

Stability of Coastal Jets

Linear Stability Calculations and Nonlinear Simulations

by

John Yawney

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Applied Mathematics

Waterloo, Ontario, Canada, 2016

© John Yawney 2016

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In this thesis, a new numerical ocean model, Tempest, has been developed for application to simple process studies of large-scale ocean dynamics. This model allows for hydrostatic, non-hydrostatic, quasi-hydrostatic, and quasi-geostrophic model options to be employed and is a rigid-lid, fully three-dimensional model that allows for two-dimensionally varying bottom topography using a terrain-following coordinate transformation. To assess the accuracy and validity of this model a number of preliminary test cases are considered. These consist of a complex linear advection test, various convection studies including those defined over one- and two-dimensionally varying bottom topography, and a series of ocean gyre tests.

Next, the stability characteristics of a barotropic and a surface intensified baroclinic coastal jet are analyzed. Barotropic jets are characterized by significant horizontal shear and can give rise to both barotropic and baroclinic modes. Furthermore, barotropic jets are often more greatly impacted by variations in the bottom topography compared to surface-intensified baroclinic jets. On the other hand, baroclinic jets have both strong horizontal and vertical shear and are representative of more commonly observed physical jets in the ocean. That being said, baroclinic jets are often more difficult to study numerically due in part to the variations in the isopycnals from the definition of the background jet structure.

To gain insights into the growth and structure of the instabilities that can arise from the perturbation of these jets both linear stability calculations and nonlinear simulations are performed. The linear stability calculations allow us to consider a wide range of parameters efficiently. The effects of prograde and retrograde topography as well as varying degrees of stratification are considered. Guided by the results from the linear stability calculations, a set of nonlinear simulations are chosen. Using both the hydrostatic and quasi-geostrophic model options, a comparison between the two sets of results are made and non-QG effects are observed. As well, the results from the linear stability calculations are validated.

Acknowledgments

I would like to thank my supervisor, Francis Poulin, for the many discussions and guidance. He has helped me gain an understanding of many concepts which were originally foreign to me and has offered many insights into the research I have conducted and results I have obtained. His patience and knowledge are truly inspiring.

I would also like to thank the members of my committee, Michael Waite, Marek Stastna, Fue-Sang Lien, and Susan Allen. I am extremely appreciative of them taking the time to review my work and for their thoughtful comments. As well, thank you to Kevin Lamb for always being available to talk and for offering great insights. The other members of the fluids group have always been helpful. A particular thank you to Sina Khani and Michael Dunphy for many inspiring conversations.

The applied mathematics administrators have helped me extensively throughout my studies and for that I would like to extend a huge thank you to them for all the work they do. In particular, Laura Frazee and Helen Warren have answered a multitude of questions, and Maureen Fraser has offered positive encouragement and support.

Outside the academic environment, I have a multitude of people to thank. There is no doubt that the person I owe the most thanks to is my wife, Elisenda Yawney. She has always been there to give me motivation and encouragement. She has been more patient, loving, and supportive than I could have ever asked for. I certainly owe my success to her. A big thanks as well to our daughter, Isabel, for being a little bundle of joy, even when that meant some sleepless nights.

I am very thankful for the support that I have received from my other family members and friends. In particular, my mother-in-law, Maria Mejia, and honorary mother-in-law, Becky Paraskevopoulos, have been extremely loving and supportive. As well, my Aunt and Uncle, Helen and George Lacko, have offered strong encouragement and words of wisdom. I owe a great deal to them. My closest friends, Ehssan Pourramezani and Mark Verschoor, have offered some great and grounding discussions.

I will close by saying a huge thank you to my mother and father, Fay Hobin and John Yawney, who have offered love and support throughout my life and studies. Father, I will always miss you and I wish you could be here for my graduation.

I dedicate this work to my wife Elisenda and our daughter Isabel.

*Your unyielding love and support fills my life with joy
and motivates me to always strive for success.*

Table of Contents

List of Tables	xii
List of Figures	xiii
1 Introduction	1
1.1 Studying the coastal oceans	3
1.2 Available ocean models	5
1.2.1 General Overview	6
1.2.2 Physical Implications	7
1.3 Thesis overview	8
2 Ocean Modelling Regimes	10
2.1 Equations of Motion	10
2.1.1 Conservation of Mass	11
2.1.2 Conservation of Linear Momentum	12
2.1.3 Rotating Reference Frame	13
2.1.4 Evolution of Tracers	14
2.1.5 Thermodynamic Equation	15
2.2 Boussinesq Model	16
2.2.1 Elliptic Problem for Pressure	17
2.3 Boundary Conditions	18

2.4	Non-Dimensionalizing the Model	20
2.4.1	Non-Dimensional Parameters	21
2.4.2	Geostrophic Scaling	21
2.5	Hydrostatic Model	24
2.5.1	Summary: Surface Pressure Solution Procedure	26
2.6	Quasi-Geostrophic Model	27
2.6.1	Leading Order Terms	27
2.6.2	Ageostrophic terms	28
3	Numerical Methods	31
3.1	C Grid Staggering	31
3.2	Terrain-Following Coordinates	33
3.3	Solution Procedure	36
3.3.1	Time-Stepping	38
3.3.2	Interpolation and Coriolis terms	41
3.3.3	Nonlinear Advection	44
3.3.4	Viscosity	50
3.3.5	Buoyancy and Deconvolution	53
3.3.6	Pressure Solve	54
3.3.7	A PETSc Aside	56
3.3.8	Density Equation	58
3.3.9	WENO	59
3.3.10	Zhang limiting	66
3.4	QG Model	68
3.5	Linear Stability Analysis	70
3.5.1	Linear Stability of the NH Model	70
3.6	Computational Efficiency	76
3.6.1	Linear Stability Analysis	76

3.6.2	MPI Splitting Efficiency	76
3.6.3	Nonlinear Model	78
3.7	Conclusions	81
4	Validation Test Results	83
4.1	WENO Nonlinear Advection Test	84
4.2	Schär Advection Test	85
4.3	Passive Tracer Test	90
4.4	Buoyancy Tests	91
4.5	Flow Down Shelf Test	97
4.6	Bowl Test	100
4.7	Wind-Driven Gyre Tests	104
4.7.1	Single Gyre Test Cases	105
4.7.2	Double Gyre Test Cases	108
4.8	Conclusions	112
5	Linear Stability Analysis	114
5.1	LSA Governing Equations	115
5.2	Barotropic Jet: PE LSA Calculations	116
5.2.1	Convergence Analyses	117
5.2.2	Reference State	119
5.2.3	The Effects of Stratification	124
5.2.4	The Effects of Topography	128
5.3	Barotropic Jet: QG LSA Calculations	132
5.3.1	Reference State	133
5.3.2	The Effects of Stratification	135
5.4	Baroclinic Jet: PE LSA Calculations	138
5.4.1	Convergence Analyses	139

5.4.2	Reference State	140
5.4.3	The Effects of Ro	143
5.4.4	The Effects of Stratification	145
5.4.5	The Effects of Topography	146
5.5	Baroclinic Jet: QG LSA Calculations	149
5.5.1	Reference State	149
5.5.2	The Effects of Ro	151
5.5.3	The Effects of Stratification	151
5.6	Conclusions	153
5.6.1	Barotropic Jet	153
5.6.2	Baroclinic Jet	155
6	Nonlinear Simulations	157
6.1	Diagnostics	157
6.1.1	Growth Rate Curves	157
6.1.2	Mechanical Energy Equation	158
6.2	Barotropic Jet	160
6.2.1	Hydrostatic Reference Simulation	161
6.2.2	Quasi-Geostrophic Reference Simulation	165
6.2.3	Hydrostatic Simulations: Effects of Stratification	167
6.2.4	Quasi-Geostrophic Simulations: Effects of Stratification	171
6.2.5	The Effects of Topography	173
6.3	Baroclinic Jet	178
6.3.1	A Hydrostatic/Non-Hydrostatic Comparison	178
6.3.2	Hydrostatic Reference Simulation	181
6.3.3	Quasi-Geostrophic Reference Simulation	185
6.3.4	Hydrostatic Simulations: Effects of Stratification	188
6.3.5	Quasi-Geostrophic Simulations: Effects of Stratification	190
6.3.6	The Effects of Topography	191
6.4	Conclusions	194

7	Conclusions and Future Work	197
7.1	Conclusions	197
7.2	Future Work	201
	APPENDICES	204
A	Maple Coeff procedure	205
B	Nonlinear Case Study	209
B.1	Defining the TestCase#.cpp File	209
B.2	Defining the Preference File	212
B.3	Running the Test	214
B.4	Analyzing the Results	215
B.4.1	Viewing Slices	215
B.4.2	Viewing Three-Dimensional Isosurface Plots	216
B.4.3	Energy and Growth Plots	217
C	LSA Configuration	218
	References	220

List of Tables

3.1	Comparing the parallel efficiency and memory usage of SuperLU_dist (SLUd) and MUMPs	76
4.1	Comparison of the L_∞ -norms for the Schär advection test using successively refined grids. For each resolution the top row represents the results at $t = 5000$ s and the bottom represents the results at $t = 10000$ s.	87
4.2	Summary of validation tests that were performed using the full NH or HY Tempest model.	113
5.1	Test of the dependence of convergence of solutions to SLEPc solver tolerances using fixed horizontal and vertical grid resolutions, $N_y = N_z = 512$	117
5.2	Test of the dependence of convergence of solutions to grid resolution and order of the finite difference approximations. For all calculations $N_y = N_z$ and the default SLEPc solver tolerance of 1e-8 was used.	118
5.3	Test of the dependence of convergence of solutions to horizontal grid resolution with a reduced vertical grid resolution of $N_z = 64$. The default SLEPc solver tolerance of 1e-8 and fourth order finite difference approximations were used.	118
5.4	Test of the dependence of convergence of solutions to varying the grid resolution in both the horizontal and vertical $N_y = N_z = N$ for the baroclinic jet.	140
5.5	Test of the dependence of convergence of solutions to varying the grid resolution in the vertical with fixed horizontal resolution $N_y = 256$ for the baroclinic jet.	140

List of Figures

3.1	Standard three-dimensional C grid staggering	32
3.2	Two-dimensional cross-section of a terrain-following grid over a Gaussian hill.	33
3.3	Cell vertex labels	36
3.4	Stability Regions for Standard RK Methods	40
3.5	Vertical biasing direction for a grid with a terrain-following transformation.	46
3.6	WENO interpolation applied in the x direction	62
3.7	Candidate stencils for the approximation of fluxes in cell I_i	63
3.8	Example MPI grid splitting in one direction and two directions.	77
3.9	Parallel efficiency test of the HY model where N represents the number of cores. The results are displayed for a fine ($256 \times 256 \times 128$) and a coarse ($128 \times 128 \times 64$) run.	79
3.10	Parallel efficiency test of the NH model where N represents the number of cores. The results are displayed for a coarse ($128 \times 128 \times 64$) run.	80
4.1	Nonlinear advection of a simple Gaussian profile using the one-dimensional Burgers equation $u_t + uu_x = 0$. The integration is performed for a total of 1 s and the results are displayed as Hovmöller diagrams.	84
4.2	Schär advection test using a grid of 150×30 and Courant number of 0.25. Subfigure (a) shows the snapshots of the density anomaly as predicted by the WENO5 and SSPRK3 schemes. Subfigure (b) shows the absolute errors for the variable topography case and subfigure (c) shows the absolute error for the flat bottom case.	88

4.3	Schar advection test using a grid of 600×120 and Courant number of 0.25. Subfigure (a) shows the snapshots of the density anomaly as predicted by the WENO5 and SSPRK3 schemes. Subfigure (b) shows the absolute errors for the variable topography case and subfigure (c) shows the absolute error for the flat bottom case.	88
4.4	A plot of the error difference between the integration with and without limiting as predicted on the coarse 150×30 grid. The differences are computed at a time $t = 10000$ s associated with the end of the integration.	89
4.5	Cross-section of the passive tracer test using the WENO5 advection schemes as predicted by (a) the HY model without Zhang limiting and (b) with Limiting. The plots are shown for the end of the integration and (c) shows the relative difference between the two approximations.	91
4.6	Simple buoyancy test. Random cooling at the surface first passes through the viscous boundary layer into the convective layer characterized by strong vertical motions (a). Upon passing through the convective layer plumes form with horizontal length scale on the order of l_{rot} (b).	94
4.7	Buoyancy test of [Jones and Marshall, 1993], showing the cross-section (a) and 3D isosurface plot (b) for the dense fluid after two days of integration at which point a series of plumes have formed as a result of the applied density flux.	95
4.8	Buoyancy test of [Jones and Marshall, 1993], showing the cross-section (a) and 3D isosurface plot (b) for the dense fluid after day 4 of integration at which point the plumes have coalesced and formed into conal structures consistent with the findings of [Maxworthy and Narimousa, 1994].	96
4.9	Buoyancy test of [Jones and Marshall, 1993], showing the surface vorticity at days (a) 0.5, (b) 2, and (c) 4 respectively.	97
4.10	Flow down a shelf problem of MITgcm [Adcroft et al., 2014] using a grid resolution of 320×60	99
4.11	Flow down a shelf problem of MITgcm [Adcroft et al., 2014] using a grid resolution of 640×120	100
4.12	Initial configuration for the bowl test case. Dense fluid is located on a shelf and a lock-exchange experiment is conducted.	101
4.13	Cross-section of the density profile at $y = L_y/3$ using Tempest for the Bowl test case, as described in [Winton et al., 1998].	102

4.14	Vorticity field as predicted in the bowl test case at the end of the integration. Subfigure (a) shows the vertically integrated vorticity while subfigure (b) presents the surface vorticity.	103
4.15	Propagating Kelvin-like waves along the Western domain boundary as observed in the density field in the bowl test case. The cross-sections are taken through the middle of the domain.	104
4.16	Single gyre, single vertical level test case.	105
4.17	500 day integration of the single gyre, single vertical level test case showing snapshots of the evolution of the vorticity from day 50 to day 500 displayed in 50 day increments.	106
4.18	500 day integration of the single gyre, 16 vertical level test case showing snapshots of the evolution of the vorticity from day 50 to day 500 displayed in 50 day increments.	107
4.19	Three-dimensional structure of the vorticity in the single gyre, 16 vertical levels test case after 350 days of integration. Subfigure (a) shows the barotropic structure, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity field.	107
4.20	Three-dimensional structure of the vorticity in the single gyre, 16 vertical levels test case after 500 days of integration. Subfigure (a) shows the barotropic structure, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity field.	108
4.21	500 day integration of the double gyre, single vertical level test case showing snapshots of the evolution of the vorticity from day 50 to day 500 displayed in 50 day increments.	109
4.22	Time mean of meridional velocity and vorticity at the midpoint of the zonal extent of the domain, integrated over the first 200 days.	110
4.23	500 day integration of the double gyre, sixteen vertical levels test case showing snapshots of the evolution of the vorticity from day 50 to day 500 displayed in 50 day increments.	110
4.24	Three-dimensional structure of the vorticity in the double gyre, 16 vertical levels test case after 300 days of integration. Subfigure (a) shows the barotropic structure, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity field.	111

4.25	Three-dimensional structure of the vorticity in the double gyre, 16 vertical levels test case after 500 days of integration. Subfigure (a) shows the barotropic structure, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity field.	111
5.1	Comparison of the growth rates of the most unstable mode for wavenumbers in the range $[0, 2]$. The results are shown for the coarse grid 10×10 matlab direct solver and the fine grid 256×64 SLEPc calculation.	120
5.2	Subplots (a)-(f) present the vertical structure of the u, v, w, ρ, p and vorticity fields for the s_0 mode with the largest growth rate for the reference barotropic jet case. In subplot (g) and (h), the horizontal structure of the across channel velocity v and the vorticity are shown.	120
5.3	Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the five most unstable modes in the reference barotropic case as predicted by the PE linear stability calculation.	121
5.4	Vertical (top row) and horizontal (bottom row) structures of the across channel velocity v for the first five modes of the reference barotropic jet case, as predicted by the PE linear stability calculations.	122
5.5	Plots of the growth rate curves for the sixteen most unstable modes in the reference barotropic case for a reduced range of k values, using the PE linear stability calculations. The sinuous and varicose barotropic and first baroclinic modes are highlighted.	123
5.6	The structure of the v_0 mode in the reference barotropic jet calculations. Subplots (a)-(f) present the vertical structure of the u, v, w, ρ, p and vorticity fields and subplots (g) and (h) present the horizontal structure of v and the vorticity.	124
5.7	Plots of the growth rates as a function of wavenumber, k , and the Burger number, Bu, for (a) the s_0 mode and (b) the s_1 mode of the barotropic jet calculations.	125
5.8	The plots of the largest growth rates for the five most unstable modes as determined from the PE linear stability calculations in the case of the barotropic jet.	125

5.9	Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the three unstable modes in the weakly stratified $Bu = 1$ case as predicted by the PE linear stability calculations for the case of the barotropic jet.	126
5.10	Vertical structure (top row) and horizontal structure (bottom row) of the v velocity for the k value associated with the maximum growth of that mode. The results are displayed for the first three modes for the case with $Bu = 1$ as predicted from the PE linear stability calculations.	127
5.11	Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the most unstable modes in the strongly stratified $Bu = 50$ case as predicted by the PE linear stability calculations for the case of the barotropic jet.	128
5.12	The vertical structures of the across channel velocity v for the first five modes of the strongly stratified barotropic jet case with $Bu = 50$, as predicted by the PE linear stability calculations.	128
5.13	Plots of the growth rates as a function of wavenumber, k , and topography parameters, a , for the s_0 mode. The figure on the left shows the full range of k values and the figure on the right shows the k value that gives rise to the largest growth rates.	130
5.14	Plots of the growth rates as a function of wavenumber, k , and topography parameters, a , for the s_1 mode. The figure on the left shows the full range of k values and the figure on the right shows the k value that gives rise to the largest growth rates.	130
5.15	Plots of the maximum growth rates (left) as a function of the topography parameters, a , for the s_2 , s_3 , and s_4 unstable mode of the barotropic jet case as predicted by the PE linear stability calculations. A summary of the maximum growth rates as a function of k and a are shown as well (right).	131
5.16	Plots of the vertical and horizontal structures for the across channel velocity v for the s_0 mode in the case of an intermediate $a = -0.5$ (a) retrograde topography. Subfigure (b) shows the same plots in the case of the intermediate $a = 0.5$ prograde topography.	132
5.17	Plots of the vertical and horizontal structures for the across channel velocity v for the s_1 mode in the case of an intermediate $a = -0.5$ retrograde topography (a). The subfigure (b) shows the same plots in the case of the intermediate $a = 0.5$ prograde topography.	132

5.18	Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the five most unstable modes in the reference barotropic as predicted from the QG linear stability calculations.	133
5.19	The horizontal structures of the across channel velocity v for the first five modes of the reference barotropic jet case as predicted by the QG linear stability calculations.	134
5.20	Plots of the growth rate curves for the eleven most unstable modes in the reference barotropic case for a reduced range of k values, using the QG linear stability calculations. The sinuous and varicose barotropic and first baroclinic modes are highlighted.	134
5.21	The plots of the largest growth rates for the five most unstable modes, as a function of Bu, determined from the QG linear stability calculations in the case of the barotropic jet.	135
5.22	Growth rate curves in the weakly stratified Bu = 1 case as predicted by the QG linear stability calculations for the case of the barotropic jet. The two modes represent the s_0 and v_0 modes.	136
5.23	Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the most unstable modes in the strongly stratified Bu = 50 case as predicted by the QG linear stability calculations for the case of the barotropic jet.	137
5.24	Plot of the maximum growth rates as a function of mode number as predicted by the PE and QG models, for the strongly stratified case of Bu = 50, when (a) Ro = 1 and (b) Ro = 0.01.	138
5.25	Comparison of growth rates using a 128×128 SLEPc computation versus the direct matlab solver on a very coarse 10×10 grid in the case of the reference baroclinic jet test.	141
5.26	Plots of the growth rate curves for the five most unstable modes in the reference baroclinic case as predicted by the PE linear stability calculations.	141
5.27	Subplots (a)-(f) present the vertical structure of the u, v, w, ρ, p and vorticity fields for the s_1 mode with the largest growth rate for the reference baroclinic jet calculation. In subplot (g) and (h), the horizontal structure of the across channel velocity v and the vorticity are shown.	142

5.28	Subplots (a)-(f) present the vertical structure of the u, v, w, ρ, p and vorticity fields for the s_2 mode with the largest growth rate for the reference baroclinic case as predicted by the PE linear stability calculations. In subplot (g) and (h), the horizontal structure of the across channel velocity v and the vorticity are shown.	143
5.29	Plots of the normalized growth rate curves for variable Ro as predicted by the PE linear stability calculations in the reference baroclinic case. Subfigure (a) and (b) show the results for the s_1 and s_2 modes.	144
5.30	Plots of the vertical and horizontal across channel velocity and vorticity fields for the s_1 mode with the largest growth rate in the case of $Ro = 0.1$ (top row) and $Ro = 0.01$ (bottom row) for the baroclinic jet case as predicted by the PE linear stability calculations.	145
5.31	Plots of the growth rates as a function of wavenumber and Bu for the s_1 mode of the baroclinic jet in the case of $Ro = 1$ as predicted by the PE linear stability calculations.	145
5.32	Plots of the vertical across channel velocity and vorticity fields, and the horizontal across channel velocity and vorticity fields in the baroclinic jet case as predicted by the PE linear stability calculations for the strongly stratified, $Bu = 50$, case.	146
5.33	Plots of the growth rates as a function of topographic height and wavenumber for the s_1 mode of the baroclinic jet case as predicted by the PE linear stability calculations. The figure on the left shows the results for a full range of wavenumbers and the figure on the right shows the cases with largest growth rates.	147
5.34	Plots of the vertical and horizontal structures for the across channel velocity for the s_1 mode in the case of (a) an extreme $a = -0.75$ retrograde topography. The subfigure (b) shows the same plots in the case of the extreme $a = 0.75$ prograde topography.	147
5.35	Plots of the growth rates as a function of topographic height and wavenumber for the s_2 mode of the baroclinic jet case as predicted by the PE linear stability calculations. The figure on the left shows the results for a full range of wavenumbers and the figure on the right shows the cases with largest growth rates.	148

5.36	Plots of the vertical and horizontal structures for the across channel velocity v for the s_2 mode in the case of an extreme $a = -0.75$ retrograde topography (a). The subfigure (b) shows the same plots in the case of the extreme $a = 0.75$ prograde topography.	149
5.37	Plots of the growth rate curves for the s_1 and s_2 modes in the reference baroclinic case as predicted by the QG linear stability calculations.	150
5.38	Plots of the vertical v (a) and vorticity (b) for the s_1 mode with the maximum growth rate as predicted by the QG case. Subplots (c) and (d) present the horizontal structures of v and vorticity respectively.	150
5.39	Plots of the growth rates as a function of wavenumber and Bu for the s_1 mode of the baroclinic jet in the case of $Ro = 1$ as predicted by the QG linear stability calculations.	152
5.40	Plots of the growth rates and wavenumbers associated with the s_1 modes with maximum growth rates for the baroclinic jet case with varying stratification as predicted by the QG calculations. Subfigure (a) shows the maximum growth rates as a function of Bu and subfigure (b) shows the associated wavenumber at each stratification. Superposed is the results from the PE calculations for comparison.	152
5.41	Plots of the vertical across channel velocity and vorticity fields and the horizontal across channel velocity and vorticity fields as predicted by the QG calculations for the strongly stratified case with $Bu = 50$	153
6.1	Plots of (a) the growth rates and (b) incremental changes in the energy and relative energy loss for the reference barotropic simulation with $Ro = 1$ and $Bu = 10$ as predicted by the Tempest HY model.	162
6.2	The structure of the perturbation vorticity field at a snapshot during the period of maximum growth. Subfigure (a) shows the vertical structure of the instability and subfigure (b) shows the horizontal structure of the instability.	163
6.3	Snapshots of the evolution of the vorticity field in the reference barotropic case as predicted by the Tempest HY model. The results are displayed in 5 day increments.	164
6.4	The three-dimensional structure of the vorticity field in the reference barotropic simulation as predicted by the Tempest HY model at day 10. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.	164

6.5	Plots of (a) the growth rates and (b) incremental changes in the energy and relative energy loss for the reference barotropic simulation with $Ro = 1$ and $Bu = 10$ as predicted by the Tempest QG model.	165
6.6	Snapshots of the evolution of the vorticity field in the reference barotropic case as predicted by the Tempest QG model. The results are displayed in 5 day increments.	166
6.7	The three-dimensional structure of the vorticity field in the reference barotropic simulation as predicted by the Tempest QG model at day 8. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.	167
6.8	Growth rate curves for the barotropic jet simulations as predicted by the Tempest HY model equations for the (a) weak $Bu = 1$ stratification case and the (b) strong $Bu = 50$ stratification case.	168
6.9	Incremental changes in the energy for the barotropic jet simulations as predicted by the Tempest HY model equations for the (a) weak $Bu = 1$ stratification case and the (b) strong $Bu = 50$ stratification case.	169
6.10	Horizontal cross-section of the density perturbation field near the surface (relative height 0.9) after 10 days of integration for the (a) weakly stratified $Bu = 1$, (b) reference $Bu = 10$, and (c) strongly stratified $Bu = 50$ cases in the barotropic jet simulations as predicted by the Tempest HY model.	170
6.11	Three-dimensional isosurface plots of the vorticity after 10 days of integration for the (a) weakly stratified $Bu = 1$ and (b) strongly stratified $Bu = 50$ cases in the barotropic jet simulations as predicted by the Tempest HY model.	170
6.12	Three-dimensional isosurface plots of the barotropic and baroclinic portions of the vorticity after 10 days of integration in the strongly stratified $Bu = 50$ cases in the barotropic jet simulations as predicted by the Tempest HY model.	171
6.13	Growth rate curves for the barotropic jet simulations as predicted by the Tempest QG model strongly stratified $Bu = 50$ case.	172
6.14	The three-dimensional structure of the vorticity field in the strongly stratified $Bu = 50$ barotropic simulation as predicted by the Tempest QG model at day 8. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.	173
6.15	The growth rate curves in the case of the barotropic jet simulation over (a) retrograde and (b) prograde topographic profiles as predicted by the Tempest HY model.	174

6.16	Plots of vertical and horizontal cross-sections of the vorticity after 10 days of integration for the barotropic jet simulation over (a) retrograde and (b) prograde topographic profiles spanning a maximum height of 50% of the domain height as predicted by the Tempest HY model	174
6.17	Plots of the incremental changes in the energy and relative energy loss in the case of the barotropic jet simulation over (a) retrograde and (b) prograde topographic profiles as predicted by the Tempest HY model.	175
6.18	The evolution of the vorticity field in the case of the barotropic jet simulation over retrograde (top row) and prograde (bottom row) topographic profiles respectively. The outputs are printed in 5 day increments.	176
6.19	The three-dimensional structure of the vorticity field in the retrograde barotropic simulation as predicted by the Tempest HY model at day 10. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.	176
6.20	The three-dimensional structure of the vorticity field in the prograde barotropic simulation as predicted by the Tempest HY model at day 10. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.	177
6.21	The three-dimensional structure of the vorticity field in the (a) retrograde and (b) prograde barotropic simulations as predicted by the Tempest HY model at the end of the integration.	177
6.22	Plots of the growth rate curves in the case of the reference baroclinic test case for a grid resolution of $128 \times 128 \times 64$ using the (a) HY and (b) NH models.	179
6.23	Plots of the incremental changes in energy and the relative energy loss in the case of the reference baroclinic test case for a grid resolution of $128 \times 128 \times 64$ using the (a) HY and (b) NH models.	180
6.24	A horizontal cross-section of the vorticity field at the final day of integration in the case of the baroclinic jet as predicted by the HY model (a) and the NH model (b)	180
6.25	Plots of (a) the growth rates and (b) incremental changes in the energy and relative energy loss for the reference baroclinic simulation with $Ro = 1$ and $Bu = 10$ as predicted by the Tempest HY model.	182

