the boundary conditions on the bounded domain. It has been shown that for certain boundary layer type flows that have an unbounded domain, such as in the viscous Orr-Sommerfeld equation, the spectrum consists of both discrete and continuous parts. Due to the way the boundary conditions are implemented it is possible for solutions with any given wavelength to exist, meaning a continuous set of solutions. Further to the boundary layer flow in an unbounded domain it is also possible to have inviscid flow in a bounded domain that exhibits continuous spectrum [19, and references therein]. In the case of an unbounded domain with a boundary layer flow, solutions with any given wavenumber may exist while still meeting the boundary condition at infinity. If solutions of any wavelength may exist it will result in a part of the spectrum with a continuous part to it. Unlike the bounded domain, resting isothermal reference state case, in which the boundary conditions demand a discrete spectrum.

The type of flow considered for the purpose of this thesis will be a boundary layer flow but with boundary conditions imposed on a finite domain; in this situation it is not clear whether a discrete only spectrum should be expected. It may be that the damping mechanism in the boundary layer can absorb dynamical waves of any length, which could result in a continuous part in the spectrum of solutions.

When only a discrete spectrum of solutions exists it is clear which solutions will be supported when the equations are discretised. When a continuous spectrum of solutions exists the discretisation of the equations will result in some discrete sampling of the continuous spectrum. In this case it is not clear that different configurations, for example when using the Lorenz or Charney-Phillips grid, will sample the continuous spectrum in the same way. If they do not sample the continuous spectrum in an equivalent manner then it would not be possible to compare discrete modes since they would all be different.

If solutions may take any wavelength and the spectrum thus contains a continuous part then it is possible to numerically search for it, for example by generating repeat solutions with increasing resolution. This would not be sufficient to prove that solutions can only take certain wavelengths and that only a discrete part of the spectrum exists. Use of the methodology of comparing low and high resolution eigenvectors will require solutions sampled from a discrete solution; this will need to be examined as part of the study.

1.7 Thesis Outline

The aim of this thesis is to investigate the spatial aspects of coupling the atmospheric boundary layer and large-scale dynamics. Questions of particular interest are:

- Does either the Lorenz grid or Charney-Phillips grid offer an overall superior configuration for accurately capturing the types of processes that can occur?
- What differences occur as a result of different ways of stretching the grid?
- How does the computational mode behave in the presence of the boundary layer mechanism?

Answering these questions in a satisfactory manner will require a methodology that allows for a systematic study of the transient behaviour in the system. The classic approach is normal mode analysis. However this is only suitable for the examination of normal systems, i.e. those whose eigenvectors of the linearised system are orthogonal. Two further points can be added in the questions that need to be addressed:

- Just how useful is the normal mode methodology?
- If normal mode analysis is found to be unusable then how viable is other possible methodology, i.e. singular values?

The project is divided into tackling the boundary layer on its own, from which a number of non trivial questions are raised, and tackling the fully coupled problem. As a result the thesis is divided into two distinct parts with chapters therein.

Part I will cover the boundary layer only problem. Chapter 2 will discuss the model used for describing the boundary layer, including the parametrisation of boundary layer turbulence. Adjustments are required to the usual form of the boundary layer equations in order to suit the methodology; this is also discussed here. Chapter 2 will also outline the fundamental differences between the types of staggering, the possible stretching of the grid spacing and the various options available for averaging the variables. The solutions of the equations are then dealt with in Chapter 3 and Chapter 4. Firstly the steady state solutions are examined in Chapter 3. Here the ability of each configuration to accurately capture the steady state structure of the solution is discussed. Chapter 4 looks at the transient part of the calculation. This includes a detailed examination of the ability of each configuration to capture the physical structure that the system supports and of the ability of the methodology to assist in the comparison. In the final chapter of Part I some conclusions for the boundary layer only case are discussed.

Part II of the thesis examines the fully coupled problem. Chapter 6 discusses the equation set required to model the full problem, so is in essence an extension of the ideas set out in Chapter 2. Chapters 7 and 8 contain the main results of the thesis, namely the answers to the questions that were posed in this introduction. Firstly a discussion is given of the dynamics only case so that the results can be extended to include a stretched grid and to test the methodology that was introduced for the boundary layer only case. A steady state which is based on the full coupled equations is generated and compared for the two types of grid staggering. Various cases are constructed to demonstrate the differences between the vertical configurations and understand physically the influence of the boundary layer on the dynamical structures and vice-versa. Chapter 9 draws together the results and offers some conclusions and motivation for further work investigating the spatial aspects of the physics-dynamics coupling problem.

Part I

Vertical Discretisations for the Stably Stratified Planetary Boundary Layer

In this first part of the thesis the simpler case where the dynamics are ignored so that system solves the boundary layer alone is studied. It is already well understood that Charney-Phillips is preferable when solving the dynamics without the boundary layer. Here an investigation is presented into whether the Lorenz grid performs better than Charney-Phillips when only the boundary layer is solved. It will be argued that the form of the boundary layer equations suggests that the Lorenz grid should be preferred, this also needs to be tested numerically to confirm. The methodology, as outlined in the introduction, will be constructed in order to test whether the Lorenz grid is indeed the favoured configuration. The equations are linearised into steady and transient parts. Once the equations are suitably posed for finding the steady state profiles the two different configurations, Lorenz and Charney-Phillips, are compared. Once results are established for the steady state part of the equations the transient part will also be examined. It is in this area of the comparison that the methodology needs to be carefully tested and examined. With a suitable methodology the ability of the Lorenz and Charney-Phillips configurations to capture the linear time evolution can be compared.

Part I is arranged as follows. Chapter 2 describes the model that will be used, Chapter 3 describes the steady state and how it is obtained and Chapter 4 discusses the transient part of the solution. Some concluding remarks are offered in Chapter 5.

Chapter 2

The Boundary Layer Model

This chapter aims to outline the model that will be used to examine the boundary layer. Firstly showing how the Navier-Stokes equations are adapted into a form suitable to modelling the boundary layer structure and then discussing the discretisation techniques that will be used.

A useful approximation for modelling the Navier-Stokes equations is the Reynolds Averaging technique. The idea behind the approximation is to consider the continuous equations as a their mean part plus fluctuations around the mean part. The equations can then be time-averaged, for a sufficient length of time the average of the fluctuating part is just equal to the mean part. The result of the averaging is that only the mean part and the products of two fluctuating terms remain. These products of fluctuating variables are called the Reynolds stresses; these are functions of fluctuating velocity and represent the effect of turbulent transport within the flow. These Reynolds stresses need to be approximated by a turbulence model, also known as a closure. The beauty of the approximation is that with the right closure these Reynolds stresses give a good representation of the characteristics of a boundary layer type flow. Rather than modelling all the small-scale viscous interactions in the full Navier-Stokes equations, a much more viable simulation, where turbulence is calculated from the resolved mean flow is obtained.

In order to be able to computationally model the boundary layer the continuous equations are solved discretely on a grid. A typical operational numerical weather

prediction model will have a vertical grid spacing increasing with height. The smallest spacing, between the ground and lowest model level, may be around 10m. The turbulent processes in the boundary layer can be sub-grid, in that the scale of motion may be smaller than the spacing between the points on the discrete grid. Without the Reynolds averaging and closure the sub-grid scale motion would not be captured; with the closure the sub-grid scales are approximated based on the large-scale mean features that are stored on the grid. This chapter goes on to demonstrate the type of closure that is commonly used for boundary layer modelling, from which it can be seen why the Lorenz grid is expected to be the preferred configuration.

Beyond the derivation of the equations the idea of using subsidence to generate a steady state is introduced, the set of test cases that will be used are listed and the types of grid spacing that will be required are introduced and discussed.

2.1 Governing Equations

The fully compressible 3D Navier-Stokes equations of fluid motion, in a rotating frame and in Cartesian coordinates, are

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} + w \frac{\partial u}{\partial z} - fv = -\frac{1}{\rho} \frac{\partial p}{\partial x} + \nabla \cdot (\nu \nabla u), \quad (2.1)$$

$$\frac{\partial v}{\partial t} + u \frac{\partial v}{\partial x} + v \frac{\partial v}{\partial y} + w \frac{\partial v}{\partial z} + fu = -\frac{1}{\rho} \frac{\partial p}{\partial y} + \nabla \cdot (\nu \nabla v), \quad (2.2)$$

$$\frac{\partial w}{\partial t} + u \frac{\partial w}{\partial x} + v \frac{\partial w}{\partial y} + w \frac{\partial w}{\partial z} = -\frac{1}{\rho} \frac{\partial p}{\partial z} + \nabla \cdot (\nu \nabla w) - g, \quad (2.3)$$

$$\frac{\partial \theta}{\partial t} + u \frac{\partial \theta}{\partial x} + v \frac{\partial \theta}{\partial y} + w \frac{\partial \theta}{\partial z} = \nabla \cdot (\nu \nabla \theta), \quad (2.4)$$

$$\frac{\partial \rho}{\partial t} + u \frac{\partial \rho}{\partial x} + v \frac{\partial \rho}{\partial y} + w \frac{\partial \rho}{\partial z} + \rho \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} \right) = 0. \quad (2.5)$$

Equations (2.1) to (2.4) represent the evolution of the horizontal (u, v) and vertical (w) components of momentum and potential temperature. Potential temperature θ is related to the other variables through the viscosity, as will become clear. Equation (2.5) represents the mass continuity equation. ρ is density and p is pressure. The coordinate system is aligned such that x represents the zonal direction, y the

meridional direction and z the vertical direction. ν is viscosity, $g = 9.80616\text{ms}^{-2}$ is gravity and $f = 1.031 \times 10^{-4}\text{s}^{-1}$ is the Coriolis parameter.

2.1.1 Reynolds Averaging

When modelling turbulence it is not necessarily feasible or useful to try to model all the scales supported within the flow. A common approach to modelling the turbulence in the atmospheric boundary layer is by using the Reynolds averaged Navier-Stokes equations. This filtering technique used in order to obtain an ‘average’ picture of the state of the flow along with a turbulent unknown part. Many of the authors that have written on the boundary layer equations tend also to work within an incompressible Boussinesq framework for the boundary layer, e.g. [28, 34, 62]. It is generally sufficient to restrict attention to incompressible flow since variation in density is relatively small across the lowest portion of the atmosphere. When deriving for an incompressible flow the Reynolds averaged equations can be found by writing all model variables as their mean plus fluctuating parts, i.e. $u = \bar{u} + u'$. Terms can then be time-averaged to gain an insight to the overall flow property. To use the technique it is assumed that for a certain period of time the fluctuating parts average to the mean flow i.e. $\overline{\bar{u} + u'} = \bar{u}$, the time must be long enough so that fluctuations average to the mean flow but not so long that the mean flow itself could be evolving. The property of the averaging is thus $\overline{u'} = 0$ and thus that $u \equiv \bar{u}$. Note however that products of fluctuating variables do not go to zero, indeed these are required for modelling turbulence.

In part II of this thesis the non-Boussinesq fully compressible equations will be required. This will allow for acoustic waves to be supported by the system, a desirable feature for testing methodology and for physical interpretation of the system. Since acoustic waves propagate very quickly they will always stand out in the spectrum. When obtaining the Reynolds averaged Navier-Stokes equations for a fully compressible flow it is not sufficient to write all model variables just as mean plus fluctuating part as this leads to terms describing fluxes of densities, preventing one from performing the necessary cancelling of terms. Instead all terms

excluding density and pressure are written as mass weighted averages, density and pressure can remain as un-weighted mean plus fluctuating part. The zonal velocity for example is now written $u = \tilde{u} + u''$, with the new mass weighted property, i.e.,

$$\tilde{u} = \frac{\overline{\rho u}}{\bar{\rho}}. \quad (2.6)$$

The most important difference to note is that when the quantity u is averaged the $\overline{u''}$ terms do not equal zero. Instead the mass weighted version will, giving that

$$\overline{\rho u''} = 0 \quad . \quad (2.7)$$

The technique for obtaining the equation set, as described also by [63], is shown here for the zonal momentum equation (2.1) and the equation for conservation of mass (2.5). First of all, terms in (2.5) are combined by the product rule then expanded into mean and fluctuating parts and the mass weighted counterparts,

$$\frac{\partial}{\partial t} (\rho + \rho') + \frac{\partial}{\partial x} (\rho + \rho') (\tilde{u} + u'') + \frac{\partial}{\partial y} (\rho + \rho') (\tilde{v} + v'') + \frac{\partial}{\partial z} (\rho + \rho') (\tilde{w} + w'') = 0. \quad (2.8)$$

Now the terms are multiplied out with the averaging applied,

$$\begin{aligned} \frac{\partial}{\partial t} (\bar{\rho} + \bar{\rho}') + \frac{\partial}{\partial x} (\bar{\rho}\tilde{u} + \bar{\rho}'\tilde{u} + \overline{\rho u''} + \overline{\rho' u''}) + \frac{\partial}{\partial y} (\bar{\rho}\tilde{v} + \bar{\rho}'\tilde{v} + \overline{\rho v''} + \overline{\rho' v''}) + \\ \frac{\partial}{\partial z} (\bar{\rho}\tilde{w} + \bar{\rho}'\tilde{w} + \overline{\rho w''} + \overline{\rho' w''}) = 0. \end{aligned} \quad (2.9)$$

Terms such as $\bar{\rho}'\tilde{u}$ equal zero since $\bar{\rho}' = 0$ by definition of the time-averaging. Terms such as $\overline{\rho u''} + \overline{\rho' u''}$ can be rewritten as $(\bar{\rho} + \bar{\rho}') \overline{u''} = \overline{\rho u''}$ which is zero by definition of the mass weighted time-averaging (2.7). The equation for the conservation of “weighted” mass is now obtained,

$$\frac{\partial \bar{\rho}}{\partial t} + \frac{\partial \bar{\rho}\tilde{u}}{\partial x} + \frac{\partial \bar{\rho}\tilde{v}}{\partial y} + \frac{\partial \bar{\rho}\tilde{w}}{\partial z} = 0. \quad (2.10)$$

Note that this shows the general form of the derivation, however the process is

shortened if equation (2.5) is written in flux form, then applying equation (2.6) leads directly to equation (2.10).

Now a similar process is required for the conservation of zonal momentum. As suggested above and following [63] it simplifies the derivation once the momentum equations are written in conservative form,

$$\frac{\partial \rho u}{\partial t} + \frac{\partial}{\partial x} (\rho u^2) + \frac{\partial}{\partial y} (\rho uv) + \frac{\partial}{\partial z} (\rho uw) - f \rho v = -\frac{\partial p}{\partial x} + T_x \quad (2.11)$$

where $T_x = \nabla \cdot (\mu \nabla u)$ and $\mu = \rho \nu$, this is the viscosity term. A good approximation for the boundary layer, as will be considered here, is the high Reynolds approximation, which says that the transfer of momentum and heat as well as any other quantities will be dominated by turbulent transport over molecular viscosity. This is due to the large scales involved with the vertical turbulent flux which may be of the order of 100m compared with a few centimetres for the molecular viscosity. Due to the complexity in deriving T_x for the fully compressible case it is excluded here, refer to [63] if requiring the full derivation.

As for the mass conservation equation each term is written as mean plus fluctuating or mass weighted equivalents,

$$\begin{aligned} \frac{\partial}{\partial t} (\bar{\rho} + \rho') (\tilde{u} + u'') + \frac{\partial}{\partial x} (\bar{\rho} + \rho') (\tilde{u} + u'') (\tilde{u} + u'') + \frac{\partial}{\partial y} (\bar{\rho} + \rho') (\tilde{u} + u'') (\tilde{v} + v'') \\ + \frac{\partial}{\partial z} (\bar{\rho} + \rho') (\tilde{u} + u'') (\tilde{w} + w'') - f (\bar{\rho} + \rho') (\tilde{v} + v'') = -\frac{\partial (\bar{p} + p')}{\partial x}. \end{aligned} \quad (2.12)$$

Equation (2.12) is now time-averaged and expanded,

$$\begin{aligned} \frac{\partial}{\partial t} (\bar{\rho} \tilde{u} + \bar{\rho}' \tilde{u} + \overline{\rho u''} + \overline{\rho' u''}) + \frac{\partial}{\partial x} (\bar{\rho} \tilde{u} \tilde{u} + 2\overline{\rho \tilde{u} u''} + \overline{\rho u'' u''} + \bar{\rho}' \tilde{u} \tilde{u} + 2\overline{\rho' \tilde{u} u''} + \overline{\rho' u'' u''}) \\ + \frac{\partial}{\partial y} (\bar{\rho} \tilde{u} \tilde{v} + \overline{\rho \tilde{u} v''} + \overline{\rho \tilde{v} u''} + \overline{\rho u'' v''} + \bar{\rho}' \tilde{u} \tilde{v} + \overline{\rho' \tilde{u} v''} + \overline{\rho' \tilde{v} u''} + \overline{\rho' u'' v''}) \\ + \frac{\partial}{\partial z} (\bar{\rho} \tilde{u} \tilde{w} + \overline{\rho \tilde{u} w''} + \overline{\rho \tilde{v} w''} + \overline{\rho u'' w''} + \bar{\rho}' \tilde{u} \tilde{w} + \overline{\rho' \tilde{u} w''} + \overline{\rho' \tilde{v} u''} + \overline{\rho' u'' w''}) \\ - f (\bar{\rho} \tilde{v} + \bar{\rho}' \tilde{v} + \overline{\rho v''} + \overline{\rho' v''}) = -\frac{\partial (\bar{p} + \bar{p}')}{\partial x}. \end{aligned} \quad (2.13)$$

Again, terms linear in fluctuations are immediately zero by definition of the time-averaging, again combining terms allows use of the property of the weighted averaging, i.e. $\overline{\rho\tilde{u}u''} + \overline{\rho'\tilde{u}u''} = \tilde{u}\overline{\rho u''} = 0$. Rewriting with the removal of all the terms that cancel gives,

$$\begin{aligned} \frac{\partial}{\partial t} (\overline{\rho\tilde{u}}) + \frac{\partial}{\partial x} (\overline{\rho\tilde{u}\tilde{u}}) + \frac{\partial}{\partial y} (\overline{\rho\tilde{u}\tilde{v}}) + \frac{\partial}{\partial z} (\overline{\rho\tilde{u}\tilde{w}}) - f\overline{\rho\tilde{v}} + \frac{\partial\overline{p}}{\partial x} \\ = -\frac{\partial}{\partial x} (\overline{\rho u''u''}) - \frac{\partial}{\partial y} (\overline{\rho u''v''}) - \frac{\partial}{\partial z} (\overline{\rho u''w''}). \end{aligned} \quad (2.14)$$

Note that derivatives expand as, for example,

$$\begin{aligned} \frac{\partial}{\partial y} (\overline{\rho\tilde{u}\tilde{v}}) &= \overline{\rho\tilde{u}} \frac{\partial\tilde{v}}{\partial y} + \overline{\rho\tilde{v}} \frac{\partial\tilde{u}}{\partial y} + \tilde{u}\tilde{v} \frac{\partial\overline{\rho}}{\partial y}, \\ &= \overline{\rho\tilde{v}} \frac{\partial\tilde{u}}{\partial y} + \tilde{u} \frac{\partial}{\partial y} (\overline{\rho\tilde{v}}). \end{aligned} \quad (2.15)$$

Applying expansion (2.15) to all of the spatial derivatives on the left hand side of (2.14) and then for the derivatives in the other equivalent momentum equations the cancelling becomes clear. Every term that appears in the conservation of mass equation also appears in every conservation of momentum equation, terms like e.g. $\frac{\partial}{\partial y} (\overline{\rho\tilde{v}})$. Since the sum of the terms in the mass conservation equation is zero the conservation of momentum equations reduce to,

$$\begin{aligned} \frac{\partial\tilde{u}}{\partial t} + \tilde{u} \frac{\partial\tilde{u}}{\partial x} + \tilde{v} \frac{\partial\tilde{u}}{\partial y} + \tilde{w} \frac{\partial\tilde{u}}{\partial z} - f\tilde{v} + \frac{1}{\overline{\rho}} \frac{\partial\overline{p}}{\partial x} \\ = \frac{1}{\overline{\rho}} \left[-\frac{\partial}{\partial x} (\overline{\rho u''u''}) - \frac{\partial}{\partial y} (\overline{\rho u''v''}) - \frac{\partial}{\partial z} (\overline{\rho u''w''}) \right], \end{aligned} \quad (2.16)$$

$$\begin{aligned} \frac{\partial\tilde{v}}{\partial t} + \tilde{u} \frac{\partial\tilde{v}}{\partial x} + \tilde{v} \frac{\partial\tilde{v}}{\partial y} + \tilde{w} \frac{\partial\tilde{v}}{\partial z} + f\tilde{u} + \frac{1}{\overline{\rho}} \frac{\partial\overline{p}}{\partial y} \\ = \frac{1}{\overline{\rho}} \left[-\frac{\partial}{\partial x} (\overline{\rho v''u''}) - \frac{\partial}{\partial y} (\overline{\rho v''v''}) - \frac{\partial}{\partial z} (\overline{\rho v''w''}) \right], \end{aligned} \quad (2.17)$$

$$\frac{\partial\tilde{w}}{\partial t} + \tilde{u} \frac{\partial\tilde{w}}{\partial x} + \tilde{v} \frac{\partial\tilde{w}}{\partial y} + \tilde{w} \frac{\partial\tilde{w}}{\partial z} + \frac{1}{\overline{\rho}} \frac{\partial\overline{p}}{\partial z}$$

$$= \frac{1}{\bar{\rho}} \left[-\frac{\partial}{\partial x} (\overline{\rho w'' u''}) - \frac{\partial}{\partial y} (\overline{\rho w'' v''}) - \frac{\partial}{\partial z} (\overline{\rho w'' w''}) \right] - g, \quad (2.18)$$

$$\frac{\partial \tilde{\theta}}{\partial t} + \tilde{u} \frac{\partial \tilde{\theta}}{\partial x} + \tilde{v} \frac{\partial \tilde{\theta}}{\partial y} + \tilde{w} \frac{\partial \tilde{\theta}}{\partial z} = \frac{1}{\bar{\rho}} \left[-\frac{\partial}{\partial x} (\overline{\rho \theta'' u''}) - \frac{\partial}{\partial y} (\overline{\rho \theta'' v''}) - \frac{\partial}{\partial z} (\overline{\rho \theta'' w''}) \right], \quad (2.19)$$

$$\frac{\partial \bar{\rho}}{\partial t} + \frac{\partial \bar{\rho} \tilde{u}}{\partial x} + \frac{\partial \bar{\rho} \tilde{v}}{\partial y} + \frac{\partial \bar{\rho} \tilde{w}}{\partial z} = 0. \quad (2.20)$$

Equations (2.16) - (2.20) represent the fully compressible Reynolds averaged Navier-Stokes equations. All the terms on the left hand side of these equations depend on the large-scale time-averaged mean part of the flow and can be considered 'known'. Terms on the right hand side include products of fluctuating variables and are 'unknown'. These unknown terms are crucial in that they represent the turbulent part of the flow, however some can be ignored due to the type of flow that is being examined. For the purposes of this thesis the fluxes need only to represent boundary layer turbulence.

This thesis will focus on the stably stratified planetary boundary layer. For the stably stratified boundary layer attention can be restricted to the case of horizontally homogeneous turbulent flux, in regions which exhibit fairly uniform terrain this is a reasonable approximation [34]. Further to this reasoning a scale argument can show this to be a good approximation, for the stably stratified boundary layer eddies will have horizontal scale of the order of a few kilometres but the shallow depth of the boundary layer will limit their vertical scale to less than a kilometre. Horizontal gradients of turbulent flux will be considerably less than vertical gradients of turbulent flux. Eddies in a convective boundary layer will stretch out vertically, eddies in a stably stratified boundary layer will stretch out horizontally. Homogeneity in gradients of horizontal turbulent flux is represented mathematically as,

$$\frac{\partial}{\partial x} (\overline{\rho u'' u''}) \ll \frac{\partial}{\partial z} (\overline{\rho u'' w''}), \quad (2.21)$$

$$\frac{\partial}{\partial y} (\overline{\rho u'' v''}) \ll \frac{\partial}{\partial z} (\overline{\rho u'' w''}). \quad (2.22)$$

The approximation made earlier that transport by viscosity is dominated by turbulent transport is represented formally as,

$$\overline{T_x} \ll \frac{\partial}{\partial z} (\overline{\rho u'' w''}). \quad (2.23)$$

One further approximation suitable for the boundary layer, and indeed much of the atmosphere, is that w is small compared with u and v so the vertical transport of vertical velocity is negligible.

$$\frac{\partial}{\partial z} (\overline{\rho w'' w''}) \ll g \quad \text{and} \quad \frac{1}{\bar{\rho}} \frac{\partial p}{\partial z}. \quad (2.24)$$

Applying these approximations and dropping the overbar notation for terms on the left hand side, since these are now the predicted variables, gives the full system of equations. The remaining flux gradients are required for modelling boundary layer turbulence.

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} + w \frac{\partial u}{\partial z} - f v + \frac{1}{\rho} \frac{\partial p}{\partial x} = -\frac{1}{\rho} \frac{\partial}{\partial z} (\overline{\rho u'' w''}), \quad (2.25)$$

$$\frac{\partial v}{\partial t} + u \frac{\partial v}{\partial x} + v \frac{\partial v}{\partial y} + w \frac{\partial v}{\partial z} + f u + \frac{1}{\rho} \frac{\partial p}{\partial y} = -\frac{1}{\rho} \frac{\partial}{\partial z} (\overline{\rho v'' w''}), \quad (2.26)$$

$$\frac{\partial w}{\partial t} + u \frac{\partial w}{\partial x} + v \frac{\partial w}{\partial y} + w \frac{\partial w}{\partial z} + \frac{1}{\rho} \frac{\partial p}{\partial z} = -g, \quad (2.27)$$

$$\frac{\partial \theta}{\partial t} + u \frac{\partial \theta}{\partial x} + v \frac{\partial \theta}{\partial y} + w \frac{\partial \theta}{\partial z} = -\frac{1}{\rho} \frac{\partial}{\partial z} (\overline{\rho \theta'' w''}), \quad (2.28)$$

$$\frac{\partial \rho}{\partial t} + \frac{\partial \rho u}{\partial x} + \frac{\partial \rho v}{\partial y} + \frac{\partial \rho w}{\partial z} = 0. \quad (2.29)$$

In part II of this thesis equations in the form (2.25)-(2.29) will be used since they are capable of supporting both the stably stratified boundary layer and the dynamics, indeed these equations represent the fully coupled problem. For the boundary layer only case the equations can be filtered further so that the boundary layer mechanisms can be examined on their own, this can be achieved in part by using the Boussinesq approximation, as outlined in the next section.

2.1.2 Boussinesq Approximation

A useful approximation, sufficient when applying equations (2.25) - (2.29) to the boundary layer is the Boussinesq approximation. The Boussinesq approximation assumes that the thermodynamic variables p and ρ can be split into hydrostatically balanced vertical reference profiles and finite amplitude perturbations away from the reference profiles, for example $\rho = \bar{\rho}(z) + \rho'$. It needs to be noted that in the literature there is a notation conflict between Reynolds averaging and Boussinesq approximation. Following [23] the overline and ' notation are used in this section to refer to reference and perturbation profiles. However in terms like $\overline{u''w''}$ the ' still refers to terms that result from Reynolds averaging. To free up notation, and avoid ambiguity, use of ' is dropped from density and pressure perturbation terms in the equations. In a situation such as the boundary layer it can be assumed that the reference profile has only small variation with height and so can be replaced by a constant reference density ρ_0 ; this is the fundamental property of the Boussinesq approximation. Note that the perturbation quantity still represents a perturbation away from the height dependent reference profile, i.e. $\rho' = \rho - \bar{\rho}(z)$. Now since density perturbations are small relative to the constant reference state ρ_0 density can be replaced by its reference state. However, since the flow characteristics are governed by buoyancy perturbations, density perturbation should be retained when it occurs in buoyancy terms. A further influence of the Boussinesq approximation is that the flow can now be considered incompressible, i.e. $\nabla \cdot \mathbf{u} = 0$. With the Boussinesq approximations equations (2.25)-(2.29) become,

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} + w \frac{\partial u}{\partial z} - fv + \frac{1}{\rho_0} \frac{\partial p}{\partial x} = -\frac{\partial}{\partial z} (\overline{u'w'}), \quad (2.30)$$

$$\frac{\partial v}{\partial t} + u \frac{\partial v}{\partial x} + v \frac{\partial v}{\partial y} + w \frac{\partial v}{\partial z} + fu + \frac{1}{\rho_0} \frac{\partial p}{\partial y} = -\frac{\partial}{\partial z} (\overline{v'w'}), \quad (2.31)$$

$$\frac{\partial w}{\partial t} + u \frac{\partial w}{\partial x} + v \frac{\partial w}{\partial y} + w \frac{\partial w}{\partial z} + \frac{1}{\rho_0} \frac{\partial p}{\partial z} = -g \frac{\rho}{\rho_0}, \quad (2.32)$$

$$\frac{\partial \theta}{\partial t} + u \frac{\partial \theta}{\partial x} + v \frac{\partial \theta}{\partial y} + w \frac{\partial \theta}{\partial z} = -\frac{\partial}{\partial z} (\overline{\theta'w'}), \quad (2.33)$$

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} = 0. \quad (2.34)$$

Note that, as mentioned earlier, when the flow is incompressible the Reynolds averaging is simplified, rather than using mass weighted terms all flow parameters can be decomposed into just mean and fluctuating part $u = \bar{u} + u'$ (Reynolds averaging notation), hence only single primed terms in the turbulent fluxes.

Recall that θ is related to the other equations through the turbulent flux terms. Also note that in the Boussinesq framework equation (2.32) may be written in terms of potential temperature, [23],

$$-g \frac{\rho - \bar{\rho}(z)}{\rho_0} = g \frac{\theta - \bar{\theta}(z)}{\theta_0}. \quad (2.35)$$

By applying the Boussinesq approximation to the system all acoustic mechanisms are filtered out by the incompressibility. In order to study the boundary layer on its own, without any dynamics, these equations need to be further simplified by assuming that the flow is homogeneous in the horizontal and that the vertical component of velocity is zero. With this level of approximation the Rossby and gravity mechanisms are also filtered out leaving the boundary layer mechanism and the inertial mechanism. Also note that the system is in hydrostatic balance. Although these approximations would be unlikely employed in an operational model, they are quite reasonable in the boundary layer and assist the study considerably by minimising the types of wave propagation that are supported.

Above the boundary layer the turbulence goes to zero and for synoptic scales inertial forces are negligible compared to Coriolis force and pressure gradient. For a steady solution here horizontal momentum equations must therefore be in geostrophic balance,

$$-f v_g = -\frac{1}{\rho} \frac{\partial p}{\partial x}, \quad (2.36)$$

$$f u_g = -\frac{1}{\rho} \frac{\partial p}{\partial y}, \quad (2.37)$$

where subscript g denotes geostrophic wind. Since all variables are independent of x and y (except pressure) the horizontal gradients of pressure are independent of height. This can be seen by differentiating the hydrostatic balance equation

$\frac{1}{\rho_0} \frac{\partial p}{\partial z} = -g \frac{\theta}{\theta_0}$ with respect to x or y , e.g. $\frac{\partial^2 p}{\partial z \partial x} = 0$. The horizontal pressure gradients can therefore be replaced with their far field values (2.36) and (2.37) throughout the domain.

To summarise, the horizontally homogeneous equations that will be used to model the boundary layer without any dynamics are,

$$\frac{\partial u}{\partial t} - f(v - v_g) = -\frac{\partial \overline{u'w'}}{\partial z}, \quad (2.38)$$

$$\frac{\partial v}{\partial t} + f(u - u_g) = -\frac{\partial \overline{v'w'}}{\partial z}, \quad (2.39)$$

$$\frac{\partial \theta}{\partial t} = -\frac{\partial \overline{\theta'w'}}{\partial z}. \quad (2.40)$$

Due to incompressibility and no vertical wind the vertical momentum equation, which now just represents hydrostatic balance, is decoupled and so does not need to be included in the solution.

2.2 Model Closure and Parametrisation

Equations (2.38) -(2.40) are the equations required for modelling the boundary layer but they do not form a closed set since the turbulent fluxes $\overline{u'w'}$, $\overline{v'w'}$ and $\overline{\theta'w'}$ are unknown. In order to model the equations a closure is required so that these unknowns can be written in terms of the mean flow. The classic approach to this is the K -diffusion closure,

$$\overline{u'w'} = \tau_x = -K_m \frac{\partial u}{\partial z}, \quad (2.41)$$

$$\overline{v'w'} = \tau_y = -K_m \frac{\partial v}{\partial z}, \quad (2.42)$$

$$\overline{\theta'w'} = \mathcal{H} = -K_h \frac{\partial \theta}{\partial z}. \quad (2.43)$$

Writing the equations in this form one can see the similarity with the original viscosity term $\nu \frac{\partial^2 u}{\partial x^2}$, the idea originates from the premise that locally eddies influence the flow in a way comparable to molecular diffusion. Indeed K_m is commonly known as the eddy viscosity and K_h the eddy heat diffusivity. This type of closure

is first order since second order fluxes are replaced with first order terms. Other first order closures exist, such as K -profile, there also exist higher than first order closures such as the turbulent kinetic energy closure where a further energy budget equation is carried, this is said to be a one and half order closure and is used in high resolution models. The focus of the work carried out here is to examine the effect of vertical staggering, since the K -diffusion closure is used in operational models and, as will be demonstrated, is sensitive to vertical staggering it is most immediately relevant. The intention here is not to compare vertical staggering for different closures nor comment on the suitability of the closure.

For simplicity equations (2.38) - (2.40) shall be written in their τ -form,

$$\frac{\partial u}{\partial t} - f(v - v_g) = -\frac{\partial \tau_x}{\partial z}, \quad (2.44)$$

$$\frac{\partial v}{\partial t} + f(u - u_g) = -\frac{\partial \tau_y}{\partial z}, \quad (2.45)$$

$$\frac{\partial \theta}{\partial t} = -\frac{\partial \mathcal{H}}{\partial z}. \quad (2.46)$$

2.2.1 Parametrisation

Now that the general form of the closure has been decided, the terms need to be linked to the mean model parameters. This will be done using the mixing length formulation [34, 54]. K_m and K_h are obtained as,

$$(K_m, K_h) = l^2 \left| \frac{\partial \mathbf{u}}{\partial z} \right| \{f_m(Ri), f_h(Ri)\}. \quad (2.47)$$

The idea behind the formulation is that when a particle is moved by turbulence it will mix with the surrounding fluid and lose the characteristics it possessed before being moved. The distance it can move before the loss is known as the mixing length l . The formulation for the mixing length as given by [54], is,

$$l = \frac{\kappa z l_\infty}{\kappa z + l_\infty}, \quad (2.48)$$

where l_∞ is the mixing length for a neutrally stratified layer and is generally ob-

tained empirically, here 20m will be used, as from an example in [54]. Numerical experimentation suggests that the value of l_∞ has relatively little influence on the overall boundary layer steady state structure and so should not influence the sensitivity to staggering, thus for simplicity it will be kept fixed for all boundary layer depths. κ is the von-Karmen constant, taken to be 0.4.

The stability functions that shall be used are the SHARP form of [41],

$$\{f_m, f_h\} = \begin{cases} a \left(\frac{1}{20Ri}\right)^2 & Ri \geq 0.1 \\ a(1 - 5Ri)^2 & 0 \leq Ri < 0.1 \\ 1 & Ri = 0 \end{cases} , \quad (2.49)$$

where $a = 1$ for f_m and $a = 1/Pr$ for f_h ; the Prandtl number Pr is a ratio of kinematic viscosity and heat diffusivity, i.e. $\frac{K_m}{K_h}$, here the neutral value is used $Pr = 0.7$ to help distinguish between f_m and f_h ; in [41] $Pr = 1$ was used. If $Ri < 0$ then the fluid is unstably stratified and a different usually non-local closure is generally required. This form of the stability functions is employed by the Met Office over the sea [11] where it is more natural to employ a fixed boundary condition on temperature, as will be crucial to the methodology subsequently presented. Ri is the Richardson number, given by,

$$Ri = g \frac{\frac{\partial \ln(\theta)}{\partial z}}{\left| \frac{\partial \mathbf{u}}{\partial z} \right|^2}. \quad (2.50)$$

Turbulence can be produced by either buoyant instability or by shear, Richardson number describes a local ratio between buoyancy and shear and so provides useful information about the onset of turbulence. There is a critical Richardson number Ri_c , when below this value the fluid is said to be turbulent, the fluid will be turbulent whenever unstably stratified. When the fluid is statically stable, as it is for the stably stratified boundary layer, whether or not the fluid is turbulent depends on how much shear there is versus how stably stratified it is. If there is

enough shear to overcome the stabilising effect of the buoyant stability the fluid can be turbulent, as is the case in the boundary layer.

2.2.2 The Surface Boundary Conditions

The parametrisation, equation (2.47), will need to be adapted at the surface, both K_m and K_h involve calculating gradients of velocity and potential temperature, furthermore gradients of K_m and K_h will be used. If not altered then values below the ground would be required. As is generally popular in numerical weather prediction, the lower boundary conditions are imposed through a roughness length, a formula based on surface layer similarity theory and the universal logarithmic wall law. Rather than for example velocity being zero at height $z = 0$ it is said to go to zero at the top of the roughness length.

The boundary conditions for the three model variables are,

$$u = v = 0 \quad \text{at} \quad z = z_{rm}, \quad (2.51)$$

$$\theta = \theta_s \quad \text{at} \quad z = z_{rh}, \quad (2.52)$$

where z_{rm} and z_{rh} are the height of the roughness length for momentum and heat; both will be taken to be 0.1m so for ease of notation both will be referred to as z_r .

The full parametrisation of, for example, τ_x away from the surface is,

$$\tau_x = -l^2 \left| \frac{\partial \mathbf{u}}{\partial z} \right| f_m(Ri) \frac{\partial u}{\partial z}. \quad (2.53)$$

In order to obtain a parametrisation of the surface stresses $(\overline{u'w'})_0 = \tau_{x0}$ and $(\overline{v'w'})_0 = \tau_{y0}$ note first that as $z \rightarrow 0$ the mixing length $l \rightarrow \kappa z$. Also note that in the lowest 10m of the boundary layer it can be assumed that fluxes are constant [34], this aids the analysis in that wind direction can be assumed to be aligned with the x -axis, giving $v \equiv 0$. Now (2.53) becomes,

$$\tau_{x0} \approx -(\kappa z)^2 \left(\frac{\partial u}{\partial z} \right)^2 f_m(Ri_b). \quad (2.54)$$

where Ri_b is the bulk Richardson number, an approximation to the gradient Richardson number and equivalent to measuring the Richardson number from z_r to a height z ; it is given by,

$$Ri_b = g(z - z_r) \frac{(\ln \theta - \ln \theta_s)}{|\mathbf{u}|}. \quad (2.55)$$

Assuming that sufficiently close to the ground, where fluxes are constant, the bulk Richardson number and thus the stability function tend to some constant, integrating (2.54) and rearranging gives that,

$$u \approx \frac{1}{\kappa} \left(\frac{-\tau_{x0}}{f_m(Ri_b)} \right)^{\frac{1}{2}} \ln \left(\frac{z}{z_r} \right), \quad (2.56)$$

$$\Rightarrow \tau_{x0} \approx - \left(\frac{\kappa}{\ln \left(\frac{z}{z_r} \right)} \right)^2 f_m(Ri_b) u^2; \quad (2.57)$$

the bracketed part of (2.57) is known as C_{mn} and is the neutral drag coefficient for momentum. Obtained in a likewise manner, the surface heat flux is,

$$\mathcal{H}_0 \approx - \left(\frac{\kappa}{\ln \left(\frac{z}{z_r} \right)} \right)^2 f_h(Ri_b) u (\theta - \theta_s). \quad (2.58)$$

In the non wind-aligned coordinate system this is,

$$\tau_{x0} = -C_{mn} f_m(Ri_b) |\mathbf{u}| u = C_m |\mathbf{u}| u, \quad (2.59)$$

$$\tau_{y0} = -C_{mn} f_m(Ri_b) |\mathbf{u}| v = C_m |\mathbf{u}| v, \quad (2.60)$$

$$\mathcal{H}_0 = -C_{hn} f_h(Ri_b) |\mathbf{u}| (\theta - \theta_s) = C_h |\mathbf{u}| (\theta - \theta_s). \quad (2.61)$$

The drag coefficients are defined as $C_m = -C_{mn} f_m(Ri_b)$, $C_h = -C_{hn} f_h(Ri_b)$.

For the purposes of this thesis the form for the surface fluxes in equations (2.59) - (2.61) is sufficient. Operational forecasting models will tend to use log-linear formulations, a useful review of the two procedures is given in [17].

2.3 Subsidence to Balance the Boundary Layer Diffusion

In the atmosphere even the very stably stratified boundary layers are not completely steady but have continuous cooling and heating through radiative processes. In idealized models it is possible to create a quasi-steady-state whereby certain flow properties continue to vary but momentum flux, potential temperature and surface potential temperature flux are constant [10, 18, 55]. Indeed situations where these conditions are satisfied in the atmosphere exist. One such example is the wintertime Antarctic boundary layer, a very shallow layer that evolves so slowly that it can be considered as steady [42, 30].

The overall aim of this project is to address the questions on how staggering and use of stretched grids affects the numerical modelling of the boundary layer, and to achieve this in a systematic way. A methodology that would provide this systematic approach is to linearise the model around a steady state and examine the parts separately, i.e. the steady state on its own and the transient modes. For this methodology to be used a completely steady state is required.

The overall structure of the quasi-steady-state of equations (2.44)-(2.46) would appear steady or evolve very slowly, though the solution is not truly steady since quantities are still evolving. Without a true steady state, where nothing is transported, it would not make physical sense to examine the transients. Although (2.44)-(2.46) are well suited for modelling the constantly changing boundary layers, found in much of the atmosphere, employing a fixed surface potential temperature is not possible. Conversely the only way to obtain a truly steady state of these equations requires a fixed surface potential temperature boundary condition along with fixed potential temperature at the top of the domain. If fixed potential temperature boundary conditions are used heat will diffuse downwards causing the potential temperature in the boundary layer to reduce until equal to the surface potential temperature throughout the depth.

It is known from the literature, as cited above, that during the Antarctic win-

ters the boundary layer can satisfy ‘steady’ conditions and that a fixed surface temperature would be a realistic approximation here due to the frozen surface but further processes must be occurring which are not accounted for in the basic boundary layer equation set. By representing these processes in the equations it would be possible use fixed boundary conditions on potential temperature and simulate a truly steady Antarctic wintertime type boundary layer. This would enable the desirable methodology to be used.

In the Antarctic strong Katabatic (drainage) winds occur persistently and these strong winds lead to some interesting phenomena. There have been a number of studies examining the Antarctic boundary layer and general climatic processes, in particular the divergent circulation that occurs as a result of the Katabatic wind in the boundary layer, [36, 56, 74, 50]. The Katabatic wind occurs as the air is cooled dramatically by the cold surface or by radiation, as the density of this air increases it moves down slope, towards the coast in Antarctica. The effect of the horizontal divergence of air in the boundary layer would be an overall reduction in potential temperature, such as seen when using a fixed surface temperature. As discussed by [36] this horizontal flow divergence must be balanced by a horizontal convergence and subsidence just above the boundary layer. Though it is not necessary to model specifically Katabatic type winds, the resulting effect they have on a boundary layer (reduction of potential temperature) would be equivalent to that which is seen with the fixed surface potential temperature and so each can be balanced by the same process. Since there is no x or y dependence in the boundary layer only model it would not be possible to produce and horizontal convergence but it is possible to produce a subsidence heating. Subsidence heating in the context of the stably stratified boundary layer has also been recently examined by [51].

A subsidence heating can be implemented by reintroducing the vertical advection of the potential temperature term, but with an imposed subsidence velocity, given by,

$$w_{\text{sub}} \frac{\partial \theta}{\partial z}. \quad (2.62)$$

The w_{sub} , known as the subsidence velocity, will be negative as it is a velocity

towards the ground, it will be zero at the surface and have absolute value increasing upwards, this can be achieved by using a tanh type function. A suitable form for this term, which is found empirically, is,

$$w_{\text{sub}} = -0.015 \tanh\left(\frac{z}{1000}\right). \quad (2.63)$$

This gives sufficient heating so as to balance the cooling of the boundary layer due to the fixed surface potential temperature whilst still giving sensible depths and structure when comparing to the boundary layers observed in the literature e.g. [6].

Subsidence velocity provides the mechanism which allows a fully steady state to be found. Further to this a radiative cooling term will be included. Consider a situation in neutral stratification with no shear, both $K_m = 0$ and $K_h = 0$, if some shear grew in that region then some eddy viscosity and eddy diffusivity would also appear, the neutral stratification would keep potential temperature from changing through the boundary layer terms but velocity would change. Then through the Coriolis it could oscillate, this type of behaviour is known as the nocturnal jet. If instead the shear grew in a region of fluid that was stably stratified the stability would damp away the shear and prevent it from oscillating. The nocturnal jet itself is an interesting phenomena but this kind of behaviour needs to be avoided since an oscillation would prevent a fully steady state being achieved. It may be that shear grows above the boundary layer, where the fluid is neutrally stratified, as a numerical artefact which could be difficult to predict. By adding a radiative cooling ensures that the whole domain is stably stratified and thus guarantees that any jet effect is damped away and the steady state properly approached. To implement the stable stratification an overall radiative cooling R_c of around 1K per day is included in the equation.

With the inclusion of both the subsidence heating and radiative cooling, equation (2.46) becomes,

$$\frac{\partial \theta}{\partial t} = -\frac{\partial \mathcal{H}}{\partial z} - w_{\text{sub}} \frac{\partial \theta}{\partial z} - R_c. \quad (2.64)$$

This subsidence warming term is not explicitly added to global forecasting models but exists, where necessary, through the normal vertical advection term. The extra terms in the energy budget due to the subsidence and radiative cooling are negligible in comparison to the overall energy in the system. Using the simplest forward finite difference in the subsidence prevents this extra term affecting the sensitivity of results on averaging. Having a steady state that can be linearised about is a highly powerful tool in terms of examining configurations of the grid and model variables.

2.4 Test Cases

Five stable boundary layers which can be used to compare any configurations of staggering with stretched and uniform grids are constructed. A typical shallow stable boundary layer may be about 100m deep, a very deep layer may be around 1000m. The five boundary layers will be in this range with different depths obtained by choosing the boundary conditions; the upper boundary condition on velocities u_g and v_g and the surface potential temperature, θ_s . The upper condition on potential temperature is fixed as $\theta_g = 308\text{K}$. The two components of geostrophic wind are chosen to be equal $u_g = v_g$. Sometimes in the literature v_g is set to zero to assist in analysis, although this is not necessary here. Later when dynamical processes are introduced it may be useful to have $u_g = v_g$ so that modes have equivalent advective properties in each direction. A deeper boundary layer will have either more shear, larger difference between upper and surface potential temperature boundary condition or both compared with a shallow boundary layer. Deeper boundary layers are less stably stratified than shallow boundary layers and have larger K_m and K_h , giving increased damping. Stratification is determined by the gradient of θ , i.e. if θ is increasing with height the regime is stably stratified. While stably stratified boundary layers have positive θ_z , unstably stratified boundary layers have a negative θ_z and neutral layers have $\theta_z = 0$. The approximate depths and boundary conditions for the five boundary layers that will be compared in the model are summarised in Table 2.1.

Notation	ug, vg (ms^{-1})	$\theta_s(K)$	Approx. depth (m)	Obukhov Length (m)
BL1	4	283	100	8.33
BL2	6	288	200	27.18
BL3	8.5	293	400	84.04
BL4	10.5	298	650	229.59
BL5	14	298	950	407.17

Table 2.1: Five boundary layers with varying depths, dependent on the boundary conditions of velocities and potential temperature.

2.5 Discretisation and Averaging

In order to solve equations (2.44) - (2.46), with their respective parametrisations, they must be discretised and solved numerically. Derivatives are solved using second order finite differencing. A vertical grid is constructed, the lowest discrete level in the model will be at z_r , the highest model level will be at 2000m, this ensures any boundary layer structure lies in the domain. The grid is constructed such that there are N z_ρ levels, so called due it being the location where density will be stored, and $N + 1$ z_w levels, so called as it will be where vertical velocity is always stored. z_ρ levels are thought of as being staggered relative to z_w , the unstaggered levels, since each z_ρ level lies between two z_w levels. The lowest z_w level is fixed at z_r and the highest z_w level fixed at 2000m the position of all other levels is dependent on the choice of spacing and the number of grid points used. The model levels can be arranged uniformly or with a stretching, this will be discussed presently. Since the equations present a highly nonlinear problem no analytical solution is known of, instead a high resolution numerical solution will be used as a ‘truth’ reference solution. For normal low resolution 11 z_w including z_r are used for the domain of 2000m, this number of grid points represents the operational resolution that was in use at the Met Office when this work was undertaken. For the high resolution run 100 z_ρ and 101 z_w grid points are used.

2.5.1 Choice of Vertical Staggering

There are two types of vertical staggering to consider: the Lorenz grid and Charney-Phillips grids. For each of the grids, the horizontal velocity variables u and v are stored at the same grid point, these are the staggered z_ρ levels. Note that ρ and w are not required in the boundary layer only formulation but are considered temporarily to demonstrate the two grid types. If considering the full dynamics then on both Lorenz and Charney-Phillips grids the vertical velocity w would be stored at the unstaggered z_w levels. Horizontal velocities u and v need to be stored together to avoid averaging in the Coriolis terms. Staggering density relative to w ensures no vertical averaging is required in $\nabla \cdot \mathbf{u}$ in the continuity equation whilst storing w at the unstaggered levels makes fulfilling the no flux boundary conditions easier.

In addition to ρ a second thermodynamic variable is required and this is chosen to be potential temperature θ . The difference between the Lorenz grid and Charney-Phillips grid is the choice of whether to place the potential temperature with horizontal velocity and density or with vertical velocity. For the Lorenz grid it is stored with density, with the Charney-Phillips grid it is stored with vertical velocity. When in the Charney-Phillips framework the unstaggered levels are sometimes given notation z_θ , since θ will be moved for the purpose of the comparison in this work the unstaggered levels shall be labelled as z_w . For the boundary layer only case, where ρ and w are not required, the Lorenz grid is effectively unstaggered since all model variables lie at the same grid points.

The configuration of the grids for both Lorenz and Charney-Phillips is shown in Figure 2.1. The z_ρ levels have integer indices, the spacing between is denoted by Δz with $\pm \frac{1}{2}$ indices, the z_w levels have $\pm \frac{1}{2}$ indices and have spacing with integer indices. The indices of the grid spacings are designated so as to match the indices of the level that they cross.

In order to calculate the Richardson number it is clear that the derivatives of velocities and potential temperature are required at the same model level. On the Lorenz grid these quantities will naturally lie on the z_w levels which is also where




Figure 2.1: The Lorenz (left) and Charney-Phillips (right) grids. z_ρ levels have integer indices and are represented by solid lines, z_w levels are the dashed lines. Δz are assigned indices by the model level that they cross, i.e. whole indices represent spacing between z_w levels.

K_m and K_h will be required in the calculation of the fluxes; no averaging will be present. When using the Charney-Phillips grid the shear will naturally lie on z_w levels while the potential temperature gradient will lie on the z_ρ levels. This means that one or the other of these quantities will need to be averaged in order to obtain the Richardson number. This is why solutions obtained when using the Lorenz grid can be expected to be more accurate than those found when using the Charney-Phillips grid.

The three basic cases to consider for the Charney-Phillips grid are,

- I. Potential temperature gradient is averaged so that Richardson number is calculated at the z_w levels.
- II. Shear is averaged so that Richardson number is calculated the z_ρ levels.
- III. Both quantities are averaged so that Richardson number is computed at z_ρ and z_w grid levels.

For the second II and third III configurations four subclasses denoted ‘a’ -‘d’ can be added. There are three stages in the computation of the shear squared,

$$\left| \frac{\partial \mathbf{u}}{\partial z} \right|^2 = \left\{ \left[\left(\frac{\partial u}{\partial z} \right)^2 + \left(\frac{\partial v}{\partial z} \right)^2 \right]^{\frac{1}{2}} \right\}^2, \quad (2.65)$$

first finding the gradients of the two components of velocity, then taking the modulus (absolute value) and finally squaring, as the Richardson number requires. The averaging can take place at any stage during this calculation.

- a. Average velocity gradients then square them in order to compute squared modulus.

$$\left| \frac{\partial \mathbf{u}}{\partial z} \right| = \left[\overline{\left(\frac{\partial u}{\partial z} \right)^2} + \overline{\left(\frac{\partial v}{\partial z} \right)^2} \right]^{\frac{1}{2}}, \quad (2.66)$$

- b. Average squared velocity gradients then compute squared modulus.

$$\left| \frac{\partial \mathbf{u}}{\partial z} \right| = \left[\overline{\left(\frac{\partial u}{\partial z} \right)^2} + \overline{\left(\frac{\partial v}{\partial z} \right)^2} \right]^{\frac{1}{2}}, \quad (2.67)$$

c. Compute the modulus and then average before squaring.

$$\left| \frac{\partial \mathbf{u}}{\partial z} \right|^2 = \overline{\left\{ \left[\left(\frac{\partial u}{\partial z} \right)^2 + \left(\frac{\partial v}{\partial z} \right)^2 \right]^{\frac{1}{2}} \right\}^2} \quad (2.68)$$

d. Compute the squared modulus and then average.

$$\left| \frac{\partial \mathbf{u}}{\partial z} \right|^2 = \overline{\left\{ \left[\left(\frac{\partial u}{\partial z} \right)^2 + \left(\frac{\partial v}{\partial z} \right)^2 \right]^{\frac{1}{2}} \right\}^2} \quad (2.69)$$

In option I Ri lies at θ levels, this is also where shear lies and thus K_m and K_h can be readily found at these levels. A further three options arise for how to compute K_h at the z_ρ levels.

- i. K_h itself can be averaged.
- ii. The stability function $f_h(Ri)$ can be averaged as well as the shear.
- iii. The Richardson number Ri can be averaged as well as the shear.

Similarly, for each of the various configurations under option II three options exist for the averaging required to obtain K_m at the z_w levels; for option III no further averaging is required for K_m and K_h .

This gives three cases under option I, twelve under option II and four under option III, a total of 19 that require consideration for Charney-Phillips, although others would be possible. With the Lorenz grid case this gives a total of 20 and for 5 depths of boundary layer a total of 100 configurations. In addition to this there are a number of possible grid spacings to consider. As the investigation progresses it should be possible to reduce the number of cases that need to be considered as it becomes apparent which of the varying options have the greatest effect. Charney-Phillips configurations are labelled as, for example II(b)-ii, denoting that Richardson number is calculated at z_ρ levels, that option b is used for calculating shear squared and that the stability function is averaged to find K_m . Under option I the bracketed lower case letter is omitted and for option III the

lower case Roman numeral is omitted. A selection of the candidate configurations are displayed in Table 2.2.

	Shear Handling	K handling
Lorenz	N/A	N/A
Charney-Phillips Option I	Required at natural location	Average K_h -i Average f_h -ii Average Ri -iii
Charney-Phillips Option II	Sub Options (a-d)	Average K_m -i Average f_m -ii Average Ri -iii
Charney-Phillips Option III	Sub Options (a-d)	At required locations

Table 2.2: Outline of a selection of choices for the boundary layer term averaging when using the Charney-Phillips grid.

Averaging in Boundary Conditions

Reconsider equations (2.44)-(2.46) but in discretised form at the lowest internal model level,

$$\left. \frac{\partial u}{\partial t} \right|_1 - f(v_1 - v_g) = - \frac{(\tau_x)_{\frac{3}{2}} - \tau_{x0}}{z_{\frac{3}{2}} - z_r}, \quad (2.70)$$

$$\left. \frac{\partial v}{\partial t} \right|_1 + f(u_1 - u_g) = - \frac{(\tau_y)_{\frac{3}{2}} - \tau_{y0}}{z_{\frac{3}{2}} - z_r}, \quad (2.71)$$

$$\left. \frac{\partial \theta}{\partial t} \right|_j = - \frac{\mathcal{H}_{j+\frac{1}{2}} - \mathcal{H}_0}{z_{j+\frac{1}{2}} - z_r}, \quad (2.72)$$

j in equation (2.72) is 1 for the Lorenz grid and $\frac{3}{2}$ for the Charney-Phillips grid. Recalling from equations (2.59)-(2.61), the discretised form for the surface fluxes are,

$$\tau_{x0} = - \left(\frac{\kappa}{\ln \left(\frac{z_1}{z_r} \right)} \right)^2 f_m(Ri_b) |\mathbf{u}_1| u_1, \quad (2.73)$$

$$\tau_{y0} = - \left(\frac{\kappa}{\ln \left(\frac{z_1}{z_r} \right)} \right)^2 f_m(Ri_b) |\mathbf{u}_1| v_1, \quad (2.74)$$

$$\mathcal{H}_0 = - \left(\frac{\kappa}{\ln \left(\frac{z_j}{z_r} \right)} \right)^2 f_h(Ri_b) |\mathbf{u}_j| (\theta_j - \theta_s). \quad (2.75)$$

subscript j in (2.75) is again 1 for the Lorenz grid and $\frac{3}{2}$ for the Charney-Phillips grid. When using the Charney-Phillips grid the $|\mathbf{u}_{\frac{3}{2}}|$ will be an averaged quantity. Now consider the discretised bulk Richardson number,

$$Ri_b = g(z_j - z_r) \frac{(\theta_j - \theta_s)}{|\mathbf{u}_j|}. \quad (2.76)$$

When using the Lorenz grid j will be 1 for all calculations of bulk Richardson number. When using the Charney-Phillips grid j will be 1 for the bulk Richardson in f_m , θ_1 will be an averaged quantity. The subscript j will be $\frac{3}{2}$ in f_h and $|\mathbf{u}_{\frac{3}{2}}|$ will be an averaged quantity.

Close to the surface the model variables u , v and θ exhibit log-like behaviour [28, 34], e.g. $u \sim \ln \left(\frac{z}{z_r} \right)$. When using the Charney-Phillips grid a straight forward averaging of these quantities may thus be inaccurate. When a curve is logarithmic a point a on that curve can be approximated as $c \ln(a)$, where c is some constant which can be pinned down using a known point. Instead of a straight forward averaging this log like behaviour can be used to approximate model variables at unknown points. For example,

$$\bar{u}_{\frac{3}{2}} = \frac{\ln \left(\frac{z_{\frac{3}{2}}}{z_r} \right)}{\ln \left(\frac{z_1}{z_r} \right)} u_1, \quad (2.77)$$

$$\bar{\theta}_1 = \frac{\ln \left(\frac{z_1}{z_r} \right)}{\ln \left(\frac{z_{\frac{3}{2}}}{z_r} \right)} (\theta_{\frac{3}{2}} - \theta_s) + \theta_s. \quad (2.78)$$

Accuracy can be further improved by extrapolating from more grid points and averaging.

2.5.2 Grid Spacing

For flows with a boundary layer such as the atmosphere the choice of grid stretching is crucial. The simplest type of grid would be one where z_w levels in the model are spaced uniformly, such that the spacing between unstagged levels Δz is simply the domain size divided by the number of staggered grid points N ; z_p levels would lie halfway between z_w levels. From a computational point of view this is somewhat inefficient when considered the underlying structure in the atmosphere. The complex topography of the Earth together with all the vegetation and buildings creates turbulent flow near the ground, the very reason that a specific boundary layer closure is used. High up in the atmosphere the flow structure has greater regularity and larger overall structure. To have a completely uniform grid throughout is to dedicate equivalent computer processing time to the large well resolved structure of the upper atmosphere as to the unresolved small scale turbulent processes close to the surface. In fact, if the resolution is so coarse that no boundary layer structure is resolved at all, then it has been shown that large errors will occur throughout the model, [39].

It is widely accepted that a grid having finer resolution close to the surface and lower resolution higher up produces a considerably more efficient and accurate way to model the atmosphere. One can construct a grid with high resolution near the ground by implementing a geometric stretching of the grid spacing, this is similar to the method in place at the Met Office. There are, however, some subtle issues, associated to the log-like behaviour of the model variables near the surface, which cause numerical solutions to converge to the true solution very slowly with uniform and geometrically stretched grids. It is crucial to the methodology that solutions converge at a reasonable rate, if they do not then there would be differences between solutions on Lorenz and Charney-Phillips for even a relatively large number of grid points such as $N = 100$. This does not mean that stretched grids should not be considered at all only that they can not be used in order to generate the high resolution reference solution. For the high resolution a transformation to a logarithmic coordinate is required in order to obtain fast enough convergence. In this

section the geometric grids are introduced and then the logarithmic transformation with more detailed reasons behind its use.

Geometric Grids

Although poor convergence is expected high ($N = 100$) and low resolution geometric grids are set up for comparison purposes. The high resolution geometric grid will be useful to examine and confirm poor convergence. Since geometric grids can be easily adapted, i.e. to bias the points closer to the ground or vice versa, it will be of interest to have low resolution geometric grids to compare with the Met Office grid. For the high resolution the minimum spacing (between the top of the roughness length and $(z_\rho)_1$) is chosen as $(\Delta zh)_0 = 100/N = 1\text{m}$, subsequent spacings are found geometrically,

$$(\Delta zh)_j = (\Delta zh)_0(1 + \alpha)^j \text{ for } j = 1, 2, \dots, 2N, \quad (2.79)$$

where (Δzh) represents the spacing between an unstaggered and a staggered level. The unknown term α can be determined by solving,

$$\frac{(\Delta zh)_0}{D} \left[\frac{1 - (1 + \alpha)^{2N-1}}{-\alpha} \right] - 1 = 0. \quad (2.80)$$

Equation (2.80) can be solved by using any root finding algorithm. Some examples of values for α are displayed in Table 2.3.

N	α
10	0.2138
100	0.0177
200	0.0088
400	0.0044

Table 2.3: Values of α for increasing resolution.

Equation (2.79) gives $2N$ grid spacings, each representing the distance between a z_ρ level and a z_w level. These are now combined to give spacings between one z_ρ level and another, the $(\Delta z)_{j+\frac{1}{2}}$ spacings, and likewise $(\Delta z)_j$ the spacing between

one z_w level and another.

$$(\Delta z)_j = (\Delta zh)_{2j-1} + (\Delta zh)_{2j}, \quad j = 1, 2, \dots, N, \quad (2.81)$$

$$(\Delta z)_{j+\frac{1}{2}} = (\Delta zh)_{2j} + (\Delta zh)_{2j+1}, \quad j = 1, 2, \dots, N - 1. \quad (2.82)$$

This produces a z_w and z_ρ stretched at the same rate. Alternatively the grid could be constructed so that the staggered levels are positioned halfway between the unstaggered, as is the case in the levels obtained from the Met Office. Following [39] a higher level of accuracy is expected when levels vary smoothly, this gives numerical derivatives with as close to second order accuracy as possible. During this study the effect of placing z_ρ halfway between z_w as opposed to a continuous stretching is investigated. The low resolution geometric grid with α as given in Table 2.3 will be compared with the Met Office grid spacing and uniform grids for computing the steady and transient solutions; results for this are presented in the next two chapters. The levels provided by the Met Office are 0.1m, 20.1m, 80.1m, 180.1m, 320.1m, 500.1m, 720.1m, 980.1m, 1280.1m, 1620.1m and 2000m for the unstaggered levels and 10.1m, 50.1m, 130.1m, 250.1m, 410.1m, 610.1m, 850.1m, 1130.1m, 1450.1m, 1810.1m for the staggered levels, levels have all had the roughness length added to them and are given accurate to one decimal place. These levels were obtained at a time when the Met Office employed a model with 38 vertical levels; it should be noted that at the time of writing an upgrade to use 70 vertical levels had occurred. The low resolution geometric grid that will be used here biases model levels slightly more toward the surface than the Met Office grid.

Logarithmic Transformation

It is well known that near the surface model variables have mean properties proportional to the log of their distance from the wall [34]. For example the logarithmic wind speed profile is given by $u \sim \ln\left(\frac{z}{z_r}\right)$. In this regime the vertical derivative approaches a singularity as $z \rightarrow 0$. The basic Taylor series expansion which leads to the centered difference scheme for the derivative of a model variable, which will

then be used in the equations for the fluxes, is,

$$\left. \frac{\partial u}{\partial z} \right|_{j+\frac{1}{2}} = \frac{u_{j+1} - u_j}{(\Delta z)_{j+\frac{1}{2}}} + \epsilon. \quad (2.83)$$

The formal order of accuracy m of the scheme is determined by examining ϵ as resolution increases and is generally determined by the leading term. For a general level the truncation error is given by,

$$\epsilon = -\frac{\Delta z^2}{6} \frac{\partial^3 u}{\partial z^3} - \frac{\Delta z^4}{120} \frac{\partial^5 u}{\partial z^5} - \dots \quad (2.84)$$

For u which is infinitely differentiable, and has properly bounded derivative, the order of this scheme would be said to be $O(\Delta z^2)$ (second order); that is to say the error in the approximation reduces like the square of the grid refinement. Close to the ground, however, the derivatives of u are not necessarily bounded, but since u is logarithmic will scale like $\frac{\partial^p u}{\partial z^p} \sim \frac{(-1)^{p+1}(p-1)!}{z^p}$. Plugging this term into the derivative in equation (2.84) means that the error ϵ near the ground would be approximated by,

$$\epsilon \approx \sum_{k=1}^{\infty} -\frac{(\Delta z)^{2k}}{(2k+1)z^{2k+1}}. \quad (2.85)$$

The increasing powers of z as $k \rightarrow \infty$ in the denominator means that the overall magnitude of successive terms in ϵ will not quickly reduce. It is not clear that the overall scheme will have second order convergence, as implied by the leading order numerator, but likely some slower rate of convergence governed by a combination of grid spacing and the influence of the derivative approaching the singularity. Although the boundary condition assumes a logarithmic variability there will be other model levels in the logarithmic region that will not have good convergence. Further it is not clear that in reducing the size of Δz the magnitude of ϵ is reduced. Take for example the uniform grid, in this case the level heights are given by $z = j\Delta z$, then equation (2.85) reduces to

$$\epsilon \approx \sum_{k=1}^{\infty} -\frac{1}{(2k+1)j^{2k+1}\Delta z}. \quad (2.86)$$

For a level j near the surface the overall error increases as Δz decreases.

The singularity is due to the derivative having a $\frac{1}{z}$ factor, which follows into the truncation error of the scheme, if the derivative were completely bounded so as to avoid this term the behaviour of the terms in ϵ would change so as to decrease as Δz decreased. A way to achieve this is to perform a coordinate transformation $\zeta = \eta(z)$, as in [8, 77]. For any model variable whose derivative is required the chain rule is applied to obtain, for example,

$$\frac{\partial u}{\partial z} = \frac{\partial \zeta}{\partial z} \frac{\partial u}{\partial \zeta}. \quad (2.87)$$

Now, provided that $\eta(z) \rightarrow \ln(z)$ as $z \rightarrow 0$ it is found that, near to the ground,

$$\frac{\partial u}{\partial z} = \frac{1}{z} \frac{\partial u}{\partial \zeta} \quad (2.88)$$

where $\frac{\partial u}{\partial \zeta}$ is a bounded derivative. Now the centred difference approximation for the derivative of u with respect to ζ gives a truncation error of,

$$\epsilon = -\frac{\Delta \zeta^2}{6} \frac{\partial^3 u}{\partial \zeta^3} - \frac{\Delta \zeta^4}{120} \frac{\partial^5 u}{\partial \zeta^5} - \dots \quad (2.89)$$

The derivatives $\frac{\partial^p u}{\partial \zeta^p}$ are bounded and so the order of the scheme is $O(\Delta \zeta^2)$, where $\Delta \zeta$ is uniform throughout the domain.

A further advantage of choosing η such that the coordinate is logarithmic near the surface is that τ_x near the surface would then be given by,

$$\tau_x \approx -\kappa^2 f_m(Ri) \left| \frac{\partial \mathbf{u}}{\partial \zeta} \right| \frac{\partial u}{\partial \zeta}. \quad (2.90)$$

In the new coordinate the Richardson number is given by,

$$Ri = g \frac{\left(\frac{\partial \zeta}{\partial z}\right)^{-1} \frac{\partial \ln \theta}{\partial \zeta}}{\left| \frac{\partial \mathbf{u}}{\partial \zeta} \right|^2} \quad (2.91)$$

in the logarithmic coordinate. The $\frac{\partial \zeta}{\partial z}$ terms have cancelled with the z^2 that appears

in the mixing length, l^2 . This gives an expression for the surface momentum flux where dependence on z only occurs through the Richardson number.

Note that the averaging as described in Table 2.2 for the Charney-Phillips configurations can be modified so that the mixing length l^2 lies at the correct level, so as to ensure the cancelling of z . For example consider the discretised version of the θ equation (2.64), which would contain the term,

$$-\frac{\partial \mathcal{H}}{\partial \zeta} \Big|_{j+\frac{1}{2}} = \frac{1}{(\Delta \zeta)} \left[(K_h)_{j+1} \frac{1}{z_{j+1}} \frac{\partial \theta}{\partial \zeta} \Big|_{j+1} - (K_h)_j \frac{1}{z_j} \frac{\partial \theta}{\partial \zeta} \Big|_j \right], \quad (2.92)$$

where K_h is averaged from z_w levels to z_ρ , for example,

$$(K_h)_j = \frac{1}{2} \left[l_{j+\frac{1}{2}}^2 \left(\frac{\partial \zeta}{\partial z} \right)_{j+\frac{1}{2}} \left| \frac{\partial \mathbf{u}}{\partial z} \right|_{j+\frac{1}{2}} f_h(Ri)_{j+\frac{1}{2}} + l_{j-\frac{1}{2}}^2 \left(\frac{\partial \zeta}{\partial z} \right)_{j-\frac{1}{2}} \left| \frac{\partial \mathbf{u}}{\partial z} \right|_{j-\frac{1}{2}} f_h(Ri)_{j-\frac{1}{2}} \right]. \quad (2.93)$$

In this example the two $\frac{1}{z}$ terms which cancel with z^2 in the mixing length come from the shear in K_h and the $\frac{\partial \theta}{\partial z}$ term which K_h multiplies. However due to the averaging the mixing length and the shear will be a half level away from the vertical gradient of θ . Instead the mixing length and the $\frac{\partial \zeta}{\partial z}$ in front of the shear are taken outside of the averaging. Both of these quantities are smooth so whether they appear inside the averaging or outside makes little difference to the configuration, especially in the high resolution run for which this coordinate system is used. Instead the averaging to find K_h is given by,

$$(K_h)_j = \frac{l_j^2}{2} \left(\frac{\partial \zeta}{\partial z} \right)_j \left[\left| \frac{\partial \mathbf{u}}{\partial z} \right|_{j+\frac{1}{2}} f_h(Ri)_{j+\frac{1}{2}} + \left| \frac{\partial \mathbf{u}}{\partial z} \right|_{j-\frac{1}{2}} f_h(Ri)_{j-\frac{1}{2}} \right]. \quad (2.94)$$

In the new coordinate system equation (2.44) is given by,

$$\frac{\partial u}{\partial t} - f(v - v_g) = \frac{\partial \zeta}{\partial z} \frac{\partial \tau_x}{\partial \zeta}. \quad (2.95)$$

The derivatives of τ with respect to ζ will be well bounded and depend on ζ rather than z ; meaning the return to finding that $\epsilon \approx O(\Delta \zeta^2)$ as $N \rightarrow \infty$.

The major disadvantage to a purely logarithmic grid is that resolution is biased

entirely in the lowest part of the domain. It may be desirable to redistribute some of the points higher up, especially if the domain is quite large. The simplest way to achieve this is by setting up a hybrid log-linear coordinate such as that given by [77],

$$\zeta = \ln \left(\frac{z + z_r}{z_r} \right) + \frac{z}{b_0}. \quad (2.96)$$

Increasing b_0 in (2.96) biases the concentration of grid points towards a more logarithmic setup. The difference between location of grid points for $N = 100$ is shown in Figure 2.2, for the log-linear grid $\beta_0 = 67.5$. Clearly the hybrid grid gives much better resolution higher in the domain whilst still providing the logarithmic behaviour close to the ground, ensuring the good convergence properties.




Figure 2.2: Distribution of points when using purely logarithmic and log-linear hybrid grids, $N = 100$.

Although the logarithmic and log-linear grids will provide good convergence properties for the high resolution runs they may not be that well suited to low resolution runs. When the number of points is quite low, as it is in global models, satisfying the logarithmic spacing near the ground may result in insufficient points higher in the domain. Later when the steady states are examined the ability of low resolution versions of the logarithmic and log-linear grids will be tested.

Chapter 3

Steady State Solution

For the proposed methodology equations (2.44), (2.45) and (2.64) are linearised about an assumed steady reference state. This is achieved by Taylor expanding variables; as an example consider a time dependent function F of variable U . The expansion around the reference state is $F(U) = F^{(r)} + \frac{dF}{dU}|^{(r)} U' + \dots$. Superscript (r) denotes steady reference state and is the part independent of time. The $'$ represents the linear departure from the reference state known as the transient, i.e. $U' = U - U^{(r)}$. An assumption of the expansion is that $|U'| \ll |U^{(r)}|$ so higher order terms can be neglected. For the model variables themselves the expansion is simple since the function is just the variable and $u = U + u'$, $v = V + v'$ and $\theta = \theta^{(r)} + \theta'$. Note that previously $'$ had been used to represent the sub-grid part of the Reynolds averaging but for any terms introduced in the remainder of Part I $'$ shall be used to denote the fluctuating or transient part of the linearisation. For the velocity components upper case will be used for the steady parts U , V and W ; for the thermodynamic terms the steady part will be denoted with a superscript (r) .

First the steady reference state part is considered by neglecting all the transient parts (i.e. $u' = v' = \theta' = 0$, $\frac{\partial}{\partial t} = 0$). Writing in K -form the steady state is thus found by solving,

$$0 = f(V - v_g) + \frac{\partial}{\partial z} \left(K_m \frac{\partial U}{\partial z} \right), \quad (3.1)$$

$$0 = -f(U - u_g) + \frac{\partial}{\partial z} \left(K_m \frac{\partial U}{\partial z} \right), \quad (3.2)$$

$$0 = \frac{\partial}{\partial z} \left(K_h \frac{\partial \theta^{(r)}}{\partial z} \right) - w_{\text{sub}} \frac{\partial \theta^{(r)}}{\partial z} - R_c. \quad (3.3)$$

There are a number of ways to compute the solution of equations (3.1)-(3.3). In the Met Office NWP the governing equations are solved using a Crank-Nicholson iteration. The simplest in terms of coding is that of false time-stepping. Since the boundary layer equations are being solved it can be assumed that any transient processes will be diffusive, this means that if the full equations are iterated forward in time they will eventually just arrive at the steady state. The form of the equations to be solved is thus just the above but with a temporal derivative of the steady part of each model variable on the left hand side. If iterated through sufficient numerical time steps, the above equations will arrive at their steady state. This method can have quite slow convergence for simple first or second order time schemes since it is limited by stability criteria, given by $\frac{K\Delta t}{\Delta z^2} < 1$, sometimes referred to as the viscous Courant number. If the resolution is fine and the boundary layer is deep, giving larger K , then the required time step may be quite small. Note also that nonlinearities in the system may reduce the time step further. It is found that false time-stepping is suitable for finding the solution of low resolution configurations where large grid spacing will allow for large time steps and so will be used for all low resolution runs. Since false time-stepping is sufficient for finding the low resolution steady states the Crank-Nicholson scheme is not considered here.

In discretised form the false time stepping of (3.1) and (3.2) is,

$$U_j^{n+1} = U_j^{n-1} + 2\Delta t \left\{ f(V_j^n - v_g) + \frac{1}{(\Delta z)_j} \left[(K_m)_{j+\frac{1}{2}} \frac{U_{j+1} - U_j}{(\Delta z)_{j+\frac{1}{2}}} - (K_m)_{j-\frac{1}{2}} \frac{U_j - U_{j-1}}{(\Delta z)_{j-\frac{1}{2}}} \right]^{n-1} \right\}, \quad (3.4)$$

$$V_j^{n+1} = V_j^{n-1} + 2\Delta t \left\{ -f(U_j^n - u_g) + \frac{1}{(\Delta z)_j} \left[(K_m)_{j+\frac{1}{2}} \frac{V_{j+1} - V_j}{(\Delta z)_{j+\frac{1}{2}}} - (K_m)_{j-\frac{1}{2}} \frac{V_j - V_{j-1}}{(\Delta z)_{j-\frac{1}{2}}} \right]^{n-1} \right\}. \quad (3.5)$$

On the Lorenz grid equation (3.3) is written as,

$$(\theta^{(r)})_j^{n+1} = (\theta^{(r)})_j^n + \frac{\Delta t}{(\Delta z)_j} \left\{ \left[(K_h)_{j+\frac{1}{2}} \frac{\theta_{j+1}^{(r)} - \theta_j^{(r)}}{(\Delta z)_{j+\frac{1}{2}}} - (K_h)_{j-\frac{1}{2}} \frac{\theta_j^{(r)} - \theta_{j-1}^{(r)}}{(\Delta z)_{j-\frac{1}{2}}} \right] - (w_{\text{sub}})_j \frac{\theta_{j+1}^{(r)} - \theta_j^{(r)}}{(\Delta z)_{j+\frac{1}{2}}} \right\}^n - R_c, \quad (3.6)$$

or on the Charney-Phillips grid as,

$$(\theta^{(r)})_{j+\frac{1}{2}}^{n+1} = (\theta^{(r)})_{j+\frac{1}{2}}^n + \frac{\Delta t}{(\Delta z)_{j+\frac{1}{2}}} \left\{ \left[(K_h)_{j+1} \frac{\theta_{j+\frac{3}{2}}^{(r)} - \theta_{j+\frac{1}{2}}^{(r)}}{(\Delta z)_{j+1}} - (K_h)_j \frac{\theta_{j+\frac{1}{2}}^{(r)} - \theta_{j-\frac{1}{2}}^{(r)}}{(\Delta z)_j} \right] - (w_{\text{sub}})_{j+\frac{1}{2}} \frac{\theta_{j+\frac{3}{2}}^{(r)} - \theta_{j+\frac{1}{2}}^{(r)}}{(\Delta z)_{j+1}} \right\}^n - R_c. \quad (3.7)$$

Note that an upwind scheme is used for the subsidence for simplicity at the lower boundary. Superscript n denotes time level and subscript j denotes the height level. A second order time scheme is required for velocity and first order for the potential temperature, this ensures stability in the scheme.

At the lowest model level the boundary conditions are included as,

$$U_1^{n+1} = U_1^n + 2\Delta t \left\{ f(V_1^n - v_g) + \frac{1}{(\Delta z)_1} \left[(K_m)_{\frac{3}{2}} \frac{U_2 - U_1}{(\Delta z)_{\frac{3}{2}}} - (C_m)_1 |\mathbf{U}_1| U_1 \right] \right\}^{n-1}, \quad (3.8)$$

$$V_1^{n+1} = V_1^n + 2\Delta t \left\{ -f(U_1^n - u_g) + \frac{1}{(\Delta z)_1} \left[(K_m)_{\frac{3}{2}} \frac{V_2 - V_1}{(\Delta z)_{\frac{3}{2}}} - (C_m)_1 |\mathbf{U}_1| V_1 \right] \right\}^{n-1}. \quad (3.9)$$

On the lowest Lorenz grid level potential temperature is implemented as,

$$(\theta^{(r)})_1^{n+1} = (\theta^{(r)})_1^n + \frac{\Delta t}{(\Delta z)_1} \left\{ \left[(K_h)_{\frac{3}{2}} \frac{\theta_2^{(r)} - \theta_1^{(r)}}{(\Delta z)_{\frac{3}{2}}} - (C_h)_1 |\mathbf{U}_1| (\theta_1^{(r)} - \theta_s) \right] - (w_{\text{sub}})_1 \frac{\theta_2^{(r)} - \theta_1^{(r)}}{(\Delta z)_{\frac{3}{2}}} \right\}^n - R_c, \quad (3.10)$$

or on the lowest Charney-Phillips grid level as,

$$\begin{aligned}
 (\theta^{(r)})_{\frac{3}{2}}^{n+1} = (\theta^{(r)})_{\frac{3}{2}}^n + \frac{\Delta t}{(\Delta z)_{\frac{3}{2}}} \left\{ \left[(K_h)_2 \frac{\theta_{\frac{5}{2}}^{(r)} - \theta_{\frac{3}{2}}^{(r)}}{(\Delta z)_2} - (C_h)_{\frac{3}{2}} |\bar{U}_{\frac{3}{2}}| (\theta_{\frac{3}{2}}^{(r)} - \theta_s) \right] - \right. \\
 \left. (w_{\text{sub}})_{\frac{3}{2}} \frac{\theta_{\frac{5}{2}}^{(r)} - \theta_{\frac{3}{2}}^{(r)}}{(\Delta z)_2} \right\}^n - R_c.
 \end{aligned}
 \tag{3.11}$$

During each iteration the values for K_m , K_h , C_m and C_h are updated using the latest values for the model variables. All model variables are sampled every 6 hours, if the maximum change, level by level, for that time period has reached a value less than either 10^{-4}ms^{-1} or 10^{-4}K then the system is considered to have reached its steady state.

3.1 Newton Iteration

In the case of the operational resolution runs the stability criteria will allow for large enough Δt , even on the geometric grid, to use false time-stepping. In fact running with up to 100 grid points is reasonable, however much more than this results in the deepest boundary layers (larger K) taking a large amount of computational time to reach the condition of stability, of the order of 10^7 iterations for 300 grid points. In addition to this the convergence of the solutions with increasing resolution is expected to be poor for the geometric and uniform grids due to the singularity in the derivative at the surface. Resolution would need to be increased to the order of thousands of grid points before negligible difference between Lorenz and Charney-Phillips would be observed. For this level of resolution generating solutions that meet the steady criteria would be highly expensive. Further to this any convergence testing would be unfeasible.

If using the logarithmic grid the convergence between the Lorenz and Charney-Phillips grid will be much faster and thus negligible difference between solutions should be achieved for a reasonable number of grid points, of the order of one

hundred. However the relatively fine spacing of the logarithmic grid spacing means that a very small time step would be required to meet the stability criteria, in turn meaning an unfeasible number of iterations required.

Instead of the false time-stepping a much more efficient method will be required. One such scheme is the Newton method, an iteration that seeks the root of a function. The method works by making a guess at the root, it then uses the derivative of the function to produce a better guess at the root and so on. Provided the original guess lies close to the root and the function does not have lots of local minima and maxima points between the guess and root then it will quickly approach the solution. The conditions of iteration are easy to envisage for say a one or two dimensional function and the same principles extend to as high a dimensional space as required; the dimension of the boundary layer model will depend on the resolution. The iteration is written,

$$\mathbf{x}_{\text{new}} = \mathbf{x} - [\nabla_{\mathbf{x}}\mathbf{F}(\mathbf{x})]^{-1}\mathbf{F}(\mathbf{x}), \quad (3.12)$$

where \mathbf{x} is a vector containing the three model variables at each model level. The vector $\mathbf{F}(\mathbf{x})$ is the right hand side of the discretised steady state equations (i.e. from equations (3.4)-(3.11)) with the latest guess for the model variables. The only complicated part of the Newton method is in obtaining the Jacobian $\nabla_{\mathbf{x}}\mathbf{F}(\mathbf{x})$. Recall however that the linearisation around the steady reference state for some arbitrary function of a model variable is the sum of the reference function and the gradient of the reference function with respect to the variable multiplied by the transient function. Entries in the Jacobian matrix are also the gradients with respect to the variable. In finding the Jacobian the matrix of transient coefficients is also obtained, this is required in the eigendecomposition when the transient part of the equations is considered later. The transient equations, and thus the Jacobian matrix, are obtained by linearising equations (2.44), (2.45) and (2.64) and leaving only terms involving first order transients, giving

$$\frac{\partial u'}{\partial t} = f v' + \frac{\partial}{\partial z} \left(K_m \frac{\partial u'}{\partial z} + K'_m \frac{\partial U}{\partial z} \right), \quad (3.13)$$

$$\frac{\partial v'}{\partial t} = -f u' = \frac{\partial}{\partial z} \left(K_m \frac{\partial v'}{\partial z} + K'_m \frac{\partial V}{\partial z} \right), \quad (3.14)$$

$$\frac{\partial \theta'}{\partial t} = \frac{\partial}{\partial z} \left(K_h \frac{\partial \theta'}{\partial z} + K'_h \frac{\partial \theta^{(r)}}{\partial z} \right) - w_{\text{sub}} \frac{\partial \theta'}{\partial z}. \quad (3.15)$$

Discretised versions of the right hand side of equations (3.13)-(3.15) can be written in matrix form, i.e. \mathbf{Ax} , where \mathbf{x} is a vector containing all transient variables at each model model. The Jacobian matrix \mathbf{A} is the matrix of transient coefficients, $\mathbf{A} = \nabla_{\mathbf{x}} \mathbf{F}(\mathbf{x})$. For use in the Newton method \mathbf{A} is used with the latest guess for the steady variables. The full expansion of the transient equations is shown in the next section.

Some care needs to be used when applying a Newton type iteration for a problem of this kind. The full stably stratified problem represents a nonlinear function. The complex nature of the underlying structure makes it necessary to assist the iteration with a good guess of the solution so that the root can be found. Setting the iteration going with a poor guess will generally lead to the solution diverging. If the complexity of the equations is reduced then it is possible to use the Newton method with simpler, less informed, initial guesses. For example the problem can be reduced to neutral conditions, ($f_m = f_h = 1$) effectively removing any dependence on $\theta^{(r)}$ in the equations. For the neutral problem it is sufficient to choose an initial guess that consists of a constant velocity profile $U = u_g$ and $V = v_g$. With these initial guesses the Newton Method will find the proper steady state profile for the neutral boundary layer. By making a gradual switch from neutral to stably stratified conditions is enough to assist the iteration in ‘homing-in’ on the root of the stably stratified boundary layer equations. The first guess for the $\theta^{(r)}$ profile, which is decoupled initially, is chosen to be exponential, $\theta^{(r)} = \theta_s + (\theta_s - \theta_g) \exp(-\frac{z}{50.0})$. The switching is achieved by placing a factor in front of every occurrence of Richardson or bulk Richardson numbers in the equation set. This factor will begin at 0 and approach 1 as the iteration moves forward, by the time the iteration ends the equations being solved will be the full stably stratified problem. As an example the factor that is used may reach 1 after around 100 iterations, although this may need to be adapted for each problem.

3.2 High Resolution Convergence

Before ensuing with the comparison of the different grid configurations at operational resolution, high resolution solutions are examined for the uniform, geometric and logarithmic grids. When comparing low resolution grids a high resolution solution will be used to place a measure on how well each candidate performs. To be sure that this measure is accurate and the methodology viable there should be negligible difference between Lorenz and Charney-Phillips high resolution solutions, when this occurs the high resolution solutions can be considered the ‘truth’ solution.

The equation set that is being modelled has some exact solution. Although this exact solution is not known analytically it would be theoretically possible to increase resolution until able to capture all scales and thus obtain an exact sampling of the true solution. If the resolution were sufficiently high then this exact sampling would be irrespective of whether using the Lorenz or Charney-Phillips grid, they would be equivalent. The convergence rate of the solution describes how quickly the solution will approach the point where Lorenz and Charney-Phillips are equivalent. If the problem exhibits poor convergence then this situation would be harder to achieve.

If numerically computing the solution of a problem with reasonably large scale structure and bounded derivatives then a stretched grid with around 100 grid points over a domain of 2000m would probably be expected to be enough to obtain good agreement between different configurations and thus achieve the truth solution. Figure 3.1 shows, for 100 grid points, the three model variables as found using the Lorenz and Charney-Phillips grids for the uniform, geometric, logarithmic and log-linear level spacing. The test is with boundary layer 1 so for each plot only the lowest 150m of the domain is shown to highlight the differences.

Studying the four rows of Figure 3.1 it is clear that the uniform grid is far from giving converged solutions. Less clear is the difference between the stretched grids. Figure 3.2 shows the absolute difference between the U s and V s predicted by the Lorenz and Charney-Phillips grids on the geometric and logarithmic grids, i.e. the




Figure 3.1: Plots U , V and $\theta^{(r)}$ on Lorenz and Charney-Phillips grids. The top row shows the uniform grid, the second row the geometrically stretched grid, the third row the logarithmic grid and the bottom row the log-linear hybrid grid. All are for boundary layer 1, with 100 grids points and with only the lowest part of the domain where interesting structure lies shown.

difference between the curves shown in the plots on columns 1 and 2 of rows 2 and 3 in Figure 3.2. Upon closer inspection it is clear that the difference for the geometric grid is approximately an order of magnitude larger than the difference in the logarithmic grid. Furthermore if the resolution is doubled the difference in

the geometric stretching plot is not significantly reduced. Certainly the solutions found using the geometric grid cannot be said to have as good convergence as the logarithmic grid.




Figure 3.2: Differences between U and V in the lowest 150m of the domain for a geometrically stretched grid (left) and logarithmic grid (right). Both are for runs with 100 grid points.

For these tests the shallowest boundary layer is used. This is the least ‘fair’ to the uniform grid, less than 10% of the grid points are within the boundary layer; for a deeper boundary layer it would have a better chance at capturing some of the near-surface structure. The log-linear hybrid grid seems to offer little to improve on the fully logarithmic grid, again however the benefits it affords will be more noticeable for deeper boundary layers. The benefits of using log-linear can be seen in Figure 3.3 which shows the $\theta^{(r)}$ field for boundary layer 5 on logarithmic and log-linear grids, clearly the solutions are closer for the hybrid grid.

The above figures point towards potential issue in using the geometric grid for the high resolution solution. As was discussed in the previous chapter the singularity in the derivative at the surface does indeed appear to be preventing good convergence properties. In order to confirm this full convergence tests need to be performed, this will also highlight the rate of convergence that can be expected and check that the logarithmic grid does indeed improve matters.

In order to perform convergence tests a number of solutions with relatively high resolution need to be obtained and so the Newton method is used for all




Figure 3.3: Solutions to the $\theta^{(r)}$ for boundary layer 5 on the logarithmic and log-linear grids.

cases, uniform, geometric and logarithmic. First a ‘very high resolution’ solution is obtained using 640 grid points and the logarithmic stretching. Convergence rates are obtained by calculating the error between the surface fluxes, τ_{x0} , τ_{y0} and \mathcal{H}_0 , at each lower resolution and the surface fluxes at very high resolution. The lower resolution runs double successively from 10 grid points to 320 grid points.

Figure 3.4 and Figure 3.5 show the convergence rates for uniform, geometric and logarithmic grids when using the Lorenz and Charney-Phillips grids respectively. Included on the plots are lines showing exact first and second order convergence rates for comparison with the actual convergence rates.

With the exception of one or two outliers it is clear from Figure 3.4 and Figure 3.5 that the logarithmic grid provides the most reliable convergence and importantly that it does on all surface fluxes simultaneously. In all three of the surface fluxes for both Lorenz and Charney-Phillips the convergence rate is somewhere between first and second order. For the other two grids there is more overall variation in the convergence rates and the magnitude of the error can be an order of magnitude larger than for the logarithmic grid. The geometrically stretched grid has better convergence for the Lorenz grid than the Charney-Phillips grid which appears first order or worse, seen by comparing the second rows of Figure 3.4 and Figure 3.5. For the geometric grid the staggered levels are smoothly stretched, it may produce even worse convergence results if the staggered levels were placed




Figure 3.4: Convergence rates for the three types of stretching using the Lorenz grid. The uniform grid is shown in the top row, geometrically stretched in the middle row and logarithmic in the bottom row. From left to right shows the dimensional errors in surface fluxes τ_{x0} , τ_{y0} and \mathcal{H}_0 against resolution. On each plot straight lines show exact first (shallower) and second order convergence.

halfway between unstaggered levels [39]. The uniform grid gives first order or worse in all except the surface heat flux and the magnitude of the error is higher than both stretched grids.

In order to obtain accurate solutions to any of the model variables throughout the depth of the boundary layer requires all three surface fluxes to be captured accurately. The logarithmic grid gives reliably faster convergence and error magnitude for all three variables. The other two grids do not and this is why larger errors are seen when comparing Lorenz and Charney-Phillips even with reasonably high resolution solutions and why increasing resolution sees limited reduction in the error.




Figure 3.5: As for Figure 3.4 but using the staggered Charney-Phillips grid.

In all future comparisons the high resolution solution will be obtained using either a logarithmic or log-linear stretching with 100 grid points.

3.3 Lorenz Configuration

Having obtained the high resolution truth solution there are now a number of low resolution configurations to consider and compare. It is first important to examine the differences between the Lorenz and Charney-Phillips staggering, establishing whether a particular configuration is favoured for obtaining the steady state. Further to the comparison between the alternate staggers is the choice for how to arrange the actual levels of the model. So far the use of uniform, geometric and logarithmic grids have been discussed in terms of their ability to obtain ‘truth’ solutions to the equations, of potential interest is also their ability to obtain the

low resolution solution. In discussion of Lorenz versus Charney-Phillips the model levels as used in a previous Met Office operational model are the most immediately relevant. These levels are obtained using a stretching method, similar to the way in which the geometric grid is constructed here. It will be interesting to compare the solutions when using the Met Office levels with those found when using the geometric grid designed here to establish if any improvement is possible with different stretching factors. It will also be interesting to compare the Met Office and geometric grids with low resolution versions of the uniform, logarithmic and hybrid of logarithmic-linear grids.

The Lorenz grid is the easiest to implement since it requires no averaging, so the comparison begins by obtaining solutions for this case. Solutions are obtained using the Met Office grid for boundary layer 1 to boundary layer 5. Figure 3.6 and Figure 3.7 show the steady reference states U , V and $\theta^{(r)}$ for the Lorenz grid, the high resolution solution has also been added to show how well the low resolution performs. The domain is always fixed at 2000m but the figures only show the part of the domain which contains the interesting boundary layer structure.

The shallowest boundary layer, boundary layer 1, is only around 100m deep; for the low resolution Met Office levels this means only two z_ρ levels lie within the boundary layer. Examining the top row plots in Figure 3.6 it is clear that this lack of resolution is fairly damaging to the solution of the velocity equations. The overall depth of the boundary layer is clearly misrepresented. Some of the jet structure in the V field is captured due to a model level coinciding with the strongest part but the jet in the U field is not captured at all. At a height of say 100m the U wind speed is predicted to be around 3ms^{-1} when in actual fact it should be around 4ms^{-1} , no errors of this magnitude occur in V . The boundary layer 1 $\theta^{(r)}$ field is quite poorly captured by the low resolution solution, with a vertical shift of up to a 20m in comparison with the high resolution. With resolution this coarse it is always unlikely that boundary layer structure will be captured, particular features will only be captured if they happen to occur near a grid point.

As the depth of the boundary layer increases and the size of the overall structure




Figure 3.6: Met Office grid solutions for boundary layer 1 to boundary layer 5 (top to bottom) with the Lorenz configuration. High resolution, shown with a dashed line, uses the logarithmic grid. This figure shows reference velocities U (left) and V (right), potential temperature $\theta^{(r)}$ is shown in Figure 3.7.

increases, more grid points come into use and the ability to capture the profiles increases. From boundary layer 2 onwards the U field is captured reasonably well, particularly in the lower part of the boundary layer, although it is not until the




Figure 3.7: Corresponding to Figure 3.6 this figure shows reference potential temperature $\theta^{(r)}$ for the Lorenz grid.

depth of boundary layer 5 that any of the jet structure can be seen in the low resolution. In the V field the deepening boundary layer also sees more of the features captured by the high resolution appearing in the low resolution. $\theta^{(r)}$ is still captured least accurately of the model variables. The overall structure is captured but has a vertical shift. More resolution at the top of the boundary layer could help to more accurately represent the depth of the boundary layer in the potential temperature field.

For the deepest boundary layer where seven of the low resolution levels fall in the boundary layer the low resolution looks reasonably close to the high resolution solution.

3.4 Charney-Phillips Option I (Ri computed at z_w levels)

Now the three Charney-Phillips configurations where Ri is computed on the z_w levels are considered. In all three cases the potential temperature gradient $\theta_z^{(r)}$ is averaged so that Ri is found at z_w levels. The three configurations come from how K_h is computed on the z_ρ levels, either average K_h I-i, average f_h I-ii or average Ri I-iii. The averaging steps taken to obtain K_m and K_h at their relevant levels are shown in Figure 3.8




Figure 3.8: The steps taken to obtain K_m and K_h with the Charney-Phillips option I grids on a section of the grid away from boundaries. The option I-ii and I-iii choices are shown in grey.

Figure 3.9 and Figure 3.10 show the Met office grid solutions for Charney-Phillips option I-i, i.e. with K_h itself averaged to the z_ρ levels. Included for comparison on the plots is the Met Office grid Lorenz solutions as shown in Figure 3.6 and Figure 3.7 and the high resolution solutions.

Despite the averaging involved with the Charney-Phillips grids, the low resolution solutions of the steady state are quite reasonable. Even the shallowest boundary layer, particularly in the velocity fields, appears to be captured quite well by the Charney-Phillips grid. The depth in the U and V fields is as accurate as the Lorenz version and all of the jet in the V field is present. The $\theta^{(r)}$ field is not as good as in the Lorenz case but results are reasonable considering the low resolution.

Despite model levels being equivalent for the velocity fields the averaging has




Figure 3.9: High (dashed) and Lorenz and Charney-Phillips option I-i (K_h itself averaged) solutions for boundary layer 1 to boundary layer 5 (top to bottom). High resolution uses the logarithmic grid, low resolution uses the Met Office stretched grid. This figure shows reference velocities U (left) and V (right), potential temperature $\theta^{(r)}$ is shown in Figure 3.10.

the effect of capturing more of the jet feature than was captured when using the Lorenz grid. For example in boundary layer 4 the jet is quite clear and there are




Figure 3.10: Corresponding to Figure 3.9 this figure shows reference potential temperature $\theta^{(r)}$ for the Charney-Phillips I-i grid.

even signs of it in boundary layer 3.

The low level structure in the U field is more accurate when using the Lorenz grid than when using the Charney-Phillips I-i grid, but vice versa in the V field. Capturing of the jet part of the velocity field is less clear cut. For the Lorenz case the strongest part of the jet in the V field is a little too high up and for the Charney-Phillips case is a little too low but there is only noticeable difference in the deeper boundary layers. Charney-Phillips does slightly better in capturing the U jet.

The Charney-Phillips I-i grid does a better job of capturing the lower part of the potential temperature field than the Lorenz grid did but the compromise is that the top of the boundary layer is smoothed. For example looking at boundary

layer 5, both grids have a vertical shift in the $\theta^{(r)}$ field, but the Lorenz grid solution has greater shift at the lower part of the boundary layer but less of a shift close to the top of the boundary layer. This may enable the Lorenz grid to provide a better approximation to the height of the top of the boundary layer, a useful quantity in boundary layer modelling [47].

Figure 3.11 and Figure 3.12 compare the three sub options I-i, I-ii and I-iii. Again the Lorenz and high resolution solutions have been added to the plots to aid comparison.




Figure 3.11: Plots of boundary layers 1, 3 and 5 for Charney-Phillips option I-i, I-ii and I-iii against Lorenz and high resolution solutions. This figure shows reference velocities U (left) and V (right), potential temperature $\theta^{(r)}$ is shown in Figure 3.12, along with a legend of entries.

Most important to note from Figure 3.11 and Figure 3.12 is that for option I-iii, where Ri is averaged, the $\theta^{(r)}$ field is captured significantly more accurately




Figure 3.12: Corresponding to Figure 3.12 this figure shows reference potential temperature $\theta^{(r)}$ for the Charney-Phillips option I grids.

than, not only the other two Charney-Phillips cases, but also the low resolution Lorenz case. This sub option will therefore be useful to include in the study of the transient evolution.

Interestingly the case where f_h is averaged seems to give solutions in the velocity field that are quite similar to the solutions found by the Lorenz grid. This similarity does not appear to extend to the $\theta^{(r)}$ field where this option offers no benefit over simply averaging K_h .

To further understand the differences between the option I configurations a test case is considered. Model variables U , V and $\theta^{(r)}$ are found by interpolating, using high order cubic interpolation, from a high resolution truth solution to the Met Office grid. The test is shown here for boundary layer 2. From these variables the Richardson number is calculated at the z_w levels by averaging the gradient of the potential temperature. The three sub options are then used to calculate the eddy diffusivity K_h and the results are shown in Figure 3.13. For each configuration the

figure shows the Richardson number, the stability function and K_h ; the height of the boundary layer is determined by seeking the point where K_m goes to zero in the equivalent high resolution solution. Note that there is a different vertical scale for each sub-plot to show the most important features and that the horizontal scale on the stability function plot is adjusted to show clearly where they go to zero.




Figure 3.13: A test case comparing the differences between the three option I Charney-Phillips configurations when calculating eddy diffusivity K_h . In the top left plot is the Richardson number, the bottom left plot shows the stability function and the right plot shows K_h .

Since there is no difference in the way the boundary conditions are implemented between options i-iii the cause for differences must be in the way the outer levels are solved. Option I-i and I-ii grids start with the same Richardson number so only two curves are shown in that plot, the green curve is averaged from the blue curve. In the stability function plot the red curve is averaged from the blue curve. Crosses on the curves note the location of the grid points. It can be seen in the figures that when Richardson number is averaged with option iii the top of the boundary layer is reduced. K_h goes to zero for option I-iii well below the location it goes to zero

for option I-i and I-ii and this is why the reduced height is observed in Figure 3.11 and Figure 3.12. This reduced height allows for the option iii grid to give better overall representation of the upper part of the boundary layer.

3.5 Charney-Phillips Option II (Ri computed at z_ρ levels)

Now steady state solutions are sought for the Charney-Phillips option II staggers to establish whether any benefit is afforded by averaging the shear instead of the potential temperature gradient. For the option II there are four times as many configurations as there were for option I, this is due to the four ways of obtaining the squared shear (a)-(d). For each option on the shear the three options i-iii can be used for obtaining K_m , giving a total of 12 choices. Figure 3.14 shows a selection from the possible steps that are required in order to obtain K_m and K_h when using option II.




Figure 3.14: As for Figure 3.8 but for the Charney Phillips option II grids. The beginning of the alternative options for finding K_m are shown in grey.

To begin with attention is restricted to the option -i case where K_m is averaged rather than f_m or Ri . Figure 3.15 and Figure 3.16 show the three reference model variables as predicted by Charney-Phillips option II(c)-i and II(d)-i. The Lorenz and high resolution solutions are also shown for comparison.

As can be seen from the figures no useful solutions are obtained when using option II(c)-i or II(d)-i. Apart from the fact that the solutions lack the correct




Figure 3.15: The three steady state model variables as found by Charney-Phillips option II(c)-i and option II(d)-i against low resolution Lorenz and high resolution. The solutions are for boundary layer 1 to boundary layer 5 (top to bottom). This figure shows reference velocities U (left) and V (right), potential temperature $\theta^{(r)}$ is shown in Figure 3.16, along with a legend of entries.

structure some of the runs failed to approach a steady solution, whether approximated through false time-stepping or the Newton method. When it became appar-




Figure 3.16: Corresponding to Figure 3.15 this figure shows reference potential temperature $\theta^{(r)}$ for the Charney-Phillips option II grids.

ent that solutions were not likely to approach steady conditions they were halted. For boundary layer 3 neither option II(c)-i or II(d)-i approached the steady state and had to be halted. For boundary layer 2 option II(d)-i had to be halted. One reason that the numerics can fail to approach steady state is due to a nocturnal jet type effect. As first explained in [9] the nocturnal jet results as the depth of the boundary layer decreases at sunset, leaving ageostrophic wind in a region of neutral stratification and so free to undergo inertial oscillation. In some of the runs the large errors resulted in big discontinuities in potential temperature giving regions of neutral stratification and in the cases that were halted an inertial oscillation was observed, implying regions of ageostrophic wind free of the K_m and K_h forcing.