6.26	Cross-sections of the vorticity plot in the case of the reference baroclinic jet as predicted by the Tempest HY model at day 20 of integration which represents a time during the period of strongest exponential growth. Subfigure (a) shows the horizontal cross-section and subfigure (b) shows the vertical cross-section at a relative height of 0.9.	183
6.27	Horizontal and vertical structures of the vorticity for the (a) s_1 mode and (b) s_2 mode of the reference baroclinic jet case as predicted by the linear stability calculations.	184
6.28	Snapshots of the evolution of the vorticity in the case of the reference baroclinic jet as predicted by the Tempest HY model. The cross-sections are taken near the surface at a relative height of 0.9. The results are displayed in 10 day increments.	184
6.29	Three-dimensional structure of the vorticity field of the reference baroclinic case as predicted by the Tempest HY model. Subfigure (a) shows the structure at day 30 and subfigure (b) shows the structure at the end of the integration.	185
6.30	Plots of (a) the growth rates and (b) incremental changes in the energy and relative energy loss for the reference baroclinic simulation with $Ro = 1$ and $Bu = 10$ as predicted by the Tempest QG model.	186
6.31	Snapshots of the evolution of the vorticity in the case of the reference baroclinic jet as predicted by the Tempest QG model. The cross-sections are taken near the surface at a relative height of 0.9. The results are displayed in 10 day increments.	187
6.32	Three-dimensional structure of the vorticity field of the reference baroclinic case as predicted by the Tempest QG model. Subfigure (a) shows the structure at day 25 and subfigure (b) shows the structure at the end of the integration.	188
6.33	Plots of (a) the growth rates and (b) incremental changes in the energy for the strongly stratified $Bu = 50$ baroclinic simulation as predicted by the Tempest HY model.	189
6.34	Isoplots of the vorticity field in the strongly stratified $Bu = 50$ baroclinic jet case as predicted by the Tempest HY model. Subfigures (a), (b), and (c) show the snapshots at days 5, 15, and 25 respectively.	189
6.35	Growth rate for the strongly stratification $Bu = 50$ baroclinic simulation as predicted by the Tempest QG model.	190

6.36	Isoplots of the vorticity field in the strongly stratified $Bu = 50$ baroclinic jet case as predicted by the Tempest QG model. Subfigures (a), (b), and (c) show the snapshots at days 5, 15, and 25 respectively.	191
6.37	The growth rate curves in the case of the baroclinic jet simulation over (a) retrograde and (b) prograde topographic profiles as predicted by the Tempest HY model.	192
6.38	Plots of the vorticity perturbations at day 0.5 of integration for the retrograde (left), flat-bottom (middle), and prograde (right) topographies in the case of the baroclinic jet simulations as predicted by the Tempest HY model.	192
6.39	Plots of the incremental changes in the energy and relative energy loss in the case of the baroclinic jet simulation over (a) retrograde and (b) prograde topographic profiles as predicted by the Tempest HY model.	193
6.40	Snapshots of the evolution of the vorticity in the case of the retrograde (top row) and prograde (bottom row) topographic profiles. Plots are displayed in 12.5 day increments starting from day 25.	194
6.41	The three-dimensional structure of the vorticity field at the end of the integration in the case of (a) retrograde and (b) prograde topographic profiles for the baroclinic jet test case as predicted by the Tempest HY model.	195

Chapter 1

Introduction

The oceans are vast and the understanding of ocean dynamics has a wide-range of implications on our lives. In particular, understanding of ocean dynamics influences decisions relating to fisheries and shipping practices [WWF-Deutschland, 2014]. As well, ocean currents play a fundamental role in nutrient transport [WWF-Deutschland, 2014, Barth et al., 2000]. Understanding the fundamental processes in the oceans allow us to make informed predictions. Insights about these processes can be gained through observations or through isolated laboratory experiments. However, to study a range of parameters we can use numerical models to simulate the evolution of oceanic flows. The mathematical representation, simulation, and interpretation of the results gained from these studies are key components of ocean modelling.

From a mathematical standpoint, the study of fluid dynamics is described by partial differential equations. These complex mathematical equations are based on a series of physical laws and, at a particular instant in time, a set of state variables can be used to describe the dynamics. These state variables have both space and time dependence [McWilliams, 2006] and can describe flows with interesting and complex variability.

Applying the techniques of numerical analysis to the specific set of partial differential equations that govern fluid flow is what is usually referred to as computational fluid dynamics. This allows us to find approximate solutions to problems that do not lend themselves to analytical techniques. Numerical techniques can be used to study a range of dynamics in the oceans or atmosphere. These numerical solutions offer us a means to try and predict the evolution of a given state. In particular, we can determine what instabilities can arise from unstable flows.

At the basin scale, ocean dynamics are $O(1000)$ km and are dominated by mostly

two-dimensional flows. However, the near surface currents are driven by winds, while the region below the pycnocline is not and this can give rise to three-dimensional motions. In contrast, on smaller scales the dynamics are highly three-dimensional and on the order of several meters. The study of the intermediate regimes, that of mesoscale $O(10-100)$ km and submesoscale $O(1-10)$ km dynamics, is particularly interesting since they play a fundamental role in the energy budget of the oceans. These motions can remain coherent for long periods of time and bridge the scales between large, basin scale dynamics and smaller, turbulent dynamics affected by dissipation.

To study the dynamical features that evolve in these regimes in greater detail we perform what are referred to as process studies. These are idealized studies of fluid flow that allow us to simulate specific flows in particular, idealized settings. In order to describe these motions accurately, it is necessary to use numerical models with computationally efficient algorithms that allow us to resolve a wide range of length and time scales.

Throughout fluid mechanics a number of questions arise, such as *What processes do we want to study?* and *What simplifying assumptions can be made to study these processes?* Non-dimensionalization can be used to guide our numerical simulations and to determine appropriate models to describe the processes we are studying. Making simplifying assumptions limit the types of motions we can describe but allow us to use simpler models that offer a reduction in computational expense.

Non-dimensional analysis can be used to describe our flow in terms of a set of non-dimensional parameters. In the inviscid NH model, these non-dimensional parameters are the Rossby number (Ro), the Burger number (Bu), and the aspect ratio (δ). These non-dimensional parameters are defined as:

$$\text{Ro} = \frac{U}{f_0 L}, \quad \text{Bu} = \frac{NH}{f_0 L}, \quad \delta = \frac{H}{L},$$

where U is our horizontal velocity scale, f_0 is the Coriolis frequency, L is our horizontal length scale, H is our vertical length scale, and N is the Brunt-Väisälä frequency. Some authors define Bu as the square of $NH/f_0 L$ however we use the convention defined above. We see that Ro is a measure of the relative magnitude of the nonlinear terms to the Coriolis terms and for a small Ro we conclude that nonlinear terms are relatively unimportant [Vallis, 2006]. Bu is a measure of the relative importance of stratification to rotation and for large Bu we conclude that the flow is dominated by stratification. The aspect ratio measures the vertical to horizontal length scales and for small aspect ratio the domain is much larger in the horizontal than the vertical. This is typically the case for large-scale oceanography.

We will be considering four dynamical models. The Boussinesq approximation is made throughout and, without further simplifying assumptions, we have the NH model. The Boussinesq approximation is ubiquitous to oceanic fluid flow and is based on the fact that the density perturbations are near negligible relative to the background density. If we consider flow for which the aspect ratio is small, and therefore the advective term in the vertical momentum equation is small [Vallis, 2006], then we can make the hydrostatic approximation which replaces the vertical momentum equation with a statement of exact hydrostatic balance. This leads to our second model called the hydrostatic (HY) model [Miller, 2007], which will be explained in section 2.5. If we include non-traditional Coriolis terms we have the quasi-hydrostatic (QH) model. If we consider the HY model, Ro is sufficiently small [Miller, 2007], and Bu is order one we can apply the quasi-geostrophic (QG) model [Charney, 1948, Charney and Stern, 1962] as explained in section 2.6.

In modelling oceanic flows through the use of computational fluid dynamics, we are often faced with technical questions such as *How computationally efficient can our code be?* and *What order of accuracy should we use?* However, with the recent advent of large, affordable computing clusters, such as those available through SciNet and SHARCNET, we are able to compute solutions on grids with increasingly finer resolutions. That being said, we still need to consider how much is being gained for extra computational costs of higher order methods.

1.1 Studying the coastal oceans

There are always interactions between the oceans and the atmosphere and these can strongly influence the global circulation in each. The oceans are forced on synoptic (planetary) length scales which are $O(1000)$ km. Two critical forces that drive the oceans are the winds from the atmosphere and the differences in solar heating across the surface of the globe. The wind-driven motion generates large circulation patterns, which are very energetic features and produce many boundary currents. In particular, large ocean gyres develop strong Western boundary currents, which owe their existence to the meridional change of the Coriolis parameter and dissipation. Some of the energy stored in these boundary currents cascades down to small length scales through nonlinear interactions and is dissipated at the very small scales.

At large scales, the dynamics of the oceans are largely governed by horizontal motions, which are nearly in geostrophic and hydrostatic balance and the Coriolis force dominates over the nonlinear advection. Therefore, we conclude that Ro is small and for $Ro \ll 1$ we

get near geostrophic balance in the horizontal. This is a balance between the horizontal pressure gradient terms and the Coriolis terms [Vallis, 2006]:

$$f\mathbf{k} \times \mathbf{u} = -\frac{1}{\rho_0} \nabla_{\text{h}} p,$$

under the traditional approximation. In the case of small aspect ratio, and therefore weak vertical advective terms, we can replace the vertical momentum equation with a hydrostatic balance relation [Vallis, 2006]:

$$\frac{\partial p}{\partial z} = -g\rho,$$

which assumes gravity is acting in the vertical.

When studying large scale motions we are often describing basin-wide circulation patterns. The horizontal aspects of these types of features can often be described using shallow water or QG models given the physical behavior of these basin-scale dynamics. When the vertical structure is critical to the dynamics it is perhaps necessary to employ the HY model [Marshall et al., 1997b], which can accurately describe some motions from the planetary scale down to the submesoscale for which the aspect ratio is small. However, it is possible that the submesoscale can contain NH effects.

In contrast, for motions on which the vertical length scale is comparable to the horizontal length scale we can potentially have turbulent three-dimensional motions which we often simulate using a fully NH model. The intermediate, mesoscale and submesoscale, length scales are very important for the oceanic energy budget, playing a fundamental role in the cascade of energy from planetary to dissipative scales. Depending on the size of the aspect ratio and the relative magnitude of the vertical velocity we can study these unbalanced motions with either an NH or HY model.

Furthermore, in submesoscale phenomena there is relatively large vertical velocities, such as during the upwelling of cold, deep, nutrient-rich ocean water in coastal areas [Ahrens, 2008] suggesting that NH effects could play a significant role in their dynamics. In these applications, NH models would be inherently more accurate. However, [Mahadevan, 2006] has shown that for much of the mesoscale and submesoscale that the dynamics appear to be hydrostatic in nature and can be accurately predicted by a HY model. However, the author does indicate that finer resolution could potentially lead to greater differences between the results predicted by NH and HY models.

We want to study the dynamics of mesoscale coastal jets that allow for the generation of submesoscale motions. Even though this is one particular example of mesoscale, and

potentially submesoscale, dynamics we hope that this may shed some light on the generation of these types of motions in other contexts. In particular, we consider barotropic and baroclinic jet studies. Barotropic jets are much simpler and are therefore a better starting point for scientific investigations. These jets can sometimes be idealized with two-dimensional shallow water and QG models. Surface-trapped baroclinic jets are often representative of coastal western boundary currents such as the *Gulf Stream*. To determine the types of instabilities that arise we can first compute the stability characteristics using a linear stability analysis. Subsequently, nonlinear simulations can be used to show the evolution of these instabilities throughout the linear and nonlinear regimes.

The benefit of combining these two approaches is that the linear theory offers a computationally efficient way to study a wide range of parameter space. In particular, we study the effects of stratification by varying Bu and also the effects of varying the topographic height of an idealized coastal shelf configuration. We solve the linear stability problem for both the barotropic and baroclinic jets, using PE and QG calculations. Then, based on the insight gained, we investigate the nonlinear simulations for a few select parameters that are of interest. The nonlinear simulations require significantly more computation time but offer more insights into the nonlinear evolution of the instabilities. In particular, it offers us a means of observing the development of mesoscale vortices and possibly submesoscale motions. The energy cascade can also be quantified using a spectral analysis but that is not done here. Instead, we focus on the initial onset of the instability.

1.2 Available ocean models

There are a significant number of ocean models that are freely available and the list is ever expanding. Some include the unstructured-grid Finite Volume Community Ocean Model (FVCOM), the Stanford Unstructured-grid Non-hydrostatic Parallel Coastal Ocean Model (SUNTANS), the MIT General Circulation Model (MITgcm), the Regional Ocean Modelling System (ROMS), the Coastal and Regional Ocean model (CROCO; built upon ROMS), the Second-generation Louvain-la-Neuve Ice-ocean Model (SLIM), and the Modular Ocean Model (MOM). These models can handle irregular geometry, parameterizations of turbulence, and a wide range of other options. This list is certainly not exhaustive. In this section we review some basic features of these models.

Although the above models can be used for process studies, in order to capture mesoscale and submesoscale dynamics in greater detail it is valuable to create a model based upon higher order methods. Particularly, higher order advection schemes allow for these motions to remain coherent for longer integrations. Other models exist that are based upon high

order methods, such as Spectral Parallel Incompressible Navier-Stokes (SPINS) solver [Subich et al., 2013]. However, this model is based on higher order spectral methods and is configured currently for only one-dimensional topography. Incorporation of two-dimensional varying bottom topography is another goal in the development of our new model, Tempest.

1.2.1 General Overview

There are a variety of models that are freely available and each has their advantages and disadvantages. Most models are based on second and third order spatial discretization and time-stepping, which tend to be more diffusive than higher order techniques. Instead of using one of these models we have worked to develop a new ocean model for idealized studies that uses high order schemes in order to resolve a wide range of grid scales. Using higher order finite difference approximations allow us to attain near spectral accuracy [Boyd, 2000] with the added benefit of computational efficiency. Using finite volume based methods allow accurate conservation properties and relatively simple formulations for cases with variable topography. Also, having the option to switch between NH, HY, and QG models can make our investigations more efficient.

The dynamical processes that we aim to investigate tend to have large Reynolds numbers and therefore have relatively little dissipation compared to flows at smaller length scales. We are very much concerned with the order of accuracy of our advective operator since that determines in large part the nonlinear iterations. Having high order of accuracy in advection implies that the amount of numerical dissipation and dispersion will be reduced and so we would expect that physical features of the flow will remain coherent for longer periods of time. In addition to a high order of accuracy, we also want to implement methods that are not prone to numerical dispersion. Therefore, the results will be less likely to give rise to numerical instabilities.

One of the most commonly used ocean models is the finite volume MITgcm model [Marshall et al., 1997b], which has a variety of advection schemes, including third and fourth order accurate methods. The issue with fourth order centered schemes, however, are their inability to accurately resolve discontinuities producing overshoots and undershoots (oscillations). That being said, the model has third order flux limiting options to prevent these oscillations. As well, there are flux limiting options that preserve positivity but these are limited to second order accuracy. The model has HY and NH capabilities allowing for comparison between the two solution procedures. There are a number of external forcing parameterizations and a wide range of potential test cases that include documentation. A number of these test cases have been used as reference tests for the Tempest model, as covered in chapter 4.

ROMS [Shepetkin and McWilliams, 2005] has similar implementation options as the MITgcm; available is second and fourth order centered schemes and the third order upwind biased method. As well, there is an option for the piecewise parabolic method of Colella and Woodward [Colella and Woodward, 1984], which can be limited via MPDATA. This limiting procedure reduces the otherwise third order (fourth order in the limit $CFL \rightarrow 0$ [Sekora and Colella, 2009]) accurate PPM method to a second order method, and preserves monotonicity and extrema but not necessarily positivity. ROMS is implemented as a HY model solver using a terrain-following grid and includes a free-surface.

The HY model is also solved by MOM. This model employs up to sixth order accurate advection schemes for tracer advection and fourth order for the nonlinear terms in the momentum equation. It uses a second order time splitting approach. MOM also implements a number of test cases and allows for complex geometrical domains. Some of the test cases from MOM were used to validate the Tempest ocean model.

There are also a number of other lower-order finite difference and finite volume methods available such as SUNTANS and FVCOM, which both currently use second order accurate discretizations. There are arguably benefits to using these approaches since lower order accurate schemes generally means that you can use a finer discretization as well as a smaller viscosity coefficient. As well, lower order methods leads to much simpler mathematical derivations and significantly less computational cost. The question of whether or not this outweighs the theoretical low accuracy is a long debated topic, especially in the cases of nonsmooth solutions [Wang et al., 2013].

At the opposite extreme, spectral based ocean models are fairly rare probably due in part to their mathematical complexity and also their computational expense. Using extremely accurate global based approximation techniques, spectral methods can give the most accurate predictions but with the highest computational cost. One such ocean model that uses spectral techniques is the Spectral Element Ocean Model (SEOM) [Iskandarani et al., 1995]. Given the global nature of the approximation operators the inversion matrix involved in the solution of the pressure equations is dense which, compared to sparse matrix inversions involved in finite difference and finite volume based ocean models, leads to computational expensive pressure solves. These dense matrices also pose limitations due to their memory constraints.

1.2.2 Physical Implications

The techniques of geophysical fluid dynamics [McWilliams, 2006] and computational fluid dynamics [Miller, 2007] offer us a means of describing the underlying physical processes

responsible for the observed phenomena in the oceans and atmosphere. GCMs are a powerful tool for the study of large-scale features. However, when studying smaller scale processes they tend to be very dissipative due, at least in part, to the low-order methods that they tend to use for discretizing the derivatives. Thus, to study the physical processes in the oceans in an idealized setting it would be beneficial to develop a model that uses higher-order accurate methods that can better resolve physical processes for a given numerical resolution. The higher-order accuracy will help to resolve the smaller length scales inherent in submesoscale dynamics. Furthermore, deriving a criteria to determine which model should be used can be difficult but by including the option to switch between the QG, HY, and NH models we can easily explore and determine when a simpler model is sufficient to answer a certain set of questions.

We do not explicitly account for sub-grid scale motions through parameterizations but instead focus on the nearly inviscid limit where possible. When dissipation is required we include viscosity, which can be seen as the simplest form of eddy parameterization. Eddy viscosity is used to represent the momentum transfer associated with nearly random motions of small fluid parcels [Vallis, 2006]. When viscosity is included we realize that we are unable to resolve the boundary layer and so generally make the free-slip assumption.

1.3 Thesis overview

This thesis will present a new numerical ocean model, Tempest, that has been developed to study simple process studies that arise in the coastal regions of the oceans. In particular, the model has been applied to barotropic and baroclinic jets and has been used to gain a better understanding of the formation of mesoscale vortices and submesoscale motions. Furthermore, the Tempest model has been designed with a set of standardized approximations including specifically the HY, NH, and QG approximations. Using the results from various simulations, we can test the validity of standard approximations to the primitive equations.

Chapter 2 describes an overview of the various ocean modelling regimes with a particular emphasis on theory. Chapter 3 describes the numerical approaches used in the development of the Tempest ocean model as well as those used in the linear stability analysis. It is also within this chapter that we briefly discuss computational efficiency of the code. Chapter 4 goes through a set of test cases, beginning with simple preliminary tests looking at specific features of the model. The chapter then continues with more classical test cases used for model validation. Next, Chapter 5 initiates the study of the barotropic and baroclinic jet cases. The result of the linear stability calculations using both PE and QG models are summarized within this chapter. Based on the predictions of the linear stability analyses,

Chapter 6 then considers specific nonlinear simulations that are of particular interest, using both the QG and HY models. The reason behind the decision to use the HY model over the NH model is also briefly described in the baroclinic jet case. Chapter 7 then touches on potential avenues for future research and development. It closes with a summary of the results from the previous chapters.

Chapter 2

Ocean Modelling Regimes

Many mathematical models consist of a system of governing equations that can be written in the language of differential equations. Many of these sets of differential equations cannot be solved analytically. Therefore, numerical techniques must be applied. In the case of fluid dynamics, the most general equations are a well-studied set of complex nonlinear partial differential equations. Due to their complexity, they can be quite expensive to solve numerically, however they can often be reduced to simpler models guided by physical assumptions.

In this chapter, we will introduce the general equations of motion with an elementary look at conservation laws. The Boussinesq approximation will next be described. Then the boundary conditions used in the solution of the equations will be discussed. Non-dimensionalization of the model will follow, which will emphasize in what parameter regime the models should be valid. Lastly, the HY and QG models will be introduced.

2.1 Equations of Motion

Under the continuum hypothesis, the infinitesimal fluid elements comprising fluid motion are assumed to form a continuum. In general, the fluid motion is described via a set of vector and scalar fields. These are the velocity \mathbf{u} , density ρ , pressure p , and other tracer quantities such as salinity and temperature. The fluid is assumed to be a continuous material that adheres to the fundamental conservation laws of physics. In particular, they respect the conservation laws of mass, momentum, and energy [Vallis, 2006].

The general form of a conservation law is [Durrant, 2010]

$$\frac{\partial q}{\partial t} + \nabla \cdot \mathbf{F}_q = S_q, \quad (2.1)$$

where q is a conserved state variable, \mathbf{F}_q represents fluxes of q , and S_q represents any sources or sinks of q . If we integrate this over a control volume V and apply Gauss' divergence theorem to the flux term then we can rewrite this in integral form as

$$\frac{\partial}{\partial t} \int_V q dV = \int_V S_q dV - \int_{\partial V} \mathbf{F}_q \cdot \hat{\mathbf{n}} dS, \quad (2.2)$$

where dV and dS represent differential volume and area elements. Both the differential (local) 2.1 and the integral (global) forms 2.2 are valuable in the study of geophysical and computational fluid dynamics.

When discussing the equations of motion, it is imperative to convey the importance of Reynolds' transport theorem:

$$\frac{D}{Dt} \int_V q(\mathbf{x}, t) dV = \int_V \frac{\partial q}{\partial t} dV + \int_{\partial V} q \mathbf{u} \cdot \hat{\mathbf{n}} dS,$$

where q represents a state variable conserved following the motion, dV and dS are differential volume and area elements. The boundary of our material volume V is denoted by ∂V . We use the notation D/Dt , following [Cohen and Kundu, 2008], to indicate that we are following a material volume. Therefore, this states that the rate of change of the sum of q throughout V is equal to the rate of change of q within V and the sum of flux of q through the boundary of V . If we instead use the divergence theorem to rewrite the surface integral on the right-hand side of Reynolds' transport theorem we get

$$\frac{D}{Dt} \int_V q(\mathbf{x}, t) dV = \int_V \left(\frac{\partial q}{\partial t} + \nabla \cdot (q \mathbf{u}) \right) dV.$$

2.1.1 Conservation of Mass

From the fundamental laws of physics we have that the mass of a material volume cannot be created or destroyed. If we label the material volume V and consider infinitesimal fluid elements of volume dV then the mass of these fluid elements is a product of their density and volume, $\rho(\mathbf{x}, t) dV$. If we integrate this quantity over the entire material volume, then we have via the fundamental physical law that

$$\frac{D}{Dt} \int_V \rho(\mathbf{x}, t) dV = 0,$$

where V is a material volume free to change in time. Applying Reynolds' transport theorem to the left-hand side we get the differential relationship for the conservation of mass often referred to as the (compressible) *continuity equation*:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0. \quad (2.3)$$

This represents a conservation law as given by the equation 2.1 for conserved state variable ρ and mass flux $\mathbf{F}_\rho = \rho \mathbf{u}$.

In the oceans, this is often replaced by the simplified incompressibility condition, $\nabla \cdot \mathbf{u} = 0$. This is an appropriate approximation since the density variations in the oceans are small, which is the underlying assumption in the Boussinesq approximation. This approximation is discussed in greater detail in section 2.2.

2.1.2 Conservation of Linear Momentum

The following derivation is a brief overview of the linear momentum equation and follows closely to the derivation found in [Vallis, 2006]. Newton's second law of motion states that the rate of change of the momentum within a control volume is equal to the sum of forces acting on this control volume. Therefore,

$$\frac{D}{Dt} \int_V \rho \mathbf{u} dV = \int_V \mathbf{F} dV.$$

Applying Reynolds' transport theorem and the divergence theorem to the left-hand side we get

$$\int_V \left(\frac{\partial}{\partial t} (\rho u_i) + \frac{\partial}{\partial x_j} (\rho u_i u_j) \right) dV = \int_V F_i dV,$$

where we have made use of indicial notation. Next, if we expand out the left hand side we obtain the following:

$$\int_V \left(\rho \frac{\partial u_i}{\partial t} + \rho u_j \frac{\partial u_i}{\partial x_j} + u_i \frac{\partial \rho}{\partial t} + u_i \frac{\partial}{\partial x_j} (\rho u_j) \right) dV = \int_V F_i dV.$$

Making use of the continuity equation 2.3, the third and fourth terms sum to zero. This leaves us with the equation

$$\int_V \left(\rho \frac{Du_i}{Dt} - F_i \right) dV = 0.$$

As this is true for any material volume, the integrand must evaluate to zero and we can express this in terms of a differential relationship in vector notation we have [Vallis, 2006]

$$\rho \frac{D\mathbf{u}}{Dt} = \mathbf{F}.$$

This states exactly how the acceleration of a fluid parcel is related to the forces acting on it. The forces \mathbf{F} can be a sum of body forces, surface forces, and tensile forces. In our case, we consider a local Cartesian coordinate system, such that the x points in the Eastward direction, y in the Northward direction, and z points upward. The body force is simply gravity $\Phi = -gz$. This force results from the fluid being placed in Earth's gravitational force field. The surface forces are due to direct contact with the surrounding medium [Cohen and Kundu, 2008]. This is comprised of two parts, the force exerted in the direction of the unit normal via pressure, p , and the viscous part due to the movement of the fluid, resulting in shear strain. This viscosity is parameterized by the viscosity coefficient ν which typically varies in the horizontal and vertical. Under the incompressibility assumption, the resulting equation is

$$\frac{D\mathbf{u}}{Dt} = -\frac{1}{\rho}\nabla p - g\mathbf{k} + \nu\nabla^2\mathbf{u}.$$

This is not an exact expression of the viscous stress, but it is generally used throughout most fluid dynamics applications [Vallis, 2006]. The viscous term is often added to other non-conservative forces and written as a force term as in [Marshall et al., 1997b]:

$$\frac{D\mathbf{u}}{Dt} = -\frac{1}{\rho}\nabla p - g\mathbf{k} + \mathbf{F}_v.$$

Writing the equations in this form also accounts for the possible differences in the horizontal and vertical viscosity coefficients.