With the exception of boundary layer 2, where option II(c)-i does a reasonable

job, there is a significant issue in the structure of the solutions. For boundary layer 1 the solution becomes negative near the ground in the U field. Other boundary layers develop large discontinuities in the solution, resulting in regions of neutral stratification and in the velocity fields shear is predicted in regions high above the boundary layer. The potential temperature field is almost captured for boundary layer 3 but has a large discontinuity. For the deepest boundary layers the potential temperature is vertically shifted almost down to the surface with huge discontinuities to the top of the boundary layer.

In all there does not appear to be any viable solutions for the Charney-Phillips II grid, all twelve configurations were considered and all found to have some degree of issue. A number of numerical tests have been considered in order to try and understand why these configurations performed badly. It is generally quite difficult to fully establish where the errors occur since they will generally occur very suddenly during in the early stages of the iteration and simultaneously in all the model parameters. There is a clear sensitivity to any error that occurs in the shear and this appears to be due to the more direct way in which shear enters into K_m and K_h . For the option I configurations the averaging is on potential temperature gradient and so this is where errors would occur. However the potential temperature gradient only enters through the Richardson number and the stability function and so likely undergoes a smoothing effect. The errors seen for option II are seen again in the next section examining the option III configurations, further discussion of the problems encountered is presented there.

3.6 Charney-Phillips Option III (Ri computed at all levels)

The final group of Charney-Phillips grids to be considered comes under option III, for these configurations Ri is calculated at all of the model levels. The number of cases to be considered is reduced by the fact that K_m and K_h are automatically found at the level at which they are required, meaning only the options on averaging

shear require consideration, these are sub options (a) to (d). Using the option III grid is equivalent to computing K_m using an option I grid and K_h using an option II grid. The steps taken for options III(c) and III(d) are shown in Figure 3.17.




Figure 3.17: As for Figure 3.8 but for the option III(c) and III(d) grids.

Plots of the low resolution solutions for the five boundary layer depths with Charney-Phillips option III(c) and III(d) are shown in Figure 3.20 and Figure 3.21. Clearly by averaging both shear and potential temperature gradient an improvement is gained on just averaging shear, as in option II. For all five boundary layers the velocity fields are considerably improved over option II, there are no regions of ageostrophic wind above the boundary layer and the predicted structure of the velocity in the boundary layer is reasonably close to the high resolution solution, although not as good as that predicted by the Lorenz or Charney-Phillips option I grids.

The results from using option III give further insight into why the option II grids performed so badly. It is clear that in averaging the shear errors can occur and that the system is quite sensitive to these errors; more so than it is to errors when averaging potential temperature gradient, implied by the good results found when using option I. For the option III grids the sensitivity to the error in averaging shear is seemingly reduced. This is due to the influence of the errors in shear as K_m is calculated and used to update U and V since under option III K_m is calculated without averaging shear. The U and V wind components are not immune from errors in shear but may feel them through the potential temperature which is more directly affected by the error through the shear averaging required to find K_h . Also

note that the predictions of potential temperature are considerably worse than the predictions of the velocity.

Consider the route that any error occurring as a result of averaging shear must take before entering the velocity terms and thus introducing further shear errors. First K_h is calculated using an averaged shear, this is then used to update potential temperature. The gradient of potential temperature gradient is calculated and then averaged, then passed into Ri , then the stability function before entering into K_m , this is then used to update velocity. The error that originates in the shear in K_h is smoothed and suppressed as it passes through this link. Conversely if K_m is calculated more directly from an averaged shear, as it is in option II, then greater sensitivity to the error is seen. Averaging shear is seen to produce discontinuities and errors in the velocity profiles, as the shear is recalculated and averaged the error becomes worse still, feeding back in continuously. A further important difference between the way K_m is calculated in option II and option III is in the way that the averaged shear enters; either directly with option II, or through the stability function via potential temperature with option III. Recall that $K_m = l^2 \left| \frac{\partial \mathbf{u}}{\partial z} \right| f_m(Ri)$, under option II the shear in K_m is averaged and so has a more direct influence, under option III it is not averaged, except in f_m . Similarly for option I the errors due to averaging potential temperature gradient can only influence the solution through the stability function.

The differences between the way averaged shear can enter the model variables in option II and option III is further highlighted in Figure 3.18 and Figure 3.19. These figures show flow diagrams of the different stages of the iteration as the model variables are updated. In both cases the velocities are stored at the z_ρ levels and potential temperature is stored at z_w levels. In Figure 3.18 there is one loop to represent that Richardson number is only computed on z_ρ levels, in Figure 3.19 there are two loops due to Richardson number being computed on both levels. Figure 3.18 shows the stages of the iteration when using Charney-Phillips option II(d)-i. For the II(d)-i configuration averaged shear is computed immediately and then used to compute Richardson number which is in turn used to compute the

stability function and then K_m and K_h , which are then used to update the model variables. Further the averaged shear enters K_m and K_h directly (dashed arrows). If an error occurs in averaging shear then u and v are quite directly affected by the error, potentially leading to even greater error. Conversely consider Figure 3.19, which shows the stages of the iteration when using the option III(d) configuration. Before any error that occurs in the averaged shear can re-enter u and v it must first travel around the K_h side of the iteration and only enters into the K_m side of the iteration through potential temperature. Any error in the shear is likely to be smoothed.




Figure 3.18: Flow diagram demonstrating the process of updating the model variables during iteration. The diagram shows the process when using the Charney-Phillips option II(d)-i grid.

Other than finding that the model is more sensitive to errors in shear than errors in potential temperature gradient it is difficult to pin down exactly where errors in shear arise. Output has been examined in the early stages of the iterations for option II and III. It seems that for some cases errors appear near the ground first then propagate throughout the domain. It is possible to see this error propagation by using the interpolated high resolution solutions to slow down the divergence of the solution. For example by holding either the K_m or K_h side of the iteration constant at values obtained from interpolating the high resolution solution or by allowing them to only vary from these below or above a certain height. It is near the surface that shear varies most strongly so it is where the greatest errors would




Figure 3.19: Flow diagram demonstrating the process of updating the model variables during iteration. The diagram shows the process when using the Charney-Phillips option III(d) grid.

be expected when averaging. A number of numerical test have been considered in an attempt to fully understand why options II and option III are found to give such poor results. It is generally found that iteration will develop significant errors after just a few iterations or very suddenly between iterations. This has made it very difficult to fully investigate the problem and establish how the apparent smoothing of errors through the stability function occurs. The presence of a computational mode in the problem was also investigated. This was done by linearising about a reference state generated for option I but using an option II transient discretisation, a description of the transient calculation follows in the next chapter. No structure with computational mode like behaviour was identified however.

The results from the configurations that have been examined here suggest that when an averaging is required it is better to avoid averaging shear. The sensitivity that the model has to errors in the predicted shear are greater than the sensitivity to errors in potential temperature gradient. Using a shear averaging may result in unpredictable behaviour in the model.

3.7 Comparison of various types of grid spacing

The previous few sections have compared the steady state results for the Met Office operational grid. It also interesting to compare the performance of the




Figure 3.20: The three steady state model variables as found by Charney-Phillips option III(c) and option III(d) against low resolution Lorenz and high resolution. This figure shows reference velocities U (left) and V (right), potential temperature $\theta^{(r)}$ is shown in Figure 3.21, along with a legend of entries.

Met Office grid against the other ways of spacing the model levels. Grids that can be considered in addition to the Met Office grid are: the uniformly spaced grid, a geometric grid, a logarithmic grid and a log-linear grid. Although the




Figure 3.21: Corresponding to Figure 3.20 this figure shows reference potential temperature $\theta^{(r)}$ for the Lorenz grid.

Met Office grid also has a geometric type stretching the one designed here has a different bias so that levels concentrate slightly closer to the ground. Two geometric grids are considered, each with the unstaggered z_w levels at the same place but with the staggered z_ρ levels placed either halfway between unstaggered levels or stretched smoothly at the same rate as the unstaggered levels. The uniform grid is useful in the sense that it is the type of grid used in previous studies such as [68] and may prove useful in later configurations, such as when examining the fully coupled problem. Results from the geometrically stretched grid will ascertain whether bringing more points close to the ground can help; this is also addressed by examining use of the logarithmic grids but with a more radical bias towards the ground.




Figure 3.22: A comparison of steady states predicted by the Lorenz grid whilst using a variety of ways of spacing the model levels. This figure shows reference velocities U (left) and V (right), potential temperature $\theta^{(r)}$ is shown in Figure 3.23, along with a legend of entries.

Figure 3.22 and Figure 3.23 shows the steady states for the five boundary layers as predicted by the Lorenz grid with the variety of grid spacing options. Figure 3.24 and Figure 3.25 shows the steady states for the five boundary layers as predicted




Figure 3.23: Corresponding Figure 3.22 this figure shows reference potential temperature $\theta^{(r)}$ for the variety of grids.

by the Charney-Phillips option I-i grid with the variety of grid spacing options. Finally Figure 3.26 shows the Lorenz grid solutions for the lowest 20m to examine the ability of each grid to capture the near surface structure.

The first and most important point to note from the various plots is that overall the Met Office grid does well at capturing the solutions. It offers a good balance across the cases and across the depth of the boundary layer. Although the grids with more bias towards the ground are slightly better for the very shallow boundary layers the Met Office grid produces the most accurate results for the deepest boundary layers

The uniform grid is never expected to perform well for this kind of problem and indeed for the shallowest boundary layer feels virtually no boundary layer effect. It




Figure 3.24: A comparison of steady states predicted by the Charney-Phillips I-grid whilst using a variety of ways of spacing the model levels. This figure shows reference velocities U (left) and V (right), potential temperature $\theta^{(r)}$ is shown in Figure 3.25, along with a legend of entries.

starts to do a better job of the deepest boundary layer but nowhere near as good as anything else.

There appears to be little difference between the two geometric grids that were




Figure 3.25: Corresponding Figure 3.24 this figure shows reference potential temperature $\theta^{(r)}$ for the variety of grids.

tested, for every boundary layer depth there is little to distinguish between them. The extra resolution near to the ground does mean that the geometric grid is able to capture the shallowest boundary layer accurately; for boundary layer 1 it is only bettered by the logarithmic grid.

Like the geometric grid the logarithmic grid does well for the most shallow boundary layer, however the concentration of grid points near to the ground means structure of deeper boundary layers is generally misrepresented. The log-linear grid is quite good for capturing the large-scale structure of the deeper boundary layers but had issue in predicting the structure of the shallower boundary layers. When using the Lorenz grid with the log-linear spacing the U profile goes negative near the surface. The bias from log to linear can be adjusted and so it may be that the




Figure 3.26: A comparison of the near surface structure as found by the Lorenz grid with the variety of grids. Shown for boundary layer 1 and boundary layer 5

log-linear grid can be made to give better representation, it is unlikely that it can improve much on the Met Office and geometric grids however.

3.8 Concluding Remarks on Steady State Solutions

For the full picture of how well a grid configuration can perform the full linearisation and transient analysis must be considered. Despite this the steady state examination has revealed useful information about the vertical configurations.

Before the steady states were compared the problem had to be modified to suit the methodology. Firstly, in order to use the methodology requires that a completely steady state is obtained, it has been shown that, with a modification, it is possible to achieve this for the boundary layer equations. In order to obtain a completely steady state requires a fixed boundary condition on the potential

temperature, the lack of imposed fluxes means subsidence warming is required to balance the overall cooling. A further requirement of the methodology is that the numerics have good convergence properties so that a ‘truth’ solution can be obtained. It has been shown that a transformation to a logarithmic coordinate is needed in order to obtain good convergence properties.

A selection of stably stratified boundary layers with steady state structures that would commonly be found in the atmosphere have been studied. In general the Lorenz grid is capable of representing the steady state more accurately than any of the Charney-Phillips grids options that were tested. The Charney-Phillips option I had the best results. Indeed for the majority of the cases examined the option II and option III configurations encountered problems. Under option I three configurations were considered, defined by the averaging in the K_h term. The case where the Richardson number was averaged had promising results, particularly in its ability to represent the potential temperature profiles.

In comparing different ways of spacing the model levels in the low resolution runs no configuration has been found to significantly outperform the Met Office grid. The geometric grids could capture the shallowest boundary layers slightly more accurately but were less accurate for the deeper boundary layers. The uniform grid gave the most inaccurate representation of the steady states, but may be useful for the deepest boundary layer if it will assist the transient methodology later. The logarithmic and log-linear grids were not found to be useful for the low resolution since the biasing meant only boundary layers at each extreme were captured well. Boundary layer 1 was captured well by the log grid but boundary layer 5 was not, conversely the log-linear grid captured boundary layer 5 well but not boundary layer 1.

Going into the transient examination it is possible to reduce the number of test cases to consider as some of the configurations have been found to be problematic in some way. Any Charney-Phillips configuration that requires the averaging of shear was prone to error and in some cases unable to approach a proper physical steady state. These configurations can be neglected going into the transient comparison

since if they are unable to capture even the steady state accurately they will be less useful for modelling the stable boundary layer. Configurations of interest moving forward are the Lorenz grid and Charney-Phillips option I-i and I-iii grids.

Chapter 4

Transient Solution

Computing the steady solution gives considerable insight into the way in which the different grid staggering can perform for this type of problem and highlights any unsuitable configurations. The model as a whole evolves in time and the ability of a particular configuration to capture the reference solution does not necessarily extend to describing its ability to represent the whole time dependent problem. In order to further understand the differences between configurations the transient part of the equations needs to be examined. The transient part of the linearisation represents the small amplitude linear departure from the reference state; combining a study of the transients with the steady state study will provide insight into how well a configuration will be able to model the full problem. Note that linearising about a reference state will not reveal information about any nonlinear temporal behaviour.

Examining the transients is a popular methodology for measuring the capabilities of a numerical model and is often done by studying the normal modes of the solution, [68, 16, 66, 67, and references therein]. Normal modes of the solution represent a set of linearly independent solutions; any discrete transient behaviour can be projected onto this set of solutions. As discussed in the introduction to this thesis forcing mechanisms can perturb the atmosphere away from balance and as restoration occurs waves are radiated. These waves that occur due to perturbations are represented by the normal modes. For the dynamics the types of waves that will

be considered are the acoustic, inertio-gravity and Rossby modes. Any model needs to be able to represent the vertical structure of the acoustic and inertio-gravity normal modes so that it can accurately represent the restoration of balance. It also needs to be able to represent Rossby modes since they are energetically dominant. The boundary layer is different to the dynamics in that modal solutions do not just radiate but will be damped, representing the boundary layer diffusion.

When examining individual normal modes the methodology is similar to that used when examining the steady states. Once some normal mode is identified in two competing low resolution configurations it is then compared to the high resolution version of that mode. Again there should be negligible difference between configurations at high resolution. The process of comparing modes is made somewhat more difficult than comparing steady states since the number of normal modes in the solution is related to the number of degrees of freedom in the system and so an order of magnitude more will exist in the high resolution solution.

The number of degrees of freedom in the system is determined by the number of prognostic variables and where they are positioned. Recall from Figure 2.1 that there are $N + 1$ z_w levels and N z_ρ levels, where N is 10 for the Met Office grid and 100 for the high resolution grid. For variables stored on the z_w levels the boundary conditions are implemented at the lowest and highest level, this gives $N - 1$ degrees of freedom for each variable. For the fully coupled problem, that is considered in Part II, prognostic variables stored on z_ρ are free to vary at all levels, giving N degrees of freedom for each variable. If a boundary condition is required on a variable stored on the z_ρ levels it is taken at the top of the roughness length or at a ‘ghost’ level above the domain. Clearly there will be one more degree of freedom when using the Lorenz grid than when using the Charney-Phillips grid due to the extra degree of freedom when storing potential temperature at the z_ρ levels. When considering the boundary layer without the dynamics it is sufficient to allow variables stored on z_ρ levels to only vary on $N - 1$ model levels so that both Lorenz and Charney-Phillips support the same number of degrees of freedom. Ensuring the same number of degrees of freedom across both Lorenz and Charney-Phillips

configurations prevents the need for the potentially arduous task of identifying the mode only supported by the Lorenz grid and removing it from any comparison. In each configuration considered there will be $3(N - 1)$ degrees of freedom. In the high resolution runs there will be an order of magnitude more modes supported and so some methodology will be required for identifying those which correspond to the modes supported by the low resolution configurations.

4.1 Transient Equations and Discretisation

Linearising (2.44), (2.45) and (2.64) about the steady reference state and ignoring terms which contain products of two or more transient terms gives,

$$\frac{\partial u'}{\partial t} = f v' + \frac{\partial \tau'_x}{\partial z}, \quad (4.1)$$

$$\frac{\partial v'}{\partial t} = -f u' + \frac{\partial \tau'_y}{\partial z}, \quad (4.2)$$

$$\frac{\partial \theta'}{\partial t} = \frac{\partial \mathcal{H}'}{\partial z} - w_{\text{sub}} \frac{\partial \theta'}{\partial z}. \quad (4.3)$$

where,

$$\tau'_x = K_m \frac{\partial u'}{\partial z} + K'_m \frac{\partial U}{\partial z}, \quad (4.4)$$

$$\mathcal{H}' = K_h \frac{\partial \theta'}{\partial z} + K'_h \frac{\partial \theta^{(r)}}{\partial z}. \quad (4.5)$$

The transient eddy viscosity and diffusivity are given by,

$$K'_{\{m,h\}} = \frac{l^2}{\left| \frac{\partial \mathbf{U}}{\partial z} \right|} \left(\frac{\partial U}{\partial z} \frac{\partial u'}{\partial z} + \frac{\partial V}{\partial z} \frac{\partial v'}{\partial z} \right) f_{\{m,h\}}(Ri) + l^2 \left| \frac{\partial \mathbf{U}}{\partial z} \right| \frac{\partial f_{\{m,h\}}}{\partial Ri} \left[g \frac{\frac{\partial \left(\frac{\theta'}{\theta^{(r)}} \right)}{\partial z}}{\left| \frac{\partial \mathbf{U}}{\partial z} \right|^2} - 2g \frac{\frac{\partial \theta^{(r)}}{\partial z}}{\left| \frac{\partial \mathbf{U}}{\partial z} \right|^4} \left(\frac{\partial U}{\partial z} \frac{\partial u'}{\partial z} + \frac{\partial V}{\partial z} \frac{\partial v'}{\partial z} \right) \right], \quad (4.6)$$

the square bracketed part of the right hand side of (4.6) represents Ri' .

The spatial discretisation of equations (4.1) to (4.3) away from the boundaries is analogous to the discretisation of the steady state part of the equations. For the

lowest model level the spatial discretisation is implemented as,

$$\left(\frac{\partial u'}{\partial t}\right)_1 = f v'_1 + \frac{\partial}{\partial z} \left[\left(K_m \frac{\partial u'}{\partial z} \right)_{\frac{3}{2}} - C_{mn} (f_m(Ri_b)|\mathbf{U}_1|u'_1 + f_m(Ri_b)|\mathbf{u}_1|'U_1 + \frac{\partial f_m}{\partial Ri_b} Ri'_b|\mathbf{U}_1|U_1) \right], \quad (4.7)$$

$$\left(\frac{\partial v'}{\partial t}\right)_1 = -f u'_1 + \frac{\partial}{\partial z} \left[\left(K_m \frac{\partial v'}{\partial z} \right)_{\frac{3}{2}} - C_{mn} (f_m(Ri_b)|\mathbf{U}_1|v'_1 + f_m(Ri_b)|\mathbf{u}_1|'V_1 + \frac{\partial f_m}{\partial Ri_b} Ri'_b|\mathbf{U}_1|V_1) \right], \quad (4.8)$$

$$\left(\frac{\partial \theta'}{\partial t}\right)_j = \frac{\partial}{\partial z} \left[\left(K_h \frac{\partial \theta'}{\partial z} \right)_{j+\frac{1}{2}} - C_{hn} \left(f_m(Ri_b)|\mathbf{U}_j|\theta'_j + f_m(Ri_b)|\mathbf{u}_j|'(\theta_j^{(r)} - \theta_s) + \frac{\partial f_m}{\partial Ri_b} Ri'_b|\mathbf{U}_j|(\theta_j - \theta_s) \right) \right] - \left(w_{\text{sub}} \frac{\partial \theta'}{\partial z} \right)_j. \quad (4.9)$$

The j in equation (4.9) is either 1 for the Lorenz grid or $\frac{3}{2}$ for the Charney-Phillips grid. For the Charney-Phillips grid all terms in equation (4.9) involving velocities also need to be averaged to the $z_{\frac{3}{2}}$ level, similarly all terms in equations (4.7) and (4.8) involving potential temperature would need to be averaged to the z_1 level. The transient bulk Richardson number is,

$$(Ri'_b)_j = g(z_j - z_r) \left[\frac{\frac{\theta'_j}{\theta_j^{(r)}}}{|\mathbf{U}_j|^2} - 2 \frac{\log(\theta_j^{(r)}) - \log(\theta_s)}{|\mathbf{U}_j|^4} (U_j u'_j + V_j v'_j) \right]; \quad (4.10)$$

again j will be either 1 for the Lorenz grid or $\frac{3}{2}$ for the Charney-Phillips grid with the relevant averaging on the terms not stored at j . Note that all Δz terms are replaced by the factor $(z_j - z_r)$ in the bulk Richardson number.

Coefficients of all the transient terms in the discrete versions of (4.1)-(4.3) can be gathered into a matrix, allowing the problem to be written as,

$$\dot{\mathbf{x}} = \mathbf{A} \mathbf{x}. \quad (4.11)$$

The transient solutions have no horizontal dependency for the boundary layer only problem and so solutions of the form $\exp(\lambda t)$ are sought, where $\lambda = \mu - i\omega$. Seeking wave like solutions leads to equation (1.2). The matrix \mathbf{A} is decomposed using the eigendecomposition to give the eigenvalues and eigenvectors. Since the problem is one of diffusion and there is no horizontal mechanism the real part of all of the eigenvalues is negative, denoting decay. Some modes may be propagating and so will have an imaginary part denoting the frequency; in general the frequencies will be small in comparison to the damping rates. A number of neutral inertial modes with frequency $\pm f$ are possible; for these modes the imaginary part of the eigenvalue dominates the real part.

4.2 Matrix Normality

The eigenvectors, or modes, of the matrix \mathbf{A} represent the transient part of the solution. Any linear time dependant behaviour may be projected onto combinations of these modes. It is not, however, always possible to easily interpret this linear behaviour by examining each individual mode. The level of interpretation will depend on the orthogonality between the modes, which depends on the form of the matrix \mathbf{A} , which in turn depends on the complexity of the equations.

For an appropriate equation set, such as that studied by [68], the matrix \mathbf{A} is normal, that is to say $\mathbf{A}^*\mathbf{A} = \mathbf{A}\mathbf{A}^*$, where \mathbf{A}^* is the conjugate transpose. When a matrix is normal its eigenvectors are all orthogonal and are referred to as the normal modes of the system. That the modes are orthogonal means that each individual mode must represent a certain scale in the problem and so can be easily interpreted and understood quantitatively. For the boundary layer only problem being examined here it is found that the matrix \mathbf{A} is far from normal. The result is that the modes are not orthogonal. When modes are not orthogonal it makes interpreting their individual contribution considerably more difficult [71]. This is demonstrated in Figure 4.1 which shows successive eigenvectors from the matrix \mathbf{A} , generated for a high resolution run of boundary layer 1.

In order to generate Figure 4.1 each eigenvector is ordered by the real part of




Figure 4.1: This figure shows the real part of u' for some successive eigenvectors using a stretched Lorenz grid to model boundary layer 1. In red is one type of behaviour for 13 successive modes, in blue shows the behaviour that the system then flips to for the next 25 modes.

its corresponding eigenvalue, so neighbouring modes have similar decay rates. If the modes were normal modes then each successive mode would have a particular scale and structure associated with it, as is clear from the figure this is not the case. Many modes in sequence can have very similar overall scale and structure.

Lack of orthogonality does not mean that modes do not represent the discrete behaviour, only that they do so in a way that makes it impossible to quantify the behaviour of the system by looking at individual modes. Further to this it is not possible to easily compare modes from different configurations. For example, the thirteen modes plotted in red in Figure 4.1 all look very similar, if the Charney-Phillips grid also supported a similar family of modes then it would not possible to know which modes corresponded to each other. The problems, due to lack of orthogonality, can be summarised as:

- It is not possible to interpret the contribution represented by an individual mode and thus know which modes are most relevant for comparison.

- It is not possible to know which modes correspond to each other across differing configurations or resolutions.

4.2.1 The form of Matrix \mathbf{A}

One may reduce the complexity of the problem in order to obtain normality in \mathbf{A} , for example by considering the Ekman layer on a uniform grid. The equations for the Ekman layer are,

$$\frac{\partial u}{\partial t} = fv + K_m \frac{\partial^2 u}{\partial z^2}, \quad (4.12)$$

$$\frac{\partial v}{\partial t} = -fu + K_m \frac{\partial^2 v}{\partial z^2}, \quad (4.13)$$

where K_m is temporarily considered constant with height. If equations (4.12) and (4.13) are discretised and written in matrix form $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}$ then it would take the form,

$$\begin{pmatrix} \dot{\mathbf{u}} \\ \dot{\mathbf{v}} \end{pmatrix} = \begin{pmatrix} \mathbf{\Psi} & \mathbf{\Phi} \\ -\mathbf{\Phi} & \mathbf{\Psi} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}. \quad (4.14)$$

Vectors \mathbf{u} and \mathbf{v} contain the model variables at each level. The matrix $\mathbf{\Phi}$ is diagonal with f on the diagonal. $\mathbf{\Psi}$ is a matrix containing the coefficients of the discretised form of $K_m \frac{\partial^2}{\partial z^2}$. If a centred difference approximation is used for the derivative and the grid is spaced uniformly then $\mathbf{\Psi}$ is a symmetric Toeplitz matrix. For this simplified model the conjugate transpose of \mathbf{A} is,

$$\mathbf{A}^* = \begin{pmatrix} \mathbf{\Psi} & -\mathbf{\Phi} \\ \mathbf{\Phi} & \mathbf{\Psi} \end{pmatrix}. \quad (4.15)$$

It then follows that,

$$\mathbf{A}\mathbf{A}^* = \mathbf{A}^*\mathbf{A} = \begin{pmatrix} \mathbf{\Psi}^2 + \mathbf{\Phi}^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{\Psi}^2 + \mathbf{\Phi}^2 \end{pmatrix}, \quad (4.16)$$

meaning that \mathbf{A} is normal.

If the grid is then stretched the matrix $\mathbf{\Psi}$ is neither symmetric or Toeplitz. As

a result of the increased complexity and the loss of symmetry in Ψ the similarity between the matrix and its conjugate transpose, as in (4.14) and (4.15), is lost. In turn the form of $\mathbf{A}\mathbf{A}^*$ becomes more complicated and different from $\mathbf{A}^*\mathbf{A}$, meaning that normality in \mathbf{A} has been lost. Further to this, in any realistic situation K_m is a function of height so a uniform grid will not necessarily provide normality in \mathbf{A} .

For the full boundary layer equations a further equation for potential temperature is required and each model variable is dependant on the other two model variables through the Richardson number. The matrix \mathbf{A} can be thought of as consisting of nine parts, not four as in (4.14), and none of the parts have the simple form that would be required to see normality in \mathbf{A} itself.

4.3 Singular Value Decomposition (SVD)

The reason that using using the eigendecomposition fails to offer easy interpretation of the behaviour of the system is due to the eigenvectors not forming an orthogonal basis. Another type of matrix decomposition is the Singular Value Decomposition (SVD), which produces left (output) and right (input) singular vectors rather than eigenvectors, and singular values rather than eigenvalues. Singular vectors always form an orthogonal basis [71].

Instead of equation (1.2) the decomposition produces solutions to an equation of the form,

$$\sigma \mathbf{u} = \mathbf{A} \mathbf{v}, \quad (4.17)$$

where \mathbf{u} is the left singular vector, \mathbf{v} is the right singular vector and σ is the corresponding singular value. The decomposition is,

$$\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*. \quad (4.18)$$

Columns of \mathbf{U} are magnitude 1 and give the output orthonormal basis \mathbf{u} , columns of \mathbf{V} also have magnitude 1 and give the input orthonormal basis \mathbf{v} . $\mathbf{\Sigma}$ is a diagonal matrix whose entries are σ .

Informally the action of a general matrix on a general vector can be thought of

as rotation and stretching. Matrix decompositions break this action down so as to represent it as a rotation followed by a stretching followed by another rotation. In the eigendecomposition the eigenvectors have the special property that they are the vectors that are only magnified when multiplied by the matrix, the first rotation is cancelled by the second. The magnification is given in the eigenvalue.

The SVD is different to the eigendecomposition. Rather than just the eigenvector basis, it has both an input and an output basis and the two have to be considered together. When multiplying singular vectors by \mathbf{V}^* , then $\mathbf{\Sigma}$ and then \mathbf{U} the rotations by \mathbf{V}^* and \mathbf{U} do not necessarily cancel each other out. Physically when the input singular vectors are multiplied by the matrix from which they are derived they are rotated and stretched to be the equivalent to the output singular vectors multiplied by the corresponding singular values. The additional property is that the output singular vector is the vector that when multiplied by the conjugate transpose of the original matrix is equal to the input singular vector multiplied by the corresponding singular value.

In the special case that the matrix is normal, Hermitian and sign-definite the rotations by \mathbf{V}^* and \mathbf{U} will cancel each other out when multiplying a singular vector. In this special case the SVD and eigendecomposition are equivalent since the singular vectors are also the vectors that are only stretched when multiplied by \mathbf{A} .

Formally the SVD can be related to the eigendecomposition by noting that,

$$\begin{aligned}\mathbf{A}\mathbf{A}^* &= \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*\mathbf{V}\mathbf{\Sigma}^*\mathbf{U}^* \\ &= \mathbf{U}(\mathbf{\Sigma}\mathbf{\Sigma}^*)\mathbf{U}^*\end{aligned}\tag{4.19}$$

and

$$\mathbf{A}^*\mathbf{A} = \mathbf{V}(\mathbf{\Sigma}\mathbf{\Sigma}^*)\mathbf{V}^*.\tag{4.20}$$

The left singular vectors are equal to the eigenvectors of $\mathbf{A}\mathbf{A}^*$ and the right singular vectors are equal to the eigenvectors of $\mathbf{A}^*\mathbf{A}$. The non-zero singular values

are equal to the square root of the non-zero eigenvalues in either case.

Due to the algorithms used to compute the SVD the singular values are automatically sorted by magnitude. Examining the corresponding singular vectors with increasing singular value magnitude, equivalent to the structures plotted in Figure 4.1, finds behaviour that changes more consistently than was found in the eigenvectors. The behaviour is more representative of the physical behaviour that would be expected. Firstly the singular vector with smallest scale corresponds to the largest singular value. Secondly the problem seen in the eigenvectors, of sequences of similar behaviour, is not seen in the singular vectors. Instead a decrease in singular vector scale with increase in corresponding singular value is seen, as might be hoped for.

Seeking solutions of the form $\exp(\lambda t)$ leads to an equation of the form $\lambda \mathbf{x} = \mathbf{A}\mathbf{x}$, solutions of this equation are not directly given by the SVD. The use of the SVD therefore depends on whether a physical interpretation can be applied to its components. The SVD is widely used in meteorological and mathematical applications, particularly in understanding the growth of errors in initialisations e.g. [49, 25, 53]. The SVD has also been used in geophysical applications e.g [72]. In previous applications the SVD is used to understand the behaviour of the matrix \mathbf{A} , where \mathbf{A} represents some dynamical system, for example a global atmospheric model. This is effectively what is required here. However, the ability of the SVD to relay useful information about the problem being examined here still needs to be explored. Ideally the SVD should represent the behaviour as it is represented by the eigendecomposition but in an orthogonal way. The extent to which interpretation of the SVD is similar to interpretation of the eigendecomposition will depend on the form of the matrix. It will be shown in the next section that by choosing an appropriate norm the form of the matrix can be manipulated to make the SVD and eigendecomposition as close to each other as possible. It is also shown that in transforming to the appropriate norm the necessary physical interpretation of the SVD is provided.

4.3.1 Energetics and The Energy-Norm

When decomposing a matrix with the eigendecomposition the form of the eigenvectors depends on the choice of norm, however the value of the eigenvalues do not. Likewise for the SVD, the singular vectors are dependent on the choice of norm. For the SVD however the singular values are also dependent on the choice of norm.

The method of SVD can be further related to the eigendecomposition by choosing the appropriate norm, effectively a transformation of the model variables. Doing this transformation should allow for easier physical interpretation of the components of the SVD. In this section the energetics of the system are examined and it is shown that by placing the model in the energy norm a link is created between the SVD and eigendecomposition. The energy norm is often used when applying SVD analysis to error growth, see for example [24, and references therein]. The link between the two decompositions is due to the matrix \mathbf{A} being normal and Hermitian when the system conserves energy and is written in the energy norm. First an equation describing the conserved energy in the system needs to be derived, it is then shown that in transforming the variables so as to describe the energy in the system the matrix of coefficients becomes normal and Hermitian.

It should be noted that the boundary layer equations that are being modelled are not energy conserving; however the basic Boussinesq equations from which they have been derived are. Despite the system not being energy conserving it is nevertheless useful to use the energy norm and provide the link between matrix decompositions. Neglecting the boundary layer terms and replacing the horizontal pressure gradient terms, the basic energy conserving Boussinesq equations, in their linearized form are,

$$\frac{\partial u}{\partial t} - fv + \frac{1}{\rho_0} \frac{\partial p}{\partial x} = 0, \quad (4.21)$$

$$\frac{\partial v}{\partial t} + fu + \frac{1}{\rho_0} \frac{\partial p}{\partial y} = 0, \quad (4.22)$$

$$\frac{\partial w}{\partial t} + \frac{1}{\rho_0} \frac{\partial p}{\partial z} = b, \quad (4.23)$$

$$\frac{\partial b}{\partial t} + wN_b^2 = 0, \quad (4.24)$$

$$\nabla \cdot \mathbf{u} = 0. \quad (4.25)$$

To aid in the derivation of the energy, equations (4.23) and (4.24) are written in terms of the buoyancy b and Brunt–Väisälä frequency N_b , these can be related back to the potential temperature using, $b = \frac{g\theta'}{\theta_0}$ and $N_b^2 = \frac{g}{\theta_0} \frac{\partial \theta}{\partial z}$. θ_0 is the reference potential temperature. Combining the above equations as, $u \times (4.21) + v \times (4.22) + w \times (4.23) + \frac{\rho_0 b}{N_b^2} \times (4.24) + p \times (4.25)$ gives the energy conservation principle for the system,

$$\rho_0 \frac{\partial e}{\partial t} + \nabla \cdot (\mathbf{u}p) = 0, \quad (4.26)$$

where $e = \frac{u^2}{2} + \frac{v^2}{2} + \frac{w^2}{2} + \frac{b^2}{2N_b^2}$. The total energy contained in the system can thus be found by computing the volume integral of e . On the discrete Lorenz grid the energy integral for the three original model variables is approximated by,

$$E = \sum_{j=1}^N \left[(\Delta z)_j \left(\frac{u_j^2}{2} \right) + (\Delta z)_j \left(\frac{v_j^2}{2} \right) + (\Delta z)_{j+\frac{1}{2}} \left(\frac{w_{j+\frac{1}{2}}^2}{2} \right) + (\Delta z)_j \left(\frac{g^2 \theta_j^2}{2\theta_0^2 (N_b)_j^2} \right) \right]. \quad (4.27)$$

Note that w is zero for the case being considered here.

Now note that equation (4.27) can be written as,

$$E = \frac{1}{2} \mathbf{x}^* \mathbf{E} \mathbf{x} \quad (4.28)$$

where $\mathbf{x} = (u, v, \theta)^T$ is the state vector, \mathbf{x}^* is its conjugate transpose and \mathbf{E} is a diagonal matrix of the form,

$$\mathbf{E} = \begin{pmatrix} \Delta z & 0 & 0 \\ 0 & \Delta z & 0 \\ 0 & 0 & \Delta z \frac{g^2}{\theta_0^2 N_b^2} \end{pmatrix}. \quad (4.29)$$

To write in the energy norm the model variables are transformed so that $\mathbf{E} = \mathbf{I}$. Introducing $\mathbf{s} = \left(\Delta z^{\frac{1}{2}} u, \Delta z^{\frac{1}{2}} v, \Delta z^{\frac{1}{2}} \frac{g}{\theta_0 N_b} \theta \right)^T$ to denote the vector of transformed

model variables the energy is then obtained as,

$$E = \frac{1}{2} \mathbf{s}^* \mathbf{I} \mathbf{s}. \quad (4.30)$$

Now it can be shown that the matrix of coefficients that corresponds to \mathbf{s} is normal. For a system that conserves energy $\frac{dE}{dt} = 0$, which written in matrix form is,

$$\frac{d}{dt} \left(\frac{1}{2} \mathbf{x}^* \mathbf{E} \mathbf{x} \right) = 0, \quad (4.31)$$

$$\implies \frac{1}{2} \left[\frac{d\mathbf{x}^*}{dt} \mathbf{E} \mathbf{x} + \mathbf{x}^* \mathbf{E} \frac{d\mathbf{x}}{dt} \right] = 0. \quad (4.32)$$

Substituting the time derivatives using (4.11) gives,

$$\frac{1}{2} [\mathbf{x}^* \mathbf{A}^* \mathbf{E} \mathbf{x} + \mathbf{x}^* \mathbf{E} \mathbf{A} \mathbf{x}] = 0, \quad (4.33)$$

$$\implies \frac{1}{2} [\mathbf{x}^* (\mathbf{A}^* \mathbf{E} + \mathbf{E} \mathbf{A}) \mathbf{x}] = 0. \quad (4.34)$$

Now if the state variables are transformed such that $\mathbf{E} = \mathbf{I}$ then the only way for equation (4.34) to be satisfied is if $\mathbf{A}^* + \mathbf{A} = 0$, or $\mathbf{A} = -\mathbf{A}^*$. If $\mathbf{A} = -\mathbf{A}^*$ the matrix is said to be skew-Hermitian. This also gives that $\mathbf{A}^* \mathbf{A} = \mathbf{A} \mathbf{A}^*$ meaning \mathbf{A} is normal. In summary, for a system that conserves energy, if the state variables are transformed so as to give the identity matrix in (4.28) then the resulting matrix of coefficients \mathbf{A} will be normal and skew-Hermitian. Although this does not guarantee a sign-definite matrix it will likely produce a situation where close similarity between the SVD and the eigendecomposition is seen, making physical interpretation of the SVD more readily available. The similarity between the SVD and the eigendecomposition for the energy conserving isothermal resting state, dynamics only system is examined in Part II.

The transformed state vector can be related to the original state vector by a matrix multiplication,

$$\mathbf{s} = \mathbf{B} \mathbf{x}. \quad (4.35)$$

Since $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}$ and $\mathbf{x} = \mathbf{B}^{-1}\mathbf{s}$ it follows that $\dot{\mathbf{s}} = \mathbf{B}\dot{\mathbf{x}} = \mathbf{B}\mathbf{A}\mathbf{x} = \mathbf{B}\mathbf{A}\mathbf{B}^{-1}\mathbf{s}$. Written in the energy norm the system is thus,

$$\dot{\mathbf{s}} = \mathbf{C}\mathbf{s}, \quad (4.36)$$

where $\mathbf{C} = \mathbf{B}\mathbf{A}\mathbf{B}^{-1}$. The matrix \mathbf{B} is just a conversion matrix, for this case it is diagonal and the entries are the square roots of the coefficients of the new model variables in (4.27), so ($\mathbf{B}^2 = \mathbf{E}$).

As for the eigenvalues the singular values have dimension of time. Once in the energy norm the singular values have a similar overall magnitude to the real part of the eigenvalues, this is examined in the next section. Leading modes are the most physically relevant modes, they are the longest lived decaying modes that have structure lying entirely within the boundary layer region. By examining the corresponding singular value and the overall structure, a leading singular vector is identified and is plotted in Figure 4.2. Despite the system not conserving energy the improved physical representation that is afforded by generating the decomposition from \mathbf{C} rather than \mathbf{A} can be seen. The figure shows the input singular vectors on the bottom row and the output singular vectors on the top row. Singular vectors computed from the energy norm form are scaled back to give just the model variables. From left to right the u' , v' and θ' parts of the singular vector are shown. For θ' the singular vector computed from the transformed variables has all of its structure lying in the boundary layer, whereas in the standard norm the behaviour can dominate above the boundary layer. For a leading boundary layer mode the structure would be expected to be dominant in the boundary layer region so clearly a simpler physical interpretation is afforded by writing in the problem in the energy norm to calculate the SVD. For these singular vectors it is mainly the θ' part that benefits from the energy norm formulation. Both velocity variables are significantly rescaled by the energy norm formulation, however both remain having similar magnitude to each other. Also note that the input and output versions of θ' are quite different in the standard norm, whereas in the energy norm the overall input and output structure is closer, making the behaviour easier to interpret.




Figure 4.2: A comparison of a leading singular vector in the standard and energy norms.

4.3.2 Form of the Spectrum

An important consideration when using this kind of methodology for comparing configurations is in the form of the spectrum. The spectrum represents the range of scales supported by the continuous system. Depending on the system that is being examined this spectrum of solutions may be either discrete, continuous or exhibit a combination of the two. For the isothermal reference state linearised Euler equations, as considered by [68], it is clear that only a discrete set of transient solutions are possible. Solutions are sinusoidal and limited by boundary conditions on a finite domain; an infinite but discrete set of possible solutions exist. When the eigendecomposition is computed for the discretised equations it will sample every solution in the discrete set up to the scale allowed for by the resolution. Among the types of fluid flow that are known to exhibit a continuous part to the spectrum is boundary layer flows in an infinite domain [19, and references therein]. That these

flows have continuous spectrum is due to there existing an infinite and continuous number of ways that transient solutions that can satisfy the boundary conditions. Any discretisation of the equations results in a sampling of the continuous spectrum; this sampling is represented in the eigendecomposition. When the discrete equations are perturbed, as they will be when comparing two different configurations or resolutions, it is possible for the sampling of the continuous spectrum to completely change.

It is not clear whether a boundary layer type flow within a bounded domain, as being examined here, should exhibit a wholly discrete spectrum or whether it can support continuous spectrum. Due to the potential difficulty presented it is beyond the scope of this thesis to determine rigorously whether or not the spectrum is wholly discrete. Indeed there are two distinct situations considered here, boundary layer on its own and boundary layer coupled to the dynamics, the form of the spectrum would need to be derived for both. If the spectrum of solutions is continuous then it would be considerably more difficult to compare configurations using this methodology.

Although it would be difficult to determine rigorously whether the spectrum is discrete or not for the stably stratified boundary layer equations if the problem is reduced to the solvable Ekman equations it can be shown that the spectrum is discrete, e.g. [34]. Of course the stably stratified boundary layer has more complicated processes occurring due to the stratification and so it is not clear that the analysis should extend from the Ekman layer to this case. The later inclusion of the dynamics further complicates matters. It could be perceived, for example, that dynamical waves of any wavelength may be absorbed or distorted by the boundary layer in a way that would result in a continuous spectrum of solutions.

It seems likely that if the coupled problem exhibits a discrete spectrum the boundary layer only problem also should. The reference state would be similar for both and the transient boundary layer terms are not altered by the presence of the extra dynamics terms. On this reasoning a more detailed examination of the form of the spectrum is postponed here but performed in Part II for the coupled case.

For the remaining studies in Part I it will be assumed that the spectrum of solutions is discrete. It is important to be aware of the possibility of continuous spectrum and this will be taken into consideration throughout the study and comparison of the boundary layer modes.

4.4 Testing Methodology

The methodology that is usually employed in order to examine the transient behaviour is the normal mode analysis. Due to the complexities contained in the boundary layer equations the mode analysis is unlikely to be able to represent easily interpretable transient behaviour in the system. Mode analysis relies on the eigenvectors forming an orthogonal basis, which, due to the stratification and grid stretching, is not the case. A decomposition that does produce vectors with an orthogonal basis is the SVD. However it is not clear the extent to which the SVD is representative of the equation's transient solution. For the eigendecomposition there is a clear relationship between the decomposition and the solution to the transient equations. Although the SVD does describe the behaviour of \mathbf{A} for the SVD there is no clear relationship between the transient solution and the decomposition. The SVD produces singular vectors, which are not the vectors which are only scaled by the matrix. Instead they exhibit a more complex property represented by a relation between input and output singular vectors. Nevertheless the two methods are related. The eigendecomposition can be considered a special case of the SVD; when the matrix is normal, Hermitian and sign-definite the two are equivalent. Numerical testing will be required to fully establish whether or not the SVD will be able to represent the transient behaviour in an easily interpretable way and thus be useful for comparing different vertical configurations. It has been argued that by writing the system in the energy norm provides further relation between the SVD and eigendecomposition and so this will be used to maximise the SVD's potential.

Due to the problems with the methodology neither decomposition should be completely relied on for understanding and comparing the transient behaviour.

Instead both must be used to complement each other. For example if a certain level of agreement exists between eigenvalues and singular values then confidence is gained in extracting information from singular values, if there is no correspondence between them it is unlikely they can be used. Similarly if the singular vectors are found to reveal behaviour similar in nature to the behaviour that would be expected in the eigenvectors then confidence is gained in their use.

With the various options on averaged terms when using the Charney-Phillips grid as well as the different ways of spacing the model levels a large number of test cases have been presented. Considering the potential difficulties in comparing transients, only those configurations that are most relevant, or are most promising, will be considered. The Met Office grid is most relevant in terms of spacing the model levels since it is taken from an operational model. The Charney-Phillips configurations that gave the most accurate results for the steady state part of the investigation were option I-i and I-iii, these were both configurations which averaged potential temperature gradient and then averaged either K_h or Richardson number. The focus of the transient investigation will thus be on these two cases, option I-i and I-iii with the Met Office stretching.

Further to the choice of vertical configuration there are five boundary layer depths to consider. Due to the limited number of grid points close to the ground it is likely to be easiest to compare singular vectors for the deepest boundary layers, where increased boundary layer structure will be resolved. The low resolution Met Office grid only has one model point inside the shallowest boundary layer meaning virtually no structure will be seen in the singular vectors; it will therefore be difficult to match low and high resolution singular vectors.

The following three sections present the transient investigation. Firstly the methodologies are examined in detail in order to determine the physical interpretation that each can offer. Secondly the high resolution solution is examined. This is to examine the convergence properties between the high resolution configurations. Finally, with all the methodology carefully examined, a comparison is performed for the configurations at operational resolution.

4.5 Physical Interpretation of Eigenvectors and Singular Vectors

In many well behaved dynamics-only cases such as that described by [68] it is possible to identify and classify the types of modes analytically. For example their system supports a definite number of acoustic, inertio-gravity and Rossby modes. For the boundary layer equations no analytical solution of the dispersion relation is currently known, instead classifying modes is achieved by examining the numerical results. By determining the value of the eigenvalue and singular value as well as examining the structure of eigenvectors or singular vectors classifying eigenmodes should be possible.

Figure 4.3 shows plots, in the complex plane, of the eigenvalues of \mathbf{A} for a logarithmic grid high resolution run of boundary layer 5. The first plot shows all of the eigenvalues, the second and third show the eigenvalues with smaller magnitude real part in more detail.

Firstly note that the real parts of all the eigenvalues are negative, meaning that all modes are damped; any positive eigenvalues would represent unsteady growing modes. The magnitude of the real part of an eigenvalue can thus be considered as describing how fast the corresponding eigenmode is damped. The eigenvalues with the largest absolute real part are the fastest damped and so should correspond to the eigenmodes with smallest scale. For these modes the evolution is dominated by diffusion rather than propagation, so the imaginary part is zero. Eigenvalues with smaller absolute real part correspond to modes that are not damped so quickly. Some of the slower damped modes have an imaginary part to the eigenvalue denoting that the corresponding eigenmode propagates, these approximately occur in conjugate pairs implying similar modes propagating in opposite directions. Generally since the real part of the eigenvalue dominates; for these it is more important that a configuration be able to accurately represent decay rather than the rate of propagation. If there were errors in propagation it would be unlikely to cause problems before the mode was damped away.




Figure 4.3: Complex plane plots of the eigenvalues for high resolution boundary layer 5. Successive plots zoom into the eigenvalues with smaller magnitude real part.

For this resolution there exist a total of 16 (8 pairs) eigenvalues whose imaginary part is very close to $\pm f$ the Coriolis parameter, these correspond to inertial modes living entirely above the boundary layer. The eigenvalues corresponding to the inertial modes can be seen at the right hand side of the lowest plot in Figure 4.3, although they all lie on top of each other. The real part of these modes is effectively zero and so they propagate with frequency f without being damped. The number of inertial modes is a consequence of the number of grid points that lie in the region above the boundary layer; if a log-linear grid were used the emphasis would shift, giving more inertial modes and fewer boundary layer modes.

Some modes having very small-scale structure are possible due to the fine spacing caused by using a logarithmic grid, these modes will be damped very fast and

so the real part of the corresponding eigenvalue has large absolute value. That different scales have different rates of damping is due to relative size of the gradients in the transient diffusion terms, structures with small-scale have larger gradients and are thus damped faster. Viewing the eigenvalues without the imaginary part can help to pick out the different types of behaviour, i.e. just by the rate at which they are damped. Figure 4.4 shows the same eigenvalues as Figure 4.3 but just the real part of the eigenvalue, denoting the decay rates of the modes.




Figure 4.4: Real part of the eigenvalues for high resolution boundary layer 5. The right plot shows the absolute value of the real part of λ with a log scaling on the y axis, the first few (purely inertial) modes, which are effectively undamped, have been omitted.

The figure shows more clearly the increasing rate at which the modes are damped, although due to the large range of damping the log-scale for the absolute value plot shows more clearly the varying behaviour. The plot to the right in Figure 4.4 shows the real part but with a log scaling, note that the first few modes, which are effectively undamped need to be omitted to benefit from the scaling. In this plot it is much clearer that there is varying behaviour in the modes. The omitted eigenvalues are the inertial type modes which lie above the boundary layer. The next set of modes with index 16 to around 22 are damped, but more slowly than others. There then exists a sequence of modes with a similar damping rate to each other (22-30). Beyond this the damping rate increases with each successive corresponding mode. Figure 4.5 shows the modes coming from the

three sections of behaviour noted in Figure 4.4; the u' part of the eigenvector is shown for corresponding eigenvalues with index 16, 27 and 75. The slower damped modes lie mainly near the top of the boundary layer and appear to be exhibiting a combination of inertial and damped processes, all have corresponding eigenvalue with an imaginary part. The eigenvalues which have similar real part correspond to modes which lie in the boundary layer only but occupy approximately the whole depth. The eigenvalues after this point, which show increasing damping rate all correspond to modes that lie inside the boundary layer but do not occupy the entire depth. The faster damping is due to the smaller scales that these modes occupy; scanning through enough successive eigenvectors the structure can be seen becoming concentrated near the ground.




Figure 4.5: u' part of the eigenvectors corresponding to eigenvalues with index number 16 (left) 27 (middle) and 75 (right). The dashed black line represents the approximate height of the boundary layer.

Obtaining and viewing the singular values and their corresponding singular vectors is somewhat easier than examining the eigenvalues and eigenvectors. Singular values are real and positive and are returned by the decomposition in order of magnitude. The downside is that the singular values are therefore incapable of describing both decay rate and frequency simultaneously. If a mode is propagating it may be possible to interpret it from the singular vector structures. Figure 4.6 shows the singular values for the high resolution run of boundary layer 5. As for

the eigenvalues the singular values have dimension s^{-1} .




Figure 4.6: Singular values for high resolution boundary layer 5. Shown with a linear scaling (left) and log scaling (right)

The first encouraging point to note from the singular values is the immediate similarity with the plot of the real part of the eigenvalues implying that some of the same physical interpretation applies. The overall damping rates as described by singular values is similar to the damping rates as described by the eigenvalues. Four types of behaviour were identified in the spectrum of eigenvalues and on close inspection these are evident also in the singular values. These regions may not necessarily be representing equivalent behaviour to the regions in the eigenvalues and this needs to be determined by examining the corresponding singular vectors. The behaviour as implied by the magnitude of the singular values can be grouped as a set of singular values all with $\sigma = f$ dividing two regions of likely damped singular vectors, with corresponding singular values greater than or less than f .

Singular vectors are now examined to establish the physical interpretation they provide. Figure 4.7 shows the singular vectors corresponding to a singular value which is equal to the Coriolis parameter. Singular vectors are normalised. The plots in the figure show that these correspond to the inertial modes; all the structure is located above the boundary layer, which lies at approximately 950m. For these modes the singular value is representing the equivalent to frequency. A mode with structure lying above the boundary layer will not be decaying since it is independent

of the damping mechanism, it will just propagate with frequency equal to the Coriolis parameter. The effect of the Coriolis is further apparent from looking at the structures of the input and output singular vectors. The input vector for v' is equal to the output vector for u' , and vice-versa (with a negative factor); also note that the θ' field contains only roundoff error and so can be neglected. Recall that an input singular vector may be rotated to match the output singular vector multiplied by the singular value. This appears to be happening here as the rotational effect of the Coriolis parameter is represented by structures in the u' direction being mapped onto the same structure in the v' direction. It is clear that these 16 singular values and their corresponding singular vectors are equivalent to the 16 inertial modes found in the eigendecomposition.




Figure 4.7: Singular Vectors corresponding to the starred singular value (index 15), here the singular value is equal to $f = 1.031 \times 10^{-4} \text{s}^{-1}$. On the top row are the left (output) singular vectors \mathbf{u} and on the bottom are the right (input) singular vectors \mathbf{v} . Here, and in similar figures, the dashed black line represents the approximate height of the boundary layer.

Neglecting the singular values that are equal to the Coriolis parameter the

sequence of corresponding singular vectors are found to have relatively smoothly varying behaviour. Although not particularly prominent in the figure the biggest change in behaviour comes at the singular value with index 37. Singular vectors with corresponding singular value having index less than 37 have structure concentrated at the boundary layer top with a varying degree of structure in θ' above the boundary layer, an example of which can be seen in Figure 4.8. The singular vectors exhibiting a mixture of structure above and below the boundary layer will likely be related to the modes whose corresponding eigenvalue had non-zero imaginary part.




Figure 4.8: Singular Vectors corresponding to the stated singular value (index 5). On the top row are the left (output) singular vectors \mathbf{u} and on the bottom are the right (input) singular vectors \mathbf{v} .

Singular vectors with corresponding singular value with index greater than 37 have structure completely within the boundary layer region. Figure 4.9 and Figure 4.10 show examples of singular vectors with only boundary layer structure. These singular vectors have structure that would be associated to the modes which are

most heavily damped. This appears to also be represented well by the magnitude of the singular value which is similar to the equivalent eigenvalue. Progressing through singular vectors with increasing corresponding singular value the structure is seen becoming more concentrated near the surface, where there is sufficient resolution to support the fastest damped small scales. For all the modes living in the boundary layer there is no clear relationship between input and output singular vectors, nor should one necessarily be expected. The transient boundary layer modes may have quite complex behaviour and evolution meaning the relationship between input and output singular vectors and singular values will be highly complex.




Figure 4.9: Singular Vectors corresponding to the stored singular value (index 37). On the top row are the left (output) singular vectors \mathbf{u} and on the bottom are the right (input) singular vectors \mathbf{v} .

The advantage to using the deepest boundary layer is that the modes contained inside the boundary layer are better resolved. Note that some structure is grid point scale, as in Figure 4.9 so not all structure will be captured by the low resolution grid. However when the comparison of modes is performed having large scale struc-




Figure 4.10: Singular Vectors corresponding to the stored singular value (index 200). On the top row are the left (output) singular vectors \mathbf{u} and on the bottom are the right (input) singular vectors \mathbf{v} .

ture will assist greatly. On the other hand, when considering a deeper boundary layer, there are a number of singular vectors where there exists structure above the boundary layer, where resolution is less. Figure 4.11 shows the singular values for the shallowest boundary layer. Clearly there are a much higher number of inertial modes supported, but more importantly a clearer pattern emerges to distinguish modes contained entirely in the boundary and those with some structure above. The same types of singular vectors that were apparent in the deeper boundary layer are identified, and across all the boundary layer case. The proportion of modes representing each type of structure depends on the depth of the boundary layer.

Throughout the spectrum it appears that the singular vectors are able to represent the kinds of structures that would be expected in the eigenvectors. It is not necessarily clear at the outset that this should be the case, however the results show that for this problem, and when writing in the energy norm, the SVD can be inter-




Figure 4.11: Singular values for high resolution boundary layer 1. Shown with a log-scaling.

preted in a useful way. The structures are similar to that found when looking at eigenvectors. Further, the physical interpretation, by magnitude of the corresponding singular value, is similar to that which would be expected, and is found when checking against the eigendecomposition. That the SVD is representing similar behaviour to the eigendecomposition increases confidence that it is revealing useful information about the transient evolution in the system. It can therefore be used as a tool for the comparison of configurations. The final step before the comparison can be performed is to check the form of the high resolution solutions and establish how dependent on the choice of configuration the results of the decomposition are.

4.6 Comparison of High Resolution Singular Vectors

Although a good physical interpretation can be attached to the SVD it is not sufficient to show that it will be useful for comparing configurations. One further

requirement is that singular vector structures representing the same behaviour are evident not only in low resolution configurations but also in all high resolution configurations.

From the work in the previous chapter, steady state solutions for Lorenz and Charney-Phillips that exhibit good convergence properties have been obtained. This does not necessarily mean a whole set of like for like singular vectors is achieved. The singular value decomposition will have a certain level of sensitivity to any perturbations that occur in the matrix from which it is derived. Since the matrices for Lorenz and Charney-Phillips will differ, due to averaging, there will be some differences in the singular vectors and singular values. Further consider the neutral drag coefficient, $C_{mn} = \left(\frac{\kappa}{\ln\left(\frac{z}{z_r}\right)} \right)^2$, which comes into the surface flux terms. For the logarithmically stretched grid the model levels get very close to the roughness length, this presents a singularity whereby C_{mn} tends to infinity as z tends to z_r . Whether the increasing C_{mn} remains balanced in the surface layer equations or serves to produce significant differences in the decompositions needs to be checked. In order to compare a low resolution Lorenz configuration against a low resolution Charney-Phillips configuration by examining a particular singular vector, the corresponding singular vector must exist in both the high resolution solutions and have negligible difference between each. Whereas in the dynamics only case [68] the best converged large-scale high resolution modes were also those picked up by the low resolution grid this is not necessarily the case here.

The nonlinearities and non-uniform grid spacing result in singular vectors which have quite complex structure. When trying to match up two singular vectors it is not possible to apply an algorithm such as zero counting to systematically identify partners. However it is possible to use the fact that the singular vectors are orthonormal. Singular value decomposition of any matrix produces unitary \mathbf{U} and \mathbf{V} , meaning rows and columns of each are magnitude 1.

For real vectors the inner product of a pair of vectors is given by the product of their norms and the angle between them. For a pair of singular vectors from the same \mathbf{U} this will thus be 1, if the inner product with itself or 0 otherwise. This

extends to complex space also, but with absolute value being closest to 1 or 0. Numerically the most efficient way to calculate the inner product is by taking the dot product of the pair of vectors.

In order to compute the inner product between a singular vector from the Lorenz configuration and a singular vector from the Charney-Phillips configuration the θ' part of the singular vector needs to be taken at the same model levels. This is done by interpolating the θ' part of the Charney-Phillips singular vector to the z_p levels, the location it is stored at with the Lorenz grid. Note that similar results were found when repeating the test while interpolating the Lorenz θ' to the Charney-Phillips grid. Using high order interpolation such as cubic will minimise any error in achieving this. To keep the new interpolated singular vector as close as possible to being a unit vector the Δz component, coming from the energy integral (4.27) is divided out before interpolating, the equivalent Δz at the new levels is then put back in.

Table 4.1 lists the singular vectors from the high resolution Lorenz solution for boundary layer 5. The number given to each singular vector is the index number of its corresponding singular value, so the singular vector '1' has the smallest singular value. For each Lorenz singular vector the inner product is computed with every singular vector from the Charney-Phillips solution. The best matched vector is then chosen as the one whose inner product is closest to 1 or -1 and these are listed in the table.

For around 95% of the singular vectors interpolated from the Charney-Phillips solution the vector norm (magnitude) was found to be between 0.8 and 1 and 80% between 0.9 and 1 (not shown in the table). The few that fell below this were all greater than 0.6 (again not shown) and all from singular vectors that would be unlikely to be captured by a low resolution grid. With the norms of the singular vectors well preserved at close to 1, the inner product should just give either ± 1 or 0, representing that the corresponding vector is either equivalent or orthogonal. The norms are not exactly preserved however so some discrepancy is expected. The data in Table 4.1 shows that there is clearly some sensitivity of the singular value

decomposition to perturbations in the matrix otherwise the table would be filled with 1s and -1s and every Lorenz vector would correspond to a single vector in the Charney-Phillips. For a number of the vectors, however, a good correlation between the two high resolution cases exists. Particularly important is that there appears to be good correlation between the Lorenz singular vectors whose index is 34 and upwards and their corresponding Charney-Phillips solution. These correspond to the modes which lie in the boundary layer and will be damped. It is of most interest to ensure that any configuration can capture the singular vectors corresponding to leading boundary layer modes accurately. These structures will be the longest lived, thus accurate representation will be important for proper evolution of the boundary layer structure. Some of the poorest correlations are in the singular vectors representing inertial oscillations, this is not a significant problem though as the modes these singular vectors correspond to would not need to feature in any comparison tests.

Using the inner product test provides an efficient way to obtain a handle on how well the high resolution configurations agree, otherwise a comparison would require visual examination of each successive Lorenz singular vector against each Charney-Phillips singular vector. It also removes some of the subjectivity in that type of approach. If two singular vectors are found to agree well by using the inner product test and are then found to also agree well when examining the structure by eye then, provided it is physically relevant, it can be confidently used to measure the ability of low resolution configurations.

Examining the correlation between Lorenz and Charney-Phillips for the high resolution using the table gives the best overall impression. Clearly some singular vectors are very well matched, shown by an inner product that is close to ± 1 , others however are not so close to ± 1 .

Since there is some discrepancy between high resolution singular vectors there is a final requirement of using the SVD analysis. The high resolution singular vectors that have been found to be well converged and physically relevant are also the singular vectors that are needed when comparing the low resolution configurations.

High Lorenz	High Charney-Phillips I	...continued from left					
1	1 (-0.954)	35	35 (0.634)	69	66 (0.848)		
2	2 (-0.941)	36	37 (-0.753)	70	73 (-0.688)		
3	3 (-0.811)	37	36 (0.688)	71	69 (-0.873)		
4	4 (0.692)	38	38 (0.631)	72	71 (0.509)		
5	5 (0.841)	39	39 (-0.709)	73	71 (-0.626)		
6	6 (0.828)	40	41 (-0.564)	74	72 (0.99)		
7	7 (-0.944)	41	41 (0.654)	75	79 (0.646)		
8	8 (-0.974)	42	42 (0.827)	76	74 (-0.861)		
9	11 (-0.583)	43	43 (-0.936)	77	75 (0.983)		
10	16 (0.591)	44	45 (-0.835)	78	82 (0.603)		
11	18 (-0.535)	45	44 (0.89)	79	78 (-0.995)		
12	9 (0.547)	46	46 (0.969)	80	77 (0.87)		
13	10 (0.553)	47	48 (0.805)	81	84 (0.511)		
14	13 (0.64)	48	47 (0.84)	82	80 (0.997)		
15	14 (-0.619)	49	49 (-0.971)	83	87 (-0.544)		
16	17 (0.596)	50	50 (0.815)	84	81 (0.844)		
17	18 (-0.634)	51	53 (-0.788)	85	83 (0.998)		
18	21 (0.552)	52	51 (-0.838)	86	90 (0.515)		
19	19 (0.562)	53	52 (0.769)	87	85 (0.71)		
20	13 (-0.546)	54	56 (0.782)	88	86 (-0.925)		
21	15 (-0.571)	55	54 (0.992)	89	90 (-0.508)		
22	20 (0.6)	56	55 (0.676)	90	89 (0.998)		
23	25 (0.659)	57	55 (-0.532)	91	88 (-0.856)		
24	23 (-0.803)	58	57 (0.985)	92	93 (0.527)		
25	24 (0.863)	59	61 (0.768)	93	91 (-0.935)		
26	26 (0.699)	60	60 (-0.928)	94	96 (-0.534)		
27	28 (-0.665)	61	59 (-0.778)	95	92 (-0.805)		
28	29 (0.665)	62	64 (-0.75)	96	94 (-0.998)		
29	30 (-0.652)	63	62 (-0.99)	97	99 (-0.547)		
30	31 (0.858)	64	67 (-0.623)	98	95 (0.846)		
31	32 (0.553)	65	63 (0.865)	99	97 (0.996)		
32	30 (-0.332)	66	65 (-0.98)	100	102 (-0.534)		
33	33 (-0.801)	67	70 (-0.581)				
34	34 (-0.917)	68	68 (-0.883)				

Table 4.1: Listed by index number up to 100, this table shows the high resolution Lorenz singular vectors from boundary layer 5 against the high resolution Charney-Phillips singular vector with whom the inner product is closest to 1 or -1. The actual inner product between those two vectors is bracketed.

That is that high resolution singular vectors that correspond to the leading low resolution singular vectors agree well across the configurations being compared at low resolution.

4.7 Comparison of Lorenz and Charney-Phillips Low Resolution

For this section when referring to Charney-Phillips it is for option I-i, later the other case that had good accuracy for the steady state, option I-iii, shall be examined. The low resolution grid is the Met Office grid.

To begin the comparison low resolution singular vectors need to be matched up using the inner product test. First the low resolution Charney-Phillips solution is interpolated onto the low resolution Lorenz grid then the inner products are computed between each Lorenz singular vector and every Charney-Phillips singular vector. The next stage of the comparison requires that the low resolution Lorenz solution is interpolated onto the high resolution Lorenz grid. Inner products can then be computed between each low resolution singular vector and every high resolution singular vector. It is then possible to use the data derived in Table 4.1 to determine whether the best matching high resolution singular vector can be used, based on whether the high resolution Lorenz and Charney-Phillips singular vectors agreed well. If high resolution vectors do not match for a particular low resolution pair then there is little reason to believe that low resolution should match. For each low resolution Lorenz singular vector Table 4.2 shows the closest low resolution Charney-Phillips singular vector as well as the closest high resolution singular vector. In order to know whether comparison of a particular singular vector is likely to be possible examine first the bracketed number in the fourth column, this value is taken from Table 4.1. If this value is close to 1 or -1 then the high resolution singular vectors should be generally well matched up. Examining the bracketed number in the second column determines how well the two low resolution singular vectors are matched up. The third column shows the inner product between low

and high resolution Lorenz configurations; that the inner product is generally further from 1 here is due to the error introduced when comparing singular vectors from different resolutions, where different scales may be captured, even in singular vectors corresponding to the same mode. In the table the singular vectors which are considered to correspond to boundary layer modes are shown with a *.

The table indicates that three starred low resolution singular vectors (15,22,24) have corresponding high resolution singular vectors where the inner product between the high resolution singular vectors is within 10% of 1 or -1. This gives three cases where singular value decompositions of Lorenz and Charney-Phillips have likely captured representations of the same eigenmode. Due to the implied similarity in the high resolution differences in the low resolution can be legitimately examined. Although the test is useful for quickly highlighting the best candidate comparisons, all the pairings that it finds should be examined by eye. There is a degree of uncertainty in the testing and it may be that matches which do not look particularly good in the results turn out to match quite well.

Before examining the singular vectors themselves it is worth examining all of the eigenvalues and singular values together as if the overall behaviour is similar then it becomes easier to interpret the singular values. The singular values are shown in Figure 4.12 and the eigenvalues are shown in Figure 4.13. The figures show singular values and eigenvalues for the low and high resolution Lorenz and Charney-Phillips grids for boundary layer 5, for all plots a log scaling has been applied, eigenvalues are thus shown as absolute value. Examining the low resolution singular values and eigenvalues which have an index number of 15 or above (corresponding to boundary layer modes) there appears to be a general trend of the Charney-Phillips grid having decay rates that are smaller. The singular values and absolute real part of the eigenvalues are smaller for Charney-Phillips in almost every case. Sorting Lorenz eigenvalues into order of magnitude does not necessarily mean that the corresponding Charney-Phillips eigenvalues would be in order of magnitude, seen by the index number in column two of Table 4.2, decay rates can be compared more accurately once two singular vectors are confirmed

Lorenz Low	Charney-Phillips Low	Lorenz High	Charney-Phillips High
1	1 (0.98)	1 (0.888)	1 (-0.954)
2	2 (0.983)	2 (0.807)	2 (-0.941)
3	3 (0.952)	6 (-0.459)	6 (0.828)
4	4 (-0.831)	4 (0.461)	4 (0.692)
5	5 (-0.772)	5 (-0.527)	5 (0.841)
6	6 (0.899)	7 (-0.695)	7 (-0.944)
7	7 (0.937)	8 (0.412)	8 (-0.974)
8	10 (0.794)	21 (-0.393)	15 (-0.571)
9	11 (0.685)	11 (-0.424)	18 (-0.535)
10	8 (0.799)	17 (0.442)	18 (-0.634)
11	9 (0.67)	12 (0.426)	9 (0.547)
12	12 (0.938)	24 (0.398)	23 (-0.803)
13	13 (0.846)	30 (-0.532)	31 (0.858)
14	14 (-0.908)	33 (0.57)	33 (-0.801)
15*	16 (-0.709)	34 (-0.739)	34 (-0.917)
16*	15 (0.672)	35 (0.621)	35 (0.634)
17*	17 (0.868)	42 (-0.461)	42 (0.827)
18*	18 (-0.898)	40 (0.494)	41 (-0.564)
19*	19 (0.964)	41 (-0.511)	41 (0.654)
20*	20 (0.764)	47 (-0.415)	48 (0.805)
21*	21 (-0.895)	48 (0.55)	47 (0.84)
22*	22 (0.99)	49 (0.46)	49 (-0.971)
23*	24 (-0.909)	61 (-0.49)	59 (-0.778)
24*	25 (-0.994)	60 (0.358)	60 (-0.928)
25*	23 (-0.806)	59 (0.272)	61 (0.768)
26*	26 (0.982)	69 (-0.527)	66 (0.848)
27*	27 (0.995)	80 (-0.406)	77 (0.87)

Table 4.2: Inner product calculations of low resolution Lorenz against low resolution Charney-Phillips option I-i and against high resolution Lorenz, values for the inner product between high resolution solutions as in Table 4.1. Listed in the first column is the index number for each low resolution Lorenz singular vector, singular vectors with a * are those that have their structure entirely in the boundary layer. The other columns show the index of the singular vector ‘most’ correlated to that Lorenz singular vector with the actual inner product shown in brackets. A value of ± 1 in the bracket would show exact correlation a value of 0 for exactly orthogonal.

as being equivalent. Nevertheless it gives some insight, the majority of starred singular vectors in Table 4.2 do have the same index in both low resolution cases.




Figure 4.12: Singular values for the low (left) and high (right) resolution Lorenz and Charney-Phillips solutions to boundary layer 5.




Figure 4.13: Absolute real part of the eigenvalues for the low (left) and high (right) resolution Lorenz and Charney-Phillips solutions to boundary layer 5.

That the Charney-Phillips grid has an overall trend of decreasing the decay rates is due to the effect of the averaging involved, which is to reduce K_m and K_h and thus reduce the rate of damping. The result of this decaying too slowly is that the corresponding mode will be longer lived in the Charney-Phillips simulation. If that mode is captured inaccurately then the resulting error will be more influential on the solution than if that mode were damped faster. The decay rates are closer between configurations in the high resolution due to the increased accuracy.

Having considered the general picture now the singular vector structures need to be examined. Scanning through the combinations of singular vectors, as predicted by the inner product test in Table 4.2, it is found that the low resolution singular vectors with index 15, 17, 22 and 25 are the most promising. 15 and 22 were identified by the test, 17 was identified as being a good match when checked by eye, as was 25. In the test 24 is found to have slightly better low and high resolution matches than 25, but when checking by eye 25 gave the better match. The left and right singular vectors for the four examples are shown in Figure 4.14 - Figure 4.17. The figures are arranged with the u' , v' and θ' component of the left (output) singular vectors on the top row and the right (input) singular vectors on the bottom row. The plots show the singular vectors after they have been transformed back to the u' , v' and θ' form.

Figure 4.14 shows the singular vectors corresponding to the Lorenz singular value with index 15. Recall that these tests are for the Met Office grid and Charney-Phillips oprion I-i. To aid in comparison only the dominant part of the singular vector is shown, whether it be the real or imaginary part. This goes for all plots of singular vectors shown in this section, unless specified.

The singular vector that is shown is an important singular vector in that it corresponds to the leading boundary layer mode and is thus the longest lived. The overall structures in the singular vectors are quite similar and it does indeed appear to be the same mode being represented in each configuration. Unfortunately there is some difference in the high resolution solutions, however there is wider difference in the low resolution solutions so there is likely some error. The magnitude of the singular value is similar across each configuration, this would be expected for a leading mode so adds further confidence that physical properties are being represented. For this singular vector the difference between low resolution singular values is quite small so a similar decay rate on the mode would be expected, in fact the singular value is slightly larger for the Charney-Phillips grid implying that it would be damped faster than it would be by the Lorenz grid. For the singular vectors the closest match is in the velocity fields, here there is similar




Figure 4.14: Boundary layer 5. The low resolution Lorenz singular vectors with index 15. Left (output) singular vectors are shown in the top row, right (input) singular vectors are shown in the bottom row. The best matched singular vectors in the low resolution Charney-Phillips and high resolution configurations are also shown. These have index as given in Table 4.2. From left to right shows the u' , v' and θ' component of the singular vector. The legend in the bottom row shows the value of the corresponding singular value for each case. Here, and in similar figures, the dashed black line represents the approximate height of the boundary layer.

structure across all configurations. In the θ' there is more difference between the low resolution cases and there is also greater difference between low and high resolution. This could be related to the singularity near the ground which would dominate in the high resolution solution.

With the exception of the right input vector for u' the Charney-Phillips grid captures these singular vectors less accurately than the Lorenz grid. Near the surface the Lorenz grid is very close to the high resolution structure in all the fields

whereas Charney-Phillips appears mostly to overestimate the magnitude. The singular values are quite close so the mode would not be expected to experience significant decay errors.

In Figure 4.15 the singular vectors that are best matched to the low resolution Lorenz singular vector with index 17 are shown. From Table 4.2 it is clear that the only other combinations of singular vectors in the high resolution that match as well as or better than those corresponding to the low resolution with index 15 correspond to the part of the spectrum with fastest decay (22, 24 and 25), two of which are plotted later. If possible it is more interesting to compare the leading singular vectors as these are longest lived and thus more important to capture. The low resolution Lorenz singular vector with index 17 is chosen due to its superior high resolution inner product compared with its immediate neighbours. From the plot it is clear that the high resolution singular vectors still match quite well, despite the reduced inner product compared with those corresponding to low resolution index 15. All configurations clearly capture the same singular vector and both low resolution cases appear to do a reasonable job. The low resolution is able to resolve quite a lot of the detail as the mode's structure is not so concentrated at the top of the boundary layer as in the previous case, this can be seen in the structure of the high resolution mode. There is again greater difference between low and high resolution in the θ' field. The Lorenz low resolution solution lies closer to the high resolution solution than Charney-Phillips in the v' vectors and the input vector for θ' . Arguably the structure is more accurate for Charney-Phillips in the u' field although there is greater difference between the high resolution u' fields. The magnitude of peaks is slightly better for Lorenz.

In Figure 4.16 the singular vectors that are best matched to the low resolution Lorenz singular vector with index 22 are shown. Comparison of this singular vector is less directly indicative of the expected performance of the low resolution configuration due the fact that the corresponding singular value is larger, meaning it is likely that it corresponds to a shorter lived mode. It will still provide further understanding of the behaviour of low resolution configurations, as will the index 25




Figure 4.15: As for Figure 4.14 but for the low resolution Lorenz singular vector with index 17.

case. The high resolution grids agree very well in the velocity fields but have quite large differences in the potential temperature. For this singular vector however, the θ structure is around an order of magnitude smaller so capturing the velocity of the corresponding mode would be more important. Actually there is little to distinguish between the structure for the low resolution configurations except near the surface where the Charney-Phillips is slightly closer to the high resolution solutions. The magnitude of the maximum and minimum points is closer to the high resolution grids for Lorenz. For this example a larger difference is observed in the singular values than was previously, this is again what would be expected for a faster less well resolved mode. The low resolution Charney-Phillips grid has smaller singular value than the low resolution Lorenz and high resolution grids implying that the decay of the corresponding mode would be too slow.




Figure 4.16: As for Figure 4.14 but for the low resolution Lorenz singular vector with index 22.

In Figure 4.17 the singular vectors that are best matched to the low resolution Lorenz singular vector with index 25 are shown. The structure by this point on the spectrum is getting quite small-scale making it increasingly harder to recognise similarities between low and high resolution solutions. To aid in comparison the singular vectors in their energy norm form, including the $(\Delta z)^{\frac{1}{2}}$ factor in (4.27), and with all still on the high resolution grid are shown in Figure 4.18. In the energy norm form there is less to differentiate between the low resolution grids showing more clearly that the singular vectors are representing the same mode. In terms of capturing the singular vector in its more natural form both the low resolution grids do a pretty poor job, mainly due to this singular vector coming from a point on the spectrum where structure may exist below the low resolution grid scale. In most aspects the Lorenz configuration outperforms the Charney-Phillips, particularly in

finding the peaks in the input vectors. There is again some difference between the singular values and again they would imply that the corresponding mode decays too slowly on the Charney-Phillips grid. This will be a short lived mode so even though neither grid captures it it should not influence the solution too dramatically.




Figure 4.17: As for Figure 4.14 but for the low resolution Lorenz singular vector with index 25.

It is clear that the SVD methodology is not as revealing as when the normal mode analysis is performed on an appropriate problem. Despite this it has been possible to use the method to make comparisons between the low resolution configurations. In all the cases that were compared the physical behaviour was similar to that which would be expected based on how the equations are understood to behave. In most of the cases the Lorenz grid was able to give better representation than the Charney-Phillips grid, either by giving better representation of the singular vector structure or singular value magnitude.




Figure 4.18: As for Figure 4.17 but with the singular vectors plotted in their energy norm form.

4.7.1 Comparison with Charney-Phillips Option I-iii

Of the cases considered when examining the steady state part of the solution the option where the Richardson number is averaged rather than K_h was found to be beneficial, this was option I-iii. The structures of the steady state potential temperature solutions found when using option I-iii were closer to the high resolution solution than the other configurations that were tested. It now needs to be examined as to whether the advantages seen for the steady state extend to the representation of the transients. Table 4.3 shows the inner product calculations for Charney-Phillips option I-iii with boundary layer 5. The first point to note is that unfortunately the high resolution option I-iii does not match the high resolution Lorenz as well as the high resolution option I-i configuration did. If it is difficult to match high resolution singular vectors it makes less sense to compare low resolution singular vectors. The agreement between low resolution singular vectors is good still.

Lorenz Low	Charney-Phillips Low	Lorenz High	Charney-Phillips High
15	15 (0.692)	34 (-0.739)	34 (-0.748)
16	16 (0.669)	35 (0.621)	34 (0.58)
17	18 (-0.746)	42 (-0.461)	42 (-0.594)
18	17 (-0.702)	40 (0.494)	39 (0.766)
19	19 (-0.945)	41 (-0.511)	40 (-0.882)
20	21 (0.75)	47 (-0.415)	48 (0.598)
21	20 (-0.794)	48 (0.55)	47 (-0.881)
22	22 (-0.986)	49 (0.46)	49 (-0.729)
23	23 (-0.774)	61 (-0.49)	58 (-0.881)
24	24 (0.87)	60 (0.358)	59 (0.987)
25	25 (-0.471)	59 (0.272)	63 (0.558)
26	26 (0.906)	69 (-0.527)	66 (0.871)
27	27 (0.96)	80 (-0.406)	77 (0.855)

Table 4.3: Inner product calculations of low resolution Lorenz against low resolution Charney-Phillips option I-iii and against high resolution Lorenz, values for the inner product between high resolution Lorenz and high resolution Charney-Phillips are shown in the final column. The table is arranged as Table 4.2 except that non-boundary layer singular vectors have been omitted.

Figure 4.19 shows the eigenvalues and singular values for Charney-Phillips option I-iii against Lorenz for the low resolution. As previously there appears a tendency for the Charney-Phillips singular values to be smaller than the Lorenz singular values. This could imply that the modes being represented are not damped quickly enough under Charney-Phillips, which would mean larger errors if the mode is not captured accurately. As previously the individual singular vectors need to be matched in order to make accurate comparison.




Figure 4.19: Absolute real part of the eigenvalues (left) and the singular values (right) for Lorenz and Charney-Phillips option I-iii, boundary layer 5.

Figure 4.20 shows the input and output singular vectors for low resolution Lorenz singular vectors with index 15. The best matched Charney-Phillips option I-iii singular vectors and high resolution singular vectors are also shown. In this instance the inner product test did not return the best low resolution match. The test found the best matched Charney-Phillips singular vector to have index 15 but the plotted Charney-Phillips singular vector has index 16, which is clearly well matched.

Clearly the option I-iii grid produces results very similar to that found for option I-i grid, plotted in Figure 4.14. Again there is a good match between both low resolution configurations as well as between low and high resolution and between high resolution. All configurations are finding the same singular vector. As was found when comparing Lorenz and option I-i there is some difference between the




Figure 4.20: As for Figure 4.14 but with option I-iii used for the low and high resolution Charney-Phillips configurations.

high resolution configurations, slightly more for these configurations than there was for option I-i. This makes it even harder to draw any solid conclusions in how well the low resolution configurations are capable of capturing the singular vector. Where there is similarity between high resolution, for example in the output v' vector, the option I-iii solution is slightly more accurate than the option I-i solution. Across all the configurations the singular values are close, however the difference between the low resolution solution is similar to the difference between high resolution so it is not possible to conclude whether the corresponding mode would decay at the wrong rate.

Figure 4.21 shows the input and output singular vectors for low resolution Lorenz singular vectors with index 17. For this Lorenz singular vector the best match that could be found in the Charney-Phillips singular vectors was that found

by the inner product test.




Figure 4.21: As for Figure 4.20 but for the low resolution Lorenz singular vector with index 17.

The overall structure of the singular vectors is quite similar but unfortunately there are quite large differences between the high resolution singular vector making complete comparison impossible. Where the high resolution vectors are similar, in the u' part of the input vector and v' part of the output vector, the Lorenz grid does slightly better than the Charney-Phillips grid. Again there is difference between singular values for the high resolution cases and for the low resolution cases. The singular value for the Charney-Phillips grid is larger than for the Lorenz grid.

When the steady states were examined there was clear benefit to option I-iii when using the Charney-Phillips grid. The structures of the steady states were more accurately represented than they were when using any other Charney-Phillips option. From the evidence that could be gathered for the transient cases it does

not seem that any additional benefit is gained by using this option. The best comparison is seen by comparing Figure 4.14 and Figure 4.20, both Charney-Phillips grids capture the same singular vector and there is no significant difference between them.

4.7.2 Shallower Boundary Layers and Alternative Grid Spacing

In the section covering the steady state the abilities of the configurations whilst using a number of boundary layer depths and grid spacings was demonstrated. In obtaining the comparison between Lorenz and Charney-Phillips for boundary layer 5 in the last section various types of spacing we experimented with in order to allow for most efficient results when employing the inner product test. As part of the experimentation the log-linear hybrid was used. While the log-linear spacing gave good results when comparing two high resolution solutions the results were less positive when comparing low resolution configurations to high resolution configurations. That results were less useful when comparing low to high resolution is likely due to the large number of model levels higher in the domain, leading to, for example, an increased number of inertial modes. Since using the logarithmic grid yielded quite useful results the ability of the hybrid grid has not been extensively investigated.

In addition to using the Met Office grid the investigation also considered other ways of spacing the low resolution grid. However it was found that little benefit was afforded by using other grid spacing in the low resolution. The uniform grid has so little resolution where important structure lies that calculating inner products becomes too error prone, also comparing singular vectors visually becomes highly subjective. Singular vectors found when using the Met Office grid have also be identified when using the geometric spacing but the structure was no closer to the high resolution than it was for the Met Office grid.

So far only boundary layer 5 has been considered. Boundary layer 5 is the deepest that is being considered and thus has the highest number of model levels within

the boundary layer region, this has been very useful in comparing singular vectors. Other boundary layer depths are now examined to increase the understanding of how the grids can capture singular vectors and establish the limitations of the methodology. Table 4.4 - Table 4.7 show the inner product calculations for boundary layer 1 to boundary layer 4, all with Charney-Phillips option I-i.

Lorenz Low	Charney-Phillips Low	Lorenz High	Charney-Phillips High
1	1 (-0.97)	1 (-0.901)	1 (-0.969)
2	2 (0.938)	4 (0.469)	4 (-0.965)
3	3 (0.971)	3 (-0.789)	3 (-0.974)
4	4 (0.972)	4 (0.685)	4 (-0.965)
5	25 (-0.654)	2 (0.722)	2 (-0.995)
6	5 (0.974)	6 (0.553)	6 (0.938)
7	6 (-0.983)	7 (0.432)	7 (0.889)
8	8 (-0.993)	75 (-0.438)	73 (0.761)
9	9 (-0.96)	17 (-0.335)	68 (-0.306)
10	11 (-0.68)	16 (0.357)	61 (0.438)
11	15 (0.467)	23 (0.371)	68 (-0.307)
12	16 (0.447)	53 (0.234)	56 (-0.34)
13	15 (-0.619)	29 (-0.307)	23 (0.512)
14	17 (0.638)	46 (-0.314)	47 (-0.335)
15	21 (-0.584)	47 (-0.304)	53 (-0.418)
16	20 (-0.515)	34 (-0.353)	29 (-0.34)
17	16 (-0.531)	53 (0.294)	56 (-0.34)
18	13 (0.559)	51 (-0.324)	54 (-0.31)
19	12 (-0.238)	49 (-0.308)	60 (-0.343)
20	19 (-0.332)	42 (-0.304)	45 (-0.401)
21	11 (-0.732)	71 (0.337)	24 (0.44)
22	22 (-0.96)	70 (0.337)	17 (-0.318)
23	23 (-0.847)	14 (0.683)	13 (-0.869)
24	24 (-0.725)	12 (-0.393)	74 (0.616)
25*	26 (-0.91)	103 (-0.417)	103 (0.819)
26*	27 (0.931)	113 (0.329)	112 (-0.633)
27*	25 (-0.001)	91 (-0.373)	94 (-0.859)

Table 4.4: Inner product calculations for boundary layer 1. Table shows low resolution Lorenz against low resolution Charney-Phillips option I-i and against high resolution Lorenz. Values for the inner product of high resolution Lorenz with high resolution Charney-Phillips are shown in final column. Low resolution Lorenz singular vectors with starred index are boundary layer singular vectors.

Lorenz Low	Charney-Phillips Low	Lorenz High	Charney-Phillips High
23*	23 (-0.805)	80 (-0.658)	80 (-0.827)
24*	24 (-0.762)	84 (0.383)	83 (0.827)
25*	26 (-0.936)	85 (0.332)	85 (0.818)
26*	25 (0.019)	83 (0.307)	84 (-0.801)
27*	27 (0.984)	100 (-0.324)	98 (0.72)

Table 4.5: As for Table 4.4 except for boundary layer 2. Only singular vectors with structure entirely inside the boundary layer are shown.

Lorenz Low	Charney-Phillips Low	Lorenz High	Charney-Phillips High
20*	20 (0.84)	59 (0.647)	59 (-0.835)
21*	25 (-0.556)	60 (-0.447)	60 (0.805)
22*	21 (-0.5)	61 (0.377)	62 (0.755)
23*	23 (0.747)	65 (0.423)	65 (-0.735)
24*	24 (-0.738)	77 (0.343)	76 (0.8)
25*	26 (-0.86)	76 (0.352)	75 (0.91)
26*	24 (0.485)	77 (-0.272)	76 (0.8)
27*	27 (0.988)	94 (-0.338)	92 (-0.699)

Table 4.6: As for Table 4.4 except for boundary layer 3. Only singular vectors with structure entirely inside the boundary layer are shown.

For a given resolution the number of singular vectors whose structure lies entirely in the boundary layer is reduced as the depth of the boundary layer decreases. For boundary layer 1 only three singular vectors are clear boundary layer structures, with the vast majority being inertial modes. At first glance the inner products seem highly inaccurate compared with the deepest boundary layer, but this is due to the difficulty in calculating inner product between inertial modes. The inner product between the high resolution singular vectors that correspond to the low resolution boundary layer singular vectors are quite good in all four boundary layer depths. Based on these findings it is worth trying to extend the comparison of the singular vectors to the shallower boundary layers.

Two cases are demonstrated, firstly in Figure 4.22, which shows the leading boundary layer singular vectors for boundary layer 1.

Lorenz Low	Charney-Phillips Low	Lorenz High	Charney-Phillips High
17*	17 (-0.855)	45 (0.581)	45 (0.913)
18*	19 (0.618)	46 (-0.593)	46 (-0.894)
19*	19 (0.698)	47 (0.553)	47 (0.722)
20*	20 (0.91)	54 (-0.534)	53 (-0.872)
21*	21 (0.86)	53 (-0.451)	54 (-0.831)
22*	22 (0.976)	55 (0.472)	55 (0.929)
23*	23 (-0.84)	62 (0.529)	61 (0.642)
24*	25 (-0.992)	67 (0.368)	66 (0.991)
25*	24 (-0.754)	63 (0.286)	65 (-0.784)
26*	26 (0.971)	74 (0.46)	72 (-0.836)
27*	27 (0.991)	85 (0.386)	83 (0.812)

Table 4.7: As for Table 4.4 except for boundary layer 4. Only singular vectors with structure entirely inside the boundary layer are shown.

A good match can be seen between high resolution vectors. Unfortunately boundary layer structure is not well enough resolved in the low resolution to find the equivalent high resolution mode and thus draw solid comparison between Lorenz and Charney-Phillips. The size of the singular values agrees with what has been seen in some of the previous cases, the Charney-Phillips grid gives slightly larger singular value in both low and high resolution. For a very shallow boundary layer it is too subjective to attempt to match singular vectors between low and high resolution by eye, only the inner product test can be used. When examining by eye a large number of high resolution singular vectors could be deemed a good match for the low resolution singular vector in question. Judging by the singular values it is difficult to be confident that the singular vectors being found are indeed the same, the high resolution singular values are an order of magnitude smaller than the low resolution singular values. In the deeper boundary layer experiments the low and high resolution singular values were much closer in magnitude.

As the boundary layer depth is increased more grid points will lie in the boundary layer region and thus the structure will be better resolved. It is only by boundary layer 4 however that a similar level of confidence is gained in the comparison as was had for the deepest boundary layer. Figure 4.23 shows an example of a




Figure 4.22: As for Figure 4.14 but for boundary layer 1 and low resolution Lorenz singular vector with index 25.

boundary layer 4 singular vector; the low resolution Lorenz singular vector has index 20. For this particular singular vector the Lorenz and Charney-Phillips grid both capture the structure accurately. The singular value as found when using Charney-Phillips grid is smaller than that found when using the Lorenz grid while the two high resolution singular values are similar.

In the following chapter a summary of results for the boundary layer only case are offered, including a summary of the findings of this chapter.




Figure 4.23: As for Figure 4.14 but for boundary layer 4 and low resolution Lorenz singular vector with index 20.

Chapter 5

Summary of Part I

The principal aim of this thesis is to understand the impact of the choice of vertical staggering when coupling the large scale dynamics to the planetary boundary layer. In particular the investigation has set out to establish differences between using the Lorenz grid and using the Charney-Phillips grid. Previously it has been shown that of these two the Charney-Phillips grid is preferable when modelling the large scale dynamics on its own. It offers optimal wave propagation and supports no computational mode. At the beginning of Part I of this thesis it was argued that the Lorenz grid should be preferable when modelling the stably stratified boundary layer on its own: when using the Lorenz grid no averaging occurs in the equations.

In order to investigate the question of whether either the Lorenz grid or the Charney-Phillips grid offers an optimal configuration for the coupled case a systematic methodology, based on linearisation of the equations, has been considered. First a steady state of the equations must be generated, linearising about this steady state gives the transients. The transient part of the expansion represents the linear time evolution of the equations. The transient evolution can be projected onto normal modes, obtained through an eigendecomposition. Implementing a linearisation of the equations representing the coupling between the dynamics and boundary layer is not necessarily trivial. In order to use it to perform the comparison between the Lorenz and Charney-Phillips grid a careful examination of the methodology has been required.

In Part I of this thesis the boundary layer has been considered on its own, without the large scale dynamics. The purpose of doing this is two-fold, to examine the suitability of the methodology and to compare the Lorenz and Charney-Phillips grid to ensure that, as expected, the Lorenz grid is the favoured configuration. If it were found that the Charney-Phillips grid represented the boundary layer more accurately it would seem unlikely anything different would be found when coupling dynamics and boundary layer.

The first stage of both implementing the methodology and examining differences between configurations is to consider the steady state part of the equations. With the equations in their basic form it was not possible to generate a fully steady state due to the downward diffusion of potential temperature. Usually this process would be balanced by imposed heat fluxes. However in order to obtain a fully steady state requires fixed surface and upper boundary conditions on potential temperature. It has been shown that a completely steady state can be reached by including a subsidence heating term to balance the diffusion. In doing this retaining a realistic profile for the model variables is still achieved.

A further issue that has required attention is in finding high resolution solutions. In order to compare configurations they must also be compared to a high resolution solution, a resolution for which there is no discernible difference between configurations. Even with a geometric spacing of the grid the convergence of the solutions was not fast enough to obtain high resolution solutions for a reasonable number of grid points. The poor convergence is due to a singularity in the derivative at the surface. It has been shown that good convergence properties can be recovered by transforming to a logarithmic coordinate. With this grid transformation it has been possible to obtain a satisfactorily high resolution solution with 100 grid points.

In terms of capturing the steady state structure the Lorenz grid performed well and overall made for the best choice. Boundary layer depths were chosen ranging from 100m to 900m. For all depths the Lorenz grid was able to represent the structure. The shallowest boundary layer had the largest errors in the low

resolution structure, due to the lack of resolution in the boundary layer region. For the deepest boundary layers the structure captured by the low resolution Lorenz grid was very close to the high resolution ‘truth’ solution. The Charney-Phillips configurations that were examined gather into configurations based on the choice of averaging for obtaining Richardson number. Option I configurations were those where Richardson number was found by averaging potential temperature gradient, option II configurations averaged shear for the Richardson number and option III configurations averaged both. Options II and III configurations were generally found to be problematic due to the sensitivity to errors in the shear that result from averaging. The shear enters the equations through the eddy diffusion K_m and diffusivity K_h , it enters directly and through the stability function. Averaging results in errors in the shear which are fed back into the model variables producing even greater errors in the shear; the result is non-physical structures in all the model variables. Under option I and II there are further options for how to obtain both K_m and K_h . One is found without further averaging but the other can be found by either averaging K itself, averaging the stability function or averaging Richardson number. For option I all three configurations worked well and were able to give accurate representation of the steady states. Except for the shallowest boundary layer the Lorenz grid gave the most accurate overall representation of the operational resolution steady states. The Charney-Phillips configurations that gave the most promising results were ones which averaged K_h itself and averaged Richardson number. The configuration that averaged the stability function did not outperform the other two option I configurations in any of the tests that were considered here.

In computing the transients a number of issues arose concerning the methodology. Linearisation expands the model variables into a reference and transient part, the transient part represents the linear time evolution. In order to obtain the structures of the discrete transients the linearisation is written as a matrix eigenvalue problem. In doing this all the possible transient structures are represented by the eigenvectors, linear evolution can then be projected onto combinations of the

eigenvectors. In order for the transient behaviour to be easily interpretable from the eigenvectors they need to form an orthogonal basis and for this the matrix must be normal [71]. For the boundary layer only problem it has been shown that the matrix can be far from normal. For the dynamics only case, as considered by [68], the matrix is normal and the eigenvectors form an orthogonal basis and are referred to as the normal modes. Due to the matrix for the boundary layer only problem being far from normal multiple eigenvectors representing similar behaviour are found, making physical interpretation very difficult. The non-normality makes it difficult to identify the modes which are of importance and therefore compare them between configurations. Instead of using the eigendecomposition to represent the transient behaviour the Singular Value Decomposition (SVD) has been considered.

The great advantage of the SVD over the eigendecomposition is that singular vectors always form an orthonormal basis. For the SVD each pair of input and output singular vectors corresponds to one singular value and this singular value gives a measure of the amount those singular vectors would be perturbed when multiplied by matrix. That the vectors are orthogonal and correspond one-to-one with singular values means it should be possible compare configurations. Indeed it has been shown that the method is quite successful for comparing configurations for the boundary layer only problem. In order to increase the ‘closeness’ of the SVD and the eigendecomposition the equations are written in the energy norm. It has been shown that for an energy conserving system the matrix of transient coefficients will be normal and Hermitian, in this situation the SVD and eigendecomposition are likely to offer very similar physical interpretation; this is discussed in detail in Part II. Although the boundary layer is not energy conserving, but dissipates energy, it is nevertheless useful to provide this link. It was shown that by writing in the energy norm the physical interpretation of the SVD results was more closely aligned with the results that would be expected in the eigendecomposition.

That the singular vectors are always orthogonal has the added benefit that an automated approach can be developed for identifying singular vectors that corre-

spond to one another in differing configurations. The inner product between two orthogonal vectors is equal to zero whereas the inner product between a vector and itself is one. For example if a singular vector in the Lorenz configuration needs to be identified in the Charney-Phillips configuration the inner products are calculated between the Lorenz singular vector and all Charney-Phillips singular vectors. It has been shown that by seeking the singular vector which has inner product closest to one or minus one it is possible to quickly identify singular vectors that correspond to the same behaviour. It has also been shown that this test can be used to match singular vectors between low and high resolution. Some inaccuracies occur in the test due to an interpolation that is required to bring variables to equivalent model levels, as this is done some error is introduced. Further, a singularity in the neutral drag coefficient has been identified which appears to affect the similarity between high resolution configurations.

For comparison of the transients the option I-i and I-iii Charney-Phillips configurations were considered as these were found to be the most promising configurations for the steady state. For all the results presented the configurations computed on the Met Office grid were compared to the configurations on the high resolution logarithmic grid. The success of the testing was dependent on the depth of the boundary layer. For deeper boundary layers, where more grid points coincide with boundary layer structure, the testing was more successful. For the shallowest boundary layers it was possible to identify similar singular vectors in both low resolution configurations but difficult to find the corresponding high resolution singular vector with any certainty. Other types of grid spacing were also tested but the results did not differ significantly from the results found for the Met Office grid. Given the difficulties with the methodology it was generally quite difficult to perform detailed comparisons between different low resolution configurations.

The first case that was presented was for boundary layer 5 and compared low resolution Lorenz versus low resolution Charney-Phillips option I-i versus high resolution cases of both configurations. Using the inner product test it was possible to identify well matched corresponding singular vectors in every configuration. The

singular vectors that were most interesting to compare were those that described boundary layer behaviour but that underwent the smallest perturbation when multiplied by the matrix. These correspond to the longest lived boundary layer modes; it is important to capture these accurately since they decay slowly and are long lived in the solution. Although these were the most interesting singular vectors to compare they were not always the easiest to compare since there could be greater difference between high resolution structures. One case that was examined was for the low Lorenz singular vector with index 22, although this is shorter lived it had much closer matched high resolution singulars than the leading boundary singular vector, index 15. Across the cases that were examined it was generally found that differences existed between the corresponding high resolution solutions making it difficult to draw confident comparisons between the low resolution configurations. Where high resolution solutions were sufficiently close it was generally found that the Lorenz grid captured the solution more accurately than the Charney-Phillips option I-i configuration.

Two sets of singular vectors were shown for the comparison of Lorenz and Charney-Phillips option I-iii configuration. The two that were shown were leading boundary layer singular vectors which were also shown for the option I-i comparison. There was increased difference between the high resolution singular vectors for this case. Where there was similarity between the high resolution singular vectors the operational resolution option I-iii solution was slightly more accurate than the operational resolution option I-i solution. In addition to the improved steady state representation this seems to be a useful configuration. The other leading singular vector shown for this configuration was not captured as well as it was for the option I-i configuration, further illustrating the difficulty in implementing this methodology.

The form of the spectrum was discussed when completing the transient part of the calculation. In the literature there exist examples of boundary layer flows that have a continuous part of the spectrum; however these have unbounded domains. It is not clear whether the stably stratified boundary layer model that is

being considered here should exhibit wholly discrete spectrum and it is beyond the scope of the work to rigorously prove either way. In Part II of this thesis a more comprehensive study of the form of the spectrum is performed. In performing the comparison between the Lorenz and Charney-Phillips grids in Part I no clearly continuous behaviour has been identified. A number of similar singular vectors have been identified across low and high resolution configurations, as well as in cases with differing averaging options for the Charney-Phillips grid and differing ways of spacing the model levels. Although the identification of similar singular vectors in different configurations is not sufficient to prove the existence of only discrete spectrum, it does increase confidence that a discrete spectrum exists.

The useful results as a whole, steady states, singular values, corresponding singular vectors and eigenvalues are showing the Lorenz low resolution grid to be more accurate than the Charney-Phillips. In most steady state runs the Lorenz matched the high resolution converged solutions better. In the singular vectors that could be examined the structure captured by the Lorenz grid was better than the structure found by Charney-Phillips. There is an overall trend in the eigenvalues and singular values suggesting that Charney-Phillips would damp the associated mode too slowly, although this was not always the case when examining individual singular vectors. It should be noted before concluding too much from the singular values there was a difference of the same order between the high resolution singular values, limiting the interpretation that can be drawn. If errors occur when capturing the structure of a mode then their error will be compounded if that mode decays slower than it should.

The results as a whole agree with the initial hypothesis that the Lorenz grid should be more accurate for the boundary layer only. This also confirms the conflict between choice of vertical arrangement: for the dynamics using the Charney-Phillips grid is more accurate, for the boundary layer using the Lorenz grid is more accurate. It now needs to be established whether or not one choice of staggering provides more accurate results when coupling dynamics and boundary layer. Of interest while examining the coupled problem will be in how well all mode struc-

tures, dynamical and boundary layer, are captured, how well the frequencies of the dynamical modes are captured and how well the decay rates of the boundary layer modes are captured. A further point of interest is the behaviour of the Lorenz grid computational mode in the presence of the boundary layer.

Part II

Comparison of Vertical Discretisations Whilst Coupling Boundary Layer and Dynamics

In this part of the thesis the full system, where the boundary layer is coupled to the large scale dynamics, is considered. From the work in Part I of this thesis, along with the findings of papers such as [68], it is now understood that when the two parts of the system are modelled independently contradictory vertical configurations would be preferable. For the boundary layer the Lorenz grid gives better representation of boundary layer modes and their rate of damping and so is preferable; for the dynamics the Charney-Phillips grid gives better dispersion properties and has no computational mode and so is preferable. What now remains to be established is how these two grid configurations will perform when used to model a situation where physics and dynamics are coupled, specifically the atmospheric boundary layer and the dynamics, where the conflict has been demonstrated.

Naively one may choose a certain configuration based on which aspects are believed to be most important to capture in the numerical model. If capturing the boundary structure accurately was considered less important than capturing accurately the frequency of the Rossby modes one would conclude that using a Charney-Phillips grid would be better. The problem with this approach, and inherent in the questions that need to be addressed here, is in understanding how the two parts of the model interact with each other. It is desirable to avoid any kind of computational mode, such as that associated to the Lorenz grid. It is difficult to predict however, exactly what happens to the computational mode when introducing the boundary layer. It could be that it damps the computational mode, meaning adverse effects may be avoided. This could be theorised by considering the strong damping mechanism that the boundary layer imposes, however, complex interactions may result in something different. This damping may also affect the lower frequency Rossby modes, it may be that they are distorted and damped considerably by the boundary layer mechanism, negating some of the benefits of using Charney-Phillips. The differences observed between how a particular mode is captured by each grid in the non-coupled case may not extend to the coupled case.