2.1.3 Rotating Reference Frame

The large-scale motions of fluid dynamics are affected by the rotation of the Earth and so the governing equations must also be cast into a rotating reference frame. We consider a local Cartesian coordinate system as above and define θ as our latitude. The local rotation axis is then defined as $\cos\theta\mathbf{j} + \sin\theta\mathbf{k}$ where \mathbf{j} and \mathbf{k} are the unit vectors in the direction of the local y and z axes respectively. If we let the rotation rate of the Earth be defined as Ω , which has the value of $7.29212\text{e-}5$ rad/s, then the Earth's total rotation vector is

$$2\boldsymbol{\Omega}_e = 2\Omega\cos\theta\mathbf{j} + 2\Omega\sin\theta\mathbf{k}.$$

At the equator, this points Northward, while at the poles this points purely in the \mathbf{k} direction. We can write this as $2\boldsymbol{\Omega}_e = f_*\mathbf{j} + f\mathbf{k}$, where f and f_* represent the traditional and non-traditional Coriolis frequencies respectively. Due to the small aspect ratio of most oceanic fluid flows, only the vertical component of the rotation vector is often retained [McWilliams, 2006]. This is called the traditional approximation and under this approximation we have that $\boldsymbol{\Omega} = \Omega\mathbf{k} = (\boldsymbol{\Omega}_e \cdot \mathbf{k})\mathbf{k} = (f/2)\mathbf{k}$, where $\boldsymbol{\Omega}$ is the local vertical component of rotation. On the other hand, if we are near the equator and NH motions are important then the non-traditional Coriolis terms may be important. Since we are primarily focusing on mid-latitude jets and domains with small aspect ratios, the traditional approximation is assumed to be valid. That being said, an option for the non-traditional approximation is included in the Tempest model.

We can further simplify our $\boldsymbol{\Omega}$ by taking a Taylor expansion about latitude θ_0 . If we assume the deviations in latitude are relatively small then we can truncate this to give us approximations to f_0 and β_0 :

$$\Omega = |\boldsymbol{\Omega}| = \Omega \left(\sin \theta_0 + \cos \theta_0 (\theta - \theta_0) + \dots \right) = \frac{1}{2} \left(f_0 + \beta_0 (y - y_0) + \dots \right).$$

Therefore, we see that:

$$f_0 = 2\Omega \sin \theta_0, \quad \beta_0 = 2\Omega \cos \theta_0 \frac{d\theta}{dy} = \frac{2\Omega \cos \theta_0}{a},$$

where a is the radius of the Earth. β_0 is equal to df/dy and captures the linear effects of variations in the latitude. This then leads to the f and β plane approximations for which we approximate the traditional Coriolis frequency f by f_0 and $f_0 + \beta_0 y$ respectively.

Transforming to a rotating reference frame the momentum equation under the traditional approximation is

$$\frac{D\mathbf{u}}{Dt} + f\mathbf{k} \times \mathbf{u} = -\frac{1}{\rho}\nabla p - g\mathbf{k} + \mathbf{F}_u,$$

where the incremental centrifugal force potential is conventionally added to the geopotential (gravitational) term [McWilliams, 2006]. If, on the other hand, we wish to retain non-traditional, f_* , terms than we can write this as

$$\frac{D\mathbf{u}}{Dt} + 2\boldsymbol{\Omega}_e \times \mathbf{u} = -\frac{1}{\rho}\nabla p - g\mathbf{k} + \mathbf{F}_u. \quad (2.4)$$

2.1.4 Evolution of Tracers

The evolution of a passive tracer concentration τ , which is a mass fraction with common units being mg/L when discussing chemical tracers or units/L when discussing biological

tracers, is written as [McWilliams, 2006]

$$\frac{\partial(\rho\tau)}{\partial t} + \nabla \cdot (\rho\tau\mathbf{u}) = \rho S_\tau,$$

where S_τ , units of $[\tau]\text{s}^{-1}$, represents non-conservative sources and sinks of τ . For example, we can have flow of freshwater into an oceanic region creating a decrease in the concentration of salinity. This can be written more compactly in the following form, using the definition of the material derivative and equation 2.3, as

$$\frac{D\tau}{Dt} = S_\tau.$$

2.1.5 Thermodynamic Equation

As the above gives a set of four equations in five unknowns other equations are required to form a closed set. These equations are statements of thermodynamic laws, while the equation of state expresses a relationship between the thermodynamic variables. In the oceans, pressure, salinity, and temperature affect the density of seawater and the conventional equation of state has the form [Vallis, 2006]:

$$\rho = \rho(p, T, S).$$

This is often simplified to a linear function of temperature and salinity, where the linearization is made around some reference value (T_0, S_0) . This linearization represents the thermal expansion of the oceans due to changes in pressure and haline contraction due to changes in salinity [Vallis, 2006], and can be expressed mathematically as

$$\rho = \rho_0 \left[1 - \alpha(T - T_0) + \beta(S - S_0) \right].$$

Analogously, we see that:

$$\frac{D\rho}{Dt} = \rho_0 \left[1 - \alpha \frac{DT}{Dt} + \beta \frac{DS}{Dt} \right].$$

The salinity is governed by the evolution equation for a material tracer and in the absence of sources or sinks the salinity is conserved. An equation for the evolution of temperature can be found from the laws of thermodynamics and under the assumption that the heat flux obeys Fourier's law with constant thermal conductivity this simplifies to the classical heat equation [Cohen and Kundu, 2008]. This then translates into the thermodynamic law for density of:

$$\frac{D\rho}{Dt} = \kappa \nabla^2 \rho.$$

Assuming an adiabatic flow we have that $\kappa = 0$ and this gives us the simplified thermodynamic equation:

$$\frac{D\rho}{Dt} = 0. \quad (2.5)$$

We are not focusing on the dependence of density on salinity and temperature, but is something that could easily be included in later updates to the model.

2.2 Boussinesq Model

Throughout this work the Boussinesq approximation will always be employed. The mean density of seawater is 1025 kg/m^3 while the variations are on the order of 1 kg/m^3 [McWilliams, 2006]. To apply this approximation we decompose our density into a background constant reference density, the background hydrostatic vertically varying density with stratification, N , and the spatially and temporally varying perturbation density. Similarly, we decompose the pressure into a hydrostatically balanced background state and a three-dimensional varying perturbation pressure:

$$\begin{aligned} \rho(x, y, z, t) &= \rho_0 + \bar{\rho}(z) + \rho'(x, y, z, t), \\ p(x, y, z, t) &= \bar{p}(z) + p'(x, y, z, t). \end{aligned} \quad (2.6)$$

The relation between the background pressure and density is given by hydrostatic balance,

$$\frac{\partial \bar{p}}{\partial z} = -g(\rho_0 + \bar{\rho}).$$

The magnitude of the background density and pressure terms are often the largest in the equations and keeping them in the numerical model is a potential source of large numerical errors. Therefore, it is better to remove these two large terms and instead only retain the perturbation pressure and density in the pressure gradient and buoyancy terms. This is especially important when applying the terrain-following coordinate approximation since it minimizes the potential errors that arise from the cross-terms in the pressure gradient. This will be discussed further in section 3.2. Rewriting the pressure gradient and gravitational terms and multiplying through by ρ , the momentum equation 2.4 becomes [Cohen and Kundu, 2008]

$$\rho \left(\frac{D\mathbf{u}}{Dt} + 2\boldsymbol{\Omega}_e \times \mathbf{u} \right) = -\nabla p' - g\rho'\mathbf{k} + \rho\mathbf{F}_u.$$

We next apply the decomposition 2.6 and divide by the constant reference density ρ_0 . After cancellation of the hydrostatically balanced terms we obtain the following equation:

$$\left(1 + \frac{\bar{\rho} + \rho'}{\rho_0}\right) \left(\frac{D\mathbf{u}}{Dt} + 2\boldsymbol{\Omega}_e \times \mathbf{u}\right) = -\frac{1}{\rho_0} \nabla p' - g \frac{\rho'}{\rho_0} \mathbf{k} + \left(1 + \frac{\bar{\rho} + \rho'}{\rho_0}\right) \mathbf{F}_u.$$

If we next make the assumption that the hydrostatic density and density perturbations are negligible relative to the constant background density in all but the buoyancy term we get the following approximation:

$$\frac{D\mathbf{u}}{Dt} + 2\boldsymbol{\Omega}_e \times \mathbf{u} = -\frac{1}{\rho_0} \nabla p' - g \frac{\rho'}{\rho_0} \mathbf{k} + \mathbf{F}_u. \quad (2.7)$$

If we apply the decomposition 2.6 to the thermodynamic equation 2.5 then we can also write this in terms of the perturbation density as

$$\frac{D\rho'}{Dt} = \frac{\rho_0 N^2}{g} w,$$

where we have used the definition of the Brunt-Väisälä frequency:

$$N = \sqrt{-\frac{g}{\rho_0} \frac{d\bar{\rho}}{dz}}.$$

This frequency describes the maximal frequency of linear gravity waves without rotation.

2.2.1 Elliptic Problem for Pressure

When solving for our state variables at a particular time in our numerical simulation, we require a means of computing the pressure field given particular velocity and density fields. To do this we take the divergence of the momentum equations 2.7 and apply the divergence-free condition. This gives us the elliptic pressure equation consistent with that of [Marshall et al., 1997a]:

$$\nabla^2 p' = -\rho_0 \nabla \cdot (\mathbf{u} \cdot \nabla \mathbf{u}) + g \nabla \cdot (\rho' \mathbf{k}) - 2 \rho_0 \nabla \cdot (\boldsymbol{\Omega}_e \times \mathbf{u}).$$

Without appropriate boundary conditions this equation remains ill-posed. Therefore, we must decide how to choose these boundary conditions on the pressure field in terms of the boundary conditions on the other fields.

In the context of the numerical model we can first update the velocities using equation 2.7 including all terms except the pressure gradient terms. These partially updated velocities can be called pseudo-updated velocities to indicate that we must still apply the pressure gradient terms to them in order to project these into a set of incompressible solutions [Chorin, 1968]. If we define \mathbf{u}^* to be our velocities, our momentum equations become

$$\frac{\partial \mathbf{u}^*}{\partial t} = -\frac{1}{\rho_0} \nabla p', \quad (2.8)$$

where \mathbf{u}^* has the definition:

$$\mathbf{u}^* = \mathbf{u}^n - 2 \Delta t \boldsymbol{\Omega}_e \times \mathbf{u}^n - g \Delta t \frac{\rho'^n}{\rho_0} \mathbf{k} + \Delta t \mathbf{F}_{\mathbf{u}}^n.$$

Written in this way this represents a forward Euler time step, however similar formulations can be used for multistage time-stepping schemes.

If we let \mathbf{u}^{n+1} represent our fully updated velocities, then we can discretize this over our sub-stage as

$$\frac{\mathbf{u}^{n+1} - \mathbf{u}^*}{\Delta t} = -\frac{1}{\rho_0} \nabla p'.$$

Taking the divergence of this, applying a volume integral, and making use of the divergence theorem gives us the following:

$$\int_{\partial V} \nabla p' \cdot \mathbf{n} dS = \frac{\rho_0}{\Delta t} \int_{\partial V} \mathbf{u}^* \cdot \mathbf{n} dS - \frac{\rho_0}{\Delta t} \int_{\partial V} \mathbf{u}^{n+1} \cdot \mathbf{n} dS.$$

However, the discrete divergence-free condition implies that the second term on the right hand side is zero and we are left with the three-dimensional elliptic problem:

$$\int_{\partial V} \nabla p' \cdot \mathbf{n} dS = \frac{\rho_0}{\Delta t} \int_{\partial V} \mathbf{u}^* \cdot \mathbf{n} dS. \quad (2.9)$$

The discrete surface integral times the gradient operator will then be inverted using appropriate boundary conditions to solve for p' . In this form and with the three-dimensional elliptic problem, the equations define the NH model.

2.3 Boundary Conditions

A necessary condition in making a system of partial differential equations that depend both on space and time well-posed is to impose both initial and boundary conditions. Boundary

conditions must be applied to all sides of the fluid domain. In the case of solid boundaries (generally lateral boundaries, rigid lid and topography below), there must be no flux of momentum across the boundary. This gives rise to the no-normal flow condition:

$$\mathbf{u} \cdot \mathbf{n}|_{\mathbf{x} \in S} = 0,$$

where S is the solid boundary.

If viscosity is included, then additional conditions must be imposed on the velocity. In this case we consider either no-slip or free-slip conditions. Under the free-slip condition, the fluid is free to turn and propagate in a direction parallel to the solid boundary. An alternative to free-slip boundary conditions are no-slip boundary conditions which imply that the velocity is zero at the boundary, $\mathbf{u}(\mathbf{x}) = \mathbf{0}$ for $\mathbf{x} \in S$. Given the transition to zero velocity, no-slip boundary conditions introduce boundary layers that have a thickness dependent on the viscosity coefficient used. In the case of small viscosity we use free-slip instead of no-slip because we are not able to resolve the boundary layers that develop. As well, our primary focus is on dynamics in the interior of the domain and so we presume that the choice of boundary conditions will not have a significant effect on the dynamics. That being said, no-slip boundary conditions have also been included in the Tempest model as a possible option.

In our case, we use a rigid lid formulation, which makes the assumption that the displacements at the surface of the ocean are negligible relative to the bathymetry [Vallis, 2006]. By imposing this boundary condition we effectively filter out the fast surface gravity waves. Including these waves would impose an extremely small time step. For tracers we have no flux boundary conditions, which imply that at solid boundaries we have no influx or outflux of material concentrations. However, in Tempest it is possible to impose an influx at a boundary as well if desired.

To determine the boundary conditions on the pressure we consider the vector form of the momentum equations 2.7 and adopt the notation

$$\mathbf{u}^* = \mathbf{u}^n + \Delta t \mathcal{G}_{\mathbf{u}},$$

where $\mathcal{G}_{\mathbf{u}}$ represents all terms except the pressure gradient terms, as in [Marshall et al., 1997b]:

$$\mathcal{G}_{\mathbf{u}} = -2\boldsymbol{\Omega}_e \times \mathbf{u}^n - g \frac{\rho'^n}{\rho_0} \mathbf{k} + \mathbf{F}_{\mathbf{u}}^n. \quad (2.10)$$

If we next consider equation 2.9 from above and apply the no normal flow condition to \mathbf{u}^n we get the following:

$$\int_{\partial V} \nabla p' \cdot \hat{\mathbf{n}} dS = \rho_0 \int_{\partial V} \mathcal{G}_{\mathbf{u}} \cdot \hat{\mathbf{n}} dS.$$

Therefore, we conclude that our boundary conditions on the pressure should be

$$\nabla p' \cdot \hat{\mathbf{n}} = \rho_0 \mathcal{G}_{\mathbf{u}} \cdot \hat{\mathbf{n}}. \quad (2.11)$$

This is the most general form of the boundary conditions because it accounts for variable normal vectors in all directions.

As a particular example, if we consider slip conditions then we have that the tangential velocity components have homogeneous Neumann boundary conditions. As well, the tracers have no flux boundary conditions. Therefore, if we consider the case for $x = 0, L_x$ then we have

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial x} = \frac{\partial w}{\partial x} = \frac{\partial \rho}{\partial x} = 0.$$

Applying all these boundary conditions as well as those implied by the no normal flow condition, $u = 0$, we have that

$$\frac{\partial p'}{\partial x} = \rho_0 (fv - f_* w).$$

Similarly, in the y direction we get the balance:

$$\frac{\partial p'}{\partial y} = -\rho_0 f u.$$

In the z direction we get hydrostatic balance and a quasi-hydrostatic balance in the case when non-traditional Coriolis terms are included in the calculation:

$$\frac{\partial p'}{\partial z} = \rho_0 f_* u - \rho' g.$$

2.4 Non-Dimensionalizing the Model

The above NH model has dimensional variables. There are a number of reasons why non-dimensional variables are valuable. In practice, one of the most important reasons is to avoid round-off errors that can be caused from large or small values [Danaila et al., 2006]. This is of particular interest when we are computing the solution to the elliptic pressure problem. However, our iterative solver checks the absolute and relative tolerance of our solution vector so it is not expected that this will be a significant concern. Another important reason to consider non-dimensionalization though is to determine the relative magnitude of the various terms in the equations of motion. This is important for determining when particular approximations are valid [Vallis, 2006].

2.4.1 Non-Dimensional Parameters

When discussing non-dimensionalization we define the problem in terms of various non-dimensional parameters. In our case, we express everything in terms of the Rossby number, Burger number, Reynolds number, and aspect ratio. All but the Reynolds number have been defined before however they are summarized here for reference:

$$\text{Ro} = \frac{U}{f_0 L}, \quad \text{Bu} = \frac{NH}{f_0 L}, \quad \text{Re} = \frac{UL}{\nu}, \quad \delta = \frac{H}{L}.$$

2.4.2 Geostrophic Scaling

We define the horizontal length scale to be L , set by the horizontal extent of the motions. In the case when we consider a barotropic or baroclinic jet, this scale is set by the jet width. The horizontal velocity scale is U and is typically equal to the absolute value of the maximum horizontal velocity of the initial state. Our vertical length scale is H and represents the depth scale of our motions, which is taken to be the decay scale of our jet. W is our vertical velocity scale and is written in terms of a geostrophic scaling as [Vallis, 2006]

$$W \sim \text{Ro} \frac{UH}{L}.$$

The UH/L is introduced by incompressibility and the Ro factor is introduced by the nearly geostrophic assumption. In this case $\text{Ro} \ll 1$ and, to leading order, the geostrophic velocities are horizontally non-divergent and the vertical velocities are negligible. For the characteristic time scale, we use an advective timescale $T \sim L/U$. Our dynamic pressure and density scaling, P and R respectively, are set by geostrophic and hydrostatic balance relations as

$$P \sim \rho_0 f_0 U L, \quad R \sim \frac{P}{gH}.$$

Non-Dimensional Variables

Using these scaling factors, the non-dimensional variables can be defined as

$$\begin{aligned} \hat{u} &= \frac{u}{U}, & \hat{v} &= \frac{v}{U}, & \hat{w} &= \frac{w}{W}, \\ \hat{x} &= \frac{x}{L}, & \hat{y} &= \frac{y}{L}, & \hat{z} &= \frac{z}{H}, \\ \hat{t} &= \frac{t}{T}, & \hat{p}' &= \frac{p'}{P}, & \hat{\rho}' &= \frac{\rho'}{R}. \end{aligned}$$

We also rescale the Coriolis frequency by defining:

$$\hat{f} = \frac{f}{f_0}.$$

Non-Dimensional Momentum Equations

Applying these non-dimensional variables to the zonal, meridional, and vertical momentum equations we get the following non-dimensional equations after simplifications:

$$\begin{aligned} \frac{D\hat{u}}{D\hat{t}} &= \frac{1}{\text{Ro}} \cdot \hat{f}\hat{v} - \frac{1}{\text{Ro}} \cdot \frac{\partial \hat{p}'}{\partial \hat{x}} + \frac{1}{\text{Re}} \left(\frac{\partial^2 \hat{u}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{u}}{\partial \hat{y}^2} + \frac{1}{\delta^2} \frac{\partial^2 \hat{u}}{\partial \hat{z}^2} \right), \\ \frac{D\hat{v}}{D\hat{t}} &= -\frac{1}{\text{Ro}} \cdot \hat{f}\hat{u} - \frac{1}{\text{Ro}} \cdot \frac{\partial \hat{p}'}{\partial \hat{y}} + \frac{1}{\text{Re}} \left(\frac{\partial^2 \hat{v}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{v}}{\partial \hat{y}^2} + \frac{1}{\delta^2} \frac{\partial^2 \hat{v}}{\partial \hat{z}^2} \right), \\ \frac{D\hat{w}}{D\hat{t}} &= \frac{1}{\text{Ro}^2 \delta^2} \left(-\hat{\rho}' - \frac{\partial \hat{p}'}{\partial \hat{z}} \right) + \frac{1}{\text{Re}} \left(\frac{\partial^2 \hat{w}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{w}}{\partial \hat{y}^2} + \frac{1}{\delta^2} \frac{\partial^2 \hat{w}}{\partial \hat{z}^2} \right), \end{aligned} \quad (2.12)$$

where the material derivative is defined as

$$\frac{D}{D\hat{t}} = \frac{\partial}{\partial \hat{t}} + \hat{u} \frac{\partial}{\partial \hat{x}} + \hat{v} \frac{\partial}{\partial \hat{y}} + \text{Ro} \hat{w} \frac{\partial}{\partial \hat{z}}.$$

Non-Dimensional Density Equation

If we apply the same non-dimensionalization to the continuity equation we get the following:

$$\frac{D\hat{\rho}'}{D\hat{t}} = \text{Bu}^2 \hat{w}. \quad (2.13)$$

Therefore, non-dimensionalization does not change the continuity equation dramatically, but does alter the coefficient on the stratification term.

Non-Dimensional Elliptic Problem for Pressure

In the discussion of the elliptic problem for pressure we first define the following:

$$\tilde{\nabla} = \left(\frac{\partial}{\partial \hat{x}}, \frac{\partial}{\partial \hat{y}}, \frac{1}{\delta^2} \frac{\partial}{\partial \hat{z}} \right), \quad \hat{\mathbf{u}}^* = (\hat{u}^*, \hat{v}^*, \text{Ro} \hat{w}^*),$$

where we have used $\hat{\mathbf{u}}^*$ to define our pseudo-updated velocity components, including all terms except those involving the pressure. Using this notation the non-dimensional analogue of equation 2.8 takes the form:

$$\frac{\partial \hat{\mathbf{u}}^*}{\partial \hat{t}} = -\frac{1}{\text{Ro}} \tilde{\nabla} \hat{p}'.$$

Upon expanding the time derivative and solving in terms of the derivatives of pressure we get

$$\tilde{\nabla} p' = \frac{\text{Ro}}{\Delta t} \hat{\mathbf{u}}^* - \frac{\text{Ro}}{\Delta t} \hat{\mathbf{u}}^{n+1}. \quad (2.14)$$

Before continuing the derivation of the pressure solve it will be valuable to have an expression for the non-dimensional divergence-free condition. Applying the above scalings to the analytical incompressibility equation we obtain the following non-dimensional divergence-free condition:

$$\frac{\partial \hat{u}}{\partial \hat{x}} + \frac{\partial \hat{v}}{\partial \hat{y}} + \text{Ro} \frac{\partial \hat{w}}{\partial \hat{z}} = 0.$$

In the discrete framework, this leads to the discrete divergence-free condition:

$$\int_{\partial \hat{V}} \hat{\mathbf{u}} \cdot \hat{\mathbf{n}} d\hat{S} = 0, \quad (2.15)$$

where $\hat{\mathbf{u}} = (\hat{u}, \hat{v}, \text{Ro} \hat{w})$.

If we now take the divergence of the equations 2.14, take the volume integral, and rewrite this as a surface integral using the divergence theorem we get

$$\int_{\partial \hat{V}} \tilde{\nabla} p' \cdot \hat{\mathbf{n}} d\hat{S} = \frac{\text{Ro}}{\Delta t} \int_{\partial \hat{V}} \hat{\mathbf{u}}^* \cdot \hat{\mathbf{n}} d\hat{S} - \frac{\text{Ro}}{\Delta t} \int_{\partial \hat{V}} \hat{\mathbf{u}}^{n+1} \cdot \hat{\mathbf{n}} d\hat{S}.$$

Applying the non-dimensional discrete divergence-free condition 2.15 to the updated non-dimensional velocities $\hat{\mathbf{u}}^{n+1}$ we get

$$\int_{\partial \hat{V}} \tilde{\nabla} p' \cdot \hat{\mathbf{n}} d\hat{S} = \frac{\text{Ro}}{\Delta t} \int_{\partial \hat{V}} \hat{\mathbf{u}}^* \cdot \hat{\mathbf{n}} d\hat{S}. \quad (2.16)$$

In this form, we see that our momentum equations 2.12, our density equation 2.13, and our elliptic problem for pressure 2.16 can all be characterized by a set of non-dimensional values that describes the nature of the flow we are studying. In particular we see that if we multiply the zonal and meridional momentum equations of 2.12 by Ro and assume that $\text{Ro} \ll 1$ we get to leading order that the motion is in near geostrophic balance. Similarly, if we multiply the vertical momentum equation of 2.12 by $\text{Ro}^2 \delta^2$ we get hydrostatic balance as our leading order.

2.5 Hydrostatic Model

Typically, the hydrostatic approximation is applied when $\delta \ll 1$, which implies weak vertical acceleration. Therefore, the hydrostatic approximation will often be used when studying mesoscale motions that are characterized by little vertical motion [Marshall et al., 1997b]. On the other hand, the HY model will break down in coastal regions where there can be significant amounts of upwelling and downwelling leading to strong vertical motions. In this case, the aspect ratio of the motion can be order one.

As well, oceanic fronts and topographic variations have been shown to produce small filament-like features in the oceans which have characteristically high vertical velocities [Mahadevan, 2006]. However, the author found that, even though characterized by relatively large vertical velocities, these features were significantly hydrostatic in nature. They persisted and were explained quite accurately by a HY model for the particular resolution that they used. Density currents and convective motions give rise to a number of NH effects and we would expect that a HY model would be of limited value in these applications [Marshall et al., 1997b].

That being said, the hydrostatic approximation is often valid in many regimes which is why it is used in a number of oceanic models [Haidvogel and Beckmann, 1998]. Furthermore, computational resources classically posed a limitation on the grid resolution in NH models. In the HY approximation, we use the following decomposition for the dynamic pressure:

$$\begin{aligned} p'(x, y, z, t) &= p_{HY}(x, y, z, t) + p_S(x, y, t) \\ &= p_S(x, y, t) - \int_H^z g\rho'(x, y, z, t) dz', \end{aligned} \quad (2.17)$$

where p_{HY} is our hydrostatic pressure component and p_S is the surface pressure, which only depends on the horizontal variables and time. The second equality arises because of hydrostatic balance

$$\frac{\partial p'_{HY}}{\partial z} = -g\rho', \quad (2.18)$$

which is an approximation to the vertical momentum equation. This approximation is valid when the vertical acceleration is small relative to the gravitational acceleration. The bottom of the domain is defined as $z = h(x, y)$ and the top is defined as $z = H$, a constant. The hydrostatic pressure is found by integrating the hydrostatic balance relation 2.18 downward from the rigid lid.