In performing the comparison of Lorenz and Charney-Phillips for the boundary

layer only case a number of problems were encountered when trying to employ the usual normal mode analysis, due to the eigenvectors being non-orthogonal. This led to the use of singular vectors which were found to give a good representation of the physical behaviour of the system and allow for individual modes to be compared by orthogonality between singular vectors. When examining the dynamics on its own with a uniform grid the normal mode analysis works as expected and can be used seamlessly for comparison of different vertical configurations. Using these techniques that have been useful while studying the boundary layer along with the normal mode analysis that has previously been used for the dynamics it should be possible to gain considerable insight into how a particular staggering will perform when coupling the two. In leading to the comparison of the Lorenz and Charney-Phillips grids the methodology will be examined in detail.

Part II is arranged as follows: Chapter 6 will outline the model that will be used for the fully coupled problem, this is essentially an extension of the model that was used for the boundary layer to include the large scale dynamical mechanisms. In order to perform a thorough comparison of the Lorenz and Charney-Phillips grids, it is first important to understand all the behaviour that is supported by the coupled model. Chapter 7 will examine the types of modes that are found in the coupled problem and the differences between these and the structures found when each part is modelled independently. This chapter will also be where the methodology is examined and the extent of the usefulness of eigenvector analysis established. The final part of this chapter will examine the Lorenz grid computational mode. Chapter 8 covers the complete comparison of the Lorenz and Charney-Phillips grid for the fully coupled problem, covering the ability of each configuration to capture both the steady and transient part of the equations. The final chapter of this thesis is Chapter 9, this will draw together all the findings of the work and outline ideas for future work.

Chapter 6

The Fully Compressible

Dynamics-Boundary Layer Model

In Part I of this thesis the full Navier-Stokes equations were approximated using the Reynolds averaging technique. This is a useful technique that allows the turbulent mechanisms in the flow to be written in terms of the mean large scale aspects that are captured on the grid. With the right closure model turbulent processes that could otherwise only be represented by capturing all relevant scales can be well approximated. The technique is popular in numerical weather prediction for modelling the atmospheric boundary layer. In Part I of this thesis the transient and steady state structures of the stably stratified atmospheric boundary layer were examined by modelling the Reynolds averaged Navier-Stokes equations. Then, by applying the Boussinesq approximation and using the f -plane, the equations were effectively filtered in order to avoid capturing any dynamical acoustic, gravity and Rossby waves.

For Part II of this thesis the fully coupled problem will be examined and this chapter describes the model that will be required in order to achieve this. The first step in deriving the model requires the return to the form of the fully compressible Reynolds Averaged Navier-Stokes equations, equations (2.25)-(2.29) in Part I. The turbulent fluxes that result from the Reynolds averaging still only need to model the turbulence that occurs in the boundary layer and so the approximations, based

on scale, allowing for neglect of horizontal gradients of turbulent flux remain in place. The boundary layer closure remains as for the boundary layer only model. The equations are,

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} + w \frac{\partial u}{\partial z} - fv + c_p \theta \frac{\partial \Pi}{\partial x} = -\frac{1}{\rho} \frac{\partial}{\partial z} (\rho \tau_x), \quad (6.1)$$

$$\frac{\partial v}{\partial t} + u \frac{\partial v}{\partial x} + v \frac{\partial v}{\partial y} + w \frac{\partial v}{\partial z} + fu + c_p \theta \frac{\partial \Pi}{\partial y} = -\frac{1}{\rho} \frac{\partial}{\partial z} (\rho \tau_y), \quad (6.2)$$

$$\frac{\partial w}{\partial t} + u \frac{\partial w}{\partial x} + v \frac{\partial w}{\partial y} + w \frac{\partial w}{\partial z} + c_p \theta \frac{\partial \Pi}{\partial z} = -g, \quad (6.3)$$

$$\frac{\partial \theta}{\partial t} + u \frac{\partial \theta}{\partial x} + v \frac{\partial \theta}{\partial y} + w \frac{\partial \theta}{\partial z} = -\frac{1}{\rho} \frac{\partial}{\partial z} (\rho \mathcal{H}), \quad (6.4)$$

$$\frac{\partial \rho}{\partial t} + u \frac{\partial \rho}{\partial x} + v \frac{\partial \rho}{\partial y} + w \frac{\partial \rho}{\partial z} + \rho \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} \right) = 0, \quad (6.5)$$

where $c_p = 1005.0$ is the specific heat capacity for dry air. Thermodynamic variables are related through the ideal gas law,

$$p = R\theta \left(\frac{p_0}{p} \right)^{-\kappa_g} \rho, \quad (6.6)$$

where p_0 is a reference surface pressure, $R = 287.05 \text{ J kg}^{-1} \text{ K}^{-1}$ is the gas constant for dry air and $\kappa_g = \frac{R}{c_p}$. A subscript g has been added to κ_g , denoting gas, to differentiate from the von-Karman constant. Recall from Part I that the form of the closure is, for example, $\tau_x = -K_m \frac{\partial u}{\partial z}$.

For the boundary layer only model it was most straightforward to work in the Boussinesq framework; it would be viable to continue to use the Boussinesq framework for the fully coupled system since only acoustic modes would be filtered out, important Rossby and inertio-gravity waves would remain. There are three factors making it desirable to be non-Boussinesq. Firstly it would otherwise require use of the generalised form of the eigenvalue decomposition, meaning use of the less well known generalised Singular Value Decomposition. Secondly, acoustic modes propagate very quickly compared to other types of waves and so are less likely to be distorted or slowed by the boundary layer mechanism. Having them included will give a part of the spectrum that is highly distinguished from the rest, useful

for analysing the methodology and examining overall transient behaviour. Thirdly the Met Office’s Unified Model is based on the fully compressible equations.

6.1 Form of the Vertical Pressure Gradient

Note that equations (6.1) to (6.5) are written in terms of Exner pressure Π . For the height based coordinate Thuburn and Woollings [68] showed that the optimal vertical discretisation for the dynamics only case was one where pressure and potential temperature were chosen for the thermodynamic variables and the Charney-Phillips grid was used. Ideally one would like to choose density ρ as one of the thermodynamic variables as it gives a more direct route to ensuring mass conservation in the system [64]. With the equations written in terms of pressure p [68] found that using ρ led only to near optimal representation of the dispersion relation; one which had no computation mode but underestimated the frequency of the Rossby modes.

As was shown by [68] the averaging required on the model variables is related to how well the frequencies of the normal modes are represented by the discretisation. When the p form of the pressure gradient is used and ρ is chosen as a prognostic variable an extra averaging step is required in the w equation compared to when p is chosen as a prognostic variable. The extra averaging when using ρ results in the slowing of the smallest scale inertio-gravity or Rossby modes.

In [64] and simultaneously in [70] it was shown that when the vertical pressure gradient term is written in terms of the Exner function the extra averaging when using ρ is avoided. The result is an optimal configuration, while still allowing ρ to be used as a prognostic variable. Pressure p can be recovered from the Exner pressure by the relation,

$$\Pi = \left(\frac{p}{p_0} \right)^{\kappa_g}. \quad (6.7)$$

6.2 Linearisation

For the coupled dynamics-boundary layer study the nature of the questions that need to be addressed are slightly different to the questions that were addressed for the boundary layer separately. Firstly the overall steady state structures will be largely the same in the coupled case, the only additional requirement is that density comes into the equations, which can be calculated through hydrostatic balance. The reference solution for the vertical velocity remains zero. The process of computing the transient part of the equation is also similar to the boundary layer only case, only with the addition of the transient vertical velocity and density.

Whereas the efforts of the Part I have been in both understanding how to properly compute the linearisation of this complex problem and in examining solutions for differences in vertical configurations, the focus of the remaining chapters shifts. The methodology is largely in place and now the task is to examine the complicated coupling interactions and then to compare the grid configurations. With this in mind, the way in which the results are presented for the coupled case will differ. In Part I two distinct chapters, covering first the computation and comparison of the steady state and then the transients, was most natural. For the coupled problem the linearisation requires only small adjustments, that are covered presently. Later in the investigation different parts of the linearisation can then be called upon as required. For example in the next chapter, when examining the coupling, the transient part of the calculation will be required. Examining the steady state is postponed until comparing Lorenz and Charney-Phillips. It will have similar overall structure in the coupled system as it had in the boundary layer problem.

6.2.1 Steady State

Capturing a suitable steady state for the whole coupled model does not differ much from the boundary layer only case. It is still sufficient (and necessary) to assume that the reference vertical wind W is zero since it is dominated by horizontal wind speeds. Coriolis parameter f is constant and all reference state variables except p are independent of x and y , making it possible to continue to employ the geostrophic

approximation for the horizontal derivatives of pressure. The only addition to the velocity equations is density in the boundary layer diffusion terms. That the reference vertical velocity is zero leads to a reference state that is in hydrostatic balance. The ρ equation becomes decoupled for the horizontally homogeneous, no vertical wind approximation. The equation for hydrostatic balance is used to compute density, this can be thought of as a diagnostic step. Note that the cooling and heating terms remain in the thermodynamic potential temperature equation so as to obtain the steady state, if it were not present then the same issues of a continually cooling boundary layer would be encountered, as discussed in Part I.

The full set of equations required to solve the steady state are obtained by leaving just the steady parts of equations (6.1), (6.2) and (6.3) and by adding in the subsidence heating and radiative cooling to (6.4),

$$0 = f(V - v_g) - \frac{1}{\rho^{(r)}} \frac{\partial}{\partial z} (\rho^{(r)} \tau_x^{(r)}), \quad (6.8)$$

$$0 = -f(U - u_g) - \frac{1}{\rho^{(r)}} \frac{\partial}{\partial z} (\rho^{(r)} \tau_y^{(r)}), \quad (6.9)$$

$$0 = -c_p \theta^{(r)} \frac{\partial \Pi^{(r)}}{\partial z} - g, \quad (6.10)$$

$$0 = -\frac{1}{\rho^{(r)}} \frac{\partial}{\partial z} (\rho^{(r)} \mathcal{H}^{(r)}) - w_{\text{sub}} \frac{\partial \theta^{(r)}}{\partial z} - R_c. \quad (6.11)$$

Reference thermodynamic variables are related through the equation of state.

$$p^{(r)} = R \theta^{(r)} \left(\frac{p_0}{p^{(r)}} \right)^{-\kappa_g} \rho^{(r)}. \quad (6.12)$$

Hydrostatic Balance

The reference state consists of thermodynamic variables that are in hydrostatic balance; the vertical pressure gradient balances the gravitational force. Equations (6.8)-(6.11) can be solved as previously by the method of false time stepping or, for the high resolution, the Newton method. Equation (6.10) can be solved during a diagnostic step (ρ diagnosed from θ). If using the false time stepping method then, by using the latest guess for the potential temperature, $\Pi^{(r)}$ that is in hydrostatic

balance is calculated, from which the pressure and density follow. $\Pi^{(r)}$ can be solved by a simple rearrangement of (6.10). i.e.,

$$\frac{\partial \Pi^{(r)}}{\partial z} = -\frac{g}{c_p \theta^{(r)}}. \quad (6.13)$$

The pressure at top boundary p_N is chosen, giving $\Pi^{(r)}$ at the top boundary. Integrating downwards using a numerical form of (6.13) $\Pi^{(r)}$ is obtained throughout the domain. Pressure is then obtained throughout by using (6.7) written in terms of reference variables and then density using (6.12). In order to approximately obtain the realistic surface pressure $p_0 = 100000\text{Pa}$, $p_N = 81000\text{Pa}$ is chosen for the top of the domain at 2000m. Note that when fixing the top boundary condition on pressure, and following the calculation in this way, the surface value will vary slightly, depending on the surface boundary condition on the potential temperature.

If using the Newton method there is no need to have a diagnostic step to calculate the hydrostatic balance, it is numerically more efficient for the condition to be satisfied automatically through the iteration. Recall from Part I that the Newton iteration is given by,

$$\mathbf{x}_{\text{new}} = \mathbf{x} - [\nabla_{\mathbf{x}} \mathbf{F}(\mathbf{x})]^{-1} \mathbf{F}(\mathbf{x}), \quad (6.14)$$

where \mathbf{x} is now the vector containing all five model variables at each model level. The function \mathbf{F} represents the right hand side of (6.8)-(6.11) with zeros for the ρ equation. $\nabla_{\mathbf{x}} \mathbf{F}(\mathbf{x})$ is the matrix of the coefficients of the transients which is described in detail in the next section.

Discretisation

Here the discretisation for the false time stepping scheme is demonstrated in order to highlight the terms that will require averaging. For the scheme the time derivative of reference model variables is put back in to the equations. The only supported transients are damped through boundary layer diffusion, allowing a steady state to be reached. For both the Lorenz and Charney-Phillips grids the horizontal

momentum equations are given by,

$$\frac{U_j^{n+1} - U_j^{n-1}}{2\Delta t} = f(V_j^n - v_g) - \left\{ \frac{1}{(\Delta z)_j \rho_j^{(r)}} \left[\bar{\rho}_{j+\frac{1}{2}}^{(r)} (\tau_x^{(r)})_{j+\frac{1}{2}} - \bar{\rho}_{j-\frac{1}{2}}^{(r)} (\tau_x^{(r)})_{j-\frac{1}{2}} \right] \right\}^{n-1}, \quad (6.15)$$

$$\frac{V_j^{n+1} - V_j^{n-1}}{2\Delta t} = -f(U_j^n - u_g) - \left\{ \frac{1}{(\Delta z)_j \rho_j^{(r)}} \left[\bar{\rho}_{j+\frac{1}{2}}^{(r)} (\tau_y^{(r)})_{j+\frac{1}{2}} - \bar{\rho}_{j-\frac{1}{2}}^{(r)} (\tau_y^{(r)})_{j-\frac{1}{2}} \right] \right\}^{n-1}. \quad (6.16)$$

Superscript n denotes time level and subscript j the spatial level, overbar, e.g. $\bar{\rho}$, denotes an averaged quantity. Options for averaging in the boundary layer terms are unchanged from Part I.

On the Lorenz grid the θ equation is,

$$\frac{[\theta^{(r)}]_j^{n+1} - [\theta^{(r)}]_j^n}{\Delta t} = \left\{ -\frac{1}{(\Delta z)_j \rho_j^{(r)}} \left[\bar{\rho}_{j+\frac{1}{2}}^{(r)} \mathcal{H}_{j+\frac{1}{2}} - \bar{\rho}_{j-\frac{1}{2}}^{(r)} \mathcal{H}_{j-\frac{1}{2}} \right] - (w_{\text{sub}})_j \frac{\theta_{j+1}^{(r)} - \theta_j^{(r)}}{\Delta z_{j+\frac{1}{2}}} - R_c \right\}^n \quad (6.17)$$

and on the Charney-Phillips grid it is,

$$\frac{[\theta^{(r)}]_{j+\frac{1}{2}}^{n+1} - [\theta^{(r)}]_{j+\frac{1}{2}}^n}{\Delta t} = \left\{ -\frac{1}{(\Delta z)_{j+\frac{1}{2}} \bar{\rho}_{j+\frac{1}{2}}^{(r)}} \left[\rho_{j+1}^{(r)} \mathcal{H}_{j+1} - \rho_j^{(r)} \mathcal{H}_j \right] - (w_{\text{sub}})_{j+\frac{1}{2}} \frac{\theta_{j+\frac{3}{2}}^{(r)} - \theta_{j+\frac{1}{2}}^{(r)}}{\Delta z_{j+1}} - R_c \right\}^n. \quad (6.18)$$

Recall that w_{sub} is negative. An analogous process of discretisation follows for use in the Newton method.

6.2.2 Transients

Leaving just the order one transient terms when linearising equations (6.1), (6.2), (6.3), (6.4) and (6.5) gives,

$$\begin{aligned} \frac{\partial u'}{\partial t} + U \frac{\partial u'}{\partial x} + V \frac{\partial u'}{\partial y} + w' \frac{\partial U}{\partial z} - f v' + c_p \theta^{(r)} \frac{\partial \Pi'}{\partial x} \\ = -\frac{1}{\rho^{(r)}} \frac{\partial}{\partial z} (\rho^{(r)} \tau'_x + \rho' \tau_x^{(r)}) + \frac{\rho'}{\rho^{(r)2}} \frac{\partial}{\partial z} (\rho^{(r)} \tau_x^{(r)}), \end{aligned} \quad (6.19)$$

$$\begin{aligned} \frac{\partial v'}{\partial t} + U \frac{\partial v'}{\partial x} + V \frac{\partial v'}{\partial y} + w' \frac{\partial V}{\partial z} + f u' + c_p \theta^{(r)} \frac{\partial \Pi'}{\partial y} \\ = -\frac{1}{\rho^{(r)}} \frac{\partial}{\partial z} (\rho^{(r)} \tau'_y + \rho' \tau_y^{(r)}) + \frac{\rho'}{\rho^{(r)2}} \frac{\partial}{\partial z} (\rho^{(r)} \tau_y^{(r)}), \end{aligned} \quad (6.20)$$

$$\frac{\partial w'}{\partial t} + U \frac{\partial w'}{\partial x} + V \frac{\partial w'}{\partial y} = -c_p \theta' \frac{\partial \Pi^{(r)}}{\partial z} - c_p \theta^{(r)} \frac{\partial \Pi'}{\partial z}, \quad (6.21)$$

$$\begin{aligned} \frac{\partial \theta'}{\partial t} + U \frac{\partial \theta'}{\partial x} + V \frac{\partial \theta'}{\partial y} + w' \frac{\partial \theta^{(r)}}{\partial z} \\ = -\frac{1}{\rho^{(r)}} \frac{\partial}{\partial z} (\rho^{(r)} \mathcal{H}' + \rho' \mathcal{H}^{(r)}) + \frac{\rho'}{\rho^{(r)2}} \frac{\partial}{\partial z} (\rho^{(r)} \mathcal{H}^{(r)}), \end{aligned} \quad (6.22)$$

$$\frac{\partial \rho'}{\partial t} + U \frac{\partial \rho'}{\partial x} + V \frac{\partial \rho'}{\partial y} + w' \frac{\partial \rho^{(r)}}{\partial z} + \rho^{(r)} \left(\frac{\partial u'}{\partial x} + \frac{\partial v'}{\partial y} + \frac{\partial w'}{\partial z} \right) = 0. \quad (6.23)$$

where transient boundary layer diffusion terms τ'_x , τ'_y and \mathcal{H}' are unchanged from Part I. In each material derivative 3 or 4 terms vanish due to either the reference $W = 0$ or reference U and V being independent of x and y . The linearised form of the ideal gas law, equation (6.6), is,

$$p' = \frac{p^{(r)}}{1 - \kappa_g} \left(\frac{\rho'}{\rho^{(r)}} + \frac{\theta'}{\theta^{(r)}} \right). \quad (6.24)$$

For Exner function the linearised ideal gas is given by,

$$\Pi' = \frac{\kappa_g \Pi^{(r)}}{1 - \kappa_g} \left(\frac{\rho'}{\rho^{(r)}} + \frac{\theta'}{\theta^{(r)}} \right). \quad (6.25)$$

Since coefficients (reference state variables) in equations (6.19) to (6.23) are independent of x , y and t , solutions can be Fourier decomposed into horizontally wavelike solutions proportional to $\exp(ikx + ily + \lambda t)$. k is the zonal horizontal wavenumber and l is the meridional horizontal wavenumber. Having set the hor-

horizontal structure reduces the calculation to a one dimensional column model, as was the form of the solution in Part I.

Inclusion of the β -effect

For the equations written in this form, where f is constant and not dependent on position, Rossby waves are unable to propagate. This poses the problem that the Rossby modes would be degenerate. It is known that the Lorenz grid supports a computational mode of zero frequency and examining it is of particular interest; without varying Coriolis it would require visual examination of mode structure to identify the computational mode, rather than looking for the mode with zero frequency. Ideally the model could be rewritten replacing f with $f = f_0 + \beta y$, giving variation with latitude (y). By doing this however, it would no longer be possible to seek solutions proportional to $\exp(ikx + ily + \lambda t)$ and hence write the problem as a one dimensional column model. Instead the technique of [68] and [78] is used. Through use of a β -effect terms which emulate latitude dependence are included in the equations for conservation of horizontal momentum, see [78] for a detailed description of the technique.

With the inclusion of the β -effect and seeking wavelike solutions for transient terms that are proportional to $\exp(ikx + ily + \lambda t)$ equations (6.19) - (6.23) become,

$$\begin{aligned} \lambda u' = & -Uiku' - Vilu' - w' \frac{\partial U}{\partial z} + fv' - \frac{ik\beta}{K^2} u' - c_p \theta^{(r)} ik \Pi' \\ & - \frac{1}{\rho^{(r)}} \frac{\partial}{\partial z} (\rho^{(r)} \tau'_x + \rho' \tau_x^{(r)}) + \frac{\rho'}{\rho^{(r)2}} \frac{\partial}{\partial z} (\rho^{(r)} \tau_x^{(r)}), \end{aligned} \quad (6.26)$$

$$\begin{aligned} \lambda v' = & -Uikv' - Vilv' - w' \frac{\partial V}{\partial z} - fu' - \frac{ik\beta}{K^2} v' - c_p \theta^{(r)} il \Pi' \\ & - \frac{1}{\rho^{(r)}} \frac{\partial}{\partial z} (\rho^{(r)} \tau'_y + \rho' \tau_y^{(r)}) + \frac{\rho'}{\rho^{(r)2}} \frac{\partial}{\partial z} (\rho^{(r)} \tau_y^{(r)}), \end{aligned} \quad (6.27)$$

$$\lambda w' = -Uikw' - Vilw' - \theta' \frac{\partial \Pi^{(r)}}{\partial z} - \theta^{(r)} \frac{\partial \Pi'}{\partial z}, \quad (6.28)$$

$$\begin{aligned} \lambda \theta' = & -Uik\theta' - Vil\theta' - w' \frac{\partial \theta^{(r)}}{\partial z} \\ & - \frac{1}{\rho^{(r)}} \frac{\partial}{\partial z} (\rho^{(r)} \mathcal{H}' + \rho' \mathcal{H}^{(r)}) + \frac{\rho'}{\rho^{(r)2}} \frac{\partial}{\partial z} (\rho^{(r)} \mathcal{H}^{(r)}), \end{aligned} \quad (6.29)$$

$$\lambda\rho' = -Uik\rho' - Vil\rho' - \rho^{(r)}(iku' + ilv') - \frac{\partial}{\partial z}(w'\rho^{(r)}). \quad (6.30)$$

Attention can be restricted to $l = 0$ since rotating the horizontal wavevector $K^2 = k^2 + l^2$ is equivalent to rescaling β .

Discretisation

The following demonstrates the discretisation of equations (6.26)-(6.30), firstly for the Lorenz grid the $\lambda u'$, $\lambda w'$ and $\lambda\rho'$ equations discretise to,

$$\begin{aligned} \lambda u'_j = & -U_j i k u'_j - V_j i l u'_j - \overline{w'}_j \frac{U_{j+1} - U_{j-1}}{(\Delta z)_{j+\frac{1}{2}} + (\Delta z)_{j-\frac{1}{2}}} + f v'_j - \frac{i k \beta}{K^2} u'_j - c_p \theta_j^{(r)} i k \Pi'_j \\ & - \frac{1}{(\Delta z)_j \rho_j^{(r)}} \left(\overline{\rho^{(r)}}_{j+\frac{1}{2}} (\tau'_x)_{j+\frac{1}{2}} + \overline{\rho'}_{j+\frac{1}{2}} (\tau_x^{(r)})_{j+\frac{1}{2}} - \overline{\rho^{(r)}}_{j-\frac{1}{2}} (\tau'_x)_{j-\frac{1}{2}} - \overline{\rho'}_{j-\frac{1}{2}} (\tau_x^{(r)})_{j-\frac{1}{2}} \right) \\ & + \frac{\rho'_j}{(\Delta z)_j \rho_j^{(r)2}} \left(\overline{\rho^{(r)}}_{j+\frac{1}{2}} (\tau_x^{(r)})_{j+\frac{1}{2}} - \overline{\rho^{(r)}}_{j-\frac{1}{2}} (\tau_x^{(r)})_{j-\frac{1}{2}} \right), \end{aligned} \quad (6.31)$$

$$\lambda w'_{j+\frac{1}{2}} = -\overline{U}_{j+\frac{1}{2}} i k w'_{j+\frac{1}{2}} - \overline{V} i l w'_{j+\frac{1}{2}} - \overline{\theta'}_{j+\frac{1}{2}} \frac{\Pi_{j+1}^{(r)} - \Pi_j^{(r)}}{(\Delta z)_{j+\frac{1}{2}}} - \overline{\theta^{(r)}}_{j+\frac{1}{2}} \frac{\Pi'_{j+1} - \Pi'_j}{(\Delta z)_{j+\frac{1}{2}}}, \quad (6.32)$$

$$\lambda\rho'_j = -U_j i k \rho'_j - V_j i l \rho'_j - \rho_j^{(r)}(iku'_j + ilv'_j) - \frac{\overline{w'}_{j+1} \rho_{j+1}^{(r)} - \overline{w'}_{j-1} \rho_{j-1}^{(r)}}{(\Delta z)_{j+\frac{1}{2}} + (\Delta z)_{j-\frac{1}{2}}}, \quad (6.33)$$

where e.g. \overline{u} represents that an averaging is required. The equations for λv_j , $\lambda\theta_j$ are similar in nature to the λu_j equation.

The horizontal velocity equations differ between the Lorenz and Charney-Phillips grids only through the boundary layer terms, with placement of variables analogous to that described in Part I. The density equation does not differ between the grid configurations. On the Charney-Phillips grid the discretisations for vertical velocity and potential temperature are,

$$\lambda w'_{j+\frac{1}{2}} = -\overline{U}_{j+\frac{1}{2}} i k w'_{j+\frac{1}{2}} - \overline{V} i l w'_{j+\frac{1}{2}} - \theta'_{j+\frac{1}{2}} \frac{\Pi_{j+1}^{(r)} - \Pi_j^{(r)}}{(\Delta z)_{j+\frac{1}{2}}} - \theta_{j+\frac{1}{2}}^{(r)} \frac{\Pi'_{j+1} - \Pi'_j}{(\Delta z)_{j+\frac{1}{2}}}, \quad (6.34)$$

and

$$\begin{aligned}
\lambda\theta'_{j+\frac{1}{2}} = & -\bar{U}_{j+\frac{1}{2}}ik\theta'_{j+\frac{1}{2}} - \bar{V}_{j+\frac{1}{2}}il\theta'_{j+\frac{1}{2}} - w'_{j+\frac{1}{2}}\frac{\theta_{j+\frac{3}{2}}^{(r)} - \theta_{j-\frac{1}{2}}^{(r)}}{(\Delta z)_{j+1} + (\Delta z)_j} \\
& - \frac{1}{(\Delta z)_{j+\frac{1}{2}}\bar{\rho}_{j+\frac{1}{2}}^{(r)}}\left(\rho_{j+1}^{(r)}\mathcal{H}'_{j+1} + \rho'_{j+1}\mathcal{H}_{j+1}^{(r)} - \rho_j^{(r)}\mathcal{H}'_j - \rho'_j\mathcal{H}_j^{(r)}\right) \\
& + \frac{\bar{\rho}'_{j+\frac{1}{2}}}{(\Delta z)_{j+\frac{1}{2}}\bar{\rho}_{j+\frac{1}{2}}^{(r)2}}\left(\rho_{j+1}^{(r)}\mathcal{H}_{j+1}^{(r)} - \rho_j^{(r)}\mathcal{H}_j^{(r)}\right). \tag{6.35}
\end{aligned}$$

The above discretised equations can be written as the matrix eigenvalue problem,

$$\lambda\mathbf{x} = \mathbf{A}\mathbf{x}. \tag{6.36}$$

Solutions or modes of the system are the eigenvectors \mathbf{x} , the eigenvalues $\lambda = \mu - i\omega$ have real part describing the decay or growth rate of the corresponding mode and imaginary part giving the frequency. Positive μ denotes growth, negative denotes decay. Positive ω denotes eastward propagating, negative denotes westward.

Transient Subsidence Heating

Previously the model had no dependency in the horizontal and was incompressible; this meant that the subsidence heating, used to balance the diffusion, could be used without concern for balancing the mass budget. If the subsidence heating was included now then some transient horizontal divergence would be required to balance the flux of mass through transient subsidence heating. This would be difficult to implement and would likely cause problems with the methodology.

The extra heating terms were needed to balance the cooling caused by the fixed surface temperature boundary condition, required to obtain the steady state. The easiest way to counteract the cooling was to introduce the subsidence term, and of course a term of this form arises in the transient part of the calculation as well. It can however, be omitted from the transient calculation without affecting the validity of the results. Although more tricky and time consuming the same steady state could have been obtained by finding an appropriate heating profile,

i.e. just a function of z that had no dependency on the model variables but had similar structure to $w_{\text{sub}} \frac{\partial \theta}{\partial z}$. Could this function have been easily found the same steady state would be obtained. Writing the equations with this form of subsidence heating would not have altered the steady state but, since not a function of a model variable, would not appear in the transients. Rather than attempting to balance the mass budget in the model it is assumed that the steady state was found with the aid of a heating profile; allowing efforts to be focused on the main points of interest, i.e. in examining the overall behaviour and comparing the configurations.

6.2.3 Energy Norm

As for the boundary layer only case it will be useful to be able to work in the Energy norm. As established previously, for an energy conserving system the matrix of coefficients in (6.36) will be normal and skew-Hermitian if the system is the written in the energy norm. That is that model parameters are linearly transformed such that the total energy in the system can be written as,

$$E = \mathbf{s}^T \mathbf{I} \mathbf{s}. \quad (6.37)$$

where \mathbf{s} is the transformed version of \mathbf{x} and \mathbf{I} is the identity matrix. As noted previously, when a matrix is normal, Hermitian (or skew-Hermitian) and sign-definite it has eigendecomposition and Singular Value Decomposition (SVD) that yield equivalent results. For the problem that has a reference state which is isothermal and has no shear, as considered by [68], energy will be conserved so the matrix will be normal and skew-Hermitian when written in the energy norm. The SVD produces real singular values, whereas a skew-Hermitian matrix has purely imaginary eigenvalues so a factor i will likely appear in the SVD. Although eigenvalues will be purely imaginary they will still be positive and negative, denoting eastward and westward propagation. This means the matrix will not be sign-definite, the impact of this in terms of interpreting the SVD, in comparison to the eigendecomposition, can be examined for the isothermal case.

The coupled problem will not conserve energy due dissipation of energy by

the boundary layer diffusion. The SVD is dependent on the norm in which the problem is written. Due to the relation for the conserved energy case, using the energy norm should achieve physical interpretation of singular vectors closest to the physical interpretation of the eigenvectors.

For a system that conserves energy,

$$\frac{DE}{Dt} + \nabla \cdot (\mathbf{u}p) = 0. \quad (6.38)$$

In order to convert to the energy norm the term E needs to be found. Equation (6.38) only holds for a system with a resting basic state and no boundary layer terms so in order to convert to the energy norm these terms must be neglected. The advective material derivative terms that exist due to shear in the reference state represent an exchange of energy between scales, rather than change of overall energy, so this approximation is reasonable. Since it is simpler to calculate energy whilst using pressure rather than density transforming the variables will also result in converting to pressure. Where required, the equation for pressure is,

$$\frac{\partial p}{\partial t} + u \frac{\partial p}{\partial x} + v \frac{\partial p}{\partial y} + w \frac{\partial p}{\partial z} + c^2 \rho^{(r)} - \rho^{(r)} g w = 0. \quad (6.39)$$

To obtain the total energy in the system a procedure analogous to that described in [66] and [67] is used, i.e. by taking,

$$u \times (6.1) + v \times (6.2) + w \times (6.3) + \frac{g\theta}{\theta^{(r)} \frac{\partial \theta^{(r)}}{\partial z}} \times (6.4) + \frac{p}{c^2} \times (6.39). \quad (6.40)$$

The speed of sound c is given by

$$c^2 = \frac{RT^{(r)}}{1 - \kappa_g}, \quad (6.41)$$

where $T^{(r)} = \Pi^{(r)}\theta^{(r)}$ is the reference temperature. Coefficients in (6.40) are chosen so as give cancelling of terms not obviously conserved. Once expanded equation

(6.40) is equivalent to equation (6.38) with,

$$E = \frac{\rho^{(r)} \mathbf{u}^2}{2} + \frac{g\rho^{(r)}\theta^2}{2\theta^{(r)} \frac{\partial\theta^{(r)}}{\partial z}} + \frac{p^2}{2\rho^{(r)}c^2}.$$

Using a simple midpoint type numerical integration the discrete energy in the system is therefore given by,

$$E = \sum_{j=1}^n (\Delta z)_j \frac{\rho_j^{(r)} u_j^2}{2} + (\Delta z)_j \frac{\rho_j^{(r)} v_j^2}{2} + (\Delta z)_j \frac{g\rho_j^{(r)} \theta_j^2}{2\theta_j^{(r)} \frac{\partial\theta^{(r)}}{\partial z} \Big|_j} + (\Delta z)_j \frac{p_j^2}{2\rho_j^{(r)} c_j^2} + \sum_{j=1}^{n-1} (\Delta z)_{j+\frac{1}{2}} \frac{\rho_{j+\frac{1}{2}}^{(r)} w_{j+\frac{1}{2}}^2}{2} \quad (6.42)$$

on the Lorenz grid, and

$$E = \sum_{j=1}^n (\Delta z)_j \frac{\rho_j^{(r)} u_j^2}{2} + (\Delta z)_j \frac{\rho_j^{(r)} v_j^2}{2} + (\Delta z)_j \frac{p_j^2}{2\rho_j^{(r)} c_j^2} + \sum_{j=1}^{n-1} (\Delta z)_{j+\frac{1}{2}} \frac{g\rho_{j+\frac{1}{2}}^{(r)} \theta_{j+\frac{1}{2}}^2}{2\theta_{j+\frac{1}{2}}^{(r)} \frac{\partial\theta^{(r)}}{\partial z} \Big|_{j+\frac{1}{2}}} + (\Delta z)_{j+\frac{1}{2}} \frac{\rho_{j+\frac{1}{2}}^{(r)} w_{j+\frac{1}{2}}^2}{2} \quad (6.43)$$

on the Charney-Phillips grid.

As the investigation progresses it will be useful to use the transient variables in their energy form, for example when establishing the energetically dominant part of a mode. So that they can be described more easily a more compact notation is used, so for example equation (6.42) is,

$$E = \frac{1}{2} \sum_{j=1}^n (\Delta z)_j E_u + (\Delta z)_j E_v + (\Delta z)_j E_\theta + (\Delta z)_j E_p + \frac{1}{2} \sum_{j=1}^{n-1} (\Delta z)_{j+\frac{1}{2}} E_\rho.$$

When working in the energy norm, instead of equation (6.36), solutions take the form,

$$\lambda \mathbf{s} = \mathbf{C} \mathbf{s}. \quad (6.44)$$

where,

$$\mathbf{C} = \mathbf{B}\mathbf{A}\mathbf{B}^{-1}. \quad (6.45)$$

Multiplying the state vector \mathbf{x} by \mathbf{B} transforms into the energy variables, i.e $\mathbf{s} = \mathbf{B}\mathbf{x}$. Vector \mathbf{s} contains pressure instead of density. The matrix \mathbf{B} will have values on the diagonal representing the coefficients in either (6.42) or (6.43), with the addition of a block corresponding to the θ part of the linearised ideal gas law (6.24), used to eliminate ρ . Once solutions are found, either in the form of eigenvectors or singular vectors, they may be scaled back to the original form to produce figures with clearer physical interpretation.

6.3 Case studies

Case studies of varying complexity can be formulated from the basic equations described so far, ranging from uncoupled isothermal to fully coupled. This gives a suite of test cases, allowing for careful examination and understanding of the types of behaviour that can occur and how the dynamics and physics interact. If the most complex fully coupled problem was set up and examined initially a significant amount of information would be generated that would be difficult to interpret and understand physically. The cases to consider are,

- Isothermal no shear uniform grid [68].
- Isothermal no shear but using a stretched grid.
- Boundary layer only case. (Covered in Part I).
- Full boundary layer plus dynamics reference state but boundary layer excluded in the transient calculation.
- Full boundary layer plus dynamics mechanisms included in both reference state and transient calculations.

The first case, where only dynamics are included is well understood and so will only be briefly revisited in this study, in part to check how well SVD matches the

eigendecomposition for a normal, skew-Hermitian, but not sign-definite, matrix. It is important to check that physical interpretation of singular vectors equates with that for eigenvectors. It is expected that introducing the stretched grid will have some effect on the solutions of the isothermal case, for example modes may only have structure where resolution is sufficient for their frequency. A grid finer in places than the uniform grid will therefore lead to supporting of some faster propagating modes. Examining this case will help to give an understanding of where modes of certain frequencies will have their dominant structure. The boundary layer only case has been examined in detail in Part I and this has provided an understanding of the decaying modes that the damping mechanism introduces, this will be used to identify these modes in the full coupled case.

The next level of complexity is one where the boundary layer mechanism is introduced into the reference state; this means the reference profiles must represent the fully coupled state. The difference between the two coupled stages will be the terms included in the transients. For both cases all the dynamics terms are included but in the first case all the terms that represent the transient boundary layer structure will be omitted. Doing this provides a set of modes that have simpler physical interpretation. It is anticipated that the complex reference state associated with the boundary layer will introduce new behaviour in the dynamical modes, for example by Doppler shifting of waves or introducing shear instability. By excluding the boundary layer at first it should be possible to quantify the full effect of the reference state first, before the full coupled problem is studied.

Chapter 7

Examination of Mode Types and Interaction

The overall aim of this chapter is to understand the interaction between the dynamics and the physics so that later the Lorenz and Charney-Phillips configurations can be effectively compared. In order to do this it is first helpful to re-visit the more well known dynamics only case. There is little new to say about the choice of grid staggering for this case since it is well established that Charney-Phillips is more suited, it does however provide the clearest insight into the types of dynamical mode. From this case the complexity is built up until an understanding of the modes that exist in the coupled case is gained. Not only does the occurrence of each type of mode need to be identified but also how the dynamics and physics interact. It will be interesting to identify which types of dynamical mode retain their physical properties after coupling and which do not and to identify how and why they change. The isothermal resting state case has so far only been investigated for a uniform grid, this can be extended to examine the effect of employing a stretched grid. Further to this the horizontal wavenumber may be altered to produce different behaviour, this is particularly relevant for coupling to the boundary layer where shorter horizontal wavelengths than that used by [68] may be relevant. For the isothermal case with uniform grid, solutions produce the normal modes; it is in this situation that eigenvectors most effectively describe the system's small

amplitude transient behaviour. As the complexity is built up the system becomes non-normal and the results obtained from the eigendecomposition depart from being the normal modes. The work of Part I showed that the system could be very far from normal for the boundary layer only, meaning eigenvectors did not form an orthogonal basis. This meant that eigenvectors could not be readily interpreted and thus could not be used to effectively compare different vertical configurations. By systematically building up the complexity of the model, as described, it should be possible to quantify the onset of the non-normality and thus the extent of the usefulness of the methodology.

7.1 Isothermal - Dynamics only Case

The following revisits the results of [68]. Here there are no boundary layer terms, the grid is uniformly spaced with $N = 20$ points over a domain 10,000m deep. The steady state is isothermal ($T^{(r)} = 250\text{K}$) and stationary ($U = V = 0$), pressure is given by $P^{(r)} = \exp\left(\frac{-gz}{RT^{(r)}}\right)$, density $\rho^{(r)}$ is obtained from the ideal gas law, the Exner function is obtained most simply as, $\Pi^{(r)} = \left(\frac{p^{(r)}}{p_0}\right)^{\kappa_g}$ and then $\theta^{(r)} = \frac{T^{(r)}}{\Pi^{(r)}}$. The horizontal wavenumber is $k = \frac{2\pi}{10^6}$, giving a horizontal wavelength of 1000km. With this level of approximation one can obtain an analytical description of the dispersion relation and normal modes for each discretisation [64].

The equations are solved with the above approximations as a matrix eigenvalue problem as before. The equations are thus written $\lambda \mathbf{x} = \mathbf{A} \mathbf{x}$, eigenvalues and eigenvectors of \mathbf{A} are obtained giving λ and \mathbf{x} . Now that there is no damping mechanism and the reference state is sufficiently simple the eigenvalues are just frequencies ω and the eigenvectors are normal modes. The normal modes are wavelike solutions, they have sinusoidal structure and vertical wavenumber m . The vertical wavenumber m can be any integer up to that allowed by the resolution multiplied by $\frac{\pi}{D}$, where D is the domain height..

Figure 7.1 and Figure 7.2 show the numerically computed eigenvalues, singular values as well as the analytical dispersion relation for the Charney-Phillips and Lorenz grids. The plots show the frequencies for a given normal mode against the

number of zeros in the p' field of that normal mode and hence figures represent the dispersion relation. The figures use number of zeros in keeping with similar plots in the previous studies, the vertical wavenumber is given by $m = \frac{n\pi}{D}$, where n is the number of zeros in the eigenvector. The three branches, from highest to lowest frequency, represent acoustic modes, inertio-gravity modes and Rossby modes. The vertical wavenumber m of a given mode can be obtained as the number of zeros in p' multiplied by $\frac{\pi}{D}$.




Figure 7.1: Frequencies of the westward propagating modes with the Charney-Phillips grid. The frequency is plotted against the number of zeros in the p' part of that mode and so shows the dispersion relation. The plot shows the frequencies as found by eigenvalues and singular values and as found analytically.

Acoustic modes have the fastest propagation and are thus the highest branch on the figure. Use of either grid staggering option will slow the fastest acoustic modes slightly. For the isothermal case the analytical frequency of an acoustic mode is given by,

$$\omega \approx \pm cm, \quad (7.1)$$

where c is the speed of sound. Acoustic waves may propagate in either an eastward or westward direction. In Figure 7.1 and Figure 7.2 only the westward modes are shown since the eastward modes have very similar behaviour.




Figure 7.2: As for Figure 7.1 but for the Lorenz grid.

The middle branch represents inertio-gravity modes; in the large m limit they tend to inertial modes, in the small m limit they tend to gravity modes. For small k , as being used here, both Lorenz and Charney-Phillips grids accurately capture the inertio-gravity part of the spectrum. Following the approximations made by [64] the analytical dispersion relation for gravity waves is given by,

$$\omega \approx \pm \left(\frac{m^2 f^2 + K^2 N_b^2}{m^2 + K^2} \right)^{\frac{1}{2}}, \quad (7.2)$$

where $K^2 = (k^2 + l^2)$ and $N_b^2 = \frac{g}{\theta} \frac{\partial \theta}{\partial z}$ is the buoyancy frequency, as used in part I. Equation (7.2) is equivalent to the dispersion relation for gravity waves in Boussinesq flow. In equation (7.2) as $m \rightarrow \infty$, $\omega \approx \pm f$ so equation (7.2) is sufficient for representing the analytical inertio-gravity dispersion relation when measuring discrete dispersion relations. As is the case for the acoustic waves, inertio-gravity waves may also propagate eastward or westward, again only the westward are considered in this case since behaviour is so similar in either direction. Note that the behaviour may not be identical between eastward and westward acoustic and inertio-gravity modes due to an asymmetry introduced by using the β -effect. In order to obtain equation (7.2) a number of approximations are made meaning that

β does not appear in the expression of the dispersion relation, however it will influence the actual dispersion relation, see the full dispersion relation in [68] or [64].

The lowest branch and therefore slowest frequency modes are the Rossby modes and here is where the greatest difference between using Lorenz and Charney-Phillips grids is found. Using the Lorenz grid will lead to significant reduction of the frequency of the large m Rossby modes, whereas when using the Charney-Phillips grid the frequency is slightly over predicted, but is closer than with the Lorenz grid. The analytical dispersion relation for the Rossby modes is given by,

$$\omega \approx \frac{-k\beta N_b^2}{K^2 N_b^2 + m^2 \left[f^2 - \left(\frac{k\beta}{K^2} \right)^2 \right]}. \quad (7.3)$$

Note that the frequencies of Rossby modes are negative so in the absence of any forcing they propagate westward.

Figure 7.1 and Figure 7.2 demonstrate one side of the argument presented for why the Charney-Phillips grid is the favoured configuration for modelling the dynamics, the other being the presence of the computational mode. Note that the computational mode does not appear on these plots due it having zero frequency. Of the types of waves supported in the atmosphere Rossby waves are generally considered most important in that they determine large scale weather patterns. It is therefore crucial that a numerical scheme can capture them accurately [34] [66]. Since the Lorenz grid is poorer at representing the frequency of these modes it is less desirable than Charney-Phillips. The other side to the argument against Lorenz, the computational mode, is discussed later. Figure 7.1 is equivalent to Figure 3 in [64]. Comparing Figure 7.2 with Figure 4 in [68] it is clear that introducing the Exner form of the vertical pressure gradient has no adverse or beneficial effect on the discrete dispersion relation when using the Lorenz grid.

Identifying Types of Dynamical Mode

In this isothermal case the type of mode is clear by comparing the numerical and analytical frequencies and by the clear branches that they appear in. Other than this a mode can be identified by examining the eigenvectors and where the dominant features lie. Later when other physical processes are introduced into the equations the frequency of the dynamical modes is likely to change, i.e. they would be Doppler shifted by background wind, leaving less distinction between the three branches as seen in Figure 7.1 and Figure 7.2. It is desirable to ensure that any mode in question can have its structure identified in order to be classified and help later identify how vertical configurations compare. Note that the acoustic modes propagate considerably more quickly so will likely retain a relatively high frequency and remain in a separate branch; this should prove to be a useful way of assessing the methodology.

To assist in the classification of a normal mode one can make use of the energy in the system (this calculation has already been implemented in order to assist in the generation of the singular value solutions to the problem). Just examining the perturbation fields can be slightly tricky as it is not necessarily clear whether a particular field dominates relative to other fields. Instead, consider the contribution to the energy in the system from each field in each eigenvector; a mode of a certain type will have highest energy in the fields associated with that motion. To convert an eigenvector to give the mode energy it can simply be multiplied by the matrix \mathbf{B} in equation (6.45). A further advantage is that this eliminates density in favour of pressure, which has been more commonly used in ascertaining mode type in previous studies [66, 67]. Note that the matrix \mathbf{B} includes $(\Delta z)^{\frac{1}{2}}$ terms from the discrete energy sum. Dividing through by $(\Delta z)^{\frac{1}{2}}$ will give clearer physical interpretation and comparison between different resolutions.

Figure 7.3 shows the mode energy variables for an acoustic mode calculated on the Lorenz grid, the figure shows the real and imaginary parts of each modal energy variable along with the corresponding eigenvalue. An acoustic mode propagates due to perturbations in pressure, i.e. as a local pressure change occurs in the

system it would be immediately restored causing the initial pressure perturbation to propagate away from the source. As this process of restoration occurs air will be required to move (from regions of high pressure to low pressure) hence acoustic waves also require perturbations to velocity. Since this is a vertical only representation of the propagation of the acoustic mode the dominant contribution to the energy would be expected in the pressure and vertical velocity field; examining the figure this can clearly be seen.




Figure 7.3: The mode energy variables for an acoustic mode in the isothermal case. The lower part of the figure shows the frequencies with the eigenvalue corresponding to the plotted mode marked with a black star.

Figure 7.4 and Figure 7.5 show the structure of the inertio-gravity waves in energy variables, for regions of the spectrum representing the small m and large m limit respectively. When a flow is rotationally and buoyantly stable, as it is here, perturbations to particles in the fluid flow will be opposed by the Coriolis force and

the buoyancy force (gravity). These forces act in the horizontal and vertical axis respectively. Consider, for example, a particle moving in a horizontal-vertical two-dimensional plane, clearly some component of the rotational and buoyant forces would act upon it, resisting any perturbation. How close the trajectory of the particle lies to the horizontal or vertical axis will determine the degree to which each force acts upon it in that axis. This explains why frequencies in this branch are seen tending to the Coriolis. As vertical wavenumber increases (corresponding to more zeros in the eigenvector) the scale of the perturbations will become dominant in the horizontal and so the restoring force will be the Coriolis and thus frequencies will lie close to Coriolis. Conversely as the vertical wavenumber decreases (corresponding to fewer zeros in the eigenvector) the scale of the perturbations will become dominant in the vertical. The restoring action of a vertical perturbation is the buoyancy force and thus the frequency of the perturbation will be closer to the buoyancy frequency. The horizontal scale of the perturbation is fixed by choosing k . Decreasing m gives a perturbation with vertically dominant scale, equivalent to the trajectory being close to the vertical axis in the two dimensional analogy. Increasing m gives a perturbation with horizontally dominant scale, equivalent to the trajectory being close to the horizontal axis in the two dimensional analogy. Note that in (7.2) as $m \rightarrow \infty$, $\omega \rightarrow \pm f$ and as $m \rightarrow 0$, $\omega \rightarrow \pm N_b$.

Figure 7.4 shows the inertio-gravity wave for smaller m . The frequency here is larger than f and the contribution to the mode energy is of the same order in horizontal velocity and potential temperature (buoyancy). Figure 7.5 shows a mode from further along the spectrum where the vertical wavenumber is larger, here the frequency is closer to f and the potential temperature energy variable is an order of magnitude smaller than the horizontal velocity energy variables.

The third type of dynamical mode that the system supports represents Rossby waves and these are the most energetically dominant dynamical modes. They evolve slowly and drive the large scale cyclones and anti-cyclones that directly affect day to day weather. It is easiest to consider Rossby modes in potential vorticity terms since their existence is directly due to a gradient in the potential vorticity.




Figure 7.4: As for Figure 7.3 but for an inertio-gravity mode.

If the potential vorticity of a fluid particle is perturbed; then due to conservation of potential vorticity, neighbouring particles are subject to that displacement also. Since potential vorticity is equivalent to the product of absolute vorticity and the stratification a perturbation to the potential vorticity will lead to perturbations in velocity and potential temperature and this is what is observed in Figure 7.6, the energy variables with largest magnitude are the horizontal velocity and potential temperature. In a state without any background wind Rossby waves propagate in the westward direction.

The energy variables that are expected to dominate for each type of dynamical mode are summarised in Table 7.1. Note that a lower case x is shown for E_p in the Rossby mode to denote that it is likely to be larger in a Rossby mode in comparison to an inertio-gravity mode, helping to distinguish between the two. These are the




Figure 7.5: As for Figure 7.3 but for an inertio-gravity mode.

findings for the isothermal case with $k = \frac{2\pi}{10^6}$ and with a $D = 10\text{km}$ domain. For different values of k and D the dominant energy variables may differ.

The number of modes that the system supports is related to the number of degrees of freedom. The number of degrees of freedom is defined by the number of model parameters and the number of grid points N that each parameter is stored on, excluding the boundary conditions. There are N z_ρ levels and $N + 1$ z_w levels but the top and bottom z_w levels are where boundary conditions are implemented. So for each parameter stored at z_ρ there are N degrees of freedom and for each parameter stored on z_w levels $N - 1$ degrees of freedom. When using the Lorenz grid there are therefore $5N - 1$ degrees of freedom. When using the Charney-Phillips grid there are $5N - 2$ degrees of freedom. The number of modes, or solutions, that the system supports is equal to the number of degrees of freedom. For the




Figure 7.6: As for Figure 7.3 but for a Rossby mode.

isothermal case with either grid the modes supported are $2N$ acoustic modes (N east N west), $2N - 2$ inertio-gravity modes ($N - 1$ east $N - 1$ west) and N Rossby modes; the extra mode supported by the Lorenz grid is the computational mode. Analytically it is found that for every internal Rossby mode there exist 2 acoustic modes and two inertio-gravity modes.