In deriving the two-dimensional surface pressure equation, it simplifies notation to define

a vertical averaging operator, which, for some particular field ψ , is

$$\overline{\psi}^z = \frac{1}{H - h(x, y)} \int_{h(x, y)}^H \psi dz'. \quad (2.19)$$

As an extension of the Leibniz theorem, we can compute the vertical average of the horizontal gradient as follows (suppressing the spatial dependence of the bottom topography):

$$\begin{aligned} \frac{1}{H - h} \nabla_h \left[(H - h) \overline{\psi}^z \right] &= \frac{1}{H - h} \nabla_h \left(\int_h^H \psi dz \right) \\ &= \frac{1}{H - h} \int_h^H \nabla_h \psi dz - \frac{\psi(h) \nabla_h h}{H - h} \\ &= \overline{\nabla_h \psi}^z - \frac{\psi(h) \nabla_h h}{H - h} \\ \Rightarrow \overline{\nabla_h \psi}^z &= \frac{1}{H - h} \nabla_h \left[(H - h) \overline{\psi}^z \right] + \frac{\psi(h) \nabla_h h}{H - h}. \end{aligned} \quad (2.20)$$

Using incompressibility we know that our vertical velocity can be diagnosed from a vertical integral of the horizontal divergence of our horizontal velocity components. That is,

$$w = - \int_H^z \nabla_h \cdot \mathbf{u}_h dz', \quad (2.21)$$

under the no normal flow condition. In the HY model, the above relationship can be used to determine the vertical velocity from the horizontal velocity components. This is in contrast to the NH model where the vertical velocity is determined by integration of the vertical momentum equation.

Next, if we take the vertical average 2.19 of the horizontal velocities, multiply the result by the domain depth, $H - h$, and take the horizontal divergence we get the following:

$$\nabla_h \cdot \left[(H - h) \overline{\mathbf{u}_h}^z \right] = \int_h^H \nabla_h \cdot \mathbf{u}_h dz' - \mathbf{u}_h(h) \cdot \nabla_h(h).$$

This has made use of formula 2.20, multiplied by $H - h$. Evaluating equation 2.21 at $z = h$ we see that the first term is just $w(h)$. However, the second term is also $w(h)$ on account of the no normal flow condition and the definition of the normal vector. Therefore, we get that:

$$\nabla_h \cdot \left[(H - h) \overline{\mathbf{u}_h}^z \right] = w(h) - w(h) = 0. \quad (2.22)$$

Following a similar approach to the derivation of the three-dimensional elliptic problem 2.9 we retain the horizontal gradient of our dynamic pressure and the acceleration term of the pseudo-updated horizontal velocity components. This leads to the following equation in which we have decomposed the pressure using equation 2.17 and we have considered only the horizontal divergence of the horizontal momentum equations:

$$\nabla_{\text{h}}(p_S + p_{HY}) = -\frac{\rho_0}{\Delta t} \mathbf{u}_{\text{h}}^{n+1} + \frac{\rho_0}{\Delta t} \mathbf{u}_{\text{h}}^*.$$

If we now vertically integrate, separate the hydrostatic and surface pressure terms, and apply the horizontal divergence we get:

$$\nabla_{\text{h}} \cdot \left[(H - b) \overline{\nabla_{\text{h}} p_S}^z \right] = \nabla_{\text{h}} \cdot \left[\frac{\rho_0 (H - b)}{\Delta t} \overline{\mathbf{u}_{\text{h}}^{*z}} \right] - \nabla_{\text{h}} \cdot \left[(H - b) \overline{\nabla_{\text{h}} p_{HY}}^z \right],$$

where the $\overline{\mathbf{u}_{\text{h}}^{n+1z}}$ term is removed by application of 2.22. Therefore, this represents a projection method in which the surface pressure is determined such that the resulting horizontal velocities respect the vertically averaged horizontal divergence-free condition 2.22. Combining the horizontal divergence terms on the right hand side this yields the following:

$$\nabla_{\text{h}} \cdot \left[(H - b) \overline{\nabla_{\text{h}} p_S}^z \right] = \nabla_{\text{h}} \cdot \left[\frac{\rho_0 (H - b)}{\Delta t} \overline{\mathbf{u}_{\text{h}}^{*z}} - (H - b) \overline{\nabla_{\text{h}} p_{HY}}^z \right]. \quad (2.23)$$

2.5.1 Summary: Surface Pressure Solution Procedure

In order to solve for the surface pressure we can update the pseudo-updated velocities as we would have done in the NH case, respecting the fact that we do not need to update the vertical velocity. Then we can update these pseudo-updated horizontal velocity components by adding to them gradients of the hydrostatic pressure:

$$\mathbf{u}_{\text{h}}^{**} = \mathbf{u}_{\text{h}}^* - \frac{\Delta t}{\rho_0} \nabla_{\text{h}} p_{HY}.$$

This effectively defines the terms in the horizontal divergence of the right hand side of 2.23, with the constant factor of $\rho_0/\Delta t$ taken outside of the brackets.

Next, the vertical average is taken, the result is multiplied by the domain depth, and a horizontal divergence is applied to get:

$$\nabla_{\text{h}} \cdot \left[(H - h) \nabla_{\text{h}} \hat{p}_S \right] = \frac{\rho_0}{\Delta t} \nabla_{\text{h}} \cdot \left[(H - h) \overline{\mathbf{u}_{\text{h}}^{**z}} \right],$$

where the averaging operator is removed from the surface pressure term since the surface pressure is independent of the vertical. Computing the volume integral of this and rewriting using the divergence theorem we arrive at the equation for the surface pressure:

$$\int_{\partial V} (H - h) \nabla_h p_S \hat{\mathbf{n}} \cdot dS = \int_{\partial V} \frac{\rho_0(H - h)}{\Delta t} \overline{\mathbf{u}_h^{**z}} \hat{\mathbf{n}} \cdot dS. \quad (2.24)$$

This formulation represents a two-dimensional elliptic problem. Therefore, this represents a significant reduction in the computational expense over the three-dimensional elliptic solve in the NH model [Marshall et al., 1997a].

2.6 Quasi-Geostrophic Model

From the discussion of non-dimensionalization of the primitive equations in section 2.4, we saw that the non-dimensional inviscid horizontal momentum equations under the traditional approximation could be written as

$$\text{Ro} \frac{D\hat{\mathbf{u}}_h}{Dt} + \hat{f} \mathbf{k} \times \hat{\mathbf{u}} + \nabla_h \hat{p}' = 0,$$

where it is possible that our Coriolis term can be written on the f -plane or β -plane.

2.6.1 Leading Order Terms

If we make the assumption that $\text{Ro} \ll 1$ and decompose our variables into dominant geostrophic components and small ageostrophic components we can rewrite the variables in the form:

$$(u, v, p') = (u, v, p')_g + \text{Ro} (u, v, p')_a,$$

where $(u, v, p')_g$ and $(u, v, p')_a$ represents the geostrophic and ageostrophic variables respectively. Therefore, written in terms of dimensional variables, we have to leading order, geostrophic balance:

$$f_0 \mathbf{k} \times \mathbf{u}_g + \frac{1}{\rho_0} \nabla_h p'_g = 0.$$

Therefore, the geostrophic flow is horizontally non-divergent:

$$\nabla_h \cdot \mathbf{u}_g = 0.$$

As well, we see that given a streamfunction defined as $\psi = p'_g / (\rho_0 f_0)$ we can write $\mathbf{u}_g = \mathbf{k} \times \nabla \psi$ [Harwood, 2005].

2.6.2 Ageostrophic terms

Geostrophic balance is a diagnostic relationship and does not tell us anything about the evolution of the flow. If we wish to consider an evolving flow that is nearly geostrophic in nature then we need to consider the next largest terms in order Ro . Our non-dimensional equations tell us that the $\text{O}(\text{Ro})$ terms contributing to the material derivative are the terms arising from the geostrophic material derivative of \mathbf{u}_g , where the material geostrophic derivative is defined as:

$$\frac{D_g}{Dt} = \frac{\partial}{\partial t} + \mathbf{u}_g \cdot \nabla.$$

In terms of dimensional variables, we see that the QG horizontal momentum equations are

$$\frac{D_g \mathbf{u}_g}{Dt} + f_0 \mathbf{k} \times \mathbf{u}_a + \beta_0 y \mathbf{k} \times \mathbf{u}_g + \frac{1}{\rho_0} \nabla_h p'_a = 0,$$

where we have used the assumption that $\beta_0 L / f_0 \sim \text{Ro}$. To leading order, the continuity equation tells us that the vertical velocity component is neglected.

Applying the curl operator to this equation we get the following, where we have made use of the definition of our streamfunction:

$$\frac{D_g}{Dt} \nabla_h^2 \psi + \frac{\partial}{\partial x} (f_0 u_a + \beta_0 y u_g) + \frac{\partial}{\partial y} (f_0 v_a + \beta_0 y v_g) = 0.$$

Through the application of the product rule and the geostrophic horizontal divergence-free condition, this leads to the following:

$$\frac{D_g}{Dt} \nabla_h^2 \psi + f_0 \nabla_h \cdot \mathbf{u}_a + v_g \beta_0 = 0. \quad (2.25)$$

However, this can be simplified further by consider the $\text{O}(\text{Ro})$ terms of the continuity equation. To $\text{O}(\text{Ro})$ we have

$$\nabla_h \cdot \mathbf{u}_a + \frac{\partial w}{\partial z} = 0.$$

Therefore, replacing the horizontal divergence of \mathbf{u}_a in equation 2.25 by the vertical derivative of the vertical velocity our equation becomes

$$\frac{D_g}{Dt} (\nabla_h^2 \psi + \beta_0 y) = f_0 \frac{\partial w}{\partial z}, \quad (2.26)$$

where we have used the fact that $v_g = D_g y / Dt$.

To determine the vertical velocity component we can consider the evolution equation for density. In our non-dimensionalization we saw that the right hand side term had a coefficient of Bu^2 so if we make the assumption $Bu \sim 1$ we get the following leading order equation written in terms of dimensional variables:

$$\frac{\partial \rho'}{\partial t} + \mathbf{u}_g \cdot \nabla \rho' = \frac{\rho_0 N^2}{g} w,$$

where we have written this in term of dimensional variables. We can use this to evaluate w as

$$w = \frac{g}{\rho_0 N^2} \left(\frac{\partial}{\partial t} + \mathbf{u}_g \cdot \nabla \right) \rho'.$$

Rewriting ρ' in terms of ψ , making use of our leading order hydrostatic balance, we get the following expression for w in terms of the streamfunction:

$$w = -\frac{f_0}{N^2} \left(\frac{\partial}{\partial t} + \mathbf{u}_g \cdot \nabla \right) \frac{\partial \psi}{\partial z}.$$

Therefore, equation 2.26 becomes

$$\frac{D_g}{Dt} \left\{ \nabla_h^2 \psi + \beta_0 y + \frac{\partial}{\partial z} \left(\frac{f_0^2}{N^2} \frac{\partial \psi}{\partial z} \right) \right\} = 0.$$

This is our quasi-geostrophic potential vorticity (QGPV) equation and if we define our QGPV as

$$q = \nabla_h^2 \psi + \frac{\partial}{\partial z} \left(\frac{f_0^2}{N^2} \frac{\partial \psi}{\partial z} \right), \quad (2.27)$$

then this can be written in the following form:

$$\frac{D_g}{Dt} (q + \beta_0 y) = 0. \quad (2.28)$$

Given a specific background flow, we can define the associated QGPV using 2.27. Then, from equation 2.28, we update this vorticity similar to how we updated the tracer fields in the HY or NH models. At each time step, or sub-stage, the streamfunction is found by inverting equation 2.27. The geostrophic velocity components can then be evaluated from the relationships:

$$\mathbf{u}_g = \mathbf{k} \times \nabla \psi.$$

Therefore, the QG model has a much simpler formulation compared to the NH and HY models but does require a three-dimensional inversion when solving for the streamfunction. That being said, the QG model can offer potential speedups over HY and NH models as it does not allow for internal waves and can possibly lead to a greater allowable time step depending on the application.

Chapter 3

Numerical Methods

In this chapter, the complete details of the non-hydrostatic (NH) Tempest model are presented. As well, the hydrostatic (HY) and quasi-geostrophic (QG) models are discussed and the governing equations for the linear stability analysis (LSA) are presented. The Chapter begins with a discussion on the staggered grid and the formulation in the presence of topography. It then touches on each aspect of the time-stepping routine. The chapter concludes with a discussion of the QG model, LSA, and an analysis of the parallel efficiency of the solvers.

3.1 C Grid Staggering

Before describing the components of the model, the basic structure of the grid is introduced. The Tempest ocean model uses a standard Arakawa C grid [[Arakawa and Lamb, 1977](#), [Marshall et al., 1997a](#)]. The velocity components are staggered in each of the coordinate directions. Therefore, for a particular cell in the domain, the associated zonal velocity is located at the center of the East interface, the associated meridional velocity is located at the center of the North interface, and the associated vertical velocity is located at the center of the top interface where East, North, and top define the positive x , y , and z directions in our local Cartesian coordinate system. The dynamic pressure and density components are located at cell centers. Using this definition, the divergence is located at the cell centers and the associated vorticity is located at the Northeast corner of the grid. This grid cell staggering is shown in figure [3.1](#).

The main reason why C grid staggering is valuable is the increased stability that staggered grids offer [[Armfield and Street, 2005](#)]. In a staggered C grid the pressure and




Figure 3.1: Standard three-dimensional C grid staggering

divergence terms are computed using local cell information as well as information from adjacent cells. In the case of an unstaggered grid all the variables are collocated, which removes the need for interpolation, however it also implies that first derivatives using standard centered schemes will be computed without including knowledge of the variables at the sample location. This represents a problem as the solution to the elliptic problem for pressure then becomes a combination of two alternating solutions [Sheu and Lin, 2003].

Also, it has been shown that C grid staggering on rectangular grids does not have spurious gravity or Rossby waves, unlike other grids [Cotter and Shipton, 2012]. That being said, there are disadvantages to the C grid. In particular, grid staggering implies that we must be conscientious of the variable locations. When approximating the Coriolis terms we must account for staggering in the variables through interpolation. When interpolating we must also be aware of potential issues with energy conservation.

Furthermore, velocities must be collocated before discretizing the nonlinear advection scheme. Coupled with the C grid, using finite volume (FV) approaches to the elliptic pressure solve and the density equation does give us advantages. In the pressure solve, the pressure gradients and velocity components are collocated, reducing some of the computational complexity associated with the elliptic solve operator. In the density equation, the upwinded cell edge approximations of the density are collocated with the velocities.

3.2 Terrain-Following Coordinates

Ocean dynamics are often strongly influenced by the underlying topography. This is particularly true in the case of coastal regions. When there are variations in topography, it is convenient to make a change of variables from the physical terrain-following coordinate to the computational (Cartesian) domain. The cells in the terrain-following coordinate system have sloping top and bottom surfaces that move with the terrain; see figure 3.2 for a demonstration of the setup in two-dimensions. These sloping cells represent the physical coordinate system. The cell interfaces in the horizontal directions remain perpendicular to a flat bottom. This implies that the normal components to the x and y interfaces are \mathbf{i} and \mathbf{j} , representing the standard basis vectors of the Cartesian coordinate system.




Figure 3.2: Two-dimensional cross-section of a terrain-following grid over a Gaussian hill.

When applying this coordinate transformation the convention is to define the physical coordinates as x, y, z and the computational coordinates as X, Y, Z . The bounds set on the physical coordinates are dictated by the (potentially two-dimensional) varying bottom topography:

$$0 \leq x \leq L_x, \quad 0 \leq y \leq L_y, \quad h(x, y) \leq z \leq H,$$

where $h(x, y)$ is the spatially varying bottom topography. The transformation from the physical to the computational space is made by applying the following transformation formulas:

$$X = x, \quad Y = y, \quad Z = \frac{z - h(x, y)}{H - h(x, y)} \in [0, 1].$$

The numerical discretization of derivatives under this transformation has truncation errors that can introduce nonzero values into the equations, where the terms may otherwise be zero. For example, the x derivative of the pressure as computed in the computational framework is

$$p_x = p_X + Z_x p_Z.$$

The truncation error in numerically approximating the right-hand side of this expression can result in a nonzero pressure gradient, where the left-hand side may otherwise be zero. This is often the case when the topography is very steep and leads to pressure gradient errors [Miller, 2007]. The pressure gradient error arises from the transformed pressure gradient being the sum of two potentially large terms, p_X and $Z_x p_Z$ shown above. The discretization error of these terms can therefore lead to large numerical errors. Further discussions can be found in [Mellor et al., 1998, Shchepetkin and McWilliams, 2003, Beckmann and Haidvogel, 1993].

However, there are a couple of ways of reducing the pressure gradient errors. Theoretically writing the Boussinesq equations in terms of smaller dynamic pressure and density components reduces this issue significantly by removing the extremely large background gradients. Furthermore, making sure to consider the hydrostatic consistency condition as explained below will reduce the errors significantly [Miller, 2007]. The hydrostatic consistency condition states that

$$|h_x| \Delta x \leq \Delta z,$$

which implies that for each horizontal step we do not move more than one vertical step. Therefore, the horizontal resolution must be sufficiently small, given a particular topographic profile and vertical resolution. When this is not the case there are other approaches that would alleviate potential errors in this coordinate system, such as interpolating to level z surface such as in [Shchepetkin and McWilliams, 2003]. This approach is not used in Tempest.

Taking the derivatives of the coordinate transformation we have the following:

$$\begin{aligned} X_x &= 1, & X_y &= 0, & X_z &= 0, \\ Y_x &= 0, & Y_y &= 1, & Y_z &= 0, \\ Z_x &= \frac{z - H}{(H - h)^2} h_x, & Z_y &= \frac{z - H}{(H - h)^2} h_y, & Z_z &= \frac{1}{H - h}. \end{aligned}$$

Applying these rules, the physical derivatives can be written in terms of the computational derivative as

$$\partial_x = \partial_X + Z_x \partial_Z,$$

$$\begin{aligned}\partial_y &= \partial_Y + Z_y \partial_Z, \\ \partial_z &= Z_z \partial_Z.\end{aligned}$$

We can also transform the second derivatives in a similar fashion:

$$\begin{aligned}\partial_{xx}^2 &= \partial_x(\partial_X + Z_x \partial_Z) \\ &= \partial_x(\partial_X) + Z_{xx} \partial_Z + Z_x \partial_x(\partial_Z) \\ &= \partial_{XX}^2 + Z_x \partial_{ZX}^2 + Z_{xx} \partial_Z + Z_x (\partial_{XZ}^2 + Z_x \partial_{ZZ}^2) \\ &= \partial_{XX}^2 + 2Z_x \partial_{XZ}^2 + Z_x^2 \partial_{ZZ}^2 + Z_{xx} \partial_Z, \\ \partial_{yy}^2 &= \partial_{YY}^2 + 2Z_y \partial_{YZ}^2 + Z_y^2 \partial_{ZZ}^2 + Z_{yy} \partial_Z, \\ \partial_{zz}^2 &= Z_{zz} \partial_Z + Z_z^2 \partial_{ZZ}^2 = Z_z^2 \partial_{ZZ}^2.\end{aligned}$$

The Laplacian operator is then defined as

$$\begin{aligned}\nabla^2 &= \partial_{xx}^2 + \partial_{yy}^2 + \partial_{zz}^2 \\ &= \partial_{XX}^2 + \partial_{YY}^2 + \left[Z_x^2 + Z_y^2 + Z_z^2 \right] \partial_{ZZ}^2 + 2Z_x \partial_{XZ}^2 + 2Z_y \partial_{YZ}^2 + \left[Z_{xx} + Z_{yy} \right] \partial_Z.\end{aligned}$$

We must also be careful in defining the area of cell interfaces, cell volumes, and normal vectors in the physical coordinate system. To define these quantities, we can use the general cross product formula for a quadrilateral to determine the area of the vertical faces. That is, using the eight vertices p_1 to p_8 as shown in figure 3.3, the area of the *bottom* face of the cell is determined as

$$A_z = \frac{1}{2} |p_3 p_1 \times p_4 p_2|.$$

We can apply the same formula for the horizontal faces or we can take the integral of the difference between the linear functions defining the top and bottom edges of these faces. Either approach gives us the formula for the west and south faces as

$$\begin{aligned}A_x &= \Delta y \Delta z_{i-1/2,j}, \\ A_y &= \Delta x \Delta z_{i,j-1/2},\end{aligned}$$

where $\Delta z_{i-1/2,j}$ and $\Delta z_{i,j-1/2}$ represents the vertical grid spacings along the center of the west and south interfaces respectively.




Figure 3.3: Cell vertex labels

Since we are considering the terrain to be represented by linear functions and the x (y) faces of the grid are separated by Δx (Δy) in the horizontal we can write an equation defining the top and bottom faces of a cell as

$$T(x, y, z) = z_{i,j,k+1/2} + ax + by,$$

$$B(x, y, z) = z_{i,j,k-1/2} + cx + dy,$$

where we consider the cell to be bounded in the horizontal by $[-\Delta x/2, \Delta x/2] \times [-\Delta y/2, \Delta y/2]$. Integrating the difference between the top and bottom faces we determine the volume of the cell to be

$$V = \Delta x \Delta y \Delta z_{i,j},$$

where $\Delta z_{i,j}$ represents the vertical grid spacing between the top and bottom faces at the center of the cell.

The normal vectors in the x and y direction are equal to the Cartesian \mathbf{i} and \mathbf{j} so we only need to compute the normal vector to the top and bottom of the cell. These vectors are computed from the cross product of the diagonals and then normalized and oriented downward by ensuring that $\mathbf{k} \cdot \hat{\mathbf{n}}_z < 0$, the z component of the normal vector in the z direction.

3.3 Solution Procedure

The Tempest ocean model is based upon a blend of finite difference (FD) and finite volume (FV) methods. To find an FD approximation to some quantity we consider a set of adjacent

points. The order of accuracy of the methods increase as we consider larger stencils. However, in the case of a centered approximation we usually have that the error term is dispersive, where the centered approximations are prone to causing oscillations in the solutions [Durrant, 2010]. Therefore, we generally use upwinding based schemes.

The FD methodology is based on a Taylor series expansion of a function about a point where we want to approximate the derivative of the function. As an example, if we wanted to build a centered approximation to the first derivative of some function $f(x)$ at the location x_0 we may use the two adjacent cells and write an approximation in terms of some unknown coefficients a, b, c as

$$f'(x_0) \approx af(x_0 - \Delta x) + bf(x_0) + cf(x_0 + \Delta x).$$

Taking a Taylor series expansion of $f(x_0 \pm \Delta x)$ about the location $x = x_0$ we can equate coefficients on powers of Δx to try and determine an approximation of highest accuracy. Such an approximation would lead to the standard second order centered approximation:

$$\frac{f(x_0 + \Delta x) - f(x_0 - \Delta x)}{2\Delta x} \approx f'(x_0) + \frac{1}{6}\Delta x^2 f'''(x_0).$$

We see that the leading order error term involves an odd order derivative and, therefore, is prone to dispersion errors. In general, if we consider an approximation through n adjacent points which may or may not include the value of the function at the sample location, then we get a set of n equations in the n unknown coefficients a, b, c, \dots

To determine general FD approximations a Maple procedure called `Coeffs` has been developed. For the complete code please refer to appendix A. This function requires as input a general expansion order k , which determines the order of the Taylor expansion for each value of the function about $x = x_0$, a set of values that indicate the stencil, and a non-negative integer value indicating the derivative of which we wish to approximate. A value of 0 implies that we wish to approximate the value of the function. The set of function location values are listed as rational values that indicate the number of Δx steps from the approximation location $x = x_0$. The `Coeffs` procedure will output the FD approximation of highest order to the value at the indicated location.

For example, if we wished to determine the FD approximation to the first derivative of the function at some location $x = x_0$ using the function values $f(x_0 - \Delta x)$, $f(x_0)$, and $f(x_0 + \Delta x)$ then we would call this procedure with the command:

> `Coeffs(6, [-1, 0, 1], 1)`

$$\textit{Expanded Approximations} = \frac{\partial}{\partial x} f(x_0, y_0) + \frac{1}{6} \left(\frac{\partial^3}{\partial x^3} f(x_0, y_0) \right) \Delta x^2 + \frac{1}{120} \left(\frac{\partial^5}{\partial x^5} f(x_0, y_0) \right) \Delta x^4$$

$$\begin{aligned}
\textit{Approximation} &= \frac{1}{2} \frac{f_{i+1} - f_{i-1}}{\Delta x} \\
\textit{Truncation Error} &= \frac{1}{6} \left(\frac{\partial^3}{\partial x^3} f(x_0, y_0) \right) \Delta x^2
\end{aligned}$$

The reason we use partial differentiation and a two variable function is that the **Coeffs** procedure is based off another procedure that computes the FD approximations to a two variable function given points in the plane adjacent to the approximation location $\mathbf{x} = (x_0, y_0)$. This procedure is called **BiCoeffs**. The calling convention for this procedure is similar but the inputs now require a total of 5 arguments. The first is the order of the two-dimensional Taylor series expansion, the second is the set of relative x locations, the third is the set of relative y locations, the fourth is the order of the x derivative, and the fifth is the order of the y derivative. Please note that if we have n indices in the set of x and y locations then we are considering n values in the approximation and not $n \times n$. If the fourth and fifth argument are both 0 we are approximating the value of the function at the sample location (x_0, y_0) .

3.3.1 Time-Stepping

There are two broad classes of time-stepping methods:

1. Multistage (Runge-Kutta) methods
2. Multistep (AB, AM, AB-AM) methods

The first type of methods require more work per time step since it requires repeated applications of the operators, while the second has a larger storage requirement since we need to request information from previous time steps. In the Tempest ocean model, we use the strong stability preserving Runge-Kutta third order (SSPRK3) time-stepping method of Shu and Osher [Shu and Osher, 1988]. A particular step of the algorithm takes the following form:

$$\begin{aligned}
q_1 &= q_n + \Delta t L(q_n), \\
q_2 &= \frac{3}{4}q_n + \frac{1}{4} \left[q_1 + \Delta t L(q_1) \right], \\
q_{n+1} &= \frac{1}{3}q_n + \frac{2}{3} \left[q_2 + \Delta t L(q_2) \right].
\end{aligned} \tag{3.1}$$

Though the computational cost is greater than lower order or multistep methods, it is not so significant as to use a method with weaker stability properties. As shown by Wicker and Skamarock [Wicker and Skamarock, 2002], a Runge-Kutta third order approach coupled with a fifth order advective scheme has stability for flows that exceed a Courant number of 1. The Courant number is defined as:

$$C = \sum_i |u_i| \frac{\Delta t}{\Delta x_i}.$$

A limitation is imposed on this Courant number in the form of the Courant-Friedrichs-Lewy (CFL) condition. This CFL condition states that the numerical solution to a particular PDE must not be independent of the data that leads to the solution of the PDE [Durrant, 2010]. Therefore, typically the CFL condition will be at least as strict as $C \leq 1$ for explicit time-stepping schemes, where stricter means that C must be less than a smaller value than 1. The condition $C \leq 1$ acts as a guide and a thorough stability analysis must be performed to determine the strict constraint. From the definition of C , we see that a stricter CFL condition implies a reduction in Δt .

Satisfying the CFL condition does not imply stability. In general, when applied to a linear problem, the RK methods have the stability regions as shown in figure 3.4 with the SSPRK3 having the same stability region as RK3. In practice, the stability limit has seen to be slightly greater than those imposed by the theoretical limit, consistent with the results found by [Sowa, 1990].