7.1.1 Capturing Modes with Singular Vectors

In Part I it was found that the matrix \mathbf{A} was far from normal and that subsequently results from the eigendecomposition were difficult to interpret, instead the SVD was considered. The problems of normality are likely to be encountered again when the fully coupled problem is examined and so it will be useful to call once again on the SVD. Before using the SVD to examine the coupled problem it will be useful

	E_u	E_v	E_w	E_θ	E_p
Acoustic Mode			X		X
I-G Mode (Small m)	X	X		X	
I-G Mode (Large m)	X	X			
Rossby Mode	X	X		X	x

Table 7.1: Summary of the expected dominant energy variables for each type of dynamical mode.

to understand how it performs for the isothermal case and thus how it represents the dynamical modes in the system. Figure 7.7 and Figure 7.8 show the singular vector representation of the acoustic mode shown in Figure 7.3 and the Rossby mode shown in Figure 7.6.




Figure 7.7: Singular vector representation of the acoustic mode shown also in Figure 7.3. The model variables as returned by the SVD are converted to show the energy variables.

The singular values reproduce the magnitude of the eigenvalues almost exactly, as seen in Figure 7.1 and Figure 7.2. However the singular vectors are not identical




Figure 7.8: Singular vector representation of the Rossby mode shown also in Figure 7.6. The model variables as returned by the SVD are converted to show the energy variables.

to the eigenvectors, indeed there remains distinction between the input right and output left singular vectors. In the case where the matrix is normal, Hermitian and sign-definite the eigenvectors would agree identically with the singular vectors [71]. In order to generate the SVD the system has been written in the energy norm. Along with the uniform grid this guarantees normality and that the matrix is skew-Hermitian, i.e. that $A = -A^*$ but not that the matrix is sign-definite. That the matrix is not sign-definite is why both input and output singular vectors do not simultaneously match the eigenvectors identically. Comparing Figure 7.7 and Figure 7.8 with Figure 7.3 and Figure 7.6 it is clear that the singular vectors produce the same overall structure as the eigenvectors and that input and output singular vectors have the same scale. However there are also two clear differences between the input and output singular vectors. Firstly there is a magnitude difference

between them that can be negative and secondly that magnitude difference includes a factor of i . For each energy variable the input singular vector can be multiplied by some factor multiplied by i to obtain the output singular vector.

That the overall structure of singular vectors is similar to the eigenvectors is due to the matrix being normal and both sets of vectors forming an orthogonal basis. Recall that writing in the energy norm produces a skew-Hermitian matrix. The property of a skew-Hermitian matrix is that the eigenvalues are purely imaginary, conversely a property of the SVD is that all singular values are real. It is a property of taking the SVD of a skew-Hermitian matrix that a multiplication by i between singular vectors must occur. If the problem were formulated to produce a Hermitian matrix the i would be absorbed. If the matrix were normal, skew-Hermitian and sign-definite the SVD could be written as

$$i\sigma\mathbf{u} = \mathbf{A}\mathbf{u}, \quad (7.4)$$

where \mathbf{u} is the input singular vector. (7.4) is effectively the same as the eigendecomposition ω is purely imaginary and $i\sigma = \omega$. The magnitude difference between input and output singular vectors cannot be avoided. It would require a sign-definite matrix to see identical behaviour between input and output singular vectors; however a sign-definite matrix has only positive or negative eigenvalues which is not the case here. Note by the similarity in Figure 7.1 and Figure 7.2 that this does not impact on the singular values.

Despite the magnitude differences between input and output singular vectors strong physical interpretation is still possible. Consider the acoustic mode in Figure 7.7, the dominant energy variable is E_w in the output (left) singular vectors and E_p in the input (right) singular vectors, consistent with the physical properties of an acoustic mode, i.e. that perturbations in w are a response to perturbations in p . The magnitude of E_w in the output vectors is equivalent to the magnitude of E_p in the input vectors. For each acoustic mode there is an eastward and westward propagating version; the frequency of the eastward propagating mode will be positive while the frequency of the westward propagating will be negative.

There is a pair of singular vectors corresponding to the pair of eigenvectors, one with E_w dominating in the output vector and E_p dominating in the input vector, as plotted, and one where the opposite is true. The pair of singular vectors correspond to the pair of eigenvectors with the same scale. However each singular vector does not represent either eastward or westward propagation, instead some combination of the two singular vectors represents eastward and westward.

For the inertio-gravity singular vectors (not plotted) behaviour similar to that seen in Part I is seen. Recall that for the boundary layer only model a set of inertial modes with frequency $\pm f$ were identified living above the boundary layer. For the inertio-gravity singular vectors in the isothermal resting state Euler equations the same behaviour is seen. The E_v part of an input singular vector has the same structure as the E_u part of an output singular vector and the E_u part of an input singular vector has the same structure as -1 times the E_v part of an output singular vector. The other important energy variable for inertio-gravity modes is E_θ ; for this energy variable the output singular vector is equal to a factor times i times the input singular vector. Again the pairs of input and output singular vectors correspond to the pairs of eigenvectors with the same scale.

The singular vectors representing the Rossby modes are also clear. All the dominant fields are concurrent with the dominant fields in the eigenvector. Recall from Table 7.1 that the dominant energy variables are expected to be E_u , E_v and E_θ . Examining Figure 7.8 it is clear that correspondence, in terms of magnitude, between one variable in the input and a different variable in the output, as seen in the acoustic and inertio-gravity singular vectors, is not seen. As seen for the potential temperature energy variable in the inertio-gravity singular vectors the input-output response can be understood in terms of the complex parts. For a Rossby singular vector the magnitude of each input variable is similar to the magnitude of the same variable in the output. For every variable the output singular vector is approximately equal to just i times the input singular vector, with a change of sign in E_v and E_w .

It is clear that strong physical interpretation and understanding is possible with

the singular vectors and that they are capable of showing the normal modes of the system. When modes occur in eastward and westward pairs, as is the case for acoustic and inertio-gravity modes, interpretation is provided through an input-output response. For a set of modes that can only propagate in one direction, as is the case for Rossby modes, the interpretation is equivalent to if the matrix were sign-definite.

7.1.2 Lorenz Grid Computational Mode

The Lorenz grid has an extra degree of freedom compared to the Charney-Phillips grid. When using the Charney-Phillips grid θ' is free to vary on $N - 1$ of the z_w levels, however on the Lorenz grid it is free to vary on all N of the z_ρ levels. This extra degree of freedom allows for a computational mode to exist. Examining the vertical momentum equation (6.3) it is clear that when using the Lorenz grid θ' will require averaging to the z_w levels, where it is multiplied by the vertical gradient of Π . Should a perturbation to θ' be a two grid wave, the averaging would result in $\theta' = 0$. This means a transient solution where θ' is a two grid wave and all other variables are zero may exist. The overall solution would be non-physical since θ' would be invisible to the dynamics. The resulting mode would have zero frequency and hence fail to propagate. This non-physical static mode could then start to interfere with other structures in the solution and lead to model errors. The computational mode is associated with the averaging in equation (6.3). Figure 7.9 shows the computational mode for the isothermal case, the two grid wave in θ' is clear. The structure of this mode is dominated by the potential temperature, however the two grid wave signal is seen also in the density field and the Π' field (not plotted).

7.1.3 Shorter Horizontal Wavelengths

So far the results have been calculated using the same horizontal wavenumber as used by [68], $k = \frac{2\pi}{10^6}$. This gives a horizontal wavelength of 1000km. This is quite large considering the increasing horizontal resolution in numerical models




Figure 7.9: This plot shows the eigenvector whose corresponding eigenvalue is zero. This is the computational mode, signified by a two grid wave in the θ field.

and relatively shallow boundary layer depth of 1km or less. Instead the horizontal wavenumber can first be increased to $k = \frac{2\pi}{10^4}$ giving a horizontal wavelength of 10km. The frequencies with this wavelength are shown in Figure 7.10 and Figure 7.11 for the Lorenz and Charney-Phillips grids.

As noted by [64] the Rossby waves for this horizontal wavenumber are represented more accurately by the discrete Charney-Phillips and Lorenz grids than in the longer wavelength case. However the inertial end of the inertio-gravity waves are represented less accurately. It can be seen that when using the Exner form of the pressure gradient, in conjunction with this horizontal scale, the Charney-Phillips configuration that has density as a prognostic variable performs as well as the Charney-Phillips configuration with pressure, i.e. Figure 7.10 against Figure 4 in [64].

That there is a difference in the frequencies of the inertio-gravity and Rossby waves is due to relative sizes of m and k . When considering the dispersion relations of both inertio-gravity and Rossby waves (7.2) and (7.3) it is clear that certain terms will dominate. For example in the inertio-gravity dispersion when K is relatively small, as it would be in the original case, any discretisation error in




Figure 7.10: Charney-Phillips: as for Figure 7.1 but with a horizontal wavelength of 10km.




Figure 7.11: Lorenz: as for Figure 7.2 but with a horizontal wavelength of 10km.

computing the buoyancy frequency will have little effect, when K is large it will, giving slowed frequencies ω . Similarly the Rossby frequency reduces to $-\frac{\beta}{k}$ for $l = 0$ and large enough k . This is approximately constant and the terms that can pick up errors from vertical discretisation become negligible, giving more accurate

representation of the dispersion relation. The Lorenz grid performs poorly against the Charney-Phillips grid for the inertio-gravity waves in this case for the same reasons that it performed badly against Charney-Phillips for the Rossby waves in the longer wavelength case. The terms that need averaging when using the Lorenz grid introduce errors into the buoyancy frequency, but whether or not this term dominates in the dispersion relation is dependent on the relative size of horizontal wavenumber in relation to the size of the vertical wavenumber.

7.2 Isothermal Case with Stretched Grid

Almost all operational models employ some kind of vertical stretching of the grid. So that results generated subsequently for the stretched grid coupled case can be fully reconciled with those generated thus far it is of interest to first check the effects of using a stretched grid on dynamical modes. This can be done for the original smaller horizontal wavenumber and the larger wavenumber. The grid is geometrically stretched with all levels stretched smoothly, rather than by placing any halfway between others.

Figure 7.12 and Figure 7.13 show the dispersion relation for the Charney-Phillips and Lorenz grids whilst employing the geometrically stretched grid that was designed in Part I. From the figures it appears that the discrete modes are gaining large inaccuracies in the dispersion of the large m waves. It has to be noted that it is not actually fair to compare the set of waves that are captured by the stretched grid with those waves captured either by the uniform grid or even those found analytically. As the wavenumber increases (and the wavelength decreases) there becomes a region in the domain where the stretched grid has insufficient spatial resolution to capture the wave. For small m both uniform and stretched grid will capture the same wave and the plots reveal that here the stretched grid has as good dispersion properties as the uniform grid. At some point along the discrete spectrum the wavelength will become smaller than the maximum grid spacing of the stretched grid and the overall wave that is captured will be different. Indeed the stretched grid will be capable of distorting waves to give much higher wavenum-

ber than the uniform grid, at least in the region of increased resolution near the ground. The frequency of waves with a given large m in Figure 7.12 and Figure 7.13 cannot be fairly compared since the waves captured by the stretched grid have much smaller wavelength, indeed the wavenumber m is height dependent for the stretched grid. When using a stretched grid it is not sensible to count zeros in the eigenvector due to the more inhomogeneous structure. Instead the x -axis refers simply to an index number. Index number increases as, depending on the type of mode, the frequency of the wave it corresponds to increases or decreases. It is not even clear which analytical wavenumber the stretched grid is trying to capture since the scale can be height dependent.

By examining the corresponding eigenvectors for the stretched grid it is clear that the scale of the modes decreases as the index number increases, so the frequency plots show the trend in the dispersion relation, but for unknown corresponding wavenumber. As vertical wavenumber m increases the frequency of acoustic waves increases, the high wavenumber modes captured by the stretched grid propagate much faster than the highest wavenumber modes captured by the uniform grid. For the inertio-gravity waves the frequency tends to the value of the Coriolis parameter as m increases, hence the frequencies appear unchanged for the stretched grid. For the slowest propagating Rossby modes the Lorenz grid gives smaller frequencies than the Charney-Phillips grid. Although there is no analytical version of the stretched grid frequency it is likely that the Lorenz is capturing the frequency less accurately than Charney-Phillips. The slowing of the frequency of the Rossby modes is due to the averaging in θ' . This in turn introduces an error in the Buoyancy frequency which occurs in (7.3). As $m \rightarrow \infty$ the Buoyancy frequency dominates in the dispersion relation and the error that it contains because of averaging causes the slowing. Having a stretched grid will not prevent this from occurring. Note that for the stretched grid case Lorenz is marginally better for the modes with index 2, 3 and 4.




Figure 7.12: Charney-Phillips: as for Figure 7.1 but with a geometrically stretched grid.




Figure 7.13: Lorenz: as for Figure 7.2 but with a geometrically stretched grid.

7.2.1 Lorenz Grid Computational Mode

The stretched grid alone will not prevent the presence of the computational mode, however it is useful just to examine the structure that it has. Figure 7.14 shows the computational mode with the stretched grid. It is unmistakably the computa-

tional mode, consisting of a two grid wave throughout the domain and having zero frequency.




Figure 7.14: Lorenz grid computational mode on the stretched grid.

7.2.2 Shorter Horizontal Wavelength

Figure 7.15 and Figure 7.16 show the dispersion relation for the Charney-Phillips and Lorenz grids respectively. The stretched grid is combined with the shorter horizontal wavelength $k = \frac{2\pi}{10^4}$. This is the furthest the simpler isothermal and no shear case can be taken towards the coupled case and so gives the closest representation of the dynamical modes than can be expected in the coupled case.

The figures highlight the combination of the effects shown in Figure 7.10, Figure 7.11, Figure 7.12 and Figure 7.13. The decrease in grid spacing afforded by the stretched grid allows for faster acoustic modes and slower Rossby and inertio-gravity modes than can be supported by the uniform grid.

The difference in the frequency of the high vertical wavenumber acoustic waves introduced by the larger horizontal wavenumber is similar for the stretched and uniform grids.




Figure 7.15: Charney-Phillips: As for Figure 7.10 but with a stretched grid.




Figure 7.16: Lorenz: As for Figure 7.11 but with a stretched grid.

The effect of the increased horizontal wavenumber was to give worse representation of frequencies at the inertial end of the inertio-gravity branch. Increasing k means that the term involving the error prone buoyancy frequency becomes important in the inertia-gravity dispersion relation. The stretched grid reduces m further increasing the noticeable error in the frequency.

The combination of better Rossby dispersion achieved with the smaller wavelength is somewhat negated by the increased m afforded by the stretched grid and again the Lorenz grid is doing poorly compared to Charney-Phillips in the very large m limit. When the larger k is combined with the uniform grid $k \gg m$ and the frequency of Rossby modes reduces to $-\frac{\beta}{k}$. When a stretched grid is used the assumption that $k \gg m$ is less accurate in the large m limit, the errors in the buoyancy frequency return into the dispersion relation. In any of these stretched grid plots it is not fair to compare the discrete dispersion relation with the analytical dispersion relation.

7.2.3 Mode Structure

While changing the horizontal wavenumber will have little effect on the structure of the isothermal modes, other than to alter how much certain fields dominate, introducing the stretched grid is likely to have quite significant effect on the structure. As mentioned previously the stretched grid will capture a largely different looking set of waves to the uniform grid. The modes whose scale is larger than the maximum spacing on the stretched grid will be captured by both but modes with smaller scale than this will differ. On the uniform grid the wavenumber will increase as integers times $\frac{\pi}{D}$, on the stretched grid the set of captured modes will also decrease in scale but cannot be easily associated with a given wavenumber since the scale may differ with height.

Understanding how the dynamical modes are captured by the stretched grid will further assist in identifying them when the whole coupled case is computed on a stretched grid. Figure 7.17 shows an example of an inertio-gravity mode. This is a mode which has relatively high wavenumber and so represents a wave whose restoring mechanism would normally be dominated by rotational forces. However the figure shows the mode captured on the stretched grid to have similar magnitude in both horizontal velocity and potential temperature; this is also found for the shorter horizontal wavelength and a uniform grid. The high wavenumber modes will have their structure dominant near the surface where the resolution is sufficient

for capturing it. Although for a longer horizontal wavelength the mode shown in Figure 7.5 has the same structure as the mode that would be captured in place of this one if using a uniform grid. Clearly the mode captured by the stretched grid in Figure 7.17 has structure that is vastly different to the overall structure of the mode captured by the uniform grid in Figure 7.5. The stretched grid acts to distort the structure of the mode, giving decreased scale near the ground where resolution is finer.




Figure 7.17: As for Figure 7.5 but the equivalent mode that is captured by the stretched grid and when using a horizontal wavelength of 10km.

Figure 7.18 shows a Rossby mode. This is for the part of the spectrum where the stretched grid and uniform grid are able to resolve the modes in an equivalent manner. Again the comparison with the uniform grid version of this mode in Figure 7.6 is across different wavenumbers. However the similarity in the overall structure

can be seen.




Figure 7.18: The Rossby mode as for Figure 7.6 but with a stretched grid and horizontal wavelength of 10km.

An important feature of the stretched grid modes, most noticeable in the large m limit, is that the mode has smaller amplitude in the higher part of the domain where the grid spacing is much larger than the mode scale. It might be expected that the grid would sample some underlying wave structure of the solution and therefore have amplitude throughout the domain but in a seemingly random way as the wave is aliased. This would be the case if a standing wave were simply sampled by a stretched grid. However the behaviour has to be understood in terms of propagation. The structure of a propagating wave will dominate in the region only where sufficient resolution exists for supporting its frequency. Many phenomena have been observed related to the propagation of waves on stretched grids. For

example, [75], where waves are found to be reflected as they propagate into regions of insufficient resolution for their wavelength. This is likely the effect that is being observed in Figure 7.17, here the inertio-gravity wave is being vertically reflected at the point where its group velocity goes to zero, which is where $\frac{\partial\omega}{\partial m} = 0$.

Capturing with Singular Vectors

Now that a stretched grid is being used for the computation the problem will be further from normal and the SVD will depart further from the eigendecomposition. Although the solutions are different, physical interpretation is still possible and this is demonstrated in Figure 7.19, which shows the same Rossby mode that is shown in Figure 7.18. Clearly the mode structure as found by the eigenvector and the singular vectors is similar. Further to this the energy variables that dominate, meridional velocity and potential temperature followed by pressure, are equivalent, although two orders of magnitude larger for the singular vectors. The relationship between input and output singular vectors that was observed for the uniform grid Rossby singular vectors is also seen here. There is multiplication by i between output and input as well as similar magnitude for the input and output of each energy variable. The only field that appears to be different in the singular vectors is E_u , this is likely due to the relative size of this field when the shorter horizontal wavelength is used. Throughout all the modes the singular vectors are found to agree well with the eigenvectors in terms of structure and dominant field. Indeed for all modes the type of wave can be classified by examining the singular vectors.




Figure 7.19: Singular vector representation of the Rossby mode as shown by eigen-vector representation in Figure 7.18. The model variables as returned by the SVD are converted to show the energy variables.

7.3 Coupled Reference State, Dynamics only in the Linearisation

The previous sections of this chapter have examined the isothermal case. This is a useful example to consider, not only because it reveals which vertical configuration is preferable for the large scale dynamics, but also because it provides a sufficiently simple set of equations to give normality in the system and therefore test the methodology in the optimal environment. Combining the ideas of Part I with the ideas discussed so far in Part II, a full understanding is gained of the types of mode that are possible when the parts of the system are dealt with separately. The challenge, before a comparison is performed of the configurations for the coupled problem, is to understand exactly how the different types of mode behave and interact with each other when coupled.

For the remainder of this chapter attention will be restricted to the Lorenz case, for now it is assumed that the modes supported by the Charney-Phillips grid will have overall behaviour and interaction that is similar to the Lorenz grid. Here the computational mode that the Lorenz grid supports can be fully examined to assess how the addition of the boundary layer influences it.

The addition of the boundary layer significantly changes the reference state of the model. Both temperature $T^{(r)}$ and horizontal velocities U and V will be dependent on height z , this will have large influence on the behaviour and structure of each mode type. Varying levels of complexity for linearising around the steady state can be designed to help break the problem down and understand the interactions. The first case considered in this section is one where only the dynamics and no boundary terms are considered in the transient part of the linearisation; the fully coupled equations are used to generate the steady state but the perturbations to all the boundary layer terms are ignored. Ignoring the boundary layer terms means that only dynamical modes will be supported in the transients, allowing their structure and evolution to be gauged in the presence of the more complex background flow.

In the next section, when the boundary layer is switched on, the behaviour of some dynamical modes will be affected. This will be represented by solutions that are damped by the boundary layer diffusion or result from the coupling. This section aims to build an understanding of how the dynamical modes are affected by the boundary layer background flow so that later the boundary layer type behaviour can be distinguished from the behaviour of the dynamics.

Note that now the boundary layer is included, the background flow varies with height, meaning the system will be further from normal. Solutions can be thought of just as eigenmodes rather than normal modes since eigenvectors will not form an orthogonal basis. For simplicity they remain being referred to as modes. Further, extra advection terms in the material derivatives means writing in the energy norm will not guarantee a normal and Hermitian matrix and so larger differences are anticipated between the eigendecomposition and the SVD for the boundary layer on in the reference state case.

The steady state is obtained by finding the solution to equations (6.8)-(6.11). The steady states for the coupled problem are quite similar to the steady states for the boundary layer only problem but with the addition of a steady profile for the density. Figures showing the structure of the coupled steady states are omitted here and shown when comparing the Lorenz configuration to the Charney-Phillips configuration in the next chapter. Some of the behaviour of the modes depends on the background flow, particularly the background shear and stratification. Overall structure of the steady states can be recalled from Figure 3.6 and Figure 3.7 on pages 90 and 91 when interpreting mode behaviour.

The transient solution comes from solving the eigenvalue form of equations (6.26)-(6.30) but with all the boundary layer terms omitted.

7.3.1 Dispersion Relation

Figure 7.20 shows the absolute imaginary part ω of the eigenvalues λ and also the singular values σ for the shallowest boundary layer, boundary layer 1. Recall from Part I that this kind of boundary layer has an approximate depth of 100m.




Figure 7.20: Absolute value of the imaginary part of the eigenvalues and the singular values for the Lorenz grid. Steady state generated for boundary layer 1, no boundary layer mechanism in the transient calculation. Met Office grid.

The grid used here is the stretched Met Office grid. In this figure the eigenvalues are ordered by descending absolute imaginary part and the index number is just a reference. Previously it was clear that figures such as Figure 7.1 represented wavenumber against frequency and therefore gave some visual representation of the dispersion relation, even if no wavenumber was directly obtainable such as when shown for a stretched grid. Although likely to have some relation it can no longer be assumed that ordering by imaginary part will automatically order by vertical wavenumber. To help determine how well the figure represents the dispersion relation, the spectrum must be sampled and the corresponding eigenvectors examined. Once branches are confirmed to represent different mode types and the mode structures have been checked for approximate wavenumber the shape of the dispersion relation can be more closely identified.

Previously only the westward propagating acoustic and inertio-gravity modes were considered, the westward set also contained the Rossby modes and the eastward counterparts had almost identical structure. The increased complexity in the background flow means differences between eastward and westward modes are more distinguished and so both will be considered.

The immediate change observed when using the boundary layer reference state

is that Rossby modes are now all eastward propagating whereas were previously all westward propagating. This is due to the wind above the boundary layer in the background flow dominating the relatively small phase velocity of the Rossby modes, Doppler shifting them to an overall eastward propagation. Similarly note that the computational mode is advected by the background flow and no longer exhibits a zero frequency but is mixed in with the other branches of mode and becomes eastward propagating. In checking structures by eye it is found that in Figure 7.20 the computational mode corresponds to the eigenvalue with index 39. Whether a given wave propagates eastward or westward depends on the type of solution, the wind speed in the background flow and in the choice of the horizontal wavenumber. For example in Figure 7.20 the horizontal wavenumber is $k = \frac{2\pi}{10^6}$ whereas in Figure 7.21 everything remains fixed except that $k = \frac{2\pi}{10^3}$, the smaller horizontal scale allows the background flow to dominate, causing all but one Rossby mode and all inertio-gravity waves to propagate eastward. Note that only acoustic modes lie in a distinct branch, other have to be classified by examining the mode structure. In following the work of [64] in the previous section the horizontal wavenumber was chosen as $\frac{2\pi}{10^4}$. This is still quite large considering the 100m deep boundary layer and so $k = \frac{2\pi}{10^3}$ is used, giving a 1km wavelength.




Figure 7.21: As for Figure 7.20 but with the larger wavenumber.

Comparing Figure 7.20 and Figure 7.21 it is clear the choice of horizontal

wavenumber is quite significant when including the boundary layer in the reference state. The frequencies for all of the inertio-gravity and Rossby solutions are within a much smaller range now, as the Doppler shifting dominates their natural frequencies. When a wave is Doppler shifted the change in frequency is inversely related to the wavelength. For the same background flow a wave with smaller horizontal wavenumber will undergo a smaller Doppler shift than a wave with larger horizontal wavenumber.

The fact that waves are being Doppler shifted so much by the background flow makes it even harder to visually represent the dispersion relation. If all the waves have similar frequencies then it becomes more likely that ordering by frequency, as is most natural to do, will produce an order unrelated to vertical wavenumber, especially since the two branches of Rossby and inertio-gravity modes are observed to be becoming entangled. Again a way to get around this, and understand to what level the figure gives a visual representation of the dispersion relation, is to examine the mode structure by eye and try to classify their type and place on the spectrum.

A further issue potentially preventing modes appearing in an order related to vertical wavenumber is due to the structure that can result due to the complex reference state. Gravity waves, for example, can only exist in regions of stable stratification, and since the background flow is close to neutrally stratified above the boundary layer, they may have the majority of their structure in the boundary layer. Meanwhile the inertial waves may have the majority of their structure above the boundary layer where rotational effects dominate over shear; despite both waves being represented by the same underlying dispersion relation it may not be possible to relate wavenumber to frequency clearly.

The branch of acoustic solutions is clear in both wavenumber cases since the frequencies of acoustic waves are relatively fast, making them less susceptible to Doppler shifting. The acoustic branch is the highest fastest frequency branch. Further to this it is likely that their structures will depend less on the background flow since they do not rely on horizontal velocity or potential temperature fields in

order to propagate. Indeed this is the only part of the frequency plot that reliably represents the dispersion relation. The conjugate pairs of eastward and westward propagating acoustic modes are clearly visible in the spectrum. That the acoustic modes are still well represented gives further confidence in the numerics and the methodology. For example it appears to exhibit a discrete spectrum and this can be checked and used for comparison. In addition it may help to measure how well the singular values can perform, if for example they were not capable of representing even this part of the spectrum it would be a strong warning against their use for the coupled problem.

7.3.2 Unstable and Decaying Modes

Now that shear is present in the reference state, it is possible to have Kelvin-Helmholtz type instability. These instabilities occur as shear can act to advect the crests of a horizontally propagating wave faster than the troughs [19]. The shear driven decay is effectively the opposite of this process. When modes are decaying or unstable it is represented by the eigenvalues having a real part, as was seen when the boundary layer was studied in Part I. Figure 7.22 shows the eigenvalues for the $k = \frac{2\pi}{10^6}$ and $k = \frac{2\pi}{10^3}$ cases plotted in the complex plane. It is clear that there now exist eigenvalues that have real part that may be positive or negative. Eigenvalues in the left half-plane correspond to damped modes, eigenvalues in the right half-plane correspond to growing modes, upper half-plane corresponds to eastward propagation and lower half-plane westward propagation.

Firstly note that the growth rate of the majority of modes is relatively small. Secondly note that it is clear from the figure that the extent of the shear induced instability or decay is dependent on the horizontal wavenumber and that modes in the shorter horizontal wavelength case can have larger growth or decay rates. The shear acts to roll up the fluid, creating the possibility of instability, if the scale of motion is similar in the vertical compared to the horizontal then the rolling up is stronger. When the horizontal wavenumber is larger the horizontal and vertical scales are closer and hence why larger real part is seen for this case.




Figure 7.22: Eigenvalues for the two horizontal wavenumber cases plotted in the complex plane. Boundary layer switched off in the transients.

As the boundary layer is switched on the solutions which correspond to the boundary layer diffusion will likely have a negative real part. It is therefore important to be able to distinguish between solutions that are damped due to the shear and those that are damped by the boundary layer.

Figure 7.23 and Figure 7.25 show the real part of the eigenvalues found for the $k = \frac{2\pi}{10^6}$ and $k = \frac{2\pi}{10^3}$ cases respectively. Figure 7.24 and Figure 7.26 show the corresponding frequency plots but with the modes with largest growth or decay shown in black. Recall that the transient solutions are proportional to $\exp(\lambda t)$ where $\lambda = \mu - i\omega$. Those plotted in black in Figure 7.24 and Figure 7.26 have $|\mu| > 10^{-8}$.

For the smaller horizontal wavenumber case the modes with largest real part appear to correspond to the inertio-gravity branch with a few Rossby modes affected. In the larger wavenumber case modes undergo stronger damping or decay and thus more modes appear in black for the same condition on the magnitude of the real part. Based on their frequencies the modes undergoing the largest growth or decay appear to be inertio-gravity and Rossby modes. However this needs to be checked by looking at the structures since frequency does not necessarily define mode type. In fact, for $k = \frac{2\pi}{10^3}$, inertio-gravity and Rossby modes are found to mix

into a single branch. The Doppler shift and general advection by the background flow is large enough to give Rossby modes with faster frequency than inertio-gravity modes.




Figure 7.23: Real part of the eigenvalues for the longer horizontal wavelength case ($k = \frac{2\pi}{10^6}$). Boundary layer 1 reference state, boundary layer switched off in the transients. Sorted by real part but absolute real part plotted




Figure 7.24: Imaginary part of the eigenvalues for the longer horizontal wavelength case. Boundary layer 1 reference state, boundary layer switched off in the transients.

The modes with fastest decay or growth in Figure 7.22 need to be identified.




Figure 7.25: Real part of the eigenvalues for the longer horizontal wavelength case $k = \frac{2\pi}{10^6}$. Boundary layer 1 reference state, boundary layer switched off in the transients.




Figure 7.26: Imaginary part of the eigenvalues for the longer horizontal wavelength case. Boundary layer 1 reference state, boundary layer switched off in the transients.

For $k = \frac{2\pi}{10^6}$ the largest growing and decaying eigenvalues have index 29 and 28 respectively in Figure 7.24. For $k = \frac{2\pi}{10^3}$ the largest growing and decaying eigenvalues have index 47 and 48 in Figure 7.26. In both horizontal wavenumber cases the modes structures are examined and found to correspond to the same largest

scale inertial modes. In both cases it is a very similar pair of modes, the dominant energy variable is horizontal velocity and all the structure is in the region above the boundary layer. For the $k = \frac{2\pi}{10^6}$ case the Doppler shift is less and it is still possible to distinguish between direction of propagation. The unstable of the pair is the eastward propagating, and the decaying version is westward propagating. Since the modes that appear most susceptible to the shear are the inertial end of the inertio-gravity modes it should be possible to distinguish between boundary layer modes and modes decayed by shear. In Part I the inertial modes were purely neutral and so unaffected by the boundary layer diffusion.

The shear instability and decay explains why the singular values do not match the frequencies so well in the shorter horizontal wavelength case, where growth and decay rates are larger. Since singular values are real the SVD will not be capable of representing propagating and growing/decaying modes in an equivalent manner to the eigendecomposition. Instead singular values may tend to represent either frequency or growth/decay rate. This will be investigated when the boundary layer is switched on in the transients. When examined it is found that for the modes corresponding to singular values and eigenvalues that do not match in Figure 7.24 the corresponding real part of the eigenvalue will be quite large. Instead of representing the frequency of the mode the singular value represents something else, which may be a combination of the growth/decay and frequency.

In this section the changes to the frequencies of the dynamical modes that are caused by the coupled reference state have been observed. However a number of questions have been raised relating to the classification of modes. In order to fully appreciate the frequency plots and generate an understanding of the dispersion properties of the system the eigenvectors clearly need to be examined. This also needs to be done in Section 7.4 when the fully coupled linearisation is considered. Of course this increases the complexity of the problem in that some modes will become damped by the boundary layer, however for many modes the structures will remain as they are and these are of less interest anyway. To avoid repetition the mode structures are examined and classified only for the later fully coupled case.

The modes that are of interest in the case examined in this section are those that are liable to be distorted and become boundary layer modes in the fully coupled problem. As these are identified it will be interesting to examine them.

7.4 Fully Coupled Linearisation

When linearising about the fully coupled reference state with only the dynamics and not the boundary layer the solutions are all dynamical structures, either acoustic, inertio-gravity or Rossby waves. When the boundary layer is turned on in the linearisation some of the solutions will correspond to the damping by the boundary layer. The efforts in this section will be directed towards identifying the different types of solutions based on their dominant structure, examining what the spectrum looks like and attempting to establish which types of modes are most likely to be influenced by the boundary layer. This also provides a useful analysis of how well the different methodologies work for classifying solutions and thus how useful they may be for comparing vertical configurations.

Figure 7.27 shows the eigenvalues in the complex plane for the fully coupled problem. Clearly the eigenvalues for the larger horizontal wavelength case have shifted so that now damped modes dominate modes driven unstable by the shear. This would imply that the boundary layer damping dominates over the shear driven decay or instability. For the shorter horizontal wavelength the influence of the boundary layer is less clear, indeed there is a mode with growth rate equivalent to the fastest damped mode in the ‘no transient boundary layer’ case.

The information in this figure needs to be broken down as before to establish the exact behaviour of the modes. Points that need to be addressed include: are the damped modes due to boundary layer diffusion or shear driven decay? If modes have been distorted by the addition of the boundary layer in the linearisation then which modes are they? How important is the horizontal wavenumber and are shorter horizontal wavelength solutions less susceptible to the boundary layer diffusion as would be suggested by Figure 7.27?

The fastest growing mode in the right hand plot ($k = \frac{2\pi}{10^3}$) of Figure 7.27 is




Figure 7.27: Eigenvalues for the two horizontal wavenumber cases plotted in the complex plane. Fully coupled transients.

still an inertial mode, as it was for the case with no transient boundary layer (Figure 7.22), however it is a slightly different inertial mode than that which was found to be the most unstable mode previously. The mode that was previously the fastest growing is now one of the fastest decaying modes. Based on findings in Part I an inertial mode would not be expected to be decayed by the boundary layer diffusion. For the boundary layer only all the inertial modes had zero real part. Further, in the previous section inertial modes were identified as being susceptible to the shear. In addition the inertial mode that is the fastest growing mode here was almost neutral when the boundary layer was switched off in the transients. This suggests increased difficulty in interpreting the $k = \frac{2\pi}{10^3}$ case; clearly there is a high degree of sensitivity.

Figure 7.28 and Figure 7.29 show the real (growth/decay) and imaginary (frequency) parts of the eigenvalues respectively for the $k = \frac{2\pi}{10^6}$ case. The most damped eigenvalues have been highlighted in black in each plot and correspond to each other. These are the modes that are most likely associated with the boundary layer diffusion mechanism since the decay rates are greater than that of any mode found when the boundary layer was switched off in the transients. The overall frequency plot is similar to that in Figure 7.20, when no boundary layer was included

in the linearisation.




Figure 7.28: Real part of the eigenvalues for the $k = \frac{2\pi}{10^6}$ case. The reference state is found using boundary layer 1 and the boundary layer is switched on in the linearisation. Modes are sorted by real part and the magnitude of the real part is plotted.




Figure 7.29: Imaginary part of the eigenvalues and the singular values for the $k = \frac{2\pi}{10^6}$ case. These eigenvalues correspond to the those in Figure 7.28 but are ordered by imaginary part here.

Figure 7.30 and Figure 7.31 show the real and imaginary parts of the eigenvalues for the $k = \frac{2\pi}{10^3}$ case. This also shows the most damped modes plotted in black. For

the same condition on the size of the real part of the eigenvalue, i.e. that it is less than -10^{-6} , there are approximately the same number modes. However there are eigenvalues with comparable magnitude that are growing and the overall magnitude of the real part of the eigenvalues is similar to that seen when no boundary layer terms were included in the transient part of the linearisation, as seen in Figure 7.25. The growth of the fastest growing mode is slightly reduced with the boundary layer switched on from 3.902×10^{-3} to 3.596×10^{-3} .

From Figure 7.28 to Figure 7.31 a general picture of how the boundary layer affects the solution is emerging. The plots of eigenvalues suggest that solutions with longer horizontal wavelength have less susceptibility to the shear, whereas solutions with shorter horizontal wavelength will have more. Even with the boundary layer switched on, modes in the $k = \frac{2\pi}{10^3}$ case remain being decayed or driven unstable by the shear. This can be seen in Figure 7.31 where modes plotted in black, denoting decay, include parts of the acoustic branch. Also, when structures were examined some fast decaying modes were found to be inertial modes. If acoustic and inertial modes are being decayed it is unlikely that it is due to boundary layer diffusion since acoustic modes have very high frequencies and inertial modes have all their structure above the boundary layer. The real part of eigenvalue for the inertial and acoustic modes when the boundary layer is switched on in the transients is of a similar magnitude to the real part that they had when the boundary layer was switched off in the transients.

It will be interesting to investigate the structures of the modes highlighted in black which share the characteristics of the modes found in the boundary layer only case. So far those plotted in black are only based on a simple criterion. The condition for highlighting in black is based on the $k = \frac{2\pi}{10^6}$ case. The condition is designed to pick out the modes which have larger decay rates when the boundary layer is switched on than the decay or growth rates of any of the modes found when the boundary layer is switched off. So the magnitude of the eigenvalues plotted in black in Figure 7.28 is larger than the magnitude of all the eigenvalues plotted in Figure 7.23. So far only the same condition has been applied to the $k = \frac{2\pi}{10^3}$




Figure 7.30: Real part of the eigenvalues for the $k = \frac{2\pi}{10^3}$ case. The reference state is found using boundary layer 1 and the boundary layer is switched on in the linearisation. Modes are sorted by real part and the magnitude of the real part is plotted.




Figure 7.31: Imaginary part of the eigenvalues and the singular values for the $k = \frac{2\pi}{10^3}$ case. These eigenvalues correspond to the those in Figure 7.30 but are ordered by imaginary part here.

figures. For the shorter horizontal wavelength it is likely to be more difficult to pin down which modes are boundary layer modes just by looking at eigenvalues. Of course for both cases the modes needs to be investigated in detail, it may be

that not all of these modes are boundary layer modes or that more than these are. To identify whether a mode is a boundary layer mode, one must examine the structure of the eigenvector, if it has structure similar to the boundary layer modes found in Part I and is damped then it is considered to be a boundary layer mode. This type of mode will be important when it comes to comparing the Lorenz and Charney-Phillips grids.

The singular values are no longer capable of representing the frequency of all the modes in either wavenumber case. This is again due to the increased damping of certain modes and the inability of the singular values to capture both the decay rate and the frequency. Whether or not the singular values and singular vectors are capable of representing the structure and behaviour of the modes in a similar way to the eigenvalues and eigenvectors can only be determined by examining the mode structures themselves.

7.4.1 Mode Structures

Longer Horizontal Wavelength ($k = \frac{2\pi}{10^6}$)

Starting with the longer horizontal wavelength case ($k = \frac{2\pi}{10^6}$), and moving along the frequency plot in Figure 7.29 from index 1 to 49, the mode structures are examined and classified. The internal acoustic modes are those that have index number on the frequency (imaginary) plots going from 1 to 18. The structure of the acoustic modes is similar between the coupled case and the isothermal case in that they consist of a dominant structure in the vertical velocity and pressure. The structure is also only dominant in regions with sufficient resolution for their frequency and vertical structure. Plots of the structures are omitted for their similarity to the isothermal examples.

The modes with imaginary index 19 and 20 in Figure 7.29 are external modes that have structure similar to the external acoustic modes found in the isothermal case. They also exist as a conjugate pair and have no significant real part and so, although the frequencies are considerably slowed by the Doppler shifting, compared to the isothermal case, they are almost certainly the external acoustic modes. The

next two modes have imaginary index numbers 21 and 22, these have dominant structure lying near the surface in horizontal velocity and potential temperature and are also highlighted in black in Figure 7.29 implying that they are likely damped by the boundary layer. All the boundary layer mode candidates will be examined together in detail later.

Figure 7.32 shows the structure in energy variables of a typical inertio-gravity mode in the coupled problem. The figure is arranged so as to show the structure of the mode along with the position of that mode in the eigenvalue spectrum, as arranged by frequency and by growth rate. The eigenvalue plots are given to show position of the corresponding mode and are equivalent to Figure 7.28 and Figure 7.29. The mode shown in Figure 7.32 is more towards the inertial end in that the structure is dominant in the horizontal velocity over the potential temperature and that $\frac{k}{m} \ll \frac{f}{N_b}$. This is also evident by the fact that the majority of the structure lies above the boundary layer, whose height is at around 100m. This was also the case in the boundary layer only problem, where a set of inertial modes with frequency close to Coriolis were found to have structure entirely above the boundary layer. The structure of the inertial modes will dominate above the boundary layer where the shear goes to zero and is dominated by rotation.

The inertial mode plotted in Figure 7.32 lies in a branch of other inertial type modes going from imaginary index 23 to 30 (with the exception of 24), these are all eastward propagating. The scale of the mode increases as frequency increases, consistent with the inertio-gravity wave dispersion relation. A set of very similar westward propagating inertial modes are found in the branch below, imaginary index going from 31 to 38 (with the exception of 36), again scale increases as frequency increases. The two modes lying in these branches that have structure not like an inertial mode have imaginary index 24 and 36. These are also modes that have large rate of damping, real index number 45 and 46 respectively, and will be examined with the boundary layer modes later. All the modes checked so far are either damped, acoustic or look like the inertio end of the inertio-gravity modes found in the isothermal case. The structure of the acoustic modes is very




Figure 7.32: A typical inertio-gravity mode in the coupled problem, the reference state is boundary layer 1 and the horizontal wavenumber is $k = \frac{2\pi}{10^6}$. The lower part of the figure shows the position of this mode in both the imaginary and real part of the spectrum.

similar to the isothermal case. The structure of the inertial modes is also quite similar except that their structure is dominant in the region above the boundary layer. When the stretched grid isothermal case was studied at the beginning of this chapter some modes at the inertial end of the inertio-gravity spectrum were found which had structure only dominant near the ground. For the fully coupled case no such modes are identified, all inertial type modes have structure throughout the domain. Modes have been redistributed to suit what is able to exist with the presence of the sheared background flow. Mode 39 is another plotted in black and has structure around the boundary layer.

The mode with imaginary index number 40 has external mode structure and so

is the most likely candidate for the Rossby external mode; this is also the mode with largest damping rate that did not meet the criteria to be plotted in black in Figure 7.29. Although this is a Rossby type mode it is likely to decay away relatively quickly in comparison to other dynamical modes. The mode with index 41 is the computational mode, the computational mode will be discussed in a subsequent section.

Modes 42 to 47 have their dominant structure in the meridional component of velocity and in the potential temperature, an example of this kind of mode is shown in Figure 7.33. This set of modes all have dominant structure similar to that found in Rossby modes in the isothermal case in that the meridional velocity perturbation dominates the zonal velocity perturbation. It also has similar magnitude in the meridional velocity compared with the potential temperature and increased pressure perturbation compared with the inertio-gravity modes. The order in which the frequency relates to vertical scale has been flipped compared with the isothermal case; now the smaller the frequency the larger the vertical scale.

All the modes 42 to 47 have a positive real part and are thus unstable modes. The larger the vertical scale, and thus the smaller the frequency the larger the growth rate of that mode. This is interesting and somewhat counter intuitive since one might expect Rossby modes to feel the effect of the boundary layer damping but instead seem to be driven unstable by the shear. However this is a very shallow boundary layer.

Mode 48 also has similar structure to the other Rossby candidate modes but is the fastest growing and thus most unstable mode. Whereas the other modes in this branch have vertical scale increasing with decreasing frequency at this mode the behaviour changes to dominate only near the surface. This is possibly a mode resulting from the coupling. Mode 49 is the most damped mode and has structure in the fields that would be expected to have structure in the boundary layer modes.

So far all the modes that were not highlighted in black have been examined. They are the acoustic modes, the inertial end of the inertio-gravity modes, and the Rossby modes. In addition to these there are 3 external modes, two acoustic




Figure 7.33: A typical Rossby mode in the coupled problem, the reference state is boundary layer 1 and the horizontal wavenumber is $k = \frac{2\pi}{10^6}$. The lower part of the figure shows the position of this mode in both the imaginary and real part of the spectrum.

and one Rossby that has quite large damping rate. There is also a mode which is unstable and has structure dominating in the same fields as the Rossby modes but has structure only near the surface. The acoustic and inertial modes behave much as they do in the isothermal stretched grid case and have been checked also to agree well with the form of these mode types when the boundary layer is switched off in the transients. Six modes have been identified (21, 22, 24, 36, 39, 49, all plotted in black) as being damped by the boundary layer. Recall that the criteria for choosing them was that the absolute real part of the eigenvalue was larger than the real part of the eigenvalues when the boundary layer was switched off in the transients. These modes have structure in the horizontal velocity and potential

temperature fields, an example of a damped mode is given in Figure 7.34. This is one of the slower damped modes and so will be of interest when comparing the Lorenz and Charney-Phillips grids.




Figure 7.34: A typical damped mode in the coupled problem, the reference state is boundary layer 1 and the horizontal wavenumber is $k = \frac{2\pi}{10^6}$. The lower part of the figure shows the position of this mode in both the imaginary and real part of the spectrum.

The damped modes have similar structure to what would be expected in the gravity modes. Having examined all the modes which are less damped, acoustic, inertial and Rossby modes were identified, the latter of which was found to be unstable. A gravity mode would be expected to have dominant structure in the boundary layer anyway since the restoring mechanism is stable stratification and the fluid is stably stratified in the boundary layer and close to neutrally stratified above. No modes that are not damped were found to have this kind of structure;

all six modes that are significantly damped have exactly this kind of structure.

Consider Figure 7.28. The six modes working from the right hand side are the damped modes. The seventh mode in, i.e. the one with real part index 43 was an external mode, working further left to modes with less damping, the two fastest frequency acoustic modes are found followed by the external acoustic modes followed by inertial modes. Any damping in acoustic and inertial modes is likely to be shear driven rather than boundary layer diffusion driven. The six modes identified originally seem therefore to be accurately describing the boundary layer modes.

For the case where the boundary layer is switched off in the transients the mode structures have also been examined and are found to have overall quite similar behaviour. The main difference is the presence of a set of gravity modes, these exist with imaginary index 28 to 31 in Figure 7.24. There are four in total. All the structure is in the lower part of the domain where the background fluid is stably stratified and dominated in horizontal velocity and potential temperature. No modes with this kind of structure, and that do not have significant negative real part, were found in the coupled case, suggesting that all four are damped into boundary layer modes. This leaves two more to be found that turn into boundary layer modes, likely slow Rossby modes. In Figure 7.24 these were highlighted in black denoting that there was a significant real part to the eigenvalues. On examining the four gravity modes closer they are also found to be the fastest decaying and growing modes, but are driven so through shear rather than the boundary layer. These are the clear candidates for being most affected by the boundary layer since the structure remains similar and all are damped with the boundary layer switched on and with a faster damping rate than when the boundary layer was switched off.

With the boundary layer switched off in the linearisation most of the modes with structure like Rossby modes have eigenvalue with negative real part, whereas when the boundary layer is switched on in the transients they all have positive real part. This suggests that the boundary layer or the sheared reference state has

some mechanism that is capable of driving the Rossby modes unstable. There is evidence of unexpected instability is dynamical waves, for example in [31] where unstable external Rossby modes are observed in the presence of dissipation. A further possible mechanism for the instability seen could be due to the way the β -effect is added to the model, although switching β to zero did not produce any significant difference to the decay rates.

Shorter Horizontal Wavelength ($k = \frac{2\pi}{10^3}$)

As has been previously discussed the shorter horizontal wavelength is physically relevant for the boundary layer, the mode structures need therefore to be examined also for this case. From Figure 7.30 and Figure 7.31 it would seem that this may be harder due to the closer clustering of the frequencies and the unexpected damping of certain modes.

For the $k = \frac{2\pi}{10^3}$ case the modes with imaginary part index going from 1 to 18 (Figure 7.31) are again found to be the acoustic modes with structure very similar to that found in the longer horizontal case. The 19th and 20th modes are also the external acoustic modes. For this case however the frequencies are closer to the other acoustic modes due to the larger k .

For the modes with index 21 up to the last, 49, classification is considerably harder than for the longer horizontal wavelength case. There are clear inertial modes from 30 to 36 and two modes with a two grid type wave in θ , discussed later in the computational mode section. The majority of the remaining modes have very similar dominant structure to each other and so it is difficult to distinguish between them. There are however, three modes at 46 to 49 which are also the fastest damped modes and have structure like the boundary layer modes in the smaller horizontal wavenumber case.

Due to increased difficulty in identifying modes in the shorter horizontal wavelength case more care will be required in identifying modes for the Lorenz versus Charney-Phillips examination. It may be that additional information is gained when comparing Lorenz and Charney-Phillips modes in the longer horizontal wave-

length case that could help with comparison in the shorter horizontal wavelength case.

Deeper Boundary Layers

Figure 7.35 and Figure 7.36 show the real and imaginary part of the eigenvalues found when the reference state is boundary layer 5. This boundary layer has a depth of around 950m. The eigenvalues whose corresponding modes have structure in the boundary layer and with significant damping rate are plotted in black.

When switching to a much deeper boundary layer the overall behaviour changes and the distribution and behaviour of modes needs to be rechecked. Firstly note that the same number of acoustic modes exist in the deeper boundary layer case, further confirming their independence from the boundary layer influence; their structures are also unaffected; these have imaginary index going from 1 to 20. Index 19 and 20 correspond to the external acoustic modes. After checking by eye it is apparent that for the deeper boundary layer there are considerably more damped modes than previously, from index 30 to 49 in the real part plot and corresponding black circles in the imaginary plot, Figure 7.36. All of the damped modes have structure lying in the boundary layer and the dominant fields are the horizontal velocity and potential temperature perturbations.

Modes with imaginary index 29 to 31 are inertial modes, imaginary index 34 is the computational mode and 35 and 36 look like inertial modes. Mode 38 is difficult to classify but looks somewhat similar to the Rossby external mode, but only in vertical velocity and pressure. Modes 47 and 48 are also inertial modes. There is clearly considerable dependence on the depth of the boundary layer. For boundary layer 1 there existed acoustic modes, inertial modes, Rossby modes and six damped boundary layer/gravity modes. In addition the Rossby modes were all unstable. For the deeper boundary layer there exist acoustic modes, inertial modes and twenty damped/gravity modes but no obvious Rossby modes except possibly the external mode. There are only four unstable modes on the figure, two of which are inertial, one of which is the computational and the only one with any significant




Figure 7.35: Real part of the eigenvalues for the $k = \frac{2\pi}{10^6}$ case. The reference state is found using boundary layer 5 and the boundary layer is switched on in the linearisation.




Figure 7.36: Imaginary part of the eigenvalues and the singular values for the $k = \frac{2\pi}{10^6}$ case. These eigenvalues correspond to the those in Figure 7.35.

growth is the external Rossby. The real part of the eigenvalue corresponding to the computational mode is 3.853×10^{-12} so it can be considered neutral.

These results would suggest that the boundary layer is indeed capable of damp- ing the Rossby modes but only when sufficiently deep, bringing increased damping, otherwise they may even be unstable. When comparing the Lorenz and Charney-

Phillips grids the shallower boundary layers may be more useful for comparing how each can capture the dynamical modes such as Rossby and inertial but the deeper boundary layers may be more useful for comparing how they capture the boundary layer modes.

Boundary layer 1 and boundary layer 5 offer the two extremes, a depth of 100m and 950m. For the shallow case structures of Rossby modes are largely unaffected by the boundary layer. As the depth of the boundary layer increases the number of clear Rossby modes decreases and the number of boundary layer modes increases. The number of damped modes with boundary layer structure is: ten for boundary layer 2, twelve for boundary layer 3 and sixteen for boundary layer 4.