The SSPRK3 method as described by the above set of equations 3.1 can be thought of as an Euler stage followed by incremental corrective stages. The bulk of the computational expense is in the first forward Euler stage. Furthermore, it is observed that each of the sub-stages of the SSPRK3 method are formed by taking linear combinations of forward Euler methods in which the coefficients multiplying each forward Euler sub-stage are positive [Durrant, 2010].

Like the forward Euler method, these methods imply that the strong-stability condition $\|q_{n+1}\| \leq \|q_n\|$ is retained by the time-stepping method when limiting on the spatial derivatives have imposed this constraint. When the strong stability methods are used, it has been found that the CFL constraint is stricter than the typical CFL constraint by a factor of $1 + \alpha$, where α is a specifically chosen parameter and $\alpha = 4$ has been suggested by [Suresh and Huynh, 1997]. Therefore, if the CFL condition is conventionally $C \leq 1$ then in application of the strong stability preserving method the suggestion by [Suresh and Huynh, 1997] is to use $C \leq 1/5$. When considering positive preservation, the limiting procedure [Zhang and Shu, 2010] does impose a much stricter CFL condition, as explained in section 3.3.10.

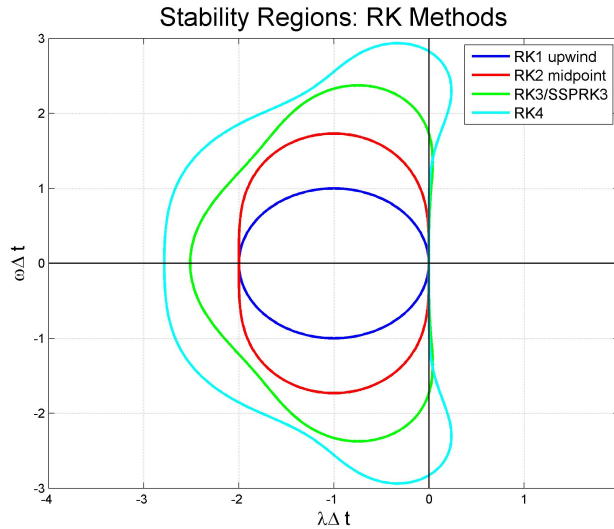


Figure 3.4: Stability Regions for Standard RK Methods

For each of the sub-stages of the SSPRK3 implementation the Tempest model follows a series of incremental steps. Therefore, we evaluate $q_1, q_2,$ and q_{n+1} as described in 3.1 using the following procedure:

1. Determine the pseudo-updated velocities using FD approximation methods. Recall from chapter 2 that the pseudo-updated velocities are the velocity components updated with all terms except for the pressure gradient terms. This update has the following steps:
 - (a) Update the velocities with the Coriolis terms, interpolating velocity components to staggered locations.
 - (b) Apply the nonlinear advection terms in the momentum equations.
 - (c) Apply viscosity to the pseudo-updated values if the viscosity coefficients are nonzero.
 - (d) Apply the buoyancy term to the vertical velocity if solving the NH equations.
 - (e) Calculate and update the horizontal velocities with the horizontal gradient of the hydrostatic pressure field if solving the HY equations.
2. Project the velocities into the set of incompressible solutions by computing the pressure and applying the pressure gradients to the velocity components. For the

elliptic problem for pressure we consider a volume integral of the elliptic problem. This is inherently a second order accurate operation due to the fact that we have point values. Writing the volume integral as a surface integral allows us to use the staggered grid to easily evaluate the divergence and accounts for the boundary conditions on the pressure.

3. If solving the HY equations, solve for the vertical velocity component using the discrete divergence-free condition.
4. Update the density field using FV methods, allowing for accurate conservation properties.

3.3.2 Interpolation and Coriolis terms

The combination of the staggering of the grid and the transformation to terrain-following coordinates means that we require the velocity components to be collocated during many of the sub-steps listed above. For example, when we are updating the velocity components to account for the Coriolis terms we must interpolate u to the location of v and vice versa. Therefore, it is valuable to explain the interpolation procedures used throughout the Tempest model.

The coefficients of the interpolation are determined from FD approximations using adjacent points. To explain the interpolation procedure we consider interpolating the v velocity components to the location of the u velocities. In a particular cell i, j, k the u velocities are located at the position $i + 1/2, j, k$ and the v velocities are located at the position $i, j + 1/2, k$. Therefore, we must interpolate the v velocity components a half step in the zonal and meridional directions. To do this we can use the `BiCoeffs` function, but we could also interpolate the v velocities to the center of the cell and then from the cell centers to the u locations.

The formula for the interpolation of v to the location of the u velocity components using the two-step approach is then defined as

$$\overline{\overline{v}}_{i+1/2,j,k}^y \approx \sum_{ii=1}^{n_x} \sum_{jj=1}^{n_y} \mathcal{C}_{ii}^x \mathcal{C}_{jj}^y v_{i+\alpha(ii),j+\beta(jj),k} \quad (3.2)$$

where n_x and n_y denote the number of values contributing to the interpolation in the x and y directions respectively and the set of values α and β are the i index offsets and j index offsets. The sets of coefficients \mathcal{C}^x and \mathcal{C}^y are the coefficients as determined by the one-dimensional interpolations in the x and y directions respectively. The overbar y denotes an average in the y direction and the overbar x denotes an average in the x direction.

For the cells in the interior of the domain, we can interpolate the v velocities to the cell centers using the standard interpolation formula:

$$\bar{v}_{i,j,k}^y \approx \frac{9(v_{i,j-1/2,k} + v_{i,j+1/2,k}) - (v_{i,j-3/2,k} + v_{i,j+3/2,k})}{16} + \frac{3}{128} \frac{\partial^4 v}{\partial y^4} \Big|_{i,j,k} \Delta y^4. \quad (3.3)$$

In a similar way, we can interpolate these cell centered values to the location of the u velocities using the following formula:

$$\bar{\bar{v}}_{i+1/2,j,k}^y \approx \frac{9(\bar{v}_{i,j,k}^y + \bar{v}_{i+1,j,k}^y) - (\bar{v}_{i-1,j,k}^y + \bar{v}_{i+2,j,k}^y)}{16} + \frac{3}{128} \frac{\partial^4 \bar{v}^y}{\partial x^4} \Big|_{i+1/2,j,k} \Delta x^4. \quad (3.4)$$

Therefore, the interpolations in the interior of the domain require sixteen adjacent v values. In the general interpolation formula 3.2 this implies that $n_x = n_y = 4$. The sets of coefficients and offsets are defined as

$$\mathcal{C}^x = \mathcal{C}^y = \left[-\frac{1}{16}, \frac{9}{16}, \frac{9}{16}, -\frac{1}{16} \right], \quad \alpha = [-1, 0, 1, 2], \quad \beta = \left[-\frac{3}{2}, -\frac{1}{2}, \frac{1}{2}, \frac{3}{2} \right].$$

In addition to interpolations in the interior of the domain, we must define appropriate interpolation stencils and coefficients when we are near solid walls. In the above case, of interpolating v to the location of the u velocities we must consider if we are near an x (East/West) or y (North/South) boundary. We consider first the case when we are near an x boundary. The u velocities on the staggered grid are defined at the East interfaces of each cell. Therefore, the u velocity associated with the first index is located at the East of the first cell and not against the Western domain boundary. As well, the u velocity located at the last index is located against the Eastern domain boundary and is set to zero in the non-periodic cases to account for no normal flow.

This means that we only need to make special considerations for the first and the second last u velocities in the x direction, located a displacement of Δx ($-\Delta x$) from the Western (Eastern) domain boundaries. At the first u velocity location we account for the boundary conditions by considering an approximation using the adjacent v velocity components as well as the unknown v value at the location of the western domain boundary. Using our `Coeffs` function we find that this interpolation is

$$\bar{\bar{v}}_{1/2,j,k}^y \approx \frac{-4\bar{v}_{-1/2,j,k}^y + 15\bar{v}_{0,j,k}^y + 10\bar{v}_{1,j,k}^y - \bar{v}_{2,j,k}^y}{20} + \frac{1}{64} \frac{\partial^4 \bar{v}^y}{\partial x^4} \Big|_{1/2,j,k} \Delta x^4. \quad (3.5)$$

If we have no-slip condition then the first term in the numerator will simply be set to 0 and we get the approximation defined as follows:

$$\overline{v}^y_{1/2,j,k} \approx \frac{15\overline{v}^y_{0,j,k} + 10\overline{v}^y_{1,j,k} - \overline{v}^y_{2,j,k}}{20} + \frac{1}{64} \frac{\partial^4 \overline{v}^y}{\partial x^4} \Big|_{1/2,j,k} \Delta x^4.$$

On the other hand, if we have free-slip conditions then we have homogeneous Neumann conditions on the boundary. In this case, we can approximate the x derivative of \overline{v}^y at the boundary location using the following FD formula:

$$\frac{\partial \overline{v}^y}{\partial x} \Big|_{-1/2,j,k} \approx \frac{-184\overline{v}^y_{-1/2,j,k} + 225\overline{v}^y_{0,j,k} - 50\overline{v}^y_{1,j,k} + 9\overline{v}^y_{2,j,k}}{60\Delta x} - \frac{5}{64} \frac{\partial^4 \overline{v}^y}{\partial x^4} \Big|_{-1/2,j,k} \Delta x^3.$$

Setting this equal to 0 and solving for the unknown $\overline{v}^y_{-1/2,j,k}$ we determine that

$$\overline{v}^y_{-1/2,j,k} \approx \frac{225\overline{v}^y_{0,j,k} - 50\overline{v}^y_{1,j,k} + 9\overline{v}^y_{2,j,k}}{184} - \frac{75}{2944} \frac{\partial^4 \overline{v}^y}{\partial x^4} \Big|_{-1/2,j,k} \Delta x^4.$$

Upon substitution into equation 3.5, we can then evaluate this interpolation in the x direction as

$$\overline{v}^y_{1/2,j,k} \approx \frac{93\overline{v}^y_{0,j,k} + 102\overline{v}^y_{1,j,k} - 11\overline{v}^y_{2,j,k}}{184} + \frac{61}{2944} \frac{\partial^4 \overline{v}^y}{\partial x^4} \Big|_{1/2,j,k} \Delta x^4.$$

We see that regardless of no-slip or free-slip condition, under this formulation we have that the truncation error in the FD approximation is fourth order accurate. More importantly, the first order error term has an even derivative and is therefore diffusive. In equation 3.2 we set $n_x = 3$, $\alpha = [0, 1, 2]$ and the coefficients \mathcal{C}^x are defined as

$$\mathcal{C}^x = \begin{cases} \left[\frac{15}{20}, \frac{10}{20}, -\frac{1}{20} \right], & \text{no-slip } x \text{ boundaries} \\ \left[\frac{93}{184}, \frac{102}{184}, -\frac{11}{184} \right], & \text{free-slip } x \text{ boundaries} \end{cases}.$$

By symmetry we can determine the indices and coefficients for the approximation of $\overline{v}^y_{N_x-3/2,j,k}$ as

$$\alpha = [-1, 0, 1], \quad \mathcal{C}^x = \begin{cases} \left[-\frac{1}{20}, \frac{10}{20}, \frac{15}{20} \right], & \text{no-slip } x \text{ boundaries} \\ \left[-\frac{11}{184}, \frac{102}{184}, \frac{93}{184} \right], & \text{free-slip } x \text{ boundaries} \end{cases}.$$

However, we must also consider when we are near y boundaries. When we consider the first, second, second last, and last approximation locations in the y direction we will have that the value of the v velocity on the y boundary will be involved in the interpolation. The v velocities on the south and north domain boundaries are equal to zero due to the no normal flow condition. Therefore, when we consider the second and second last interpolations we have that the formula is the same as the interior cells, equation 3.3, except with the coefficients of the v values associated with the boundaries set to zero.

When we consider interpolating the v velocities to the cell centers in the case of the first row of cells in the y direction then we must alter the interpolation, using a biased approximation. After considering the approximation and accounting for no normal flow we find that

$$\bar{v}_{i,0,k}^y \approx \frac{15v_{i,1/2,k} - 5v_{i,3/2,k} + v_{i,5/2,k}}{16} - \frac{5}{128} \frac{\partial^4 v}{\partial y^4} \Big|_{i,0,k} \Delta y^4.$$

Therefore, in equation 3.2 we have that $n_y = 3$ and the offsets and coefficients in the y direction are

$$\beta = \left[\frac{1}{2}, \frac{3}{2}, \frac{5}{2} \right], \quad \mathcal{C}^y = \left[\frac{15}{16}, -\frac{5}{16}, \frac{1}{16} \right].$$

In the case of the last row of cells in the y direction we must also consider a biased approximation symmetrical to the case for the first row. Therefore, in the approximations involving $\bar{v}_{i,N_y-1,k}^y$ we have $n_y = 3$ and

$$\beta = \left[-\frac{5}{2}, -\frac{3}{2}, -\frac{1}{2} \right], \quad \mathcal{C}^y = \left[\frac{1}{16}, -\frac{5}{16}, \frac{15}{16} \right].$$

We use the above interpolation in the treatment of the Coriolis terms. For the Coriolis terms, we can consider the f or β plane approximations. In both cases, within the Tempest model, the definition of the coefficients is the same with β set to zero when we are consider the f plane approximation. We can also opt to include non-traditional terms for which f_* would be defined as nonzero.

3.3.3 Nonlinear Advection

As part of the pseudo-update procedure the nonlinear advection terms are evaluated. We use a FD procedure for these nonlinear advection terms because the velocities are defined as point values. The nonlinear terms are applied in advective form because of the

computational simplicity this poses on a staggered grid. As well, this method lends itself well to upwind biasing, which is a component of the WENO method we have chosen to use.

To describe the procedure in greater depth let us consider the update of the nonlinear advection terms in the u momentum equation. For a particular u velocity component located at a general location $(i + 1/2, j, k)$ our update takes the form

$$u^* = u + \Delta t \left[u \left(\frac{\hat{u}_{i+1} - \hat{u}_i}{\Delta X} + Z_x \frac{\hat{u}_{k+1/2} - \hat{u}_{k-1/2}}{\Delta Z} \right) + \overline{\overline{v}}^{yx} \left(\frac{\hat{u}_{j+1/2} - \hat{u}_{j-1/2}}{\Delta Y} + Z_y \frac{\hat{u}_{k+1/2} - \hat{u}_{k-1/2}}{\Delta Z} \right) + \overline{\overline{w}}^{zx} \left(Z_z \frac{\hat{u}_{k+1/2} - \hat{u}_{k-1/2}}{\Delta Z} \right) \right],$$

where we have only labeled the indices that vary from the sample location $(i + 1/2, j, k)$. We have also suppressed the superscripts indicating the initial time of the sub-stage. The hatted values represent the value of the numerical flux function at the indicated index locations.

Upwind Biasing

Upwind biased approximations are typically used, in particular the weighted essentially non-oscillatory (WENO) scheme, for which we have positive and negative approximations to each of the numerical flux function values at each interface of the current cell. The WENO scheme is beneficial due to its essentially non-oscillatory behaviour. We choose the positive or negative numerical flux based on the sign of the advecting velocity components, $u, \overline{\overline{v}}^{yx}, \overline{\overline{w}}^{zx}$ respectively. In the case of the cross terms in the horizontal derivatives we consider the direction bias as dictated by the sign of the advecting velocity multiplied by the associated vertical normal vector component. This is described in figure 3.5.

We see that if the x component of the *downward-facing* normal vector, $\mathbf{i} \cdot \hat{\mathbf{n}}_z$, is positive and the u velocity is positive that the direction of the flux biasing will be downward, or from the positive k direction. Similarly, if u and $\mathbf{i} \cdot \hat{\mathbf{n}}_z$ are both negative then the direction will be from the positive k direction as well. When u and $\mathbf{i} \cdot \hat{\mathbf{n}}_z$ are of opposite signs then we have that the flux direction is from the negative k direction. Therefore, we have the following:

$$\hat{u}_{k+1/2} = \begin{cases} \hat{u}_{k+1/2}^- & \text{if } (u \mathbf{i}) \cdot \hat{\mathbf{n}}_z \leq 0 \\ \hat{u}_{k+1/2}^+ & \text{if } (u \mathbf{i}) \cdot \hat{\mathbf{n}}_z > 0 \end{cases}.$$

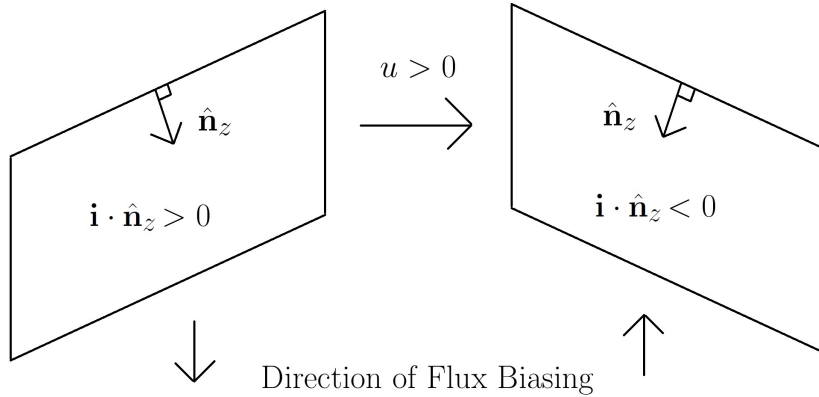


Figure 3.5: Vertical biasing direction for a grid with a terrain-following transformation.

Ghost Cell Extensions

When determining the flux values in a given direction we use adjacent point values. When we are near a boundary, some of these adjacent values may lie outside the domain. Cells that lie outside the domain are called ghost cells and when we are determining an approximation to the ghost values we must account for the boundary conditions. A commonly used approach is to compute the values of the ghost cells from even and odd extensions depending on if we are applying a homogenous Neumann or homogenous Dirichlet condition respectively [LeVeque, 1998, Guy and Fogelson, 2005]. As an example, under this approach we could extend the u velocity at the South domain boundary by taking

$$u_{i+1/2,-j,k} = \begin{cases} u_{i+1/2,j-1,k} & \text{free-slip } y \text{ boundaries} \\ -u_{i+1/2,j-1,k} & \text{no-slip } y \text{ boundaries} \end{cases}, \quad j = 1, \dots, h,$$

where h is the width of the ghost cell region. By considering a Taylor series argument these extensions can be shown to be second order accurate. In Tempest, the same idea is used but the approximations involve more known values of the function. The approach is similar to that used in section 3.3.2.

To demonstrate this approach, we consider the example of extending the u velocity across the south boundary. When we consider the $\partial_Y u$ term in the update of the u velocity we require three adjacent points from the north and south of the sample location. If we consider the first u velocity in the y direction then we must consider approximations to the three ghost elements to the south. In this particular case the sample location is $(i + 1/2, 0, k)$ and the ghost elements are located at $(i + 1/2, -j, k)$ for $j = 1, 2, 3$. Since this interpolation is unidirectional we will drop the i and k indices for simplicity.

To devise a fourth order approximation to the ghost elements we use the first three points in the y direction plus the potentially unknown value of the function at the boundary. Using this approach the approximate values of the ghost cells are

$$\begin{aligned}
u_{-1} &\approx \frac{48u_{-1/2} - 45u_0 + 15u_1 - 3u_2}{15} + \frac{1}{8} \frac{\partial^4 u}{\partial y^4} \Big|_{-1} \Delta y^4, \\
u_{-2} &\approx \frac{64u_{-1/2} - 90u_0 + 40u_1 - 9u_2}{5} + \frac{3}{2} \frac{\partial^4 u}{\partial y^4} \Big|_{-2} \Delta y^4, \\
u_{-3} &\approx 32u_{-1/2} - 50u_0 + 25u_1 - 6u_2 + \frac{25}{4} \frac{\partial^4 u}{\partial y^4} \Big|_{-3} \Delta y^4,
\end{aligned} \tag{3.6}$$

where the index location $j = 0$ refers to the center of the first cell. Therefore, $u_{j-1/2} = u_{-1/2}$ refers to the u velocity value at the South interface of the first cell.

In the case of no-slip conditions we simply set $u_{-1/2} = 0$ in the equations above and we get the fourth order approximations. In order to maintain fourth order accuracy in the case of free-slip conditions we must determine an approximation to $u_{-1/2}$ that is at least fourth order accurate. To do this we consider the approximation to the first derivative at the boundary location using the same stencil as above. This gives the following approximation:

$$\frac{\partial u}{\partial y} \Big|_{-1/2} \approx \frac{-184u_{-1/2} + 225u_0 - 50u_1 + 9u_2}{60 \Delta y} - \frac{5}{64} \frac{\partial^4 u}{\partial y^4} \Big|_{-1/2} \Delta x^3.$$

Setting this equal to 0 we get our fourth order approximation to $u_{-1/2}$:

$$u_{-1/2} \approx \frac{225u_0 - 50u_1 + 9u_2}{184} - \frac{75}{2944} \frac{\partial^4 u}{\partial y^4} \Big|_{-1/2} \Delta x^4.$$

Evaluating the equations 3.6 we see that the fourth order extensions to the ghost elements under free-slip boundary conditions are

$$\begin{aligned}
u_{-1} &\approx \frac{21u_0 + 3u_1 - u_2}{23} + \frac{1}{23} \frac{\partial^4 u}{\partial y^4} \Big|_{-1} \Delta y^4, \\
u_{-2} &\approx \frac{-54u_0 + 104u_1 - 27u_2}{23} + \frac{27}{23} \frac{\partial^4 u}{\partial y^4} \Big|_{-2} \Delta y^4, \\
u_{-3} &\approx \frac{-250u_0 + 375u_1 - 102u_2}{23} + \frac{125}{23} \frac{\partial^4 u}{\partial y^4} \Big|_{-3} \Delta y^4.
\end{aligned}$$

We observe that the truncation error terms are again diffusive.

We can use analogous formulas to extend the u velocities across y and z boundaries, the v velocities across x and z boundaries, and the w velocities across x and y boundaries. The only cases that are different are for the approximation of ghost values of the u, v, w velocities across x, y, z boundaries respectively. To demonstrate these cases consider the u velocity being extended beyond the East domain boundary. The u velocity located at the East boundary is at a index location $(N_x - 1/2, j, k)$ and the three adjacent ghost cells to the East are located at $(N_x + i, j, k)$ for $i = (1/2, 3/2, 5/2)$.

The approach is the same but we note that, as a result of the grid staggering, the first ghost cell in this direction is located a distance of Δx from the boundary as opposed to $\Delta x/2$. Therefore, we must reevaluate the formula for the ghost cell approximations. Suppressing the j, k indices, we determine the formulas in this direction to be

$$\begin{aligned} u_{N_x+1/2} &\approx -6u_{N_x-3/2} + 4u_{N_x-5/2} - u_{N_x-7/2} + \frac{\partial^4 u}{\partial x^4} \Big|_{N_x+1/2} \Delta x^4, \\ u_{N_x+3/2} &\approx -20u_{N_x-3/2} + 15u_{N_x-5/2} - 4u_{N_x-7/2} + 5 \frac{\partial^4 u}{\partial x^4} \Big|_{N_x+3/2} \Delta x^4, \\ u_{N_x+5/2} &\approx -45u_{N_x-3/2} + 36u_{N_x-5/2} - 10u_{N_x-7/2} + 15 \frac{\partial^4 u}{\partial x^4} \Big|_{N_x+5/2} \Delta x^4. \end{aligned}$$

After obtaining the ghost cell values using the above methods we can proceed with our approximation techniques as we would for a cell in the interior of the domain.

WENO FD Formulation

Since the u velocity components are defined as point values we define a sliding average function such that the average of this function over the cell is equal to this point value [Qiu and Shu, 2011]. For example, given the point value of the u velocity we assume that the u velocity of cell (i, j, k) can be expressed in terms of some other function $u^*(x, y, z)$ as [Shu, 1991]

$$u_{i+1/2,j,k} = \frac{1}{\Delta x} \int_{x_i}^{x_{i+1}} u^*(x, y, z) dx,$$

from which we conclude that

$$\frac{\partial u}{\partial x} \Big|_{i+1/2,j,k} = \frac{u_{i+1}^* - u_i^*}{\Delta x}.$$

Therefore, in order to determine an approximation to this derivative we need to choose $\hat{u}_{i,j,k} \approx u_{i,j,k}^*$ and $\hat{u}_{i+1,j,k} \approx u_{i+1,j,k}^*$. It is not that easy to approximate u^* since we do not know the function. However, we do have the cell average values of the function, which are the u velocity point values. In this way, the FD and FV WENO procedures are equivalent but give rise to different approximations. In the FD context, we obtain a high order accurate approximation to the derivative of u at the sample location. On the other hand, in the FV context we obtain high order accurate approximations to the point values of the state variable for which we have cell average values.

It is also important to note that applying the technique in this way requires an equally spaced grid, otherwise the method will be no more than second order accurate [Shu, 1991]. To show this we can consider an example approximation using this technique. If we consider the approximation of $\partial_x u$ in the u momentum equation we can evaluate approximations to the interface values given point values in the interior. The standard fifth order upwind approximation gives the flux formulas:

$$\begin{aligned}\hat{u}_{i+1}^+ &\approx (2u_{i+7/2} - 13u_{i+5/2} + 47u_{i+3/2} + 27u_{i+1/2} - 3u_{i-1/2}) / 60 \\ \hat{u}_{i+1}^- &\approx (2u_{i-3/2} - 13u_{i-1/2} + 47u_{i+1/2} + 27u_{i+3/2} - 3u_{i+5/2}) / 60 \\ \hat{u}_i^+ &\approx (2u_{i+5/2} - 13u_{i+3/2} + 47u_{i+1/2} + 27u_{i-1/2} - 3u_{i-3/2}) / 60 \\ \hat{u}_i^- &\approx (2u_{i-5/2} - 13u_{i-3/2} + 47u_{i-1/2} + 27u_{i+1/2} - 3u_{i+3/2}) / 60\end{aligned}$$

If we expand the first flux formula using a Taylor expansion we find that

$$\frac{1}{60} (2u_{i+7/2} - 13u_{i+5/2} + 47u_{i+3/2} + 27u_{i+1/2} - 3u_{i-1/2}) = u_{i+1}^+ - \frac{1}{24} \frac{\partial^2 u}{\partial x^2} \Big|_{i+1} \Delta x^2.$$

Therefore, this is a second order accurate approximation to the edge value of u , however if we expand the flux approximations above we find that

$$\frac{\hat{u}_{i+1}^\pm - \hat{u}_i^\pm}{\Delta x} = \frac{\partial u}{\partial x} \Big|_{i+1/2} \pm \frac{1}{60} \frac{\partial^6 u}{\partial x^6} \Big|_{i+1/2} \Delta x^5,$$

and, if we were to use alternate positive and negative approximation, we would obtain the following:

$$\frac{\hat{u}_{i+1}^\mp - \hat{u}_i^\pm}{\Delta x} = \frac{\partial u}{\partial x} \Big|_{i+1/2} \mp \frac{1}{30} \frac{\partial^5 u}{\partial x^5} \Big|_{i+1/2} \Delta x^4.$$

We conclude that we obtain an approximation that is one order higher by considering both positive or negative fluxes. Also, the first order term is dissipative instead of dispersive,

which is beneficial for numerical stability. Therefore, we consider either the FD derivative using both positive fluxes or both negative fluxes as determined based upon the sign of the advecting velocity component as described previously in this section.