7.4.2 Comments on the Methodology

From examining the mode structures a picture is starting to develop for how the methodology performs. When looking at the boundary layer on its own, without any dynamics, problems were encountered associated with the normality of the matrix of transient coefficients. The matrix was so non-normal that comparing eigenvectors became impossible as any orthogonality between them was lost. By adding in the dynamics this problem is somewhat alleviated. It does not return the matrix to normality, and in checking inner products it is evident that eigenvectors are not completely orthogonal, however it is much better than for the boundary only case. Indeed as eigenvectors are examined it is clear that some physical interpretation can be attached to the eigenvectors; the issue that was seen previously where lots of simultaneous eigenvectors had almost identical structure is not seen here, even when resolution is increased.

Building up the complexity in the way done here enables one to gauge also the ability of the singular values and singular vectors to represent the behaviour of the system in a way that can be interpreted straightforwardly. For the isothermal resting case, even in the non-normal stretched grid case, all the modes could be easily identified in the singular vectors and the frequencies were well represented by the singular values. Part of the reason that frequency can be represented well

by singular values is because the modes are all neutral. Singular values are purely real while eigenvalues will be purely imaginary, although the problem could be rewritten so that eigenvalues were purely real. As the reference state is changed the difference between the singular values and eigenvalues is increased. Eigenvalues become complex and can be positive and negative; singular values remain real and positive. In this more complex situation the two decompositions will inevitably have to represent the behaviour in different ways. The boundary layer reference state introduces shear instability and Doppler shifting, producing eigenvalues with significant imaginary as well as real part. As the boundary layer depth is increased K_m and K_h increase so the rate of damping and the number of damped modes increases. For boundary layer five the difference between singular values and the imaginary part of the eigenvalues is increased due to the larger damping. In Part I the singular values were quite similar to the real part of the eigenvalues, this was due to the eigenvalues having dominant real part. Here the imaginary and real part of the eigenvalues are similar in magnitude and so it is not clear which part the singular value will be closest to, if any. It becomes more difficult to identify a mode type by looking only at the singular value due to the large range of values they can take, from representing the high frequencies of the acoustic modes to the damping rate of boundary layer modes. Instead the singular vectors themselves need to be examined. Physical interpretation of the dynamical modes through the singular vectors is still possible, as can be seen in Figure 7.37 and Figure 7.38. These figures show how the modes shown in Figure 7.32 and Figure 7.33 are represented by the singular vectors.

The overall structure looks quite similar to that given in the eigenvectors and indeed the behaviour identified in the isothermal resting state singular vectors occurs here. For example in Figure 7.37 is the inertial mode as shown in eigenvectors in Figure 7.32, apart from the similar structure it is clearly an inertial mode since the input singular vector for E_v has equivalent structure to the output singular vector for E_u and the negative version vice-versa, exactly as would be expected due to the findings of earlier in this chapter and in Part I. This increased informa-




Figure 7.37: The singular vectors with closest structure to the inertial mode plotted in Figure 7.32. $k = \frac{2\pi}{10^6}$, boundary layer 1.

tion can help in distinguishing inertial modes from Rossby modes in the shorter horizontal wavelength case where dominant structure does not change significantly. The Rossby mode in Figure 7.38 is also clearly identifiable, having the dominant structures in the same fields as found in all previous cases. A further point to note is that the relationship between input and output vectors that was seen for the isothermal case is still seen. In the dominant energy variables E_v , E_θ and E_p the output singular vectors include a multiplication by i from the input singular vectors.




Figure 7.38: The singular vectors with closest structure to the Rossby mode plotted in Figure 7.33. $k = \frac{2\pi}{10^6}$, boundary layer 1.

7.4.3 Mode Tracking: Motivation

A number of modes have been identified that become more quickly damped when the boundary layer is switched on. It would be useful if these modes could be tracked as the boundary layer is gradually switched on so that the form of the mode most susceptible to the boundary layer could be established. It would also be useful if the computational mode could be tracked. In the shallow boundary layer it is quite likely that the only modes that become boundary layer modes start off as gravity modes since no near neutral gravity modes were identified with the transient boundary layer switched on for this boundary layer depth. For the deeper boundary transients off case both Rossby and gravity modes appear to turn into boundary layer modes and it would be useful to be able to distinguish between

them. Being able to track modes is potentially problematic however. If there is a continuous part to the spectrum then completely different modes will exist in each case, boundary layer off and boundary layer on.

As discussed in Part I the form of the spectrum is an underlying issue that requires careful consideration when using this type of methodology. It is clear that the isothermal Euler equations can only have a discrete spectrum; there is no viscosity and therefore only waves with vertical wavenumber like integer times $\frac{\pi}{D}$ can match the imposed boundary conditions; there are infinitely many but only discrete possibilities. It has also been argued in Part I that, when solving the boundary layer only case, there is only a discrete spectrum. If both individual components support only a discrete spectrum it seems likely that, when coupled, the spectrum should also be discrete, but care should still be exercised. It is possible that the boundary layer reference state but with boundary layer transients switched off could have continuous spectrum. One may hypothesise that when coupled equations are used a continuous part to the spectrum could exist. For example if say the boundary layer mechanism were capable of absorbing certain dynamical modes of any wavelength. Due to the complexity of the fully compressible equations it would be beyond the scope of this work to derive rigorously whether the spectrum consists only of discrete values or contains a continuous part. Some tests will be implemented to search for a continuous part of the spectrum. From examining the mode structures by eye no continuous part has been identified. Similar physical behaviour is seen across the range of cases that were examined, all of which can be understood quantitatively.

Given the possibility of having a continuous spectrum it may not be wise to try and track modes that appear to turn into boundary layer modes but it should be possible to confirm which modes do not feel the boundary layer. In addition to the tracking of modes, the form of the spectrum is fundamental for the task of comparing Lorenz and Charney-Phillips. If a continuous spectrum is possible then there is no reason to believe that the two grids should even capture equivalent modes.

The most obvious test to look for a continuous spectrum is to examine the behaviour as resolution increases. If the modes that are captured during a low resolution run are not captured at all by the high resolution run then it would imply a continuous spectrum. Having all the modes captured by the high resolution would be necessary for a discrete spectrum but not sufficient to prove that only discrete spectrum exists. Initially the eigenvalues themselves were checked with increasing resolution but since the eigenvalues are within a very close range, i.e. all the inertial modes have almost identical frequency, this was not found to be helpful. The alternative, and the basis of the test used here, is to visually compare eigenvectors across different resolutions to ensure all low resolution solutions are captured in the high resolution, this is potentially a highly arduous and ambiguous process. The results so far show the eigenvector methodology to be working quite well, although eigenvectors are not orthogonal they are not as far from being orthogonal as in Part I. Computing inner products can save a lot of time when comparing eigenvectors from different resolutions and remove some of the ambiguity. It is not sufficient to rely on this test however because of the non-normality in the problem, any matching eigenvectors returned by the test need to be checked by eye.

The whole process can be sped up considerably though. Using the geometrically stretched grid that was constructed in Part I two resolutions of 20 and 30 grid points are compared. The eigenvectors in the low resolution are interpolated to the high resolution grid. A matrix is constructed whose entries are the inner product between each eigenvector in the low resolution (rows) and each eigenvector in the high resolution (columns), the column with the entry closest to 1 or -1 gives the high resolution eigenvector most similar to the low resolution one. Table 7.2 shows the high resolution eigenvector most closely resembling each low resolution eigenvector, a tick in the table denotes that when checking by eye the eigenvectors gave best agreement. For this particular example boundary layer 4 was used in the reference state, but the boundary layer was switched off in the transients.

It is clear from the table that the agreement is good and that all modes that are found in the low resolution case are also found in the high resolution case;

n = 20	n = 30		...continued from left					
	Orth Test	By Eye						
1	5	✓	35	56	✓	69	103	✓
2	6	✓	36	55	✓	70	107	✓
3	9	✓	37	58	✓	71	118	✓
4	10	✓	38	57	✓	72	110	✓
5	11	✓	39	60	✓	73	112	✓
6	12	✓	40	59	✓	74	113	✓
7	16	✓	41	61	✓	75	114	✓
8	15	✓	42	63	✓	76	115	✓
9	20	✓	43	64	✓	77	116	✓
10	19	✓	44	65	✓	78	117	✓
11	21	✓	45	66	✓	79	118	✓
12	22	✓	46	68	✓	80	119	✓
13	26	✓	47	71	✓	81	123	✓
14	25	✓	48	73	✓	82	122	✓
15	29	✓	49	75	✓	83	124	✓
16	30	✓	50	73	✓	84	125	✓
17	32	✓	51	77	✓	85	126	✓
18	31	✓	52	77	✓	86	128	✓
19	36	✓	53	80	✓	87	131	✓
20	35	✓	54	81	✓	88	132	✓
21	38	✓	55	83	✓	89	133	✓
22	37	✓	56	85	✓	90	135	✓
23	42	✓	57	85	✓	91	137	✓
24	41	✓	58	88	✓	92	138	✓
25	44	✓	59	89	✓	93	140	✓
26	43	✓	60	91	✓	94	141	✓
27	48	✓	61	93	✓	95	143	✓
28	47	✓	62	92	✓	96	144	✓
29	50	✓	63	96	✓	97	146	✓
30	49	✓	64	98	✓	98	148	✓
31	52	✓	65	99	✓	99	149	✓
32	51	✓	66	97	✓			
33	54	✓	67	99	✓			
34	53	✓	68	103	✓			

Table 7.2: Comparison of the modes found when using a geometrically stretched grid with 20 and 30 grid points, boundary layer 4 reference state but with boundary layer switched off in transients. A tick denotes that when checking by eye the mode in each configuration agrees.

no continuous part to the spectrum is evident. The test used to generate the results in the table has been repeated for a number of different cases, covering boundary layer switched on, boundary switched off, uniform grids, the Met Office grid and for different boundary layer depths. The results were reasonably good and in most cases the test returned the best match for around 90% of the low resolution eigenvectors. For the majority of individual modes for which the test failed, a candidate making a good match could be found by eye. The methodology of comparing eigenvectors works better for some cases than others. The table shows a boundary layer off case. Here every eigenvector could be identified by the test and agreed when checked by eye. When the boundary layer is switched on the test does not always produce complete success, particularly for the boundary layer modes which are likely further from orthogonal. It was also found that the methodology worked better when the two chosen resolutions were quite close, again due to orthogonality issues; when there are more modes to choose from a number may have inner product close to 1 or -1. When a number are found close to 1 or -1 a few may need to be checked by eye before finding the most appropriate match.

7.4.4 Mode Tracking: Methodology

Given the lack of evidence for a continuous spectrum after extensive testing it seems that tracking may be able to produce some useful information. Tracking as the boundary layer terms come into the problem is done by altering a coefficient in front of the transient boundary layer terms gradually from 1 to 0, effectively a matrix perturbation problem. The problem in doing this is that the matrix perturbation is unbounded since the boundary layer terms include second order derivatives. Even placing a factor $\ll 1$ in front of the transient boundary layer term may result in significant change in overall behaviour. This means that standard techniques for matrix perturbation problems are generally not useful. Some matrix perturbation procedures look to calculate the change in a particular eigenvalue, for example by looking for the eigenvalue after the perturbation lying closest to the location of that particular eigenvalue before the perturbation. This would likely be difficult

for this problem since again, even a coefficient $\ll 1$, compared to 0, in front of the transient boundary layer terms could move eigenvalues a relatively long way in the complex plane. Recall that for $k = \frac{2\pi}{10^3}$ the fastest growing mode in the boundary layer off case became the fastest decaying mode when the boundary layer was switched on. If eigenvalues move as much as that for a relatively small perturbation it could not be tracked by looking for the closest eigenvalue after perturbation. Given the success of calculating inner products of eigenvectors this is used instead to attempt to track modes as the boundary layer is switched on. It seems more natural to perform the tracking from boundary layer switched off to on (0 to 1) but since ‘singular’ behaviour will likely occur when going from 0 to any amount of boundary layer switched on the tracking is performed from on to off.

A simple iterative procedure starting with the boundary layer switched on (coefficient = 1) is considered. The coefficient is reduced, to say 0.9, and inner products are calculated between the two cases, coefficient = 1 and coefficient = 0.9. Before the coefficient can be reduced further towards boundary layer off, two conditions must be found to hold. The conditions that must hold are that i. for the best matched modes the inner product must be sufficiently close to 1 or -1 and ii. that modes correspond one-to-one. If the conditions do not hold then the original reduction in the coefficient is iteratively reduced until they do. So for example, if after the initial reduction to 0.9 one of these tests does not hold then the step would be halved to give 0.95. If the reduction is halved approximately 25 times and the conditions still do not hold then the test is likely to fail. Due to the simple nature of the process it will not have success in every situation, and has issues reaching 0 for deep boundary layers, where modes are further from being orthogonal. Generally if the test failed it was as the coefficient in front of the boundary layer transients got close to 0.

Figure 7.39 shows the values of the eigenvalues as the matrix is perturbed so that the boundary layer goes from on to off in the transients. Completely on is the right hand side of the plots, completely off at the left. The figure is for the longer horizontal wavelength case with boundary layer 1 for the reference state and

the Met Office grid. Using the results of the test the eigenvectors are examined to check which retain their structure as the boundary layer is switched from on to off. Eigenvector structures are checked by eye along each line in Figure 7.39, if the structure is equivalent at either end of the line then that line is plotted in blue, if the structure is clearly different then the line is plotted in red. This is combined with a check that each mode with different structures at either end can not be matched with another eigenvector.




Figure 7.39: The imaginary (left) and real (right) parts of the eigenvalues as the boundary is switched from on to off in the transient calculation. Boundary layer 1 reference state and the Met Office grid. Modes found when the boundary layer is switched on that are not found when the boundary layer is switched off are plotted in red.

The red lines in the figure do not necessarily describe which mode with the boundary layer off turns into each boundary layer mode. There may not be smooth behaviour as the tracking occurs, especially going from no boundary layer transients to ‘any’ boundary layer transients. The tracking could therefore incorrectly show a relation between modes. However it is reasonably safe to deduce which modes stay approximately the same. In Figure 7.40 the imaginary part of all the eigenvalues at the boundary layer off end of Figure 7.39, i.e. the same as Figure 7.24, are plotted. The modes which did not stay approximately the same, i.e. corresponding to red lines in Figure 7.39, are plotted with a black asterisk. When examining the structure of these boundary layer off modes the four together in the middle (28 to 31) are all found to look like gravity modes. The two modes with index 48 and 49

have dominant fields corresponding to what would be expected in Rossby modes but with structure dominant near the ground.




Figure 7.40: The imaginary part of the eigenvalues with the boundary layer switched off in the transient calculation. Modes found when the boundary layer is switched on that are not found when the boundary layer is switched off are plotted in black. Boundary layer 1 reference state, $k = \frac{2\pi}{10^6}$ and the Met Office grid.

The results of the tracking confirm the results of simply examining modes. For this boundary layer 1 case with $k = \frac{2\pi}{10^6}$ a set of six boundary layer type modes were found, however no gravity modes were found, but gravity modes were identified in the boundary layer off case implying they had all been damped. Only four gravity modes were identified in the boundary layer off case implying that other types of mode were also damped by the boundary layer. The results of the tracking would imply that the two slowest Rossby modes in the boundary transients off case have been damped by the boundary layer and turned into boundary layer modes, since neither could be found in the coupled modes.

7.5 Lorenz Grid Computational Mode

As discussed previously the Lorenz grid has one too many degrees of freedom, resulting in a spurious computational mode. Within the isothermal framework the computational mode could be identified by a two grid wave existing throughout the

domain. The structure is a two grid wave in θ' as this is the kind of structure that when averaged in the vertical momentum equation would be zero. It is a major disadvantage to the Lorenz grid that it supports a computational mode since it has been shown to manifest itself as non-physical properties in the model, such as the baroclinic instability as discussed in [4].

It is of interest to understand what happens to the computational mode when the complexity is increased to include a sheared background flow and boundary layer mechanism. Figure 7.41 shows the computational mode for the boundary layer 1 case, with dynamics only in the linearisation, the grid is the Met Office grid. The two grid wave in θ' and ρ' is clear but now only exists in the region above the boundary layer.




Figure 7.41: This figure shows the computational mode for the boundary layer 1 background flow but with the boundary layer transient terms switched off.

In addition to the two grid signal in θ' and ρ' the two grid signal also appears in u' , v' and w' . In the region above the boundary layer the flow is in geostrophic balance, $fv = c_p\theta\frac{\partial\Pi}{\partial z}$ and $-fu = c_p\theta\frac{\partial\Pi}{\partial z}$. The signal in θ' results in a similar signal in ρ' through Π' which in turn results in a signal in the velocities through geostrophic balance. The plot reveals that, as expected, shear alone in the background flow is not sufficient to suppress the computational mode throughout the domain. The

shear does suppress the structure of the computational mode in the shear region though.

That the shear suppresses the computational mode can be understood in terms of the speed of propagation of waves at different heights. If shear exists in the background flow then waves at different heights in the shear region will propagate at different rates, thus the vertical wavelength will only be twice the grid spacing at certain instants. A schematic to further demonstrate this is shown in Figure 7.42. A snap shot of three theoretical θ' waves at different heights on theoretical grids are shown in a vertical two grid wave configuration, a peak in the highest wave coincides with a trough on the wave below it. Immediately after the theoretical point shown in the figure the three waves would cease to exist in a two grid wave configuration, the peak of the highest wave would have moved faster than the trough of the wave below it and so on.




Figure 7.42: Two grid wave schematic. Under shear waves at different heights travel at different velocities.

This argument would suggest that having any amount of shear in the fluid might alleviate issues associated to the computational mode; this certainly would be beneficial since throughout the atmosphere there is always some amount of shear. Reconsider the schematic in Figure 7.42; although the waves will propagate

at varying rates, at some point they will realign to produce a two grid wave.

Figure 7.43 shows the computational mode for the fully coupled problem, where the boundary layer transient terms are switched back on. Again the computational mode remains, existing as a two grid type wave everywhere above the boundary layer. It is unchanged by the addition of transient boundary layer terms. In the full coupled problem the strong diffusion mechanism in the boundary layer is not sufficient for suppressing the computational mode throughout the domain. This has been checked for all the boundary layer depths and found to hold unanimously.




Figure 7.43: This figure shows the computational mode for the boundary layer 1 background flow, this is the full coupled problem where the transient boundary layer transient terms are switched on.

7.6 Chapter Summary

This chapter has aimed to build up an understanding of the types of modal solutions that are captured by the coupled problem. To build on the knowledge gained from examining the boundary layer only case in Part I, the dynamical modes have been examined. The case that allows for the most concise representation of the dynamical modes is the isothermal resting reference state, solved on a uniform grid;

the solutions for which have been examined by previous authors. By building up from the isothermal case a picture of how the solutions behave for the coupled problem can be constructed. First the solutions were sought for the isothermal resting state on a stretched grid, this is important since using a stretched grid offers the best overall representation of the atmosphere in a numerical model. As the grid is stretched the set of modes that are captured changes. The largest scale modes will be captured in an equivalent manner, however some smaller scale modes, sampled from a part of the spectrum that the uniform grid cannot support, will change. These waves will have all their structure in the region of the grid which has sufficient resolution for supporting their frequency. In addition to stretching the grid, solutions for a shorter horizontal wavelength were computed. As resolution of atmospheric models increases, how the behaviour of waves at shorter wavelengths is captured becomes an important question. In addition to this the shorter horizontal wavelengths are of a similar scale to the depth of the boundary layer. For the isothermal case the shorter horizontal wavelength meant better discrete representation of the Rossby frequency but worse representation of the inertio-gravity frequency.

An intermediate case where the dynamics and physics are coupled in the reference state but then only dynamics are included in the transient state was considered before the fully coupled problem. This allows one to build up a picture of how the dynamical modes are distorted and altered by the added complexity in the reference state. Once the reference state includes shear there is considerably more sensitivity to the choice of the horizontal wavenumber. For smaller wavenumbers the frequencies of the three types of dynamical mode types are closer to that seen in the isothermal no shear case. The larger horizontal wavenumber leads to a more condensed dispersion relation where Rossby and inertio-gravity waves appear to inhabit a single branch. The big difference seen as the wavenumber changes is due Doppler shifting by the background wind, causing the frequencies to be distorted from their natural range.

With the understanding gained from the three simpler test cases - isothermal

resting reference state on a uniform grid, isothermal resting reference state on a stretched grid and fully coupled reference state but boundary layer excluded from the transients - the fully coupled case was considered. This is the most complex case that will be considered and the case for which the difference between Lorenz and Charney-Phillips is of most interest. This chapter has largely been concerned with examining the coupling. However one important result regarding the comparison of the Lorenz and Charney-Phillips grids has been obtained in examining the structure and behaviour of the Lorenz grid computational mode. As outlined in the introduction to this thesis, spurious behaviour can result from the presence of a computational mode and this makes the Lorenz grid a dangerous choice when designing a numerical model. If the boundary layer's damping mechanism had been found to damp the computational mode throughout the domain it could have been argued that the Lorenz grid could be more useful than previously thought.

In addition to investigating the computational mode the interaction of dynamical and boundary layer modes was investigated. The dependency on the boundary layer depth is clear. The deeper the boundary layer the more boundary layer modes that are supported and the more potential problems in using the methodology. The investigation has revealed that gravity and Rossby modes can be expected to be most susceptible to the boundary layer damping. These are the kinds of modes present when the boundary layer is switched off but distorted in some way when switched on.

Going forward into the comparison of Lorenz and Charney-Phillips grids both horizontal wavenumber cases may be useful. The original and smaller wavenumber of $k = \frac{2\pi}{10^6}$ produces results which are potentially easier to interpret since the three types of dynamical mode are more evident by differences in the frequencies. The larger wavenumber case $k = \frac{2\pi}{10^3}$ may cause problems with the methodology if solutions are harder to interpret. Nevertheless it provides a more appropriate length scale for coupling with a shallow boundary layer and is relevant given the increasing resolution in current weather and climate models. Although it is harder to distinguish between dynamical modes in this case the boundary layer modes

were found reasonably easily and so a comparison of how each grid captures these may still be possible.

In addition to the two different wavelengths there are still the different boundary layer depths to help with comparison. Deeper boundary layers support more boundary layer modes and shallow boundary layers support more dynamical modes. If wanting to compare how well Lorenz and Charney-Phillips capture the Rossby frequency, for example, then using a shallower boundary layer will help. If wanting to compare the slowest damped boundary layer modes a deeper boundary layer may be more useful.

In addition to understanding the structures and behaviour of the various modes in the problem, the work in this chapter has sought to examine the extent to which the normal mode type methodology can be useful. From Part I it is clear that having the boundary layer on its own causes problems in terms of comparing eigenmodes. The matrix is far from normal and eigenvectors become far from orthogonal. In this chapter, by building up the complexity, it seems that the usefulness of the eigendecomposition is somewhat recovered. There appears to be ‘enough’ dynamics in the problem to counteract the issues found previously, particularly for shallow boundary layers. This should improve the ability of the eigendecomposition to reveal differences between the configurations, the subject of the subsequent chapter. For the boundary layer only it was found that the usefulness of the systematic methodology afforded by the eigendecomposition could be somewhat recovered by examining the singular vectors of the matrix. By examining the isothermal case a situation where the singular values produce very similar results to the eigenvalues is obtained. As the complexity of the problem is increased the SVD departs further from the eigendecomposition. However the results shown in this chapter suggest that the physical behaviour is still well represented by the SVD. Along with the results that were obtained in Part I it does seem that the singular values and singular vectors can be used to understand the behaviour of the system. In the next chapter, when configurations are compared, it will be investigated whether the decomposition can be useful for comparison purposes. It seems likely, based

on the results of this chapter, that for deeper boundary layers the eigenvectors will be further from orthogonal than in the shallower boundary layers. Here singular vectors will be needed for the comparison.

Chapter 8

Comparison of Lorenz and Charney-Phillips Grids

The aim of this chapter is to compare the Lorenz and Charney-Phillips configurations for the fully coupled dynamics-boundary layer problem. It is now well understood that when modelling the dynamics of the atmosphere the Charney-Phillips grid is preferable, in that it gives optimal wave dispersion properties. On the other hand the Lorenz grid has the benefit of being more suitable for conservation, however the Lorenz grid has one too many degrees of freedoms and as a result supports a computational mode in the dynamics. Further to this the Lorenz grid is the preferable configuration for modelling the stably stratified boundary layer. It was hypothesised during the introduction to this thesis that the boundary layer may be capable of distorting the computational mode, for example by slowing and diffusing it into a decaying boundary layer mode. This would be highly beneficial for the Lorenz grid. Although the computational mode has been identified as still existing in the coupled problem it will still be useful to compare how well all the other mode structures are captured by each grid.

As was the testing order for the boundary layer only, the comparison can be divided up into examining first the steady state structure and then the transient structure. Computing the steady states for the coupled problem is similar to the boundary layer only problem and so differences in overall structure are not antic-

ipated. The transient structures for the fully coupled problem are considerably more complex than for the boundary layer only and so require more careful consideration. In the previous chapter the transient solutions for the Lorenz grid were examined, this was done to provide insight into the coupling mechanisms, establish particular modes of interest and check the methodology. In doing this cases with properties that can aid in the comparison were established, based on boundary layer depth and horizontal wavenumber. In conjunction with the array of cases considered in Part I this leads to a huge number of potential configurations, these can be summarised as,

- Three options on averaging to obtain K_h in boundary layer terms when using Charney-Phillips.
- Five grid spacing options, uniform, geometric, Met Office, logarithmic and log-linear grids.
- Two options for horizontal wavenumber $k = \frac{2\pi}{10^6}$ or $k = \frac{2\pi}{10^3}$.
- Five depths of boundary layer, 1-5.

When examining the transients it is neither possible nor particularly interesting to demonstrate a comparison of every mode for every case that was considered. Instead just those configurations likely to yield interesting results are considered. For the low resolution runs the Met Office type stretching is most relevant since it is similar to that used in operational models; attention can first be restricted to using this grid. In terms of choosing model parameters it was found that the smaller wavenumber could be more useful in terms of identifying and interpreting modes. In part I it was found that the way in which K_h was obtained, i.e. by averaging either Richardson number, stability function or K_h itself, made some difference. For the coupled steady state all three are considered, from there the choice can be narrowed down. This just leaves the choice of boundary layer depth. The shallower the boundary layer the more dynamical modes that are supported and the deeper the boundary layer the more damped boundary layer modes will

be supported. Initially both shallowest and deepest will be considered in order to draw comparisons of each type of mode.

8.1 Steady State Comparison

The steady state equations to be solved are given in equations (6.8) - (6.11). The only addition to these equations in comparison to those used to solve for the boundary layer only is the density in the boundary layer diffusion terms. The density itself is computed through a hydrostatic balance diagnostic step.

8.1.1 High Resolution Steady States

As previously, the low resolution Lorenz and Charney-Phillips solutions are compared to a high resolution solution. For the methodology to be viable there should be negligible difference between Lorenz and Charney-Phillips for the high resolution; again this can only be achieved for a relatively small number of grid points by employing a logarithmic stretching near the surface. For the coupled case the log-linear grid is used for the high resolution to ensure good representation of the dynamical modes whilst allowing for good convergence properties. In computing the high resolution steady states the Newton method is again used with success. Recall from Part I that in order to use the Newton method a smooth switching between neutral and stable stratification is employed, this provides the iteration with a suitable guess for the solution. Using the Newton method has the added advantage that it requires calculating the Jacobian matrix which is also used in finding the transients. Note that, although not shown, using the logarithmic coordinate transform is found to produce good convergence between Lorenz and Charney-Phillips grid solutions. High resolution solutions from Lorenz and Charney-Phillips grids agree well; when a high resolution solution is used in comparisons it can be considered to have been obtained from either.

The high resolution density is shown in Figure 8.1 for all boundary layer depths. Density varies relatively smoothly in the boundary layer and over a relatively small

range. It can be seen in the figure that the density near the surface varies, this is due to the way in which the boundary conditions are implemented. Different cases of boundary layer depth were chosen by setting the surface boundary condition for potential temperature. In addition an upper boundary condition is set on pressure to calculate hydrostatic balance. The difference in surface potential temperature results in different pressure and therefore density at the surface. Note that boundary layer 5 and boundary layer 4 have similar density in the boundary layer due to their boundary conditions for potential temperature being equivalent; the different depths between these two cases is due to different boundary conditions on geostrophic wind.




Figure 8.1: High resolution steady state densities for all boundary layer depths.

It is clear that density varies quite smoothly in the boundary layer although has a sharp feature near the top of the boundary layer. It is therefore anticipated that little difference will be found between the high resolution coupled steady state profiles and the high resolution boundary layer only profiles, except for possibly near the top of the boundary layer for shallower boundary layers. The density inside the vertical gradient terms in equations (6.8), (6.9) and (6.11) will approximately cancel with the division by density in front. The similarity between the high resolution coupled and boundary layer steady states can be seen in Figure 8.2. Clearly the high resolution velocity and potential temperature are similar in each

case. This also provides a useful check on the numerics used in the Newton method.




Figure 8.2: Lorenz grid high resolution steady states for boundary layer 5. Shown for the boundary layer only against the coupled problem.

8.1.2 Operational Resolution Steady States

The steady states are now computed for the low resolution Met Office grid. For the low resolution there may be increased differences between the boundary layer only profiles and the coupled profiles, larger spacing introduces greater errors when averaging densities. Figure 8.3 and Figure 8.4 show the steady states as computed using the Lorenz and the three Charney-Phillips option I configurations, against high resolution. For all three option I configurations Richardson number Ri is computed at z_w levels. The option I-i configuration is the one where K_h at z_ρ levels is obtained by averaging K_h itself, option I-ii is when the stability function and shear are averaged to z_ρ levels and option I-iii is when the shear is also averaged but the stability function is obtained from Ri averaged to z_ρ levels. Recall the schematic for the averaging when using Option I in Figure 3.8 on page 92.

Figure 8.3 and Figure 8.4 show the equivalent to Figure 3.9 to Figure 3.12 in Part I. It is clear from the figures that inaccuracies that occurred in finding the low resolution steady state profiles for the boundary layer only problem extend to the coupled problem. The shallowest boundary layers are the least well represented by the low resolution grids. In all four fields the lack of grid points close to the top of




Figure 8.3: Steady States for the Fully Coupled Problem. Figure shows horizontal velocities U (left) and V (right), boundary layer 1 to boundary layer 5 (top to bottom).




Figure 8.4: Steady States for the Fully Coupled Problem. Figure shows thermodynamic variables $\theta^{(r)}$ (left) and $\rho^{(r)}$ (right), boundary layer 1 to boundary layer 5 (top to bottom).

the boundary layer causes a vertical shift in the structure. Generally speaking the operational resolution Charney-Phillips option I-i configuration gives a better representation of the high resolution solution for the coupled problem than it did for the boundary layer only problem. This is most clear in the velocity fields U and V for boundary layer 5, where Charney-Phillips option I-i is better than Lorenz. For the boundary layer only the Lorenz grid gave the better representation of horizontal velocity in boundary layer 5. The Charney-Phillips option I-i configuration is also found to outperform the Lorenz grid for the shallower boundary layers. For the potential temperature the Charney-Phillips I-i configuration captures the structure near the ground more accurately than the Lorenz grid which was also the case for the boundary layer only. Now the Lorenz grid does a worse job in capturing the higher part of the boundary layer and so there is less difference between the two. As was found for the boundary layer only there is a clear advantage gained by using Charney-Phillips option I-iii. It produces the most accurate representation of the reference thermodynamic variables near the top of the boundary layer for all boundary layer depths. It is also representing the U component of velocity more accurately than the other grids for the shallowest boundary layer. On the other hand the option I-iii configuration gives the worst representation of the thermodynamic variable in the lowest part of the boundary layer, which is independent of the boundary condition implementation. There is little advantage to be gained in using option I-ii, which gives the worst representation of the thermodynamic variables in all cases.

The Charney-Phillips option I-i and I-iii configurations capture the overall structure of the velocity fields most accurately. Option I-i does marginally better, particularly lower in the domain. The Lorenz grid captures the highest part of the velocity least accurately but does outperform option I-iii nearer the surface. Option I-iii is least accurate for the velocities near the surface as well. Density does not vary greatly in the boundary layer and is generally captured most accurately by Charney-Phillips options I-i and I-iii configurations. For the potential temperature the option I-iii has almost no vertical shift near the top of the boundary layer giv-

ing it best overall representation of the structure. However over approximately the lowest 20% of the boundary layer depth option I-iii is least accurate for potential temperature, option I-i is most accurate here. As was found for the boundary layer only option I-i and I-iii are the most promising Charney-Phillips configurations. The main difference now is that they mostly outperform the Lorenz grid, whereas previously the Lorenz grid outperformed all the Charney-Phillips configurations.

The only difference between the coupled equations and the boundary layer only equations is the additional diagnostic step in which density in hydrostatic balance is calculated. It is due to this step that the Lorenz grid loses the benefit it gains by requiring no averaging in Richardson number. Recall that the hydrostatic balance is given by,

$$\frac{\partial \Pi^{(r)}}{\partial z} = -\frac{g}{c_p \theta^{(r)}}. \quad (8.1)$$

The density ρ is obtained from the quantities in hydrostatic balance. Performing a numerical integration of (8.1) gives Exner pressure which leads to giving pressure and density. To ensure that the derivative is second order requires that θ is stored halfway between Π , i.e. at z_w levels. This is the natural place for θ when using the Charney-Phillips grid but θ needs to be averaged when using the Lorenz grid. Although averaging introduces some error, it is smaller than the error that would be introduced by using a first order accurate scheme for the derivative. The averaging required when using the Lorenz grid causes less accuracy in the density and as a result the steady states are less accurate than those captured by the Charney-Phillips grid.

8.2 Coupled Transient Comparison

Examination of the steady states revealed almost no benefit in using the Lorenz grid in comparison to the Charney-Phillips grid for the coupled problem. Recall that for the boundary layer only the Lorenz captured the steady states most accurately. These findings would imply that errors in the Lorenz grid averaging to compute the dynamics may dominate the errors in the Charney-Phillips grid averaging to

compute the boundary layer terms. This needs to be examined for the linear time dependent part of the equations by examining the transients. For the steady state Charney-Phillips options I-i and I-iii appear to be the most useful configurations since they both outperformed the Lorenz and option I-ii configurations. Neither option is preferred outright however, option I-i is generally more accurate in the lower part of the boundary layer and option I-iii is more accurate for the upper part of the boundary layer. For the transient comparison Charney-Phillips option I-i is considered since it almost always outperformed the Lorenz grid for the steady states. Charney-Phillips option I-iii also had some promising results for the steady state, however was often outperformed by the Lorenz grid near the surface.

Before examining any transients the computational mode can be identified in the Lorenz configuration and removed. Since it will not be captured by the Charney-Phillips grid it does not need to be included in any comparison. The process of identifying the computational mode, so as to remove its eigenvector and the corresponding eigenvalue, can be sped up considerably by automated counting of zeros in the eigenvectors. The counting is done from the top of the domain to the approximate top of the boundary layer. Unless using the deepest boundary layer with low resolution the vector with the most zeros above the boundary layer will generally be the computational mode. Any plots of frequencies or growth rates that are shown henceforth will have had the computational mode removed.

Broadly the equations support two types of modes: dynamical and boundary layer. From the study of the coupling in Chapter 7, as well as the findings of Part I, it is clear that the behaviour of these modes is considerably different; both physically and in terms of the methodology. The eigenvalues corresponding to the dynamical modes generally have dominant imaginary part and the structure of the modes can be dominant throughout the domain. Conversely the boundary layer modes have eigenvalues with significant real part to them and have structure dominating in the boundary layer region.

For these two general types of mode there appears to be two regimes for the methodology. For the dynamical modes the eigendecomposition performs well, even

in the coupled problem. It is likely that eigenvectors can be used to compare the performance of configurations for dynamical modes. Conversely the problems of non-normality that affected the results of the boundary layer only study are likely to still occur when examining the damped boundary layer modes.

Due to the differences in both the structures of the solutions and the effectiveness of the methodology the comparison of the transients is divided into examining the dynamical modes and the boundary layer modes separately.

8.2.1 Examination of Dynamical Modes

In this section the ability of the Lorenz and Charney-Phillips grids to capture the dynamical modes is compared.

Boundary Layer Switched off in the Linearisation

With the computational mode removed, any eigenvectors and corresponding eigenvalues found when using the Lorenz grid should also be found when using the Charney-Phillips grid. However this will be dependent on how well the methodology works and this depends on the configuration and boundary layer depth being considered. The most clear results in terms of the methodology were afforded by using smaller wavenumber $k = \frac{2\pi}{10^6}$ and with the boundary layer switched off in the transients. There are some modes where eigenvalues have a real positive (growing) or real negative (decaying) part, but the imaginary part generally dominates. This is a useful case to check first in the comparison as it allows clear examination of the dynamical modes and the frequencies they should have. Figure 8.5 shows the Lorenz and Charney-Phillips option I-i frequencies (imaginary part of the eigenvalue) for this wavenumber with boundary layer 1 in the reference state but the boundary layer switched off in the transients.

It is clear from the figure that the low resolution Lorenz and Charney-Phillips grids give similar frequencies for the majority of the dynamical modes, as was found for the dynamics only isothermal case. The only significant difference is in the two slowest propagating Rossby modes which propagate considerably slower




Figure 8.5: Boundary layer 1 reference state, boundary layer switched off in the transients. Frequencies as found by the Lorenz grid versus the frequencies as found using the Charney-Phillips option I-i configuration.

for the Lorenz grid; this is also what was found for the isothermal case. When the modes were examined in the previous chapter these two Rossby modes were found not to exist in the fully coupled problem but had been replaced by boundary layer modes. Although it is a useful check this demonstrates why, when comparing the two configurations, the fully coupled equations must be considered. Modes that only exist in the boundary layer off case are of little interest for comparison since they would not form part of the solution for an overall model with these parameters. For the remainder of this chapter, and the comparison between the grids, attention will be restricted to the fully coupled linearisation. In restricting the attention to fully coupled only, comparison of modes that do not actually exist in the coupled problem is avoided.

Fully Coupled: Mode Frequencies

So far the methodology is found to work effectively for comparing modes in the coupled problem. However it is not always straightforward to compare low resolution modes with high resolution modes. Techniques such as checking orthogonality between high and low resolution will be less likely to work between eigenvectors.

Any vectors deemed orthogonal need to be checked by eye. For some modes checking orthogonality is found to work well, for others many high resolution modes may still need to be checked by eye before finding a suitable match. Obtaining a high resolution frequency for every low resolution mode would be highly arduous.

Figure 8.6 shows the imaginary part of the Lorenz and Charney-Phillips eigenvalues for the fully boundary layer on problem with boundary layer 1 and $k = \frac{2\pi}{10^6}$. In order to produce more comparable results the six most damped modes have been omitted. For the most damped modes the real part of the eigenvalue is important. If the frequency is not captured accurately it is not particularly damaging since the modes will damp away before errors in propagation become detrimental to the model. With the most damped modes taken out of the spectrum the Lorenz and Charney-Phillips grid can be seen producing very similar results. All frequencies are virtually identical.




Figure 8.6: Boundary layer 1 reference state, boundary layer switched on in the transients and horizontal wavelength $k = \frac{2\pi}{10^6}$. The most damped modes are excluded from the plot.

It should be noted that the Lorenz and Charney-Phillips grids do not necessarily damp the same dynamical modes at the same rate. For example if the number of damped modes in Figure 8.6 is increased to include the seven fastest damped modes then the Charney-Phillips grid solution gains a mode in the inertio-gravity branch while the Lorenz grid gains a solution in the Rossby branch. Having examined

different test cases there does not appear to be a general rule for the order in which modes are damped by the boundary layer. Tracking the modes that become damped when the boundary layer is switched on is likely to be inaccurate.

Figure 8.6 showed a comparison of the frequencies for the clear dynamical modes, it is apparent that relatively little difference occurs between the grids. If greater difference were evident it would have been useful to obtain the corresponding high resolution frequency for each mode. In the next section the ability of each grid to capture the structure of these dynamical modes is examined.

Figure 8.7 shows the frequency of the modes found using the deepest boundary layer, boundary layer 5, the wavenumber is $k = \frac{2\pi}{10^6}$. As before the most strongly damped modes have been omitted from the plot since their decay rate is the important and dominating factor. For a boundary layer of this depth the majority of undamped modes are the acoustic modes with index 1 to 20 on Figure 8.7. These are followed by six inertial modes for which the frequencies found by Lorenz and Charney-Phillips agree well. In terms of capturing the propagation of the dynamical modes little difference exists between the two grid configurations, as found for boundary layer 1.




Figure 8.7: Boundary layer 5 reference state, boundary layer switched on in the transients and horizontal wavelength $k = \frac{2\pi}{10^6}$. The most damped modes are excluded from the plot.

The other boundary layer depths (not plotted) have also been examined and

the results for boundary layer 1 and boundary layer 5 extend to all cases. For this longer horizontal wavelength case little difference is seen between the frequencies of the non strongly damped dynamical modes. Given the similarity of the frequencies of the dynamical modes it is clear that overall both grids are either good or both grids are poor. As was seen in the comparison of the coupling, for this horizontal wavelength it is the slowest Rossby modes that are captured less accurately, which are also the modes damped by the boundary layer.

In Chapter 7 it was established that for the shorter horizontal wavelength ($k = \frac{2\pi}{10^3}$) it is the inertio-gravity waves that are slowed by using the Lorenz grid. As also established in Chapter 7, when performing the mode tracking, inertio-gravity modes are not the types of mode that are expected to be damped by the boundary layer. Therefore differences are expected between the frequency of the dynamical modes for the $k = \frac{2\pi}{10^3}$ case.

Figure 8.8 shows the comparison of frequencies for the Lorenz and Charney-Phillips grids for boundary layer 1 and $k = \frac{2\pi}{10^3}$. The shorter horizontal wavelength case was considerably harder to interpret than the longer horizontal wavelength case, as was found when examining the coupling in the previous chapter. As well as a number of growing modes the branches of inertio-gravity and Rossby waves become entangled by the Doppler shifting. In Figure 8.8 the Lorenz modes are plotted by decreasing frequency, doing this for Charney-Phillips configuration does not necessarily result in matching corresponding modes however. Instead each Lorenz mode has to be compared to all Charney-Phillips modes to find the best matching mode. For low resolution this is most effectively achieved by examining them by eye. The modes that are omitted from Figure 8.8 are either due to there being no suitable match between Lorenz and Charney-Phillips or because they are quickly decaying boundary layer type modes.

It is clear from the figure that with this shorter horizontal wavelength that there still remains clear difference between the frequencies of the dynamical modes, as expected. For the dynamics only case with this wavenumber it was found that inertio-gravity modes were slowed by Lorenz. Here there are some modes that




Figure 8.8: Boundary layer 1 reference state now with the shorter horizontal wavelength $k = \frac{2\pi}{10^3}$, boundary layer switched on in the transients. Modes which were not captured by both grids are removed.

have higher frequency with Lorenz compared to Charney-Phillips and some that have lower frequency. In the next section when dynamical modes are identified and examined it will be useful to examine these differences.

Fully Coupled: Mode Structures

When the coupling mechanisms were examined it was found that the proportion of dynamical to boundary layer modes was dependent on the depth of the boundary layer. Deeper boundary layers have a stronger damping mechanism and are capable of damping all of the Rossby and gravity waves; shallow boundary layers have weaker damping and can only support a few strongly damped modes. Further to this, the arrangement of the model levels must be considered. If more model levels are present in the boundary layer region it will lead to a larger number of damped modes being supported; if more levels are placed above the boundary layer it will lead to a larger number of dynamical modes being supported. For example if the logarithmic grid is used for the high resolution then more boundary layer modes will be supported, if the log-linear grid is used instead the ratio will shift to support more dynamical modes. The log-linear grid is used for the high resolution to give a

balance, ensuring sufficient resolution above the boundary layer for capturing the full structure of larger scale dynamical modes, whilst allowing for approximately second order convergence of high resolution solutions. Using the log-linear grid has the added benefit of capturing fewer boundary layer modes, which are far from orthogonal and thus less suited to the eigendecomposition methodology.

When the boundary layer is included in the reference state but not in the transients and $k = \frac{2\pi}{10^6}$ differences remain in the frequencies of the Lorenz grid and Charney-Phillips grid Rossby modes. This is expected from looking at the isothermal case. However the modes which are most significantly slowed by Lorenz, compared to Charney-Phillips, are the slowest propagating modes, which tend not to be present as Rossby modes in the coupled problem. Although the rate of propagation of dynamical modes in the coupled problem is similar between grids it is also important to check how well the structures of the modes are captured. For all cases examined the acoustic mode structures are found to be captured equally accurately by both grids in comparison to high resolution. Two good approximations can be made for acoustic modes, firstly that they are independent from any steady or transient boundary layer structure and secondly that acoustic modes do not depend on the transient θ' structure, the handling of which is the difference between the grids. Due to these approximations it is sufficient to understand the differences between the grids for the acoustic modes by examining the dynamics only case [68], plots comparing their structures are omitted here. The inertio-gravity and Rossby mode structures will depend on the boundary layer and thus they still need to be examined. The configuration which allows for the clearest comparison of dynamical modes is the shallow boundary layer 1 case with $k = \frac{2\pi}{10^6}$.

Figure 8.9 shows the same inertio-gravity mode that was shown in Figure 7.32 for the Lorenz grid. The figure shows the mode energy variables as captured by the Lorenz grid, the Charney-Phillips grid and by the 100 grid point log-linear grid. For the log-linear 100 grid point model there is no noticeable difference between the Lorenz and Charney-Phillips grids for this mode and so it is considered the actual solution. In order to give clear plots only the dominant part of the complex

eigenvector is shown, generally either the real or imaginary part of the mode is represented at once and so it is sufficient to compare only this in order to check how well the mode is captured. The frequencies and growth/decay rates of the plotted mode, as captured by each grid, are shown in the legend of the eigenvalue plot. The imaginary and real parts of the eigenvalue are ordered separately and the absolute value part plotted.




Figure 8.9: Shows the same inertio-gravity mode that was shown in Figure 7.32 except now against Charney-Phillips I-i and the high resolution run. The values of the frequency and growth rate are given for each case in the eigenvalue legend. Note that the two eigenvalue plots are ordered separately and the index numbers do not correspond to one another.

It is clear from the figure that with this configuration, i.e. boundary layer 1 and $k = \frac{2\pi}{10^6}$, the two different staggering options produce results that are quite close. This is a mode that has dominant structure in the E_u and E_v fields and so it is important that these are captured accurately, since Lorenz and Charney-

Phillips configurations do not differ in the Coriolis parts of the equations it is expected that they should produce similar results. The high resolution solution demonstrates how the inertial mode can only exist above the boundary layer. Also note how the structure of the high resolution mode is clearly bunched up above the boundary layer, this can be understood through a WKB type approximation. Rearranging the inertio-gravity wave dispersion relation given in equation (7.2) on page 198 gives,

$$m \approx \left(\frac{K^2 (N_b^2 - \omega^2)}{\omega^2 - f^2} \right)^{\frac{1}{2}}. \quad (8.2)$$

Regions where N_b^2 is larger will give larger m . In the boundary layer N_b^2 is largest and it decreases upwards, resulting in larger m lower in the domain. This height dependent wavenumber was seen in the low resolution mode structures when comparing the coupling in the previous chapter but it is much more noticeable when they are compared with the high resolution solution. Above the boundary layer both grids capture the wavelength of the mode accurately but slightly more accurately higher up.

In the high resolution solution there is a sharp jump to zero in E_u and E_v at approximately the height of the boundary layer which is not resolved by either low resolution grid. In trying to represent this sharp feature the Charney-Phillips grid produces some overshoot, rather than just going to zero the variable changes sign then goes to zero. The Charney-Phillips grid also produces a larger upward vertical shift than the Lorenz grid for the location of the sharp jump.

The E_θ part of this mode is also captured reasonably well by both grids, but is again more accurate higher in the domain where the structure is smoother. Again there is a sharp jump to zero at the top of the boundary layer but this time it is captured more accurately by the Charney-Phillips grid. E_θ is approximately zero in the boundary layer region for the high resolution and Charney-Phillips grid whereas has a large positive spike with the Lorenz grid.

All of the inertio-gravity modes are captured reasonably accurately by the two grids. This has also been confirmed for other boundary layer depth test cases.

Figure 8.10 shows the Rossby mode that was shown in Figure 7.33, again the

Lorenz grid is compared to the Charney-Phillips grid and the high resolution solution. Low resolution results are similar to those found for the inertio-gravity modes and the overall structure of the mode is captured well. As was found for the inertio-gravity modes there is high resolution structure at the top of the boundary layer that cannot be resolved by the low resolution grids. The Charney-Phillips grid captures this structure near the top of the boundary layer more accurately. Again the results extend to other Rossby modes captured in this and other test cases. Both Lorenz and Charney-Phillips find this mode to be slowly growing, the rate of growth is similar for each and similar to the high resolution. The frequency of this mode is also well captured by the low resolution grids.




Figure 8.10: Shows the same Rossby mode that was shown in Figure 7.33 except now against Charney-Phillips I-i and the high resolution run. The values of the frequency and growth rate are given for each case in the eigenvalue legend.

The results found for boundary layer 1 have been checked for other boundary

layer depths and similar results have been found across the cases. For the dynamical modes supported in the coupled equations, with $k = \frac{2\pi}{10^6}$, there is little to distinguish between Lorenz and Charney-Phillips. Both grids can capture the frequencies and growth rates well in comparison to the high resolution and the overall structure of the modes is also captured well.

In Figure 8.7 the frequencies in the $k = \frac{2\pi}{10^3}$ case were examined. This horizontal wavelength is of a similar order to the depth of the boundary layer and thus the depth of the region of shear. As a result the flow is more susceptible to shear influence, for example the shear driven decay and instability that have been observed. Rather than the system supporting just modes whose behaviour can be understood in terms of the dynamics or boundary layer mechanisms considerably more complex behaviour is encountered. This means that the behaviour of the different modes is much harder to interpret. Although these difficulties exist it has been possible to go some way to comparing the Lorenz and Charney-Phillips grids for this case. Figure 8.11 shows the low resolution Lorenz and Charney-Phillips grid versions of the mode with index 27 in Figure 8.7.

The mode plotted in Figure 8.11 was identified as being of interest since a significant frequency difference is observed between the low resolution Lorenz and Charney-Phillips grids. The Rossby and inertio-gravity modes appear in a single branch of frequencies for the $k = \frac{2\pi}{10^3}$ meaning the only way to identify the mode type is to examine the structure. However, as is clear from the figure, it is not straightforward to identify this mode type; the dominant energy variables do not coincide with those that would be expected to dominate in any particular dynamical mode. The singular vectors can assist in identifying this mode. Recall that there are certain relationships between the input and output singular vectors for each type of mode. For example an acoustic mode would be represented by having dominant input energy variable as pressure and dominant output energy variable as vertical velocity, or vice-versa. Singular vectors with structure very similar to that seen in Figure 8.11 can be identified. For these singular vectors the relationship between input and output is the same as that which was seen for the dynamical




Figure 8.11: A comparison of Lorenz and Charney-Phillips for a dynamical Rossby mode for the $k \frac{2\pi}{10^3}$ case. The plotted mode corresponds to the mode with index 27 in Figure 8.7.

Rossby modes, implying that this is in fact a Rossby mode. In contrast, for this wavenumber differences in frequencies would be expected for the inertio-gravity modes.

From Figure 8.11 it is clear that the high resolution modes are not sufficiently close to draw any solid conclusion about the behaviour of the Lorenz and Charney-Phillips grids. First of all the structure in the boundary layer region is completely different. The high resolution Charney-Phillips mode has no structure at all in the boundary layer whereas the high resolution Lorenz version has significant structure. Towards the top of the boundary layer the high resolution Lorenz grid captures a two-grid like structure. This large difference in structure in the boundary layer also appears to be the cause for the large phase differences in the E_v and E_w

fields. It is not possible to determine from this experiment which of the high resolution solutions, if either, is most accurate; indeed one would expect them to be considerably closer considering the large number of grid points in the boundary layer region. The high resolution solutions are further from each other than the low resolution solutions, suggesting non-physical behaviour is occurring, especially given the two-grid like structure. The frequencies of the high resolution solutions are very similar, however the Lorenz grid solution has much larger decay rate than found for the other solutions. The structure in the boundary layer with the high resolution Lorenz grid extends to the low resolution solution of E_v and E_p . If the Lorenz grid is producing spurious behaviour it is also seen in the low resolution. The low resolution Lorenz grid under predicts the frequency and the low resolution Charney-Phillips grid over predicts the frequency. It is difficult to establish whether the low resolution structure is captured more accurately because of the differences in the high resolution and because the low resolution versions of the mode are actually closer than the high resolution versions.