3.3.4 Viscosity

After applying the Coriolis and nonlinear advection terms, we apply viscosity to the pseudo-updated velocity components. This viscosity is applied via an implicit formulation and the splitting of the steps in this way does introduce numerical errors associated with the time-stepping. However, when we consider the nearly inviscid limit for which $\nu_H, \nu_V \ll 1 \text{ m}^2/\text{s}$ the errors associated with this splitting will presumably be small.

The viscosity operators are defined on a five-point fourth order FD stencil. This level of accuracy is used to allow for potentially larger viscosity coefficients and because the computational complexity of the implicit solve of the viscosity is negligible compared to that of the pressure, described below. We can determine the limitation on the fourth order explicit scheme at interior elements by considering the standard fourth order approximation formula:

$$\begin{aligned}
u_{i,j,k}^{n+1} = & u^n - \nu_H \Delta t \left[\frac{u_{i-2,j,k}^n - 16u_{i-1,j,k}^n + 30u_{i,j,k}^n - 16u_{i+1,j,k}^n + u_{i+2,j,k}^n}{12\Delta x^2} \right] \\
& - \nu_H \Delta t \left[\frac{u_{i,j-2,k}^n - 16u_{i,j-1,k}^n + 30u_{i,j,k}^n - 16u_{i,j+1,k}^n + u_{i,j+2,k}^n}{12\Delta y^2} \right] \\
& - \nu_V \Delta t \left[\frac{u_{i,j,k-2}^n - 16u_{i,j,k-1}^n + 30u_{i,j,k}^n - 16u_{i,j,k+1}^n + u_{i,j,k+2}^n}{12\Delta z^2} \right].
\end{aligned}$$

Applying a Von Neumann stability analysis we determine the amplification factor is

$$A = 1 - \frac{1}{12} \left[d_x \mathcal{F}(\Delta x) + d_y \mathcal{F}(\Delta y) + d_z \mathcal{F}(\Delta z) \right],$$

where

$$\begin{aligned}
\mathcal{F}(x) &= 2 \cos(2x) + 30 - 32 \cos(x), \\
d_x &= \frac{\nu_H \Delta t}{\Delta x^2}, \quad d_y = \frac{\nu_H \Delta t}{\Delta y^2}, \quad d_z = \frac{\nu_V \Delta t}{\Delta z^2}.
\end{aligned}$$

In order for the strong stability condition $|A| \leq 1$ to hold we must have the following:

$$-2 \leq -\frac{1}{12} \left[d_x \mathcal{F}(\Delta x) + d_y \mathcal{F}(\Delta y) + d_z \mathcal{F}(\Delta z) \right] \leq 0,$$

where the right inequality holds since $d_x, d_y, d_z, \mathcal{F}(\cdot) \geq 0$. Each of the $\mathcal{F}(\cdot)$ is bounded above by 64 so the left inequality leads to the stability condition:

$$\nu_H \frac{\Delta t}{\Delta x^2} + \nu_H \frac{\Delta t}{\Delta y^2} + \nu_V \frac{\Delta t}{\Delta z^2} \leq \frac{3}{8}.$$

Therefore, we could use an explicit time-stepping approach for the viscosity term however this imposes a limitation on the time step proportional to the grid spacing squared. For sufficiently small viscosity coefficients or in potentially coarse grid simulations explicit time-stepping is justified as the constraint could be more lenient than that imposed by the CFL condition. However, for general applicability we use implicit methods.

If we write the above using the Backward Euler time-stepping method then we obtain the amplification factor:

$$A = \left\{ 1 + \frac{1}{12} \left[d_x \mathcal{F}(\Delta x) + d_y \mathcal{F}(\Delta y) + d_z \mathcal{F}(\Delta z) \right] \right\}^{-1},$$

which is always less than or equal to 1. Therefore, the fourth order centered FD scheme with implicit time-stepping is unconditionally stable.

The formulation for the implicit solve uses a backward Euler formulation:

$$\mathbf{u}^{n+1} = \mathbf{u}^n + \Delta t \mathcal{L} \mathbf{u}^{n+1},$$

where \mathcal{L} is the Laplacian operator. Solving this for \mathbf{u}^{n+1} gives us the following:

$$\mathbf{u}^{n+1} = (\mathbf{I} - \Delta t \mathcal{L})^{-1} \mathbf{u}^n = \mathcal{D}^{-1} \mathbf{u}^n,$$

where $\mathcal{D} = \mathbf{I} - \Delta t \mathcal{L}$ is the diffusion operator and \mathbf{I} is the identity matrix. The \mathcal{D} matrix is initialized to \mathbf{I} and the Δt times the \mathcal{L} matrix is added to the identity at the end of the matrix assembly. The \mathcal{L} matrix is assembled in the following way:

$$\mathcal{L} = \partial_{XX}^2 + \partial_{YY}^2 + (Z_x^2 + Z_y^2 + Z_z^2) \partial_{ZZ}^2 + 2(Z_x \partial_X + Z_y \partial_Y) \partial_Z + (Z_{xx} + Z_{yy}) \partial_Z.$$

Each application of this implicit viscosity operator requires a three-dimensional matrix solve at each sub-stage of an SSPRK3 step. However, the computational expense imposed by these solves is, in practice, minor when compared to the pressure solve since the viscosity typically only introduces small changes to the velocity fields. Using the default solver settings above and the default tolerances the implicit viscosity solve typically requires between 1-3 iterations, whereas the pressure solve generally requires between 70-100.

In defining the FD approximations to the derivatives, we can use similar approaches to the ones described in the interpolation and nonlinear advection sections above when we are near domain boundaries. We will not cover this in the same level of detail here but consider the ∂_{XX}^2 and ∂_{YY}^2 terms involved in the viscosity applied to the u velocity. All the other second derivatives involve similar coefficients and stencils. These can be determined by reorienting the below approximations and/or by symmetry. Deriving the FD approximations to fourth order accuracy and applying the no normal flow boundary condition we get the following:

$$12\Delta X^2 \left. \frac{\partial^2 u}{\partial X^2} \right|_{i+1/2} = \begin{cases} -20u_{i+1/2} + 6u_{i+3/2} + 4u_{i+5/2} - u_{i+7/2} & i = 0 \\ 16u_{i-1/2} - 30u_{i+1/2} + 16u_{i+3/2} - u_{i+5/2} & i = 1 \\ -u_{i-3/2} + 16u_{i-1/2} - 30u_{i+1/2} + 16u_{i+3/2} - u_{i+5/2} & i = 2, \dots, N_x - 4 \\ -u_{i-3/2} + 16u_{i-1/2} - 30u_{i+1/2} + 16u_{i+3/2} & i = N_x - 3 \\ -u_{i-5/2} + 4u_{i-3/2} + 6u_{i-1/2} - 20u_{i+1/2} & i = N_x - 2 \\ 0 & i = N_x - 1 \end{cases}.$$

In determining the approximation to the $\partial_{YY}^2 u$ term we must account for either free-slip or no-slip conditions. In the case of no-slip conditions, we determine the approximations

$$\left. \frac{\partial^2 u}{\partial Y^2} \right|_j = \begin{cases} (-25u_j + 10u_{j+1} - u_{j+2})/5\Delta y^2 & j = 0 \\ (175u_{j-1} - 280u_j + 147u_{j+1} - 10u_{j+2})/105\Delta y^2 & j = 1 \\ (-u_{j-2} + 16u_{j-1} - 30u_j + 16u_{j+1} - u_{j+2})/12\Delta y^2 & j = 2, \dots, N_y - 3 \\ (-10u_{j-2} + 147u_{j-1} - 280u_j + 175u_{j+1})/105\Delta y^2 & j = N_y - 2 \\ (-u_{j-2} + 10u_{j-1} - 25u_j)/5\Delta y^2 & j = N_y - 1 \end{cases}$$

and in the case of free-slip conditions we have the following:

$$\left. \frac{\partial^2 u}{\partial Y^2} \right|_j = \begin{cases} (-25u_j + 26u_{j+1} - u_{j+2})/23\Delta y^2 & j = 0 \\ (335u_{j-1} - 669u_j + 357u_{j+1} - 23u_{j+2})/264\Delta y^2 & j = 1 \\ (-u_{j-2} + 16u_{j-1} - 30u_j + 16u_{j+1} - u_{j+2})/12\Delta y^2 & j = 2, \dots, N_y - 3 \\ (-23u_{j-2} + 357u_{j-1} - 669u_j + 335u_{j+1})/365\Delta y^2 & j = N_y - 2 \\ (-u_{j-2} + 26u_{j-1} - 25u_j)/23\Delta y^2 & j = N_y - 1 \end{cases}.$$

3.3.5 Buoyancy and Deconvolution

When solving the NH equations we must add the buoyancy forcing to the vertical momentum equation. The density is defined as cell averages across a particular cell. To determine an approximation to the density point value at the location of the vertical velocity we would have to apply a deconvolution operator to obtain accuracy greater than second order. However, lower order methods are often applied in the vertical since the vertical velocity and grid spacing are generally much smaller than the horizontal [Ullrich and Jablonowski, 2012]. As the vertical velocity increases in magnitude, which can happen in some submesoscale processes at fronts or in coastal waters, this could pose a greater constraint on the accuracy.

To derive the second and fourth order approximations we recognize that we have cell-averaged density quantities. To account for cell staggering we considering the following:

$$\begin{aligned} \bar{\rho}_{i,j,k}^{xyz} + \bar{\rho}_{i,j,k+1}^{xyz} = \frac{1}{\Delta x \Delta y \Delta z} & \left[\int_{x_i - \Delta x/2}^{x_i + \Delta x/2} \int_{y_j - \Delta y/2}^{y_j + \Delta y/2} \int_{z_k - \Delta z/2}^{z_k + \Delta z/2} \rho dz' dy' dx' \right. \\ & \left. + \int_{x_i - \Delta x/2}^{x_i + \Delta x/2} \int_{y_j - \Delta y/2}^{y_j + \Delta y/2} \int_{z_k + \Delta z/2}^{z_k + 3\Delta z/2} \rho dz' dy' dx' \right]. \end{aligned}$$

We define an antiderivative to $\rho(x, y, z)$, $R(x, y, z)$, such that $\partial_{xyz}^3 R = \rho$. Evaluating the antiderivative at the bounds of the integration and expanding out the terms in a three-dimensional Taylor series about the location of w , $(x_i, y_j, z_k + \Delta z/2)$ this yields the following:

$$\begin{aligned} \bar{\rho}_{i,j,k}^{xyz} + \bar{\rho}_{i,j,k+1}^{xyz} = 2 \rho_{i,j,k+1/2} & + \frac{\Delta x^2}{12} \partial_{xx}^2 \rho_{i,j,k+1/2} + \frac{\Delta y^2}{12} \partial_{yy}^2 \rho_{i,j,k+1/2} \\ & + \frac{\Delta z^2}{3} \partial_{zz}^2 \rho_{i,j,k+1/2} + O(\Delta x^4, \Delta y^4, \Delta z^4). \end{aligned}$$

Therefore, we can evaluate the point value of the density field at the location of the vertical velocity by solving the above for $\rho_{i,j,k+1/2}$, which yields:

$$\rho_{i,j,k+1/2} \approx \frac{1}{2} (\bar{\rho}_{i,j,k}^{xyz} + \bar{\rho}_{i,j,k+1}^{xyz}) - \frac{\Delta x^2}{12} \partial_{xx}^2 \rho_{i,j,k+1/2} - \frac{\Delta y^2}{12} \partial_{yy}^2 \rho_{i,j,k+1/2} - \frac{\Delta z^2}{6} \partial_{zz}^2 \rho_{i,j,k+1/2}. \quad (3.7)$$

This is a fourth order accurate approximation to the density point value at the location of the vertical velocity. It is obvious from this that the second order approximation is the average of the cell average values from the adjacent cells above and below.

As an alternative to this we interpolate the cell averaged values to the edge using the fourth order approximation in the vertical:

$$\bar{\rho}_{i,j,k+1/2}^{xy} = -\frac{1}{12}\bar{\rho}_{i,j,k-1}^{xyz} + \frac{7}{12}\bar{\rho}_{i,j,k}^{xyz} + \frac{7}{12}\bar{\rho}_{i,j,k+1}^{xyz} - \frac{1}{12}\bar{\rho}_{i,j,k+2}^{xyz} + O(\Delta x^2, \Delta y^2),$$

which can be determined from equation 3.7. The reasoning behind this is that we usually use an interpolation in the vertical and then a two-dimensional deconvolution operator to the interface area averaged quantity [Ullrich and Jablonowski, 2012]. This deconvolution operator is required to obtain high-order accurate point values at the center of the interfaces. As the vertical velocity is typically very small relative to the horizontal velocities, we consider the above approximation in the application of the buoyancy terms.

3.3.6 Pressure Solve

The pressure solve takes one of two forms depending on if we are solving the HY or NH equations. In the HY case the first step is to compute the hydrostatic pressure which we do by integrating the hydrostatic equation in the vertical. The density point values are defined at the location of the vertical velocity components using the interpolation above and then multiplied by $-g$, where g is the gravitational constant. This forms the solution vector. We then invert the vertical derivative matrix to determine the hydrostatic pressure at the cell centers.

The hydrostatic pressure gradient force is then added to the horizontal velocity components. The updated horizontal velocity components are then vertically integrated and scaled. The discrete horizontal surface integral is applied to these values to form the right hand side of the surface pressure equation 2.24 as defined in chapter 2. Recall that this surface pressure equation is

$$\int_{\partial V} (H - h) \nabla_h p_S \hat{\mathbf{n}} \cdot dS = \int_{\partial V} \frac{\rho_0(H - h)}{\Delta t} \mathbf{u}_h^{**z} \hat{\mathbf{n}} \cdot dS, \quad (3.8)$$

where \mathbf{u}_h^{**} represents our pseudo-updated horizontal velocities including the horizontal gradient of our hydrostatic pressure. To form the discrete horizontal surface integral we apply the discrete divergence operator to the horizontal derivative operators, which are multiplied by the domain depth.

For computational efficiency the derivative operators in this case are defined using second order accurate methods. The pressure solve in both the HY and NH cases are computationally expensive and using second order approximations creates a much sparser

elliptic problem. This is presumed to offer a significant speedup over fourth order discretizations. In addition, the discrete divergence operators are defined as the interface areas times the normal component which inherently assumes a midpoint rule for the average interface area quantities. In order to be truly fourth order accurate we would have to use fourth order derivative approximations and a fourth order divergence operator, which would involve a convolution procedure. This convolution divergence operator would then be applied to the fourth order FD derivative approximations with possible cross-terms from the transformation. This would create a fairly dense matrix and slow down the method significantly. These reasons have guided our choice of lower order approaches.

Upon solving for the linear algebraic system, the horizontal velocities are then updated with the horizontal gradient of the surface pressure. The vertical velocity is then determined by inverting the discrete divergence-free condition written in terms of a surface integral:

$$\int_{\partial V} \mathbf{u} \cdot \hat{\mathbf{n}} dS = 0.$$

Therefore, the vertical velocity in the hydrostatic limit is by definition second order accurate. However, in the hydrostatic approximation we anticipate that the vertical velocity will be orders of magnitude smaller than the horizontal velocity components. The vertical velocity solve takes the form:

$$(w^{n+1}A)_t - (w^{n+1}A)_b = (u^{n+1}A)_e - (u^{n+1}A)_w + (v^{n+1}A)_n - (v^{n+1}A)_s,$$

where the superscript $n + 1$ denotes the updated velocities at the end of the current stage of the SSPRK3 time-stepping procedure and the subscripts denote the evaluation locations. For example, the $(u^{n+1}A)_e$ term denotes the u velocity evaluated at the East interface of the cell multiplied by the East interface area. The above equation is solved by inverting the left hand side under the no normal flow conditions applied at the top and bottom domain boundaries.

In the case of the NH pressure solve we have a similar formulation. The vertical velocity is defined as a pseudo-updated variable and the buoyancy term is applied as described above. The matrix involved in the elliptic solver is defined as the discrete surface integral times the discrete gradient term.

$$\int_{\partial V} \nabla p' \cdot \hat{\mathbf{n}} dS = \frac{\rho_0}{\Delta t} \int_{\partial V} \mathbf{u}^* \cdot \hat{\mathbf{n}} dS. \quad (3.9)$$

The boundary conditions on pressure were defined by equation 2.11 of Chapter 2. This balance was written as

$$\nabla p' \cdot \hat{\mathbf{n}} = \rho_0 \mathcal{G}_{\mathbf{u}} \cdot \hat{\mathbf{n}},$$

where \mathcal{G}_u represents all terms except the pressure gradient terms and was defined in equation 2.10. After determining the dynamic pressure from the above three-dimensional elliptic problem 3.9 we update all the velocity components, thereby projecting the solution into a divergence-free space [Chorin, 1968].

In both the HY and NH models we are solving a linear system of the form $\mathcal{A}\mathbf{x} = \mathbf{b}$, where in the HY model $\mathbf{x} = p_S$ and in the NH model $\mathbf{x} = p'$. The construction of the matrix \mathcal{A} is performed at the beginning of a simulation and stored for use throughout the time steps. To perform this construction we first define divergence and derivative matrices. The divergence matrices are defined as the $\hat{\mathbf{n}} dS$ (midpoint integration rule) in the formulas 3.8 and 3.9. The derivative matrices are the FD approximations to the staggered pressure derivatives. In the case of the HY model, these matrices are also multiplied by the domain depth at the interface location to account for the $H - h$ scaling factor. We then multiply the derivative matrices by our divergence matrices to form our \mathcal{A} matrix.

3.3.7 A PETSc Aside

In sections 3.3.4 and 3.3.6 we discussed solving linear systems. For this task the Portable, Extensible Toolkit for Scientific Computation (PETSc) [Balay et al., 1997, Balay et al., 2015] has been used extensively. PETSc interfaces directly with C++ programs and has a wide assortment of efficient and fully parallelized routines for solving linear and nonlinear equations. In Tempest, the linear solver routines of PETSc are used. In addition, matrix assembly is all done through PETSc using a wide assortment of their available linear algebra routines, such as matrix-matrix multiplication and matrix-vector multiplication.

Within Tempest, there are specific functions dedicated to matrix assemblies and linear solver definitions. Each of the solvers, including the implicit viscosity solver, vertical velocity solver, HY elliptic pressure solver, and NH elliptic pressure solver, are defined in their own routine. A matrix assembly in PETSc takes the form:

```
Mat A;
MatCreate(MPI_COMM_WORLD, &A);
MatSetSizes(A, nLocal, nLocal, nGlobal, nGlobal);
MatSetType(A, MATMPIAIJ);
MatSeqAIJSetPreallocation(A, 0, d_nnz);
MatMPIAIJSetPreallocation(A, 0, d_nnz, 0, o_nnz);
```

This creates a sparse parallel matrix which is described by the matrix type MATMPIAIJ. In the creation call we can use MPI_COMM_WORLD or PETSC_COMM_WORLD since the PETSC_COMM_

WORLD is an alias to MPI_COMM_WORLD. The sizes nLocal and nGlobal refer to the local size of the nLocal×nLocal sub-matrix on the current processor and nGlobal×nGlobal is the total size of the matrix. By far the most important step is the preallocation of the matrix which is handled by the last two commands above. Properly preallocating can often reduce the amount of time spent in the matrix assemblies by upward of 50 times [Balay et al., 2015].

After the matrix has been properly defined the elements are set by iterating over all local elements of the current processor. The MatSetValues routine of PETSc is then used to set the nonzero elements. The preallocations must agree with the indices of the elements set in the MatSetValues routine or efficiency of the construction will be compromised.

Once the final linear solver matrix is defined then we can set it to a scalable linear solver (KSP) environment. This is done using the following commands:

```
KSPCreate(MPI_COMM_WORLD, &K);
KSPSetOperators(K, A, A);
KSPSetType(K, KSPGMRES);
KSPSetFromOptions(K);
KSPSetUp(K);
KSPSetTolerances(K, RTol, ATol, PETSC_DEFAULT, PETSC_DEFAULT);
```

By default, the solvers in Tempest are set to use the generalized minimal residual method (GMRES) method, which has been defined via the KSPSetType command above. This method appears to offer stable solutions to a wide range of problems. However, a huge benefit to PETSc is that by calling the KSPSetFromOptions function we can fix any of the solver options at the command line, without having to recompile any portion of the code. An example call could then take the form:

```
mpirun -n N ./Tempest ./Prefs/my_pref -ksp_type gmres -pc_type bjacobi
```

The values RTol and ATol define the relative and absolute tolerance of the solver and are used to define the convergence condition [Balay et al., 2015]:

$$\|\mathbf{r}\|_2 < \max(\text{Rtol}\|\mathbf{b}\|_2, \text{Atol}),$$

where $\mathbf{r} = \mathcal{A}\mathbf{x} - \mathbf{b}$. In Tempest, these can be set in the particular preference file associated with the test we are considering. To facilitate quick convergence the iterative solver is given the initial guess defined by the solution at the previous stage of the time-stepping approach. It is possible that the solution can converge to this solution instead of an updated solution

but this has never occurred in our applications using at least the default tolerances of 1e-8 set in Tempest.

PETSc also uses preconditioning techniques with the default parallel preconditioner being the block jacobi method, which uses an incomplete LU factorization on each of the individual processors. Another great option is available via the external package HYPRE. This preconditioner offers a multi-grid approach which appears to offer convergence in a smaller number of steps but is slower per step. Generally the time for convergence using either HYPRE or default preconditioning appears to be nearly the same. As HYPRE has not shown to have any particular advantage over the default preconditioner the default preconditioner is used for most application. HYPRE does however appear to give more stable results in the case of the QG model.

3.3.8 Density Equation

At each sub-stage of the SSPRK3 time-stepping method we first update the velocity components as described previously. The density is then updated using an FV approach. The FV formulation is effective at conserving scalar quantities such as fluxes and tracers, and therefore ensures that mass is conserved. Under the NH or HY model we have the density update equation:

$$\frac{D}{Dt}(\rho(z) + \rho') = 0.$$

By including the stratification term in this way we can apply our FV approach to $\rho(z) + \rho'$ and ensure that the treatment of the sum of these terms is conservative. If we instead include it on the right hand side then we could potentially have a source or sink of mass caused from numerical errors in interpolating w to the ρ locations. We can write the above as

$$\frac{\partial}{\partial t}(\rho(z) + \rho') = -\nabla \cdot [(\rho(z) + \rho') \mathbf{u}],$$

where we have used incompressibility.

If we now include $\rho(z)$ implicitly within ρ' we can write this compactly as:

$$\frac{\partial \rho'}{\partial t} = -\nabla \cdot (\rho' \mathbf{u}),$$

If we take the cell average of this equation we get an FV update equation:

$$\frac{1}{V_{i,j,k}} \int_V \frac{\partial \rho'}{\partial t} dV = -\frac{1}{V_{i,j,k}} \int_V \nabla \cdot (\rho' \mathbf{u}) dS.$$

Rewriting this in terms of a surface integral, by application of the divergence theorem, and using the definition of a cell averaged value we get

$$\frac{\partial \bar{\rho}'}{\partial t} = -\frac{1}{V_{i,j,k}} \int_{\partial V} (\rho' \mathbf{u}) \cdot \hat{\mathbf{n}} dS,$$

where the overbar denotes a volume averaged quantity:

$$\bar{\rho}' = \frac{1}{V_{i,j,k}} \int_V \rho' dV.$$

This can be discretized in time as

$$\bar{\rho}'^{n+1} = \bar{\rho}'^n - \frac{\Delta t}{V_{i,j,k}} \int_{\partial V} (\rho' \mathbf{u})^n \cdot \hat{\mathbf{n}} dS,$$

where $\bar{\rho}'^{n+1}$ refers to our cell-averaged density values at the time at the end of the sub-stage update. Any source or sink terms, $S_{\rho'}$, can be added to this as a non-conservative volume integrated quantity:

$$\bar{\rho}'^{n+1} = \bar{\rho}'^n - \frac{\Delta t}{V_{i,j,k}} \int_{\partial V} (\rho' \mathbf{u})^n \cdot \hat{\mathbf{n}} dS + \frac{\Delta t}{V_{i,j,k}} \int_V S_{\rho'} dV. \quad (3.10)$$

For example, when we are considering an influx of cold dense fluid applied as a surface forcing we would represent this as a positive $S_{\rho'}$, thereby increasing the mass of the system. Following the update of ρ' to ρ'^{n+1} we redefine ρ' by removing the $\rho(z)$ terms, as these are constant in time.

3.3.9 WENO

In the absence of sources or sinks equation 3.10 can be written in terms of flux functions as

$$\bar{\rho}'^{n+1} = \bar{\rho}'^n - \frac{\Delta t}{V_{i,j,k}} \left[\hat{F}_{i+1/2} - \hat{F}_{i-1/2} + \hat{G}_{j+1/2} - \hat{G}_{j-1/2} + \hat{H}_{k+1/2} - \hat{H}_{k-1/2} \right],$$

where the flux functions are defined as follows:

$$\begin{aligned} F_{i\pm 1/2} &= \int_{\partial A_{i\pm 1/2}^x} (\rho' u) dS, \\ G_{j\pm 1/2} &= \int_{\partial A_{j\pm 1/2}^y} (\rho' v) dS, \\ H_{k\pm 1/2} &= \int_{\partial A_{k\pm 1/2}^z} (\rho' \mathbf{u} \cdot \hat{\mathbf{n}}_z) dS. \end{aligned} \quad (3.11)$$

The interface areas are denoted A^x, A^y, A^z in the x, y, z directions respectively and dS is a differential area element.

To interpolate the density values to the edge we predominantly use the WENO scheme. To obtain formal fourth order accuracy in this FV formula the area integrals of 3.11 must be approximated using a tensor product of quadrature rules. Using a WENO approximation we can evaluate the point value of these flux functions in one of two ways. Both result in multiple WENO approximations required per cell. In the first approach, the general one-dimensional WENO approach is applied iteratively and the evaluations are made along Gaussian quadrature points in each direction. The alternative is to form a large multidimensional function and apply the same methodology as that of the one-dimensional WENO scheme. These approaches are explained in great detail for the two-dimensional case in [Shi et al., 2002].

In general, higher than second order accurate FV methods are not usually applied in practice [Shu, 1991] due to the computational expense associated with their very large stencils. In [Wang, 2002], they quote that a fourth order FV scheme in three dimensions is estimated to require a stencil size of at least 120. However, there are more recent approaches that apply the deconvolution and convolution principles which could offer more computationally feasible approaches [Buchmüller and Helzel, 2014]. On the other hand, FD approximations can be applied in multiple dimensions by simply applying a dimensional splitting approach [Jiang and Shu, 1996], however FD methods need not be conservative [Mueller, 2015].

In our particular model, we have considered the second order FV approach, thereby allowing us to apply the one-dimensional WENO approach in a dimensional splitting framework. Under this approach, we simply approximate the cell interface averaged values via the midpoint rule. The FV approach is beneficial for the mass conservation equation since, by design, the method is conservative; the outflux from one cell is the influx to another. Even though the FV approach used in Tempest is theoretically second order accurate using higher order interpolation methods has been shown to yield more accurate results numerically and physically by [Ullrich and Jablonowski, 2012]. Furthermore, it has been shown that low order advection schemes can produce significant errors in vertical advective terms in the density equation [Allen et al., 2003].

Due to the grid cell staggering we use a similar approach to the one used in the nonlinear advection. The cell-centered density value is WENO interpolated to the interfaces and then multiplied by the collocated velocity components to form the flux. The choice of whether or not to use the positive or negative flux is based upon the sign of the velocity component located at the edge. In the vertical we consider the sign of the velocity component multiplied

by the associated normal vector component as shown in figure 3.5 and explained in section 3.3.3. Under this approach it is possible to use a positive flux at one interface and a negative flux at the opposite, which leads to at best fourth order accuracy as shown in the nonlinear advection section 3.3.3.

To establish the ideas behind the WENO scheme we apply the procedure to the x component of the above equation. Written in a semi-discrete form we consider the evolution equation defined as follows:

$$\frac{\partial \rho'}{\partial t} = -\frac{1}{V_{i,j,k}} \left[(\hat{\rho}' u A^x)_{i+1/2} - (\hat{\rho}' u A^x)_{i-1/2} \right].$$

This is written using the midpoint rule.

Instead of applying WENO to the flux $\mathbf{F}(\rho') = \rho' \mathbf{u}$ we use the underlying interpolation procedure, which is at the heart of the WENO method [Shu, 2009]. Applying the interpolation procedure gives us approximations to ρ' at the boundaries. The ENO and WENO methods are both based on the underlying approximation polynomials using the candidate stencils surrounding a specific cell. In the case of a FV formulation, where cell averaged values are known, we consider the candidate polynomial to have the same average value over each of the cells in the stencil. This leads to a set of equations from which we can determine the coefficients of the underlying polynomial.

The ENO method uses these polynomials to determine the smoothest stencil choice as based on analysis of the primitive of the approximation polynomials. Where ENO uses one smoothest stencil, WENO uses all the candidate stencils to form a weighted approximation to the edge value using weights based on the smoothness of the approximation polynomials [Zhang et al., 2011]. By comparison, the upwinding method of the equivalent order of the WENO scheme uses the union of the intervals comprising all the stencils. However, the upwind scheme does allow for oscillations in the solution.

We know the cell averaged values:

$$\rho'_{i,j,k} = \frac{1}{V_{i,j,k}} \int_{V_{i,j,k}} \rho' dV = \frac{1}{\Delta x} \int_{-\Delta x/2}^{\Delta x/2} \left(\frac{1}{\tilde{A}_x} \int_{\tilde{A}_x} \rho' dS \right) dx,$$

where $\tilde{A}_x = V_{i,j,k}/\Delta x$. Therefore, we see that what we are approximating is the average density over the x face. The reason we get second order accuracy when we multiply the resulting approximation by the interface area is that the adjacent cells in the y and z directions also determine approximations to the x face average value. Therefore, we are applying the midpoint rule at the interface. The idea is shown in figure 3.6, where the curve at the x face is the true distribution of the x face average densities.




Figure 3.6: WENO interpolation applied in the x direction

To demonstrate the WENO procedure we consider the interpolation of the function $r(x)$ which is defined as

$$r(x) = \frac{1}{\tilde{A}_x} \int_{\tilde{A}_x} \rho dS.$$

We have suppressed our y and z coordinate because we are applying the midpoint rule. For this function we know the average quantities in the x direction, the cell averaged density values. Therefore, we consider a WENO approximation at cell I_i . In this case there are three candidate stencils which contain three cells each as shown in figure 3.7. It should be noted that the WENO procedure applied to the cell-averaged values in the stencils shown will give the negative flux at the interface located at $i + 1/2$ and the positive flux at the interface located at $i - 1/2$. To determine the positive and negative fluxes of interface $i + 1/2$ and $i - 1/2$ respectively we can apply the same procedure to the stencils centered around the adjacent cells.

Dropping the j, k subscripts on ρ' , we know the average value of the function for each of the intervals shown in figure 3.7 defined as:

$$\rho'_n = \frac{1}{\Delta x} \int_{I_n} r(x) dx, \quad n = i - 2, \dots, i + 2.$$

From these averages, we can construct a polynomial such that the average value of the polynomial over each interval is equivalent to these values. For each of the stencils




Figure 3.7: Candidate stencils for the approximation of fluxes in cell I_i

$\{S_{i-1}, S_i, S_{i+1}\}$ we can construct a quadratic to fit the samples. Integrating this quadratic over each of the intervals we can determine the unknown coefficients of the quadratic function. We can then use these polynomials to evaluate an approximation at each of the endpoints of the interval. If we consider the evaluation at the right end point we get the following approximations given the three stencils:

$$\begin{aligned}\hat{r}_{i+1/2}^- \Big|_{S_{i-1}} &= \frac{1}{3}\rho'_{i-2} - \frac{7}{6}\rho'_{i-1} + \frac{11}{6}\rho'_i, \\ \hat{r}_{i+1/2}^- \Big|_{S_i} &= -\frac{1}{6}\rho'_{i-1} + \frac{5}{6}\rho'_i + \frac{1}{3}\rho'_{i+1}, \\ \hat{r}_{i+1/2}^- \Big|_{S_{i+1}} &= \frac{1}{3}\rho'_i + \frac{5}{6}\rho'_{i+1} - \frac{1}{6}\rho'_{i+2}.\end{aligned}$$

In the above $\hat{r}_{i+1/2}^- \Big|_{S_i}$ represents the standard third order upwind-biased flux. If, on the other hand, we use a union of all the intervals and form a fourth degree polynomial then we get the approximation:

$$\hat{r}_{i+1/2}^- \Big|_S = \frac{1}{30}\rho'_{i-2} - \frac{13}{60}\rho'_{i-1} + \frac{47}{60}\rho'_i + \frac{9}{20}\rho'_{i+1} - \frac{1}{20}\rho'_{i+2},$$

where $S = S_{i-1} \cup S_i \cup S_{i+1}$. This represents the standard fifth order upwind-biased flux.

Forming a linear combination of the above three point stencil approximations we observe that:

$$\hat{r}_{i+1/2}^- \Big|_S = \frac{1}{10} \hat{r}_{i+1/2}^- \Big|_{S_{i-1}} + \frac{3}{10} \hat{r}_{i+1/2}^- \Big|_{S_i} + \frac{6}{10} \hat{r}_{i+1/2}^- \Big|_{S_{i+1}}.$$

The weights $1/10, 3/10, 6/10$ are called the optimal weights of the WENO approximation and are usually denoted C_k [Jiang and Shu, 1996]. In regions where ρ' are particularly smooth the WENO weights will tend to these values. It should be noted that each of the three point stencil approximations give third order accurate approximations to $r(x)$ and the five point stencil approximation gives a fifth order accurate approximation to $r(x)$ at the boundary.

To show that the large stencil yields a fifth order accurate approximation to $r(x)$ at the boundary we summarize the analysis of [Shu, 1991, Shu, 1998]. From our $r(x)$ values we can define a primitive function as:

$$R(x) \equiv \int_{-\infty}^x r(x) dx.$$

From the $r(x)$ values we find the edge values of this primitive function. From these edge values, if we find an approximation to $R(x)$ as $\hat{R}(x)$ we see that on the five point stencil we have six potential edge values. By standard FD approximations we have that $\hat{R}'(x)$ is equal to $R'(x) + O(\Delta x^5)$. Taking the average value of $\hat{R}'(x)$ over the interval I_i we see that:

$$\begin{aligned} \frac{1}{\Delta x} \int_{I_i} \hat{R}'(x) dx &= \frac{1}{\Delta x} [\hat{R}(x_{i+1/2}) - \hat{R}(x_{i-1/2})] \\ &= \frac{1}{\Delta x} [R(x_{i+1/2}) - R(x_{i-1/2})] \\ &= \frac{1}{\Delta x} \int_{I_i} r(x) dx = \rho'. \end{aligned}$$

Therefore, we see that over the interval I_i that $\hat{R}'(x)$ approximates $R'(x) = r(x)$ to fifth order accuracy and $\hat{R}'(x)$ is just our $\hat{r}(x)$ approximation polynomial.

The idea behind the WENO scheme is to form a weighted combination of the smaller stencil approximations to determine an approximation that is high-order but does not allow oscillations in the solutions. Therefore, we determine weights ω_k such that:

$$\hat{r}_{i+1/2}^- = \sum_{k=0}^2 \omega_k \hat{r}_{i+1/2}^- \Big|_{S_{i-1+k}}. \quad (3.12)$$

One such option is to choose $\omega_k = C_k$ but this will just give the fifth order upwind method above. Combining the fifth order upwind method with 3.12 we can write this in an alternate form as

$$\hat{r}_{i+1/2}^- = \hat{r}_{i+1/2}^- \Big|_S + \sum_{k=0}^2 (\omega_k - C_k) \hat{r}_{i+1/2}^- \Big|_{S_{i-1+k}}.$$

From this formulation, it was concluded by [Jiang and Shu, 1996] that the weights need to be chosen such that $\omega_k = C_k + O(\Delta x^2)$.

Extensive numerical testing by [Jiang and Shu, 1996, Liu et al., 1994] led to the WENO weights being defined as:

$$\omega_k = \frac{\alpha_k}{\sum_{k=0}^2 \alpha_k}, \quad k = 0, 1, 2, \quad (3.13)$$

where α_k are defined as:

$$\alpha_k = \frac{C_k}{(\varepsilon + \beta_k)^2}, \quad k = 0, 1, 2. \quad (3.14)$$

with a suggested value of $\varepsilon = 1\text{e-}6$ used to avoid division by zero [Jiang and Shu, 1996].

The β_k in the above formula are called smoothness indicators and are implemented in Tempest in accordance with that proposed in [Jiang and Shu, 1996]. The definition of these smoothness indicators are:

$$\beta_k = \sum_{l=1}^2 \int_{I_i} \Delta x^{2l-1} \left[\hat{r}_{i+1/2}^- \Big|_{S_{i-1+k}} \right]^2 dx, \quad k = 0, 1, 2. \quad (3.15)$$

This represents the sum of the L_2 -norms of all the derivatives of the particular interpolation polynomial $\hat{\rho}_{i+1/2}^-$ over the interval of interest, I_i [Jiang and Shu, 1996]. Given the underlying polynomials we can find the discrete β_k and apply these formulas at each time step.

On the other hand, if the underlying function is not smooth it has been shown by [Henrick et al., 2005, Borges et al., 2008] that ω_k must be chosen such that $\omega_k = C_k + O(\Delta x^3)$. A computationally efficient way to choose the coefficients with this level of error was presented by [Borges et al., 2008]. The authors define this in a modified WENO method they call the WENO-Z scheme. This scheme is comparable to the above formulation but they define a value:

$$\tau_5 = |\beta_0 - \beta_2|.$$

They then define new smoothness indicators β_k^z as:

$$\beta_k^z = \left(\frac{\beta_k + \varepsilon}{\beta_k + \tau_5 + \varepsilon} \right), \quad k = 0, 1, 2,$$

and the weights:

$$\alpha_k^z = \frac{C_k}{\beta_k^z}.$$

This method is not currently used in Tempest but is a simple extension.

WENO FV Summary

The WENO approach is summarized as:

1. Compute the edge approximations under the various stencils, $\hat{r}_{i+1/2}^- \Big|_{S_{i-1+k}}$.
2. Store the optimal weight coefficients, C_k .
3. Compute the smoothness parameters, β_k , according to equation 3.15.
4. Compute the α_k values using equation 3.14 given the smoothness parameters and optimal weight coefficients.
5. Normalize the α_k values so that their sum is one, thereby defining ω_k as in equation 3.13.
6. With these weights compute the non-oscillatory approximation to the edge value using equation 3.12.

3.3.10 Zhang limiting

Although the WENO scheme discussed in the previous section produces high-order accurate approximations in smooth regions while producing non-oscillatory results around discontinuities in the solution, there is no guarantee that the updated values will preserve positivity. For physical reasons, we often want to preserve the positivity of tracers since negative concentrations are unphysical [Kent et al., 2012]. This is not a major concern for velocity or variations in density since both fields can take on negative values.

The approach that we will be using is an approach that is a modification of a limiting scheme originally introduced by Liu and Osher [Liu and Osher, 1996] and later modified by Zhang and Shu [Zhang and Shu, 2010, Zhang et al., 2011]. The approach is modified to limit only the minimum value strictly enforcing positivity. The maximum is not forced and it is possible for new extrema to be produced. The method is equivalent to that used in the recent paper [Du et al., 2015]. In all cases the standard underlying approach is to find a θ such that the modified edge approximation

$$\tilde{r}_{i+1/2,j,k} = \theta_{i,j,k} (\hat{r}_{i+1/2,j,k} - \rho'_{i,j,k}) + \rho'_{i,j,k}, \quad (3.16)$$

will preserve positivity. The $\hat{r}_{i+1/2,j,k}$ represents the approximation prior to being limited. These edge approximations can be found via the WENO method of the previous section, or any other approximation technique.

Typically these edge approximations will represent the negative or positive flux and both should be limited according to this procedure. In this way, the $\theta_{i,j,k}$ associated with the downwind cell adjacent to the approximation location should be used to limit the particular flux. The $\theta_{i,j,k}$ is the limiting parameter and is defined in such a way as to preserve only positivity. The formula for this value is used according to that in [Du et al., 2015]:

$$\theta_{i,j,k} = \min \left\{ \left| \frac{\rho'_{i,j,k}}{\rho'_{i,j,k} - m_{i,j,k}} \right|, 1 \right\},$$

where $m_{i,j,k}$ represents the minimum value of the underlying approximation functions over cell $I_{i,j,k}$.

However, we do not know the value of $m_{i,j,k}$. To this end, [Zhang and Shu, 2010] demonstrated that the minimum value of the function could be determined by considering a fourth degree Hermite type polynomial such that the average of the polynomial over cells I_{i-1}, I_i, I_{i+1} agree with $\rho'_{i-1}, \rho'_i, \rho'_{i+1}$ and the point values of this function equals the interpolated values at the boundary. The extension to multiple dimensions is to take tensor products of these polynomials.

However, it was later shown by [Zhang et al., 2011] that a much simpler approach is to assume a Hermite type polynomial, $P(x, y, z)$ and then use the mean value theorem to devise a function value such that $m_{i,j,k}$ can be evaluated as the minimum of this value and the approximated endpoint values. The benefit to this approach is that it extends trivially to multiple dimensions. This value, in three-dimensions, under the midpoint rule is:

$$P_{i,j,k} = \frac{6}{5} \left[\rho'_{i,j,k} - \frac{1}{36} \left(\hat{r}_{i+1/2,j,k}^- + \hat{r}_{i-1/2,j,k}^+ + \hat{r}_{i,j+1/2,k}^- + \hat{r}_{i,j-1/2,k}^+ + \hat{r}_{i,j,k+1/2}^- + \hat{r}_{i,j,k-1/2}^+ \right) \right],$$

where the $\hat{r}_{i+1/2,j,k}^-, \hat{r}_{i-1/2,j,k}^+, \hat{r}_{i,j+1/2,k}^-, \hat{r}_{i,j-1/2,k}^+, \hat{r}_{i,j,k+1/2}^-, \hat{r}_{i,j,k-1/2}^+$ represent the WENO approximations to the edge averaged values in each of the coordinate directions. We can then take the minimum value over the cell as:

$$m_{ijk} = \min \left\{ P_{i,j,k}, \hat{r}_{i+1/2,j,k}^-, \hat{r}_{i-1/2,j,k}^+, \hat{r}_{i,j+1/2,k}^-, \hat{r}_{i,j-1/2,k}^+, \hat{r}_{i,j,k+1/2}^-, \hat{r}_{i,j,k-1/2}^+ \right\}.$$

Each of these WENO approximations can then be limited via the formula 3.16. In the strictest case, the CFL condition imposed by this limiting procedure is that the Courant number be no greater than $1/(12n)$ where n is the number of dimensions.

3.4 QG Model

The QG model is much less complicated than the NH and HY primitive equation models previously discussed. The computation of updated fields in the QG model takes only two steps. First, we advect the potential vorticity forward in time, similar to the density update equation mentioned previously. To ensure that the potential vorticity is conserved we apply an FV approach. However, the potential vorticity is assumed to be point values and so this FV formulation is second order accurate. The QGPV equation was derived in Chapter 2 and is

$$\frac{D_g}{Dt} (q + \beta y) = 0,$$

where the QGPV is defined as

$$q = \nabla_h^2 \psi + \frac{\partial}{\partial z} \left(\frac{f_0^2}{N^2} \frac{\partial \psi}{\partial z} \right).$$

We consider the special case where the boundary conditions are periodic in the x direction and either periodic or free-slip boundary conditions are used in the y direction. Combining the no normal flow condition and the free-slip condition we have the boundary conditions on the PV field at the y boundaries are homogeneous Dirichlet:

$$q = 0, \quad y = 0, L_y.$$

This is a specific choice and other boundary conditions could be chosen; the chosen boundary conditions need not be true in general. Boundary conditions are used when we are extrapolating our potential vorticity field to the ghost cells at the domain boundaries. In the advection, we do not require vertical boundary conditions since the advection operators are only applied in the horizontal.

After the vorticity is updated the new geostrophically balanced velocity components are updated using a two-step procedure. The first step is to solve for the streamfunction using the definition of our QGPV:

$$q = \nabla_h^2 \psi + \frac{\partial}{\partial z} \left(\frac{f_0^2}{N^2} \frac{\partial \psi}{\partial z} \right).$$

To do this we define the Laplacian-like operator

$$D = \nabla_h^2 + \frac{f_0^2}{N^2} \frac{\partial^2}{\partial z^2}, \tag{3.17}$$

assuming that N is constant. We then have the linear problem $\psi = D^{-1}q$, which is solved via PETSc's KSP environment. This problem is well-posed under appropriate boundary conditions. To determine these boundary conditions, we know that the vertical velocity must vanish at the top and bottom of the domain and from the density equation this gives us $D\rho'/Dt = 0$. Therefore, if we assume that the buoyancy is zero at the top and bottom we have that $\psi_z = 0$ and by virtue of $D\rho'/Dt = 0$ we have the following boundary condition on the streamfunction:

$$\psi_z = 0, \quad z = 0, H.$$

In terms of boundary conditions in the y direction, we assume homogeneous Dirichlet conditions on the perturbations in the streamfunction, which are the so-called natural boundary conditions. In terms of the total PV we must have that $v = \psi_x = 0$ implying that the streamfunction is constant at each vertical level but not necessarily the same constant. Coupled with the homogeneous Dirichlet conditions on the perturbations, we apply non-homogeneous Dirichlet conditions on the streamfunction:

$$\psi = C(z), y = 0 \quad \psi = D(z), y = L_y,$$

recognizing the fact that our constant streamfunction at the North and South walls are not necessarily equal and can vary with depth.

Upon solving equation 3.17 for ψ we can then determine our geostrophic velocity components from this streamfunction using the relationships $u_g = -\psi_y, v_g = \psi_x$. The streamfunction is defined collocated with q , at cell centers. The x and y derivatives are then staggered at the location of u and v respectively. These are then interpolated to the location of v and u .

The differentiation followed by interpolation is the equivalent to defining a derivative over the staggered mesh, which can be verified via two-dimensional Taylor series expansions. These operators can be chosen to be either second or fourth order accurate and is set in the preference file. The computation of these derivatives are performed for every sub-stage of every time step so in order to perform them efficiently we store PETSc matrices `matDxPsi`, `matDyPsi` and use PETSc's `MatMult` function to compute the derivatives.

In terms of time-stepping, the QG model also uses the SSPRK3 time-stepping scheme. This allows consistency with the approaches used in the HY and NH models. As well, WENO5 is the default choice for updating the vorticity equation and SSPRK3 and WENO5 are ubiquitous.

3.5 Linear Stability Analysis

An unstable jet will typically undergo exponential growth. During the early periods of exponential growth, in which the perturbations are small relative to the background state, linear theory can accurately predict the growth rates. LSA gives us a means of computing the linear stability characteristics for a flow over a wide range of parameters relatively efficiently. We consider the linear stability of various jet configurations and our model equations reflect this setup.

3.5.1 Linear Stability of the NH Model

We consider a model with bottom topography that varies in the across channel direction and a background jet that varies in both the across channel and vertical, y and z , directions. These types of motions can be described using the NH model. To derive the linearized equations we decompose our variables into a temporally constant background state and small amplitude perturbations around this state that are permitted to grow in time. We consider background states that can vary with depth and across the channel and perturbations around this background state defined in terms of the computational coordinates X, Y, Z . The perturbed fields then take the form:

$$\begin{aligned}u &= U(y, z) + u'(X, Y, Z, t), \\v &= 0 + v'(X, Y, Z, t), \\w &= 0 + w'(X, Y, Z, t), \\\rho &= R(y, z) + \rho'(X, Y, Z, t), \\p &= P(y, z) + p'(X, Y, Z, t).\end{aligned}$$

In the above set of equations we define the background state in terms of the physical coordinates and take exact derivatives of these terms. We apply the terrain-following transformation equations to the perturbations and so defined them in terms of the computational coordinates.

If we assume that the perturbations are small relative to the background state then we can neglect terms involving the multiplication of two or more perturbation terms. This process is called linearization. Upon substituting the expressions for the perturbed fields into our above primitive equations, linearizing, and canceling of the background balanced

states we get the linear perturbation equations:

$$\begin{aligned}
u'_t + Uu'_X + (U_y - f)v' + (f^* + U_z)w' &= -\frac{1}{\rho_0}p'_X, \\
v'_t + Uv'_X + fu' &= -\frac{1}{\rho_0}(p'_Y + Z_y p'_Z), \\
w'_t + Uw'_X - f^*u' &= -\frac{1}{\rho_0}(Z_z p'_Z) - g\frac{\rho'}{\rho_0}, \\
\rho'_t + U\rho'_X + v'R_y + w'R_z &= \frac{\rho_0}{g}N^2w', \\
u'_X + (v'_Y + Z_y v'_Z) + (Z_z w'_Z) &= 0.
\end{aligned}$$

Since the channel is periodic in the x (X) direction, we can assume the perturbations will have a normal decomposition in x and in time. This normal mode solution has the form:

$$\left[u', v', w', \rho', p' \right] = \Re \left\{ \left[\hat{u}, \hat{v}, \hat{w}, \hat{\rho}, \hat{p} \right] (Y, Z) e^{i(kX - \omega t)} \right\},$$

and upon substitution into the above linear equations, we get the following set of normal mode equations:

$$\begin{aligned}
i\omega\hat{u} &= ikU\hat{u} + (U_y - f)\hat{v} + (f^* + U_z)\hat{w} + \frac{ik}{\rho_0}\hat{p}, \\
i\omega\hat{v} &= ikU\hat{v} + f\hat{u} + \frac{1}{\rho_0}(\hat{p}_Y + Z_y\hat{p}_Z), \\
i\omega\hat{w} &= ikU\hat{w} - f^*\hat{u} + g\frac{\hat{\rho}}{\rho_0} + \frac{1}{\rho_0}(Z_z\hat{p}_Z), \\
i\omega\hat{\rho} &= ikU\hat{\rho} + R_y\hat{v} + \left(R_z - \frac{\rho_0}{g}N^2 \right)\hat{w}, \\
0 &= ik\hat{u} + (\hat{v}_Y + Z_y\hat{v}_Z) + (Z_z\hat{w}_Z).
\end{aligned}$$

In matrix form this yields the generalized eigenvalue problem:

$$Ax = B\lambda x,$$

where

$$A = \begin{bmatrix} ikU & U_y - f & f^* + U_z & 0 & \frac{1}{\rho_0} ik \\ f & ikU & 0 & 0 & \frac{1}{\rho_0} (\partial_Y + Z_y \partial_Z) \\ -f^* & 0 & ikU & \frac{g}{\rho_0} & \frac{1}{\rho_0} (Z_z \partial_Z) \\ 0 & R_y & R_z - \frac{\rho_0}{g} N^2 & ikU & 0 \\ ik & \partial_Y + Z_y \partial_Z & Z_z \partial_Z & 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad (3.18)$$

$x = [\hat{u}, \hat{v}, \hat{w}, \hat{\rho}, \hat{p}]^T$ and $\lambda = i\omega$. For a given perturbation field α we have that:

$$\begin{aligned} \alpha(X, Y, Z, t) &= \Re \{ \hat{\alpha} e^{i(kX - \omega t)} \} \\ &= [\hat{\alpha}_r(Y, Z) \cos(kX) - \hat{\alpha}_i(Y, Z) \sin(kX)] e^{\omega_i t} \cdot \Re \{ e^{-i\omega_r t} \} \end{aligned}$$

where subscript r and i denote the real and imaginary parts of the variable. Therefore, the eigenvalue with the largest growth rate corresponds to the eigenvalue with the largest ω_i value. In the above eigenvalue problem we are solving for the eigenvalues $\lambda = i\omega = -\omega_i + i\omega_r$. Therefore, to maximize ω_i we seek the eigenvalues λ such that $\lambda_r = -\omega_i$ is a minimum. In other words, we are seeking the eigenvalue with the smallest real part. The corresponding eigenvector will give the spatial structure of the mode, computed as $\hat{\alpha}_r(Y, Z) \cos(kX) - \hat{\alpha}_i(Y, Z) \sin(kX)$.

Consistent with our nonlinear model, the variables are discretized using a standard C grid. This means that for a given x slice, the $\hat{u}, \hat{\rho}$, and \hat{p} are all located at cell centers, while the \hat{v} and \hat{w} are located at the North and top edges of the cells. If we consider a grid that is sub-divided into $N_y \times N_z$ cells then we are considering $N_y \times N_z$ $\hat{u}, \hat{\rho}, \hat{p}$ values, $(N_y + 1) \times N_z$ \hat{v} values, and $N_y \times (N_z + 1)$ \hat{w} values.

However, accounting for no normal flow conditions at the top/bottom and North/South boundaries of the domain we can remove the top/bottom rows of \hat{w} and the North/South rows of \hat{v} from consideration since these represent zero rows. That is, for each of the second row sub-matrices we will remove the rows that pertain to North/South boundaries and for each of the third row sub-matrices we will remove the rows that pertain to the top/bottom boundaries. Similarly, for each of the second (third) column sub-matrices we will remove the columns pertaining to the North/South (top/bottom) boundaries.

We must consider the interpolation of the fields in sub-matrices $A_{1,2}, A_{1,3}, A_{2,1}, A_{3,1}, A_{3,4}, A_{4,2}$, and $A_{4,3}$ to account for the grid staggering. Note that when referencing our

sub-matrices we consider row and column indexing beginning at one. Furthermore, in sub-matrices $A_{5,2}$ and $A_{2,5}$ interpolation is also required when computing the Z components of the derivatives in the terrain-following coordinate system.

Using standard FD interpolations the no normal flow boundary condition is accounted for trivially (by setting the required values to 0). The exception to this is in sub-matrix $A_{3,4}$ for which we need to interpolate ρ to the w locations under the no flux boundary condition. We have only considered w for the interior cell tops (excluding the top and bottom boundary values since $w = 0$ at these locations) so we must consider fourth order FD interpolations at the first cell top and last cell bottom to include the information about the boundary.