The difficulties encountered for this mode are encountered when comparing any non-acoustic mode for the $k = \frac{2\pi}{10^3}$ case.

8.2.2 Examination of Damped Modes

It is clear from the steady states and the dynamical modes that there exists differences between the Lorenz and Charney-Phillips configurations. However the differences are not as clear cut as they were when the dynamics and boundary layer were considered separately. The Lorenz grid captured the boundary layer steady states more accurately however the coupled steady states are generally captured more accurately by the Charney-Phillips grid. When the dynamics was studied on its own the Charney-Phillips grid gave optimal wave propagation. For the coupled problem the extent to which the Charney-Phillips grid outperforms the Lorenz grid in terms of propagation is reduced; the main differences occurred for the slowest Rossby modes, however these are distorted into boundary layer modes and so growth rates become important. From the individual dynamical modes that

could be examined both the Lorenz and Charney-Phillips grids were able to capture the structure accurately. To complete the comparison the damped boundary layer modes are examined. These present a more challenging problem in that the methodology of eigendecomposition is not expected to work as well as it does for the dynamical modes. However, it is worth first attempting the eigendecomposition to establish whether or not it can be useful.

Boundary Layer 1

The approximate rates of damping for the boundary layer modes can be seen in the real part of the eigenvalues, plotted in Figure 8.9 and Figure 8.10. The real part of the eigenvalues μ represents the growth or decay of the mode. The right hand end of the plot shows the damped boundary layer modes and the left hand end shows the growing modes. There appears to be quite good agreement between the two grids for the most damped modes. However this needs to be examined in detail. Figure 8.12 shows the six most damped boundary layer modes for Lorenz and Charney-Phillips for boundary layer 1, $k = \frac{2\pi}{10^6}$. Each Lorenz mode is matched to a Charney-Phillips mode by checking orthogonality. By this test the ordering of the damping rates is found to be the same for each, i.e. the fastest damped Lorenz mode is also the fastest damped Charney-Phillips mode. The eigenvectors corresponding to these modes will generally not be orthogonal so the check is likely inaccurate; examining Figure 8.12 by eye it does seem to have found the best matches though. That the orthogonality test has any success is likely due to the small number of boundary layer modes supported by the system. Whether or not these are indeed the same modes is difficult to determine for such low resolution. A comparison of the actual growth rates for the six modes plotted Figure 8.12 is provided in Table 8.1.

There are clear differences between Lorenz and Charney-Phillips, there are differences in the structures of the modes and by Table 8.1 the Charney-Phillips grid appears to be damping most modes more slowly than Lorenz. The problem now is in obtaining the equivalent high resolution modes so it can be determined whether




Figure 8.12: The E_u and E_θ parts of the six most damped Lorenz and Charney-Phillips modes for the boundary layer 1 configuration. The index corresponds to ordering by real part so index 48 is the fastest damped.

Index	μ (Lorenz)	μ (Charney-Phillips)
43	-4.434×10^{-5}	-1.706×10^{-5}
44	-1.080×10^{-4}	-9.093×10^{-5}
45	-1.268×10^{-4}	-1.054×10^{-4}
46	-7.923×10^{-4}	-6.893×10^{-4}
47	-1.590×10^{-3}	-8.123×10^{-4}
48	-1.602×10^{-3}	-2.334×10^{-3}

Table 8.1: The rate of damping μ for the most damped modes in the boundary layer 1, $k = \frac{2\pi}{10^6}$, test case. Index corresponds to the modes plotted in Figure 8.12.

or not a particular configuration gives the better overall representation. The orthogonality test was employed which finds a ‘closest’ mode but even when checking by eye it is not clear if it is indeed the same mode. Given the lack of orthogonality it is not clear that two configurations should capture the same mode anyway. To improve matters the deeper boundary layers can be used, this way there are more grid points in the boundary layer region and larger scale structure to compare.

Boundary Layer 5

Figure 8.13 shows the real part of the eigenvalues for the boundary layer 5 test case, the modes that were omitted from the frequency plot are shown in black. As was being suggested in Table 8.1, for the shallowest boundary layer, the figure points towards the decay rate found when using the Charney-Phillips grid being smaller than the Lorenz grid. This was also found when examining the boundary layer only case in Part I where boundary layer modes were damped too slowly when using the Charney-Phillips grid. If the Lorenz grid does indeed find the decay rate more accurately then it would suggest a disadvantage in using the Charney-Phillips grid. Modes that are not damped quickly enough could produce more inaccurate results, particularly if the structures of those modes are also captured inaccurately.




Figure 8.13: Boundary layer 5 reference state, boundary layer switched on in the transients. The most strongly damped modes which were omitted from Figure 8.7 are shown in black.

Again although there appears to be faster damping of boundary layer modes by Lorenz this needs to be confirmed by comparing the structures to check that modes with the same index are indeed the same mode. Each pair of low resolution modes then needs to be found in the high resolution case to check what kind of structure and damping rate it should have. Further to this it has to be checked that each mode in the Lorenz grid high resolution agrees and is present in the

Charney-Phillips grid high resolution.

In the study in Part I the most interesting part of the boundary layer mode spectrum was the slowest of the damped modes. They are longer lived in the solution meaning they will have the largest influence in the solution. Now that the dynamics are included the rate of damping has to be considered against the time scales of the dynamics. The leading boundary layer modes have slowest damping rates, however if their decay rate is negligible in comparison to the frequencies of the dynamical modes then capturing their decay rate is less important. The fastest damped boundary layer modes are still less important in that they decay quickly, giving errors in capturing them less chance to impact the solution. So for the coupled problem the boundary layer modes that are most important to capture are those that are neither damped too quickly nor too slowly in comparison to the dynamical time scales. From Figure 8.7 the frequencies of the dynamical mode are in a range from approximately 10^{-4} to 10^{-5} so it is the boundary layer modes with decay rates in this range that are also important.

For the deepest boundary layer, there were approximately 21 modes that had the fastest damping and also had structure associated with the boundary layer. These have index 27 to 48 in Figure 8.13. Using the eigenvalues and eigenvectors the structure of the modes with decay rates on a similar time scale to frequencies of dynamical modes are investigated. These are the modes with real index near to 30. Figure 8.14 shows the Lorenz grid mode with real index 31 plotted against the Charney-Phillips grid mode with index 33, this is one of the most important boundary layer modes. Indexes correspond to Figure 8.13. The Charney-Phillips grid mode could be identified immediately using the orthogonality test. Both real and imaginary parts of the energy variables are plotted separately since there is greater difference between them for E_θ , ensuring a proper comparison.

Examining the structures in Figure 8.14 it seems highly likely that the grid configurations are both capturing similar modes. The horizontal velocity energy variables are almost identical. Although there are differences in E_θ this is somewhat expected given the different ways in which θ' is handled and the overall structure




Figure 8.14: Shows the damped Lorenz mode with real index 33 against Charney-Phillips I-i counterpart. Boundary layer 5. The values of the frequency and growth rate are given for each case in the eigenvalue legend.

looks quite similar for each. For this mode the Charney-Phillips grid predicts larger decay rate than predicted by the Lorenz grid. As for boundary layer 1, problems are encountered when trying to compare these modes to a high resolution version.

In order to attempt to compare the performance of the Lorenz and Charney-Phillips grids for the boundary layer modes a large number from different test cases were considered. For the boundary layer depth boundary layer 4 and boundary layer 5 depths were considered as these had the most promising results in Part I. For the high resolution solution both logarithmic and log-linear grids were considered and for the low resolution the Met Office and geometric grid were tried. Both horizontal wavenumber cases were also considered. Across these cases it was usually possible to identify modes which looked quite similar across the low resolution

solutions but very difficult to be sure of having found matches in the high resolution. As a result of the difficulties in matching low resolution to high resolution it has not been possible to use the eigendecomposition to perform a comprehensive comparison of the performance of the Lorenz and Charney-Phillips grids for capturing the boundary layer modes.

That the eigendecomposition is not useful is somewhat expected given the problems encountered throughout this study. Firstly the lack of orthogonality between modes becomes more problematic as resolution increases and more modes are included. The lack of orthogonality does not so severely prevent interpretation as it did in Part I. Despite this improvement however, when performing the orthogonality test there will be many potential pairings. Further to this the scale in the high resolution modes makes it difficult to compare structures; there may be significant small scale detail in the high resolution mode as was seen in Part I. This means that even when scanning through by eye it is difficult to believe with any confidence that the modes being compared are equivalent. The final issue that makes the whole process tricky is due to the near singularity in the surface layer formulation meaning potential differences, even between high resolution configurations.

Although it is difficult to determine the high resolution versions of boundary layer modes it is possible get reasonable matching between low resolution versions. In addition to the mode plotted in Figure 8.14 there are another three modes living in the interesting part of the spectrum which appear to be captured by both low resolution Lorenz and Charney-Phillips grids. The real part index and corresponding decay rates and frequencies of these modes are highlighted in Table 8.2.

The values of the decay rates in the table are somewhat inconclusive. Two modes are damped more slowly when using Charney-Phillips grid; two are damped faster. Generally the decay rates are quite close between each grid. For the frequencies there is also little difference, suggesting that corresponding modes result from similar dynamical modes being distorted by the boundary layer, and in an equivalent manner. These are the only clear cut results that can be obtained for

Real Index		Decay Rate μ		Frequency ω	
Lorenz	Charney-P	Lorenz	Charney-P	Lorenz	Charney-P
29	28	-1.052×10^{-5}	-4.346×10^{-6}	1.980×10^{-5}	1.999×10^{-5}
30	31	-1.640×10^{-5}	-1.928×10^{-5}	8.394×10^{-5}	7.554×10^{-5}
31	33	-2.561×10^{-5}	-2.645×10^{-5}	3.911×10^{-5}	4.260×10^{-5}
32	32	-3.028×10^{-5}	-2.089×10^{-5}	1.743×10^{-5}	1.776×10^{-5}

Table 8.2: The rate of damping μ and frequency ω for some of the least damped boundary layer modes in the boundary layer 5 test case.

the interesting part of the spectrum for this case. When other modes are examined there is less similarity between Lorenz and Charney-Phillips and it is difficult to be sure whether modes are equivalent. Based on these few results obtainable for boundary layer 1 and boundary layer 5, little difference is seen between Lorenz and Charney-Phillips grids. Decay rates of the boundary modes with the most relevant time scale appear to be quite similar.

8.3 Comparing Transients Using the SVD

Since the eigenvectors do not appear to offer much assistance in terms of comparing boundary layer modes, the singular values and singular vectors are next considered. For the boundary layer only they were found to be quite useful and for the cases considered in Chapter 7 they were shown to be able to represent the physical properties of the coupled system. If singular vectors that represent the boundary layer structures can be compared then it will give some insight into how the different configurations capture the structure of boundary layer modes. The possible difficulties in using the SVD is in the understanding gained from the singular values. They have been found to give a good insight into the corresponding decay rates or frequency of modes but only in the decoupled cases, where either real or imaginary part of the eigenvalue dominates throughout. Interpretation of the singular values for the coupled problem is not always clear.

Figure 8.15 shows a boundary layer type mode in the low resolution Lorenz singular vectors, parameters are for boundary layer 5 and $k = \frac{2\pi}{10^6}$. Although it

is not necessarily clear from the singular value that this is a damped mode the structure looks similar to the leading damped modes found in the eigenvectors. Using a test of orthogonality, as done in Part I, the corresponding low resolution Charney-Phillips representation of this singular vector is found, this is also shown in the figure. As was found for the eigenvectors with this kind of structure, the E_u and E_v parts are captured well by both grid configurations, but there is significant difference in the E_θ part.




Figure 8.15: A boundary layer type mode in the boundary layer 5, $k = \frac{2\pi}{10^6}$, test case, represented in the singular vectors the corresponding singular value is shown below. The figure shows the low resolution Lorenz and Charney-Phillips representations of this singular vector.

As was performed for the eigenvectors a large selection of test cases are used to try and compare the performance of the Lorenz and Charney-Phillips grids for capturing boundary layer singular vectors. The problem again is in finding a

high resolution solution from which conclusions about which low resolution grid is better can be drawn. When the boundary layer was studied on its own it was not always straightforward to find corresponding high resolution singular vectors. Generally some conclusions were possible by seeking the corresponding high resolution singular vector using the orthogonality test and checking that it was represented equivalently in both high resolution cases. This does not seem to be the case for the coupled problem. The whole process is more ambiguous and even with extensive searching across the test cases it has not been possible to find sufficient similarity between high resolution singular vectors corresponding to low resolution modes in the interesting part of the spectrum. This is possibly due to the increased number of mode types in the problem. Previously all modes were boundary layer type and so with enough testing some could be found. The problem of the near singularity also seems to be causing problems, high resolution Lorenz and Charney-Phillips singular vectors were found which agreed well in the velocity fields but poorly in the potential temperature. Recall that it is in the surface potential temperature flux that the large differences in C_{mn} and C_{hm} occur. When studying the boundary layer it was also possible to increase the confidence in finding a match by looking at the singular value. However it is more difficult to assume the singular value will be similar between grids now. Further, the fact that the singular value is harder to interpret means it is not possible to be sure that the singular vectors being examined represent modes lying in the interesting part of the spectrum. Two modes that have similar decay rates may have quite different singular values. If modes with similar decay rates represent distortions of very different dynamical modes they are likely to have different frequencies and thus different singular values. Any kind of tracking to establish which mode has been distorted is likely to be highly inaccurate.

SVD For Comparing Dynamical Modes

Although the SVD was not found to work well for the boundary layer type modes it can be effective when examining dynamical modes. Figure 8.16 shows the low

resolution Lorenz and Charney-Phillips singular vectors representing the largest scale acoustic mode. This is the kind of mode that is captured accurately by both Lorenz and Charney-Phillips and this can be seen in the dominant parts of this singular vector, input pressure and output vertical velocity.




Figure 8.16: The largest scale acoustic mode shown in the singular vectors. The Lorenz and Charney-Phillips representations are compared with a high resolution truth solution.

Further to the ability of both low resolution grids to capture the mode it is also easy to find the high resolution version, also shown in the figure. In order to find the corresponding singular vector in other configurations orthogonality is checked. All the singular vectors representing other behaviour will have inner product close to zero while, provided the unit norm is retained by appropriate Δz weighting, the inner product with corresponding singular vector is close to the unit circle (absolute value 1).

Clearly for this type of mode the SVD methodology works very well. Performing the SVD comparison of a dynamical mode is useful for testing the numerical process of finding corresponding singular vectors and thus checking the methodology. However the eigendecomposition also works well for these cases so is generally of less interest as a tool for comparison. Unfortunately the results for the boundary layer modes, for which the SVD is needed, are nowhere near as good.

At the beginning of this Chapter a number of other test cases were considered, namely different ways of obtaining K_h and different ways of stretching the grid. However these are the kinds of changes that would be most likely to influence the ability to capture the boundary layer modes, hence these are the most difficult to study making any conclusive results difficult to achieve. A number of these additional cases were tried for the transient comparison but were not found to improve matters regarding the comparison of boundary layer modes so no in depth discussion is presented.

8.4 Chapter Summary

Comparison of the Lorenz and Charney-Phillips grids has been shown for the fully coupled boundary layer and dynamics. First the steady state was examined, and then the transients for two cases of horizontal wavenumber, $k = \frac{2\pi}{10^6}$ and $k = \frac{2\pi}{10^3}$. For the $k = \frac{2\pi}{10^6}$ case results were presented for the shallowest and deepest boundary layers, boundary layer 1 and boundary layer 5. It has been difficult to generate highly conclusive results and that is due to the complex nature of the flow that is being considered and due to the inability of the methodology to relay easily interpretable results.

It is found that the Charney-Phillips grid gives the best overall representation of the steady state structures. This is in contrast to the boundary layer only model, for which the Lorenz grid is found to represent steady states most accurately. The Charney-Phillips configurations that gives the most promising results are those that average K_h or Ri , option I-i and I-iii. Neither wholly outperforms the other; option I-i gives better near surface structure and option I-iii gives better structure

near the top of the boundary layer. The Lorenz grid will sometimes perform better than option I-iii near the surface but is generally unable to outperform option I-i. This in mind the Charney-Phillips option I-i was used to perform the transient comparison.

Of the two horizontal wavenumber cases, the $k = \frac{2\pi}{10^6}$ case is most easy to interpret. The behaviour of the system divides clearly into the three types of dynamical mode and the boundary layer modes. From examining the dynamics only case it is found that the Lorenz grid reduced the frequencies of the slowest, smallest scale, Rossby modes; this is also identified as occurring for the linearisation about a sheared reference state. When the boundary layer and dynamics are coupled certain dynamical modes are distorted and become decaying boundary layer type modes. Through tracking (in Chapter 7) these have been identified as being the slowest Rossby modes and the gravity modes, the number of each depends on the strength of the damping and thus the depth of the boundary layer. Due to the distortion into boundary layer modes the only dynamical modes which are not quickly decaying have their frequencies captured equally well by the Lorenz and Charney-Phillips grid.

Results for the $k = \frac{2\pi}{10^3}$ case are considerably harder to interpret. The flow has strong instabilities and the distinction between different types of dynamical mode by frequency is not possible. Certain low resolution dynamical modes in the coupled problem have been identified as having different frequencies between Lorenz and Charney-Phillips, unlike in the $k = \frac{2\pi}{10^6}$ case. Although the structures of these mode types agree well in the low resolution it has been difficult to identify corresponding high resolutions that look well converged. It is therefore difficult to establish whether or not either Lorenz or Charney-Phillips captures the modes more accurately. Some of the structure in the high resolution Lorenz grid mode is different to what would be expected and to what is found in a similar mode in the $k = \frac{2\pi}{10^6}$ case. This would imply possible spurious or unstable behaviour being produced by the Lorenz grid and its associated averaging.

Comparing the ability of the Lorenz and Charney-Phillips grids to capture the

boundary layer modes has not been possible. The boundary layer modes are far from orthogonal and thus results from the eigendecomposition are not easy to interpret and cannot be used to compare grid configurations. For the boundary layer only study in Part I the SVD was found to be quite useful for comparing boundary layer modes. This was also tried here but could not be made to work. The SVD is found to be well suited for examining the boundary layer on its own or the dynamics on its own, or for examining the dynamics in the coupled problem. When it is used to examine boundary layer modes in the coupled problem it is firstly very difficult to interpret what the singular value represents and secondly find corresponding high resolution boundary layer singular vectors.

Chapter 9

Overall Summary

In this chapter all of the results and findings of this thesis are summarised. The extent to which the posed questions have been answered is examined. Areas where it has not been possible to completely fulfil the initial aims will be identified and discussed. Some ideas for further work will be outlined at the end of the chapter.

The overall aim, as set out at the beginning of the study, was to determine whether either the Lorenz or Charney-Phillips vertical configurations offered the greater accuracy when coupling large scale dynamics to the sub-grid scale physics. In the previous work of Thuburn and Woollings [68] it was shown that certain configurations are favoured for the large scale dynamics. When they compared Lorenz and Charney-Phillips it was Charney-Phillips that was favoured. In their study Thuburn and Woollings were able to implement a systematic methodology, based on normal mode analysis. The aim of this thesis was to use a methodology like this to first confirm that the Lorenz grid is better for certain physics processes and secondly to examine the fully coupled problem.

The case that was chosen for the study was the coupling between the large scale Rossby, inertio-gravity and acoustic waves and the sub-grid scale motions of the stably stratified planetary boundary layer. These three dynamical wave types (represented by the normal modes) are those supported by the inviscid Euler equations of motion. For the coupled dynamics-boundary layer case there is a clear conflict in which grid configuration to choose: for the optimal representation of the

dynamical modes the Charney-Phillips grid is optimal [68] whereas the Lorenz grid allows for the boundary layer equations to be written without any averaging. Determining whether the Lorenz grid would be optimal for the planetary boundary layer was part of this study. An inherent question in the study is in determining the nature of the Lorenz grid computational mode when coupling dynamics and physics.

Using the normal mode methodology, as in [68], is appealing since it allows for a highly systematic approach and reveals for what scales a particular configuration performs better. In this method the time dependent part of the system is decomposed into a set of linearly independent small amplitude oscillations called modes. Linear time dependent motion can then be represented by combinations of these modes. Examining all of these structures independently to see which are most accurately captured can offer a very clear insight into why a configuration performs better. The alternative would otherwise be to code up the full problem with certain initial conditions and compare model output after certain time. This only offers insight into the performance of the configurations for a very specific situation, dependent on the initial conditions. The approach can be quite ambiguous and provide little general insight into how configurations can be expected to perform. However, the advantage of using a time integration is that the whole nonlinear response would be considered, whereas normal mode analysis can only measure the linear response. The overriding issue that needs to be considered if attempting to use the normal mode analysis is that for the methodology to work seamlessly requires a normal system. For the fully coupled problem it is not always clear whether this suitable system can be achieved.

Due to the complexities in applying the normal mode analysis the focus of the thesis has not only been on answering the central question over choice of vertical configuration but also in analysing the methodology and establishing its suitability for this kind of problem. Further to this the actual physical interactions in the coupled problem are not simple and have required careful examination. This has been required in order to understand and effectively compare configurations but

has also revealed some interesting behaviour in itself.

The discussion of the outcome of this work can be divided into examining the effectiveness of the methodology and examining the results, i.e. which configuration is better. The next two sections contain the discussion of these two parts of the work respectively.

9.1 Methodology

In order to use eigenmode analysis the system is linearised around a steady reference state and the system of transient equations is written in matrix form $\lambda \mathbf{x} = \mathbf{A} \mathbf{x}$. The matrix \mathbf{A} contains the coefficients of the transient model prognostic variables and will depend on the steady state. Modes of the system \mathbf{x} are the eigenvectors of the matrix of coefficients and the eigenvalue λ corresponding to the eigenvector describes the frequency and growth or decay rate of the mode.

The issues that were encountered first were in obtaining a proper steady state of the system. Firstly the form of the boundary layer terms required modification to allow a steady state to be reached and secondly a coordinate transform was required in order to get good convergence with increasing resolution so that a ‘truth’ solution could be obtained in reasonable computational time. That the boundary layer terms require modification is due to the imposed boundary conditions on potential temperature. In most realistic situations the planetary boundary layer is far from steady but evolves continuously as heating and cooling occurs. In operational models this is represented by imposed fluxes, for example at the surface, and by the diffusion mechanism. In order to use the proposed methodology it would be insufficient to represent flux terms as this would never result in a completely steady state; effectively leaving the diffusion unbalanced. The result is that the boundary layer continuously loses potential temperature, reducing it to the value of the surface boundary condition throughout the domain. This reduction in potential temperature would normally be balanced by the natural fluxes at the top of the boundary layer and at the surface. In order to get around this modelling artefact and obtain a steady state but with a realistic θ profile a subsidence heating was

imposed. This downward advection of potential temperature balances the diffusion while allowing a steady state to be reached. With the right chosen parameters it will also produce realistic stable boundary layer steady state profiles. A further radiative cooling term was added to the potential temperature equation to give stable, rather than neutral, stratification above the boundary layer. Having stable stratification will damp any spurious creation of shear above the boundary layer, which could otherwise prevent a steady state being reached.

Differences in solutions found when using Lorenz and Charney-Phillips grids occur due to averaging of prognostic variables. At operational resolution these differences can be quite significant. In order to quantify the differences between the two configurations they need to be compared to the high resolution solution. A high resolution is such that there is no noticeable difference between solutions found when using Lorenz and Charney-Phillips. Near to the surface the velocity and potential temperature in the system behave in a logarithmic way and this results in derivatives of those profiles approaching a singularity as the grid is refined. This singularity reduces the convergence rate of the scheme and so it is not possible to obtain good convergence. With a geometric type grid stretching it would require of the order of one thousand grid points to obtain suitably similar Lorenz and Charney-Phillips results and obtain the high resolution solution. Given the large number of test cases to consider this would require an unrealistic amount of computational effort. In order to avoid the singularity a coordinate transform was performed to give a logarithmic grid stretching. In doing the transform a close to second order convergence was recovered and a high resolution solution was obtained for one hundred grid points. Due to the small grid spacing near the surface with a high resolution logarithmic grid the false time-stepping scheme, that was used for the low resolution steady state generation, could not be used, due to the small time step that would be necessary. Instead the Newton method was used successfully by transitioning smoothly from neutral to steady conditions in the boundary layer terms. This is an appealing method as it allows for rapid generation of steady state profiles and generates the Jacobian matrix, also used in obtaining the tran-

sient part of the calculation. Implementing a Newton method thus requires little extra coding, ideal if testing lots of cases.

If the matrix of coefficients \mathbf{A} is normal then the eigenvectors form an orthogonal basis and represent normal modes of the system. These normal modes are well understood and can be interpreted with relative ease. When the matrix is far from normal the eigenvectors do not form an orthogonal basis and so are not normal modes, instead they are thought of as eigenmodes. These eigenmodes still represent physical behaviour of the system but can be very similar to each other and are thus considerably harder to interpret than when looking at normal modes. When the boundary layer is included in the equations, whether with all the dynamics or not, the matrix will not be normal. Although it is not normal two levels of interpretation depending on whether dynamics are included or not are seen.

When the boundary layer was studied on its own it was found that the matrix was far from normal and the usefulness of the eigenvectors was completely lost. This was discussed in the previous summary chapter at the end of Part I. When the dynamics were included in the equations the usefulness of the eigendecomposition was somewhat recovered. It has been possible to interpret the behaviour and compare the Lorenz and Charney-Phillips grids for the dynamical modes using the eigendecomposition. It has also been possible to identify the boundary layer modes and compare low resolution versions of boundary layer modes. However the non-normality caused problem when trying to compare boundary layer modes from low resolution solutions with boundary layer modes from high resolution solutions. For the boundary layer only it was not possible to even compare low resolution eigenvectors so an improvement is certainly gained in the fully coupled problem.

A method that was employed to try and avoid the issue of orthogonality was the SVD. This is a method similar to the eigendecomposition and indeed the two are related for certain situations. Although this method does not offer the same interpretation as the eigendecomposition it will always produce singular vectors which form an orthogonal basis. The method was used mainly for the boundary layer only case where lack of orthogonality between eigenvectors was most prevent-

ing of physical interpretation. Although the method was used with a degree of success for this case it was found to be less useful for the fully coupled problem. Singular vectors representing boundary layer structure were identified in the low resolution configurations but, as was found for the eigenvectors, it was not possible to match those up with high resolution singular vectors, at least not to the same level of accuracy that was possible for the boundary layer only case. With the fully coupled problem more different mode types have to be simultaneously captured by the decomposition and this hinders the SVD. When only boundary layer modes are captured the chance of finding corresponding modes is increased by having more to choose from.

A further issue that hinders the finding of corresponding high resolution modes is the presence of a singularity in the surface formulation of the boundary layer terms. In the neutral drag coefficient there is a term in the denominator tending to zero for height levels close to the roughness length. For the logarithmic grid, for which levels can get very close to the roughness this singularity appears to impact the solution. The surface formulations seem unable to balance with the singularity for such bunched up height levels.

In addition to the problems of orthogonality between eigenvectors the form of the spectrum was considered. Comparison of modes across configurations requires that only discrete spectrum is supported. If continuous spectrum is supported then it is entirely feasible that different configurations could sample from different parts of the continuous spectrum. The form of the spectrum was carefully examined and no evidence was found for the presence of continuous spectrum. When dealing with a complex problem like that being considered here however, continuous spectrum can not be completely ruled out. Ideally a rigorous derivation would be formulated to show the form of the spectrum for the equations, though this is beyond the scope of this project.

In all, the methodology does not work seamlessly. Unfortunately the system can be far from normal meaning eigenvectors become difficult to interpret and compare between configurations. Instead a selection of workarounds have been devised and

used to varying degrees of success. Importantly it has allowed for some comparison of grid configurations to be made using a more systematic approach. Additionally it has provided good insight into the limitation of the methodology and thus the non ideal situations where it may still be useful. For example if one wished to study the dynamical modes in any coupled problem but was less interested in the modes of the physics the eigendecomposition could still be quite useful.

9.2 Vertical Configurations

It is well understood from the literature that the Charney-Phillips grid is the preferred choice for modelling the dynamics, at least in terms of obtaining optimal wave propagation and in that it supports no computational mode. This has been further confirmed in this work. In Part I and discussed in the summary therein, it was shown that the Lorenz grid is the preferred choice for the stably stratified planetary boundary layer. It seems likely from the discrete equations that the Lorenz would be preferred since no averaging occurs. This was confirmed through studying the steady state and the singular vectors and singular values of the linearised system.

When the equations include both dynamics and physics the difference between the two grids is quite clear for the steady states but less so for the transients. The steady states of the full coupled problem are captured more accurately by the Charney-Phillips grid than the Lorenz grid. In particular the option I-i and option I-iii configurations performed well, these are where either eddy diffusivity or Richardson number is averaged. For the boundary layer only the Lorenz grid had performed better for the steady states, however when density is included the benefits afforded by the lack of averaging in the diffusion terms is outweighed by the need to satisfy hydrostatic balance and the associated averaging.

When [68] and later [64] investigated the dynamics only case they found the Charney-Phillips grid to be preferred due to two key factors. Firstly the Lorenz grid was shown to support a computational mode and secondly the Charney-Phillips grid was shown to give optimal representation of the dispersion of the Rossby and

inertio-gravity waves. The Lorenz grid significantly slows either Rossby or inertio-gravity modes, depending on the horizontal wavenumber. The size of the horizontal wavenumber determines whether buoyancy frequency terms either dominate or do not in the dispersion relations. The associated averaging associated with this term when using the Lorenz grid results in slowing of frequencies when this term dominates. When horizontal wavenumber is small the buoyancy terms dominate in the Rossby dispersion relation, when it is large buoyancy terms dominate in the inertio-gravity dispersion. In Part II of this thesis it was shown that the same modes are slowed in these two horizontal wavenumber regimes when using the stretched grid. However it was not possible to determine an exact dispersion relations for the stretched grid due to the increased complexity in the mode structures. Modes with higher frequency are found to be supported by the stretched grid and these have structure only on regions with sufficient resolution to support their increased frequency. As waves propagate into regions that cannot support their frequencies they can be reflected; evidence of this kind of behaviour has been observed for the dynamical modes on a stretched grid.

The dynamical modes that were found to be much better represented by the Charney-Phillips grid did not always exist in the coupled problem. For example in the longer horizontal wavelength case the Lorenz grid slowed the Rossby modes in the large m limit, i.e those with shortest vertical scale. However these are the modes most susceptible to the boundary layer diffusion. In the cases tested for this horizontal wavelength no coupled case Rossby modes were found to have considerable difference between Lorenz and Charney-Phillips. It is concluded that some of the benefits afforded by using the Charney-Phillips grid regarding frequencies are lost once the boundary layer is included. Conversely for the shorter horizontal wavelength case it was found that dynamical modes did have differences in frequencies as predicted by the Lorenz and Charney-Phillips grid. These would be expected to be inertio-gravity modes since these are slowed by the Lorenz grid in the dynamics only case. However some Rossby modes were also identified as having frequency differences. It was generally more difficult to analyse the larger

horizontal wavenumber case since modes were more sensitive to the Doppler effect making it hard to distinguish between mode branches. It was also not possible to find converged high resolution modes.

As discussed in the methodology section of this chapter it was difficult to generate conclusive results regarding the boundary layer type modes. It was possible to compare some of the leading boundary layer structures across the Lorenz and Charney-Phillips configurations when at the operational resolution but not to compare with the high resolution versions. Caution is taken in concluding too much from this part of the results given the issues with methodology. Some modes were found to be damped faster by Charney-Phillips compared to Lorenz and some were found to be damped slower. However there were not unanimous results across the test cases. For the boundary layer only the Lorenz grid gave better predictions for the steady states, for the coupled equations the Charney-Phillips grid gave better representation of the steady states. This in mind some improvement in the representation of coupled boundary layer modes may be expected when using the Charney-Phillips grids.

The Lorenz grid computational mode was examined for two cases, one where the boundary layer was switched off in the transients and one where the boundary layer was switched on. The case where it is switched off appeared to reveal that shear suppresses the structure of the computational mode. This may not necessarily always be the case as the effect of the shear will just be to realign horizontal waves away from a two-grid wave configuration, it could just as well act to align back into a two-grid wave configuration. When the resting isothermal reference state is considered the computational mode has zero frequency and appears as a two-grid wave in θ throughout the domain. When the boundary layer is included and the reference state has shear the computational mode structure is altered. Instead of having zero frequency it has frequency of the same order as the other dynamical modes as it is advected by the background flow. It can also have structure in the other fields due to geostrophic balance above the boundary layer. This methodology is not particularly well suited to examining the computational mode further, to

establish the influence of having shear throughout the domain. The isothermal case of [68] was extended to include shear throughout the domain in the reference state; however for the cases tried continuous parts of the spectrum were identified. For the fully coupled problem the computational mode was identified in all cases. As with the boundary layer switched off case, the structure of the computational mode was suppressed only in the boundary layer region. All the results imply that the boundary layer is not capable of distorting the computational mode as might have been hoped. This means the main downside of using the Lorenz grid remains. This also agrees with the results of [81]. In [81] behaviour associated to the Lorenz grid computational mode was identified in a hurricane model, formulated with a well mixed boundary layer.

9.3 Concluding Discussion

Despite the difficulties encountered in attempting to apply this kind of methodology to this kind of problem some significant insight has been gained. The boundary layer only case was examined and it has been shown that the Lorenz grid is the preferred configuration. In doing this a fully steady state boundary layer has been obtained along with good convergence of the high resolution solutions. The singular vector analysis has been employed and found to be useful in the situation where eigenvectors are not.

In studying the coupled problem understanding has been gained of the coupling processes that occur and which types of dynamical modes are replaced by boundary layer modes. The Lorenz and Charney-Phillips configurations have been examined for the coupled problem and it has been shown that, for the dynamical modes at least, there is little to distinguish between the two grids. Gaining a full insight into the accuracy of the boundary layer modes has not been possible but it seems likely that the Charney-Phillips grid will do better than when no dynamics are included. The Charney-Phillips option I-i configuration outperformed the Lorenz grid for the coupled steady states. It was also found that the computational mode supported by the Lorenz grid configuration is only suppressed in the boundary layer, but

not everywhere in the domain as had been originally hoped. Given these findings it seems unlikely that, neglecting conservation properties, much benefit would be gained by using the Lorenz grid for a model covering the large scale dynamics and the stably stratified planetary boundary layer.

From the work undertaken here a number of potential issues have arisen that could not be covered within the scope of this project but might prove of interest for further study. Firstly it would be interesting to further investigate whether the spectrum is discrete, continuous or consists of both discrete and continuous parts. Second, some interesting reflection type behaviour was identified in the dynamical modes when using the stretched grid. It may be possible to identify the types of discrete modes that can be expected to be captured by a given stretched grid by considering the boundary condition at the point of reflection. This would give a further understanding of the use of normal modes with a stretched grid. Although the aim of this thesis was to use a systematic methodology to examine differences between grid staggerings it may further complement the findings by performing time integrations of the full model with different initial conditions.

This study has used the stably stratified planetary boundary layer for the physics. This configuration gives the most obvious conflict in the choice of grid staggering. It would be interesting to try adding in other physics which may benefit from different choices, for example cloud processes or gravity wave drag. One could test the methodology for these cases and also consider coupling between different sub-grid models, e.g. radiation and convection.

Bibliography

- [1] D. S. Adamson, S. E. Belcher, D. J. Hoskins, and R. S. Plant. Boundary-layer friction in midlatitude cyclones. *Quart. J. Roy. Meteorol. Soc.*, 132:101–124, 2006.
- [2] A. Arakawa and C. S. Konor. Vertical differencing of the primitive equations based on the Charney-Phillips grid in hybrid $\sigma - p$ vertical coordinates. *Mon. Wea. Rev.*, 124(3):511–528, March 1996.
- [3] A. Arakawa and V. R. Lamb. Computational design of the basic dynamical processes of the UCLA general circulation model. *Meth. Comput. Phys.*, 17:173–265, 1977.
- [4] A. Arakawa and S. Moorthi. Baroclinic instability in vertically discrete systems. *J. Atmos. Sci.*, 45(11):1688–1708, June 1988.
- [5] R. J. Beare. Boundary-layer mechanisms in extratropical cyclones. *Quart. J. Roy. Meteorol. Soc.*, 133:503–515, 2007.
- [6] R. J. Beare, M. K. Macvean, A. A. M. Holtslag, J. Cuxart, I. Esau, J-C. Golaz, M. A. Jimenez, M. Khairoutdinov, B. Kosovic, D. Lewellen, T. S. Lund, J. K. Lundquist, A. McCabe, A. F. Moene, Y. Noh, S. Raasch, and P. Sullivan. An intercomparison of large-eddy simulations of the stable boundary layer. *Boundary-Layer Meteorol.*, 118(2):1573–1472, February 2006.
- [7] A. Beljaars. The parametrization of the planetary boundary layer. ECMWF Course Notes, May 1992.

- [8] A. C. M. Beljaars, J. L. Walmsley, and P. A. Taylor. A mixed spectral finite-difference model for neutrally stratified boundary-layer flow over roughness changes and topography. *Boundary-Layer Meteorol.*, 38:273–303, 1987.
- [9] A. K. Blackadar. Boundary layer wind maxima and their significance for the growth of nocturnal inversion. *Bull. Am. Met. Soc.*, 38:283–290, 1957.
- [10] R. A. Brost and J. C. Wyngaard. A model study of the stably stratified planetary boundary layer. *J. Atmos. Sci.*, 35(8):1427 – 1440, August 1978.
- [11] A. R. Brown, R. J. Beare, J. M. Edwards, A. P. Lock, S. J. Keogh, S. F. Milton, and D. N. Walters. Upgrades to the boundary-layer scheme in the Met Office numerical weather prediction model. *Boundary-Layer Meteorol.*, 128(1):117–132, July 2008.
- [12] J. G. Charney and N. A. Phillips. Numerical integration of the quasi-geostrophic equations for barotropic and simple baroclinic flows. *J. Meteor.*, 10:71–99, 1953.
- [13] M. J. P. Cullen, T. Davies, M. H. Mawson, J. A. James, S. C. Coulter, and A. Malcolm. *Numerical Methods in Atmosphere and Ocean Modelling*, volume The André J. Robert Memorial Volume, chapter An Overview of Numerical Methods for the Next Generation UK NWP and Climate Model, pages 425–444. Canadian Meteorological and Oceanographical Society, 1997.
- [14] M. J. P. Cullen and J. A. James. A comparison of two different vertical grid staggerings. In *Proc. 10th Conference on Numerical Weather Prediction*, pages 38–40. Amer. Meteor. Soc., 1994.
- [15] M. J. P. Cullen and D. J. Salmond. On the use of a predictor-corrector scheme to couple the dynamics with the physical parameterizations in the ECMWF model. *Quart. J. Roy. Meteorol. Soc.*, 129:1217–1236, 2003.
- [16] T. Davies, A. Staniforth, N. Wood, and J. Thuburn. Validity of anelastic and other equation sets as inferred from normal-mode analysis. *Quart. J. Roy. Meteorol. Soc.*, 129(593):2761–2775, December 2003.

- [17] F. Delsol, K. Miyakoda, and R. H. Clarke. Parameterized processes in the surface boundary layer of an atmospheric circulation model. *Quart. J. Roy. Meteorol. Soc.*, 97:181–208, 1971.
- [18] S. H. Derbyshire. Nieuwstadt’s stable boundary layer revisited. *Quart. J. Roy. Meteorol. Soc.*, 116:127–158, 1990.
- [19] P. G. Drazin and W. H. Reid. *Hydrodynamic Stability*. Cambridge University Press, 1984.
- [20] M. Dubal, N. Wood, and A. Staniforth. Analysis of parallel versus sequential splittings for time-stepping physical parameterizations. *Mon. Wea. Rev.*, 132(1):121–132, January 2004.
- [21] M. Dubal, N. Wood, and A. Staniforth. Mixed parallel-sequential-split schemes for time-stepping multiple physical parameterizations. *Mon. Wea. Rev.*, 133(4):989–1002, April 2005.
- [22] M. Dubal, N. Wood, and A. Staniforth. Some numerical properties of approaches to physics-dynamics coupling for NWP. *Quart. J. Roy. Meteorol. Soc.*, 132:27–42, 2006.
- [23] D. R. Durran. *Numerical Methods for Wave Equations in Geophysical Fluid Dynamics*. Springer, December 1998.
- [24] M. Ehrendorfer. The total energy norm in a quasigeostrophic model. *J. Atmos. Sci.*, 57(20):3443–3451, October 2000.
- [25] B. F. Farrell. Small error dynamics and the predictability of atmospheric flows. *J. Atmos. Sci.*, 47:1193–1206, 1990.
- [26] M. S. Fox-Rabinovitz. Computational dispersion properties of horizontal staggered grids for atmospheric and ocean models. *Mon. Wea. Rev.*, 119(7):1624–1639, July 1991.

- [27] M. S. Fox-Rabinovitz. Computational dispersion properties of vertically staggered grids for atmospheric models. *Mon. Wea. Rev.*, 122(2):377–392, February 1994.
- [28] J. R. Garrett. *The Atmospheric Boundary Layer*. Cambridge University Press, 1992.
- [29] C. Girard and Y. Delage. Stable schemes for nonlinear vertical diffusion in atmospheric circulation models. *Mon. Wea. Rev.*, 118(3):737–745, March 1990.
- [30] D. Handorf, T. Foken, and C. Kottmeier. The stable atmospheric boundary layer over an Antarctic ice sheet. *Boundary-Layer Meteorol.*, 91:165–189, 1999.
- [31] I. Held, R. T. Pierrehumbert, and R. Lee Panetta. Dissipative destabilization of external rossby waves. *J. Atmos. Sci.*, 43(4):388–396, 1986.
- [32] D. Holdaway, J. Thuburn, and N. Wood. On the relation between order of accuracy, convergence rate and spectral slope for linear numerical methods applied to multiscale problems. *Int. J. Numer. Meth. Fluids.*, 56(8):1297–1303, December 2008.
- [33] A. Hollingsworth. A spurious mode in the “Lorenz” arrangement of ϕ and T which does not exist in the “Charney-Phillips” arrangement. Tech. Memo. 211, ECMWF, Reading, UK, 1995.
- [34] J. R. Holton. *An Introduction to Dynamic Meteorology*, volume 88 of *International Geophysics Series*. Elsevier Academic Press, fourth edition, 2004.
- [35] Y-J. G. Hsu and A. Arakawa. Numerical modeling of the atmosphere with an isentropic vertical coordinate. *Mon. Wea. Rev.*, 118(10):1933–1959, October 1990.
- [36] I. N. James. The Antarctic drainage flow: Implications for hemispheric flow on the southern hemisphere. *Antarctic Sci.*, 1(3):279–290, 1989.
- [37] E. Kalnay. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 2003.

- [38] E. Kalnay and M. Kanamitsu. Time schemes for strongly nonlinear damping equations. *Mon. Wea. Rev.*, 116(10):1945–1958, 1988.
- [39] E. Kalnay-De-Rivas. On the use of nonuniform grids in finite-difference equations. *J. Comput. Phys.*, 10:202–210, 1971.
- [40] J. Kent. Folding timescales for atmospheric Lagrangian surfaces. MSc. dissertation., University of Exeter, College of Engineering, Mathematics and Physical Sciences, EX4 4QF, UK, September 2006.
- [41] J. C. King, W. M. Conolley, and S. H. Derbyshire. Sensitivity of modelled Antarctic climate to surface and boundary-layer flux parameterizations. *Quart. J. Roy. Meteorol. Soc.*, 127:779–794, 2001.
- [42] J. C. King, T. A. Lachlan-Cope, R. S. Ladkin, and A. Weiss. Airborne measurements in the stable boundary layer over the larsen ice shelf, Antarctica. *Boundary-Layer Meteorol.*, 127(3):413–428, June 2008.
- [43] Z. Kuang. The norm dependence of singular vectors. *J. Atmos. Sci.*, 61(23):2943–2949, December 2004.
- [44] R. Laprise. The Euler equations of motion with hydrostatic pressure as an independent variable. *Mon. Wea. Rev.*, 120(1):197–207, May 1991.
- [45] L. M. Leslie and R. J. Purser. A comparative study of the performance of various vertical discretization schemes. *Meteorol. Atmos. Phys.*, 50:61–73, 1992.
- [46] S-J. Lin. A “vertically Lagrangian” finite-volume dynamical core for global models. *Mon. Wea. Rev.*, 132(10):2293–2307, October 2004.
- [47] A. P. Lock, A. R. Brown, M. R. Bush, G. M. Martin, and R. N. B. Smith. A new boundary layer mixing scheme. part i: Scheme description and single-column model tests. *Mon. Wea. Rev.*, 128(9):3187–3199, September 2000.
- [48] E. N. Lorenz. Energy and numerical weather prediction. *Tellus*, 12:364–373, 1960.

- [49] E. N. Lorenz. A study of the predictability of a 28 variable atmospheric model. *Tellus*, 17:321–333, 1965.
- [50] G. Mastrantonio, R. Ocone, A. Pellegrini, and G. Fiocco. Sodar observations of the Antarctic boundary layer in a deglaciated area: Preliminary results. *Il Nuovo Cimento*, 12(3):589–597, 1990.
- [51] J. D. Mirocha and B. Kosovic. A large-eddy simulation study of the influence of subsidence on the stably stratified atmospheric boundary layer. *Boundary-Layer Meteorol.*, 134:1–21, 2010.
- [52] F. Molteni, R. Buizza, T. Palmer, and T. Petroliagis. The ECMWF ensemble prediction system: Methodology and validation. *Quart. J. Roy. Meteorol. Soc.*, 122:73–119, 1996.
- [53] F. Molteni and T. Palmer. Predictability and finite-time instability of the northern winter circulation. *Quart. J. Roy. Meteorol. Soc.*, 119:269–298, 1993.
- [54] P. Mote and A. O’neill. *Numerical Modeling of the Global Atmosphere in the Climate System. Proceedings of the NATO Advanced Study Institute, Castelvecchio Pascoli, Italy, May 25-June 5 1998.*, volume 550 of *NATO ASI Series C: Mathematical and Physical Sciences*. Dordrecht, Kluwer Academic Publishers, first edition, 2000.
- [55] F. T. M. Nieuwstadt. A model for the stationary, stable boundary layer. In J.C.R. Hunt, editor, *Turbulence and Diffusion in Stable Environments*, New Series Number 4, pages 149–179, Oxford University Press, Walton Street, Oxford, OX2 6DP, March 1983. The Institute of Mathematics and its Applications Conference Series, Oxford University Press.
- [56] T. R. Parish. On the role of Antarctic katabatic winds in forcing large-scale tropospheric motions. *J. Atmos. Sci.*, 49(15):1374–1385, 1992.
- [57] S. Rapaka, S. Chen, R. J. Pawar, P. H. Stauffer, and D. Zhang. Generalized stability theory. part i: Autonomous operators. *J. Fluid Mech.*, 609:285–303, 2008.

- [58] E. K. Schneider. An inconsistency in vertical discretization in some atmospheric models. *Mon. Wea. Rev.*, 115(9):2166–2169, September 1987.
- [59] A. Staniforth, A. A. White, N. Wood, J. Thuburn, M. Zerroukat, E. Cordero, T. Davies, and M. Diamantakis. Unified Model documentation paper. Unified Model Documentation Paper 15, Met Office, Met Office, FitzRoy Road, Exeter, EX1 3PB, United Kingdom, September 2006.
- [60] A. Staniforth, N. Wood, and J. Côté. Analysis of the numerics of physics-dynamics coupling. *Quart. J. Roy. Meteorol. Soc.*, 128:2779–2799, 2002.
- [61] A. Staniforth, N. Wood, and J. Côté. A simple comparison of four physics-dynamics coupling schemes. *Mon. Wea. Rev.*, 130(12):3129–3135, December 2002.
- [62] R. B. Stull. *An Introduction to Boundary Layer Meteorology*. Kluwer Academic Publishers, 1988.
- [63] J. C. Tannehill, D. A. Anderson, and R. H. Pletcher. *Computational Fluid Mechanics and Heat Transfer*. Computational and Physical Processes in Mechanics and Thermal Sciences. Taylor and Francis, second edition edition, April 1997.
- [64] J. Thuburn. Vertical discretizations giving optimal representation of normal modes: Sensitivity to the form of the pressure-gradient term. *Quart. J. Roy. Meteorol. Soc.*, 132(621):2809–2825, October 2006.
- [65] J. Thuburn and A. Staniforth. Conservation and linear rossby-mode dispersion on the spherical C Grid. *Mon. Wea. Rev.*, 132(2):641–653, February 2004.
- [66] J. Thuburn, N. Wood, and A. Staniforth. Normal modes of deep atmospheres. i: Spherical geometry. *Quart. J. Roy. Meteorol. Soc.*, 128(584):1771–1792, December 2002.

- [67] J. Thuburn, N. Wood, and A. Staniforth. Normal modes of deep atmospheres. ii: f - F -plane geometry. *Quart. J. Roy. Meteorol. Soc.*, 128(584):1793–1806, December 2002.
- [68] J. Thuburn and T. Woollings. Vertical discretizations for compressible Euler equation atmospheric models giving optimal representation of normal modes. *J. Comput. Phys.*, 203(2):386–404, March 2005.
- [69] T. Tokioka. Some considerations on vertical differencing. *J. Meteor. Soc. Japan*, 56:98–111, April 1978.
- [70] M. D. Toy and D. A. Randall. Comment on the article “vertical discretizations for compressible Euler equation atmospheric models giving optimal representation of normal modes”. *J. Comput. Phys.*, 223(1):82–88, April 2007.
- [71] L. N. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, 1997.
- [72] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll. Hydrodynamic stability without eigenvalues. *Science*, 261:578–584, 1993.
- [73] A. Untch and M. Hortal. A finite-element scheme for the vertical discretization of the semi-Lagrangian version of the ECMWF forecast model. *Quart. J. Roy. Meteorol. Soc.*, 130:1505–1530, 2004.
- [74] M. R. van den Broeke, N. P. M. van Lipzig, and E. van Meijgaard. Momentum budget of the east Antarctic atmospheric boundary layer: Results of a regional climate model. *J. Atmos. Sci.*, 59:3117–3129, 2002.
- [75] R. Vichnevetsky. Wave propagation and reflection in irregular grids for hyperbolic equations. *App. Num. Math.*, 3:133–166, 1987.
- [76] N. P. Wedi. The numerical coupling of the physical parameterizations to the “dynamical” equations in a forecast model. ECMWF Tech Memo 274.
- [77] W. Weng and P. A. Taylor. On modelling the one-dimensional atmospheric boundary layer. *Boundary-Layer Meteorol.*, 107:371–400, 2003.

- [78] A. A. White. *A View of the Equations of Meteorological Dynamics and Various Approximations*. Large-Scale Atmosphere-Ocean Dynamics I: Analytical methods and numerical models. J. Nordbury, I. Roulstone, Eds., Cambridge University Press, 2002.
- [79] D. Williamson. Time-split versus process-split coupling of parameterizations and dynamical core. *Mon. Wea. Rev.*, 130(8):2024–2041, August 2002.
- [80] N. Wood, M. Diamantakis, and A. Staniforth. A monotonically-damping second-order-accurate unconditionally-stable numerical scheme for diffusion. *Quart. J. Roy. Meteorol. Soc.*, 133:1559–1573, 2007.
- [81] H. Zhu and R. K. Smith. Effects of vertical differencing in a minimal hurricane model. *Quart. J. Roy. Meteorol. Soc.*, 129:1051–1069, 2003.