Using a direct eigenvalue solver for this problem requires a lot of computational memory in two or higher dimensions and therefore one is restricted to rather coarse grids [Ménèsguen et al., 2012]. Instead, iterative techniques can be used to exploit the sparsity of the matrix, allow finer resolutions to be used. In an iterative approach we are restricted to finding a subset of modes but we can often find the most unstable modes, which are the ones of greatest physical interest for our purposes. In order to accelerate convergence, we use our coarse approximations as an initial guess in our iterative solver.

Therefore, given a particular test problem, we first use Matlab’s direct solver to compute the largest growth rates across the entire wavenumber spectrum for a rather coarse grid. For efficiency, we use a $N_y \times N_z = 10 \times 10$ grid for this direct solver. A summary of the wavenumbers and growth rates are then outputted and read into our SLEPc code. The SLEPc routines then compute more accurate approximations to the eigenvalues and eigenvectors of the eigenvalue problem across the entire wavenumber spectrum.

SLEPc

The Scalable Library for Eigenvalue Problem Computation (SLEPc) [Hernandez et al., 2005] is open-source software used for the computation of eigenvalue problems of this nature. It accepts either generalized or non-generalized problems. It was built as an extension to PETSc. SLEPc can also interface with a number of external software packages, such as MUMPS and SuperLU, which are particularly useful in the computation of solutions to eigenvalue problems that require a spectral transformation.

SLEPc natively does not handle zeros rows in the B matrix very well (it replaces them with the equivalent rows of the identity matrix), so we need to use external software packages to compute solutions to the above eigenvalue problems. The zero rows in B lead to infinite eigenvalues, which in the context of an iterative solver can pose issues for convergence. Both MUMPS and SuperLU_dist (the MPI parallelized version of SuperLU) are valid options, however MUMPS appears to give better convergence results. From

our experience, SuperLU_dist sometimes predicts spurious modes. We do have to be careful when considering the MUMPS solver, since it has been argued that it is not as versatile as the SuperLU_dist solver. However, we can allot a greater memory allocation to the MUMPS factorization array to solve errors that arose for certain wavenumbers [Amestoy et al., 2000, Amestoy et al., 2001].

For the problems we have considered, given an initial guess σ we treat this problem using a spectral transformation technique that is part of the SLEPc package. In this technique the original eigenvalue problem $Ax = B\lambda x$ is transformed to $Fx = \nu x$ where F is defined as $(A - \sigma B)^{-1}B$. The shift followed by the inversion is why this method is often called the shift and invert method [Bai et al., 2000]. The eigenvalues of the original problem can be solved from σ and ν as [Roman et al., 2015]

$$\lambda = \sigma + \frac{1}{\nu}.$$

The closer the initial guess σ is to the solution the more likely the method will converge rapidly.

Linear Stability Analysis of the QG Model

As an alternative we can describe our flows using the QG model. Our jet profile is oriented so that the flow is West to East so we can define a background state in terms of a streamfunction dependent on y . If we consider a baroclinic jet then the streamfunction will depend on z as well. Therefore, in the more general case we have a jet with background:

$$\psi = \Psi(y, z), \quad u = U(y, z) = -\Psi_y, \quad q = Q(y, z) = \Psi_{yy} + \frac{f_0^2}{N^2}\Psi_{zz}$$

assuming constant stratification and where a subscript refers to partial differentiation with respect to the subscripted variable. Perturbations to this basic state are considered in the following form:

$$\begin{aligned} \psi &= \Psi(y, z) + \psi'(x, y, z), \\ q &= Q(y, z) + q'(x, y, z). \end{aligned}$$

If we now consider an f -plane approximation and substitute the above perturbed solutions into the QGPV equation 2.28 we get, upon linearization, the following equation:

$$q'_t - \Psi_y q'_x + \psi'_x Q_y = 0,$$

or, in terms of our background velocity $U(y, z)$,

$$q'_t + U q'_x + \psi'_x Q_y = 0. \tag{3.19}$$

We assume a normal mode solution with the following form:

$$[q', \psi'] = \Re \left\{ [\hat{q}, \hat{\psi}](y, z) e^{i(kx - \omega t)} \right\}.$$

Upon substitution into the linear equation 3.19, and cancellation of the exponential factor, we get the eigenvalue problem:

$$i\omega \hat{q} = ikU\hat{q} + ik\hat{\psi}Q_y.$$

We can use the QGPV relation to write

$$\hat{q} = (\partial_{yy} - k^2) \hat{\psi} + \frac{f_0^2}{N^2} \partial_{zz} \hat{\psi},$$

under the assumption that N is constant. Using this relationship, our normal mode equation becomes

$$\begin{aligned} i\omega \left[(\partial_{yy} - k^2) \hat{\psi} + \frac{f_0^2}{N^2} \partial_{zz} \hat{\psi} \right] &= ikU \left[(\partial_{yy} - k^2) \hat{\psi} + \frac{f_0^2}{N^2} \partial_{zz} \hat{\psi} \right] + ik\hat{\psi}Q_y, \\ \frac{\omega}{k} \left[(\partial_{yy} - k^2) + \frac{f_0^2}{N^2} \partial_{zz} \right] \hat{\psi} &= \left[U (\partial_{yy} - k^2) + U \frac{f_0^2}{N^2} \partial_{zz} + Q_y \right] \hat{\psi}, \end{aligned}$$

where we divided through by ik in the last step. Written in this way, the matrices have strictly real entries.

In matrix form, this can be written as the generalized eigenvalue problem:

$$A\hat{\psi} = B\lambda\hat{\psi},$$

where $\lambda = \frac{\omega}{k}$ and:

$$B = \left[(\partial_{yy} - k^2) + \frac{f_0^2}{N^2} \partial_{zz} \right], \quad A = \left[U (\partial_{yy} - k^2) + U \frac{f_0^2}{N^2} \partial_{zz} + Q_y \right] = [UB + Q_y]. \quad (3.20)$$

In terms of boundary conditions, we know that the vertical velocity must vanish at the top and bottom of the domain. Therefore, written in terms of the streamfunction, we have:

$$\hat{\psi}_z = 0, \quad z = 0, H.$$

At the horizontal y boundaries, we must have that the streamfunction is constant since we have that $\hat{v} = \hat{\psi}_x = 0$ and x is periodic. For simplicity we choose a constant value of zero. Assuming a periodicity in x in our normal mode solution we must have:

$$\hat{\psi} = 0, \quad y = 0, L_y$$

in order for the x derivatives to be 0 at these locations.

3.6 Computational Efficiency

3.6.1 Linear Stability Analysis

To test the efficiency of the linear stability code, we consider the case of a barotropic jet profile. We consider a range of twenty-one k values such that $k \in [0.5, 1.3]$ where the k values have been standardized by the width of the jet, set to 10 km. For this particular profile $k = 0.9$ was determined to be the most unstable k value so we test k values in a neighborhood about this particular value. We consider a grid with $N_y \times N_z = 512 \times 128$. This grid is finer than our reference grid which is $N_y \times N_z = 256 \times 64$ and this is to allow for more computation time to surpass. It is also for this reason that we consider a range of wavenumbers. If we considered a coarser grid or a subset of k values than it may be the case that more time will be spent in matrix assembly than computing the solutions to the eigenvalue problems. The results for this efficiency test, using the SuperLU_dist (SLUd) and MUMPS external packages are shown in table 3.1.

# cpus	SLUd run-time	SLUd RAM	MUMPS run-time	MUMPS RAM
1	1h12m20	3.8GB	11m11s	2.0GB
2	36m20	3.9GB	7m26s	2.6GB
4	38m25s	5.6GB	5m56s	3.4GB
8	$\gg 5$ h	6.9GB	4m22s	5.0GB
16	--	11.4GB	3m15s	8.1GB

Table 3.1: Comparing the parallel efficiency and memory usage of SuperLU_dist (SLUd) and MUMPS

We observe from these results that the SuperLU_dist package does not converge as quickly for our particular test problem, and MUMPS seems significantly more efficient. For this reason, we have decided to use MUMPS as the default external package. Furthermore, MUMPS appears to use significantly less memory resources. Though the MUMPS code does exhibit a small speed-up from four to sixteen cores, four cores have been used in our LSA calculations. The four core simulation also requires less than half the memory when compared to the sixteen core simulation.

3.6.2 MPI Splitting Efficiency

It is valuable to highlight how the MPI communications and, in particular, the MPI grid splitting can affect the efficiency of the parallelization. Since the numerical approximations

depend on adjacent cell information the boundary cells of a particular processors block of data must gather information from adjacent processors. The boundaries of the local block are then expanded to hold the information of these non-local cells. These bordering regions are called the halo regions and the particular elements of the halo region are called halo elements.

MPI processor splitting can be unidirectional or bidirectional as shown in figure 3.8. In the unidirectional case, the periodic halo elements are only transferred when the domain is set to periodic in the unsplit direction. Unidirectional splitting can be performed in either the x direction or y direction and is set by the `XProcs` (`YProcs`) option in the preference files. The Tempest model does not allow for variable amounts of data per processor so it is important that the number of x (y) processors divides the number of x (y) elements.

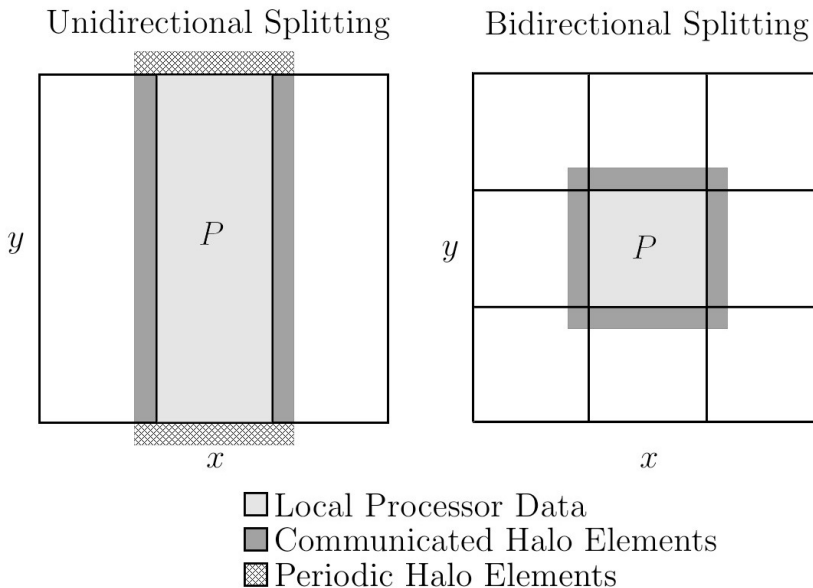


Figure 3.8: Example MPI grid splitting in one direction and two directions.

Under the simplest communication model the total time spent in message passing is determined by the number of messages passed times the latency plus the total number of elements to be processed [Rabenseifner, 2010]. If we assume the same number of processors and transfer directions, which is the case in general, then the measure of the efficiency of bidirectional splitting can be measured by the amount of halo elements per processor. The less halo elements involved in the communication the more efficient the splitting.

If we assume a periodic domain and consider a grid of $N_x \times N_y$ elements then if we split the domain into P processors in one direction only (without loss of generality we will consider the x direction) we have that the total exchange vector size will be determined by the area of the shaded regions in the left figure of figure 3.8. If the width of the halo region is h , then each of the two dark shaded regions has an area of $h N_y$, while each of the lightly shaded regions has an area of $h(N_x/P + 2h)$. Therefore, the total exchange vector size will be

$$T_1 = 2h N_y + 2h \left(\frac{N_x}{P} + 2h \right) = 2h \left(\frac{N_x}{P} + N_y \right) + 4h^2.$$

If we instead split the domain bidirectionally we will consider P_x and P_y processors in the x and y directions respectively. In this case, the total size can be determined by the total area of the shaded regions in the right figure of figure 3.8. The total exchange vector size is determined to be:

$$T_2 = 2h \left(\frac{N_x}{P_x} + \frac{N_y}{P_y} \right) + 4h^2.$$

To fairly compare the two types of splitting we consider the case for which $P = P_x P_y$. The bidirectional splitting is at least as efficient as the unidirectional splitting when $T_2 \leq T_1$. If we make the assumption that $N_x = N_y$ then this holds for:

$$P_x P_y - P_y - P_x + 1 = (P_x - 1)(P_y - 1) \geq 0,$$

which holds for all $P_x, P_y \geq 1$. We conclude that, as a general rule, when bidirectional splitting can be used, it should be used.

3.6.3 Nonlinear Model

Hydrostatic Simulation

To test the efficiency of the nonlinear model, both the HY and NH model options were considered. In the HY case, the test case considered is that of a perturbed baroclinic jet for 1000 time steps with and without topography. Two discretizations were considered; a coarser grid of $128 \times 128 \times 64$ cells and a finer grid of $256 \times 256 \times 128$ cells. When considering the MPI splitting, the convention suggested by PETSc is to ensure that there are at least 10,000 elements per processor, with 20,000 or greater being ideal [Balay et al., 2015]. The coarser resolution case allows for $\sim 16,000$ unknowns per processor in the 64 processor case, while the finer resolution case allows for over 100,000.

A summary of the results are shown in figure 3.9. We can see in this figure that the two-dimensional splitting almost always outperforms the one-dimensional splitting, as predicted by our communication calculation from the MPI section 3.6.2. As well, the fine resolution computations significantly outperform the coarser runs. For this reason, it is recommended that when using Tempest at least 20,000 unknowns are allotted per processor. That being said, under two-dimensional splitting with 64 processors and a fine grid, the parallel efficiency for the flat bottom and hyperbolic tangent topographic profile are 94.11% and 92.7% respectively.

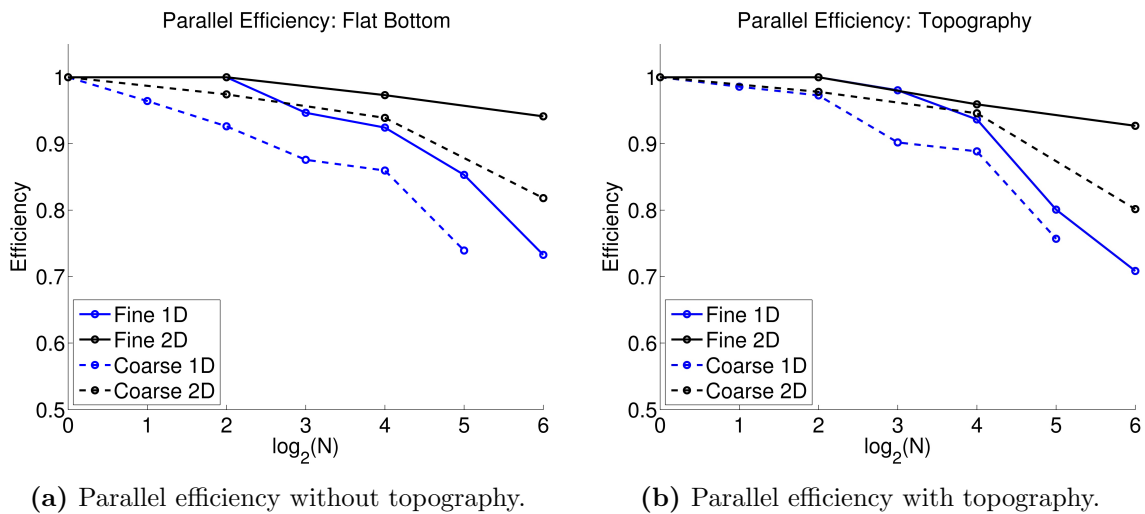


Figure 3.9: Parallel efficiency test of the HY model where N represents the number of cores. The results are displayed for a fine ($256 \times 256 \times 128$) and a coarse ($128 \times 128 \times 64$) run.

The efficiency does drop pretty significantly in the case of one-dimensional splitting with 64 cores but this implies that our local grid is only four elements wide. The halo region is also four elements wide so at each exchange we are transferring double the amount of data. This presents a significant bottleneck.

In terms of run times the reference 4 core simulations required 10.8 and 10.5 hours for the cases without and with topography respectively. These times do not account for matrix assembly but rather only the 1000 integration steps of the test case. The topography appears to have very little effect on the run time of these tests during the early stages of integration. It is expected that the run time will be affected by the topography at later integrations when the topography generates more dynamical features. However, this is not expected to affect the efficiency since the number and size of communications remains

constant throughout the integration.

Non-hydrostatic Simulation

In the case of the NH simulation we consider the same baroclinic jet simulation on the coarser grid with resolution $128 \times 128 \times 64$. The results from these efficiency calculations are shown in figure 3.10. We again observe that the two-dimensional splitting technique leads to better parallel efficiency relative to the one-dimensional splitting.

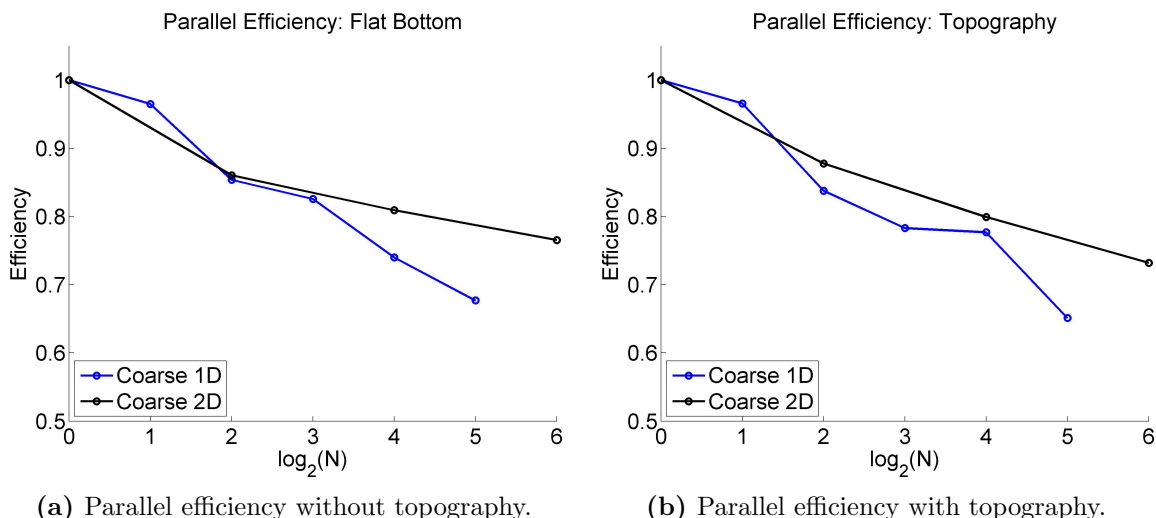


Figure 3.10: Parallel efficiency test of the NH model where N represents the number of cores. The results are displayed for a coarse ($128 \times 128 \times 64$) run.

In addition, we observe that the NH model does not perform as well as the HY model in terms of efficiency. The 64 core two-dimensional splitting for the tests with and without topography gave rise to 73.19% and 76.54% efficiency respectively. These results are slightly worse than those of the coarse HY model which gave efficiencies of 80.15% and 81.82% respectively. The causes for the observed differences in efficiency are unclear, however the significant difference between the HY and NH models is the two-dimensional versus three-dimensional elliptic pressure solves respectively. Therefore, it would seem that the NH three-dimensional solve does not scale as well as the HY two-dimensional solve. PETSc documentation [Balay et al., 2015] does indicate that increasing the number of processors does usually cause an increase in the number of iterations. Therefore, in the case of the

more complex NH solve it may be that the number of iterations dramatically increases compared to the increase in the number of iterations observed in the HY model.

From these tests, a number of conclusions can be made:

- At least 20,000 elements should be considered per processor.
- The HY model yields much better efficiency over the NH model.
- Two-dimensional splitting should be used when possible.

3.7 Conclusions

In this chapter, the full details of the Tempest ocean model have been introduced. The model is loosely based off of code that was originally provided by Paul Ullrich [Ullrich and Jablonowski, 2012]. From the original provided code, the model was modified to consider WENO5 advection schemes (as well as other advection options), higher order density terms, SSPRK3 time-stepping, various model options, fully three-dimensional studies, and two-dimensional varying bottom topography. As well, the code was updated to incorporate parallelized input/output operations through the use of the PETSc library.

By allowing for various model options, the Tempest ocean model allows for the efficient simulation of motions on various length scales, particularly through the use of the HY and QG models. Comparing HY and QG results to those obtained from NH simulations will allow users of the Tempest ocean model to gain a better understanding of NH and non-QG effects. Furthermore, empirical studies could provide information on when HY and QG models could give valid results, even when outside their theoretical regimes.

Through the use of high order advection schemes, the Tempest ocean model is able to simulate mesoscale and submesoscale dynamics with a higher degree of accuracy. The mesoscale and submesoscale structures remain coherent for longer periods of time. As well, the use of stable advection schemes allow for longer time integrations. The non-oscillatory behaviour of the advection schemes also suppress oscillations produced by the terrain-following transformation.

Terrain-following coordinates offer a way to model two-dimensional varying bottom topography. By incorporating FD and FV methods where they are most valuable, the Tempest ocean model is able to efficiently calculate terms implicitly even in the presence of topography. The calculations do require more time in the presence of topography than in the case of a flat bottom, however the parallel efficiency remains relatively unaffected in the presence of variable topography. The terrain-following coordinate system is prone to

small-scale, oscillatory errors in the presence of steep topography. Though these can be controlled by increasing grid resolution or by consider viscid equations, it is important to note this potential source of error.

The Tempest model currently does not allow for barotropic and baroclinic mode splitting. Nor does it allow for variable grid spacing or time-stepping, though variable time-stepping should be fairly easy to implement in a future update.

Chapter 4

Validation Test Results

To validate the Tempest ocean model we have considered a range of preliminary tests. The tests chosen are based on examples used in widely accepted models, particularly the MITgcm and MOM, which have provided valuable reference results and analysis. A number of test cases for ocean model development are available, however one particular challenge in selecting tests for Tempest is to find examples that have been defined on a Cartesian grid. This is another reason for selecting the following set of examples.

The first test is a simple nonlinear advection test showcasing the abilities of WENO to handle problems that could potentially lead to oscillations. The results are shown in section 4.1. The second test is called the Schär advection test and is considered in section 4.2. It has been chosen because conventionally the terrain-following coordinate system performs very poorly for this test case. The next test, section 4.3, is a coarse resolution simulation considering the advection of a passive tracer initially defined by non-negative values. We apply the Zhang limiting technique to this problem to demonstrate how we can ensure positivity. Following this test case, we consider two convection based studies in section 4.4.

Following the convection tests, we wanted to gauge the ability of Tempest to handle topographic features. For this, we considered two tests in particular. The first was a two-dimensional shelf test covered in section 4.5. The second test was a bowl test covered in section 4.6. Lastly, the horizontal advection scheme and the viscosity terms are tested using single and double gyre problems which are listed in section 4.7.

4.1 WENO Nonlinear Advection Test

This first test is a simple nonlinear advection test showcasing the abilities of WENO to handle problems that could potential lead to oscillations. The WENO method is compared to the fifth order upwinding and fourth order centered differencing methods. This test case considers the simple one-dimensional Burgers equation:

$$u_t + uu_x = 0.$$

The initial distribution of u is that of a Gaussian perturbation of the form

$$u(x, 0) = \exp\left(-\frac{(x - x_0)^2}{2 s_x^2}\right),$$

where $x_0 = 1/4$, $s_x = 1/20$, and the domain is defined on a unit interval. The integration is performed for a total of 1 s and the Courant number is set to 0.1. The grid is defined with a resolution of 100 equally spaced cells. As the maximum velocity occurs at the peak of the Gaussian profile, it is expected that the Gaussian will lean forward until a shock occurs. Upon the formation of the shock, it will travel forward as time passes and numerical dissipation will result in gradual decrease in the peak.

The results are shown in figure 4.1 as Hovmöller plots over the course of the integration. Subfigures 4.1a, 4.1b, and 4.1c show the results for the fifth order WENO scheme (WENO5), the fifth order upwind scheme (UW5), and the fourth order centered scheme




Figure 4.1: Nonlinear advection of a simple Gaussian profile using the one-dimensional Burgers equation $u_t + uu_x = 0$. The integration is performed for a total of 1 s and the results are displayed as Hovmöller diagrams.

(CD4) respectively. We see that the WENO5 scheme performs better than the UW5 scheme, and that both UW5 and WENO5 perform much better than the CD4 scheme even for this simple test case. It is clear that in both the UW5 and CD4 cases that oscillations develop in the solutions, which are suppressed by the non-oscillatory behavior of the WENO5 method. This is why we believe that the WENO scheme should do a better job in resolving ocean dynamics in comparison to these other two high order methods.

4.2 Schär Advection Test

As another preliminary test case, we consider a two-dimensional linear advection test. This test was originally introduced by [Schär et al., 2002] and begins with a bubble defined to the West of a set of strongly varying hills. Constant wind is applied and the bubble is advected across the domain, passing over the hills. Analytically the bubble should undergo a simple translation, however error is introduced by the terrain-following transformation and by the numerical errors associated with the advection scheme. The choice of this particular test was motivated by the fact that the terrain-following transformation conventionally performs very poorly in response to the strong variation in the topography [Schär et al., 2002, Li et al., 2014]. The reason is that the topography undergoes a significant amount of fluctuations over a short period and the gradients of the topography are very strong.

We attempted this test for two reasons. The first is that we wanted to test the validity of our FV density equation and in particular verify the theoretical accuracy with topography. Secondly, we wanted to determine to what extent our solution is affected by the terrain-following coordinate system. This test is defined in an atmospheric setting however it is a simple linear advection test and the definition of the domain is purely for reference to the original test paper [Schär et al., 2002].

The initial density profile is defined as

$$\rho(x, z) = \begin{cases} \cos^2\left(\frac{\pi}{2}r\right), & r \leq 1 \\ 0, & \text{elsewhere} \end{cases},$$

where r is

$$r = \sqrt{\left(\frac{x - x_0}{A_x}\right)^2 + \left(\frac{z - z_0}{A_z}\right)^2}.$$

Consistent with the definition from [Schär et al., 2002, Li et al., 2014] we use $A_x = 25$ km, $A_z = 3$ km. We set the initial center of the density anomaly to $(x_0, z_0) = (25, 9)$ km with

the domain defined as

$$0 \text{ km} \leq x \leq 150 \text{ km}, \quad 0 \text{ km} \leq z \leq 15 \text{ km}.$$

For a consistent comparison we initially use the same grid spacing of $\Delta x = 1000$ m and $\Delta z = 500$ m.

The velocity profile is set to

$$u(z) = u_0 \begin{cases} 1, & z_2 \leq z \\ \sin^2 \left(\frac{\pi}{2} \frac{z - z_1}{z_2 - z_1} \right) & z_1 \leq z \leq z_2, \\ 0 & z \leq z_1 \end{cases}$$

with $u_0 = 10$ m/s, $z_1 = 4$ km, and $z_2 = 5$ km. For all simulations the Courant number is fixed at 0.25, which, for the initial case, enforces a time step of 25 s. The simulations are run to an end-time of 10000 s. This implies that analytically the density anomaly will be translated 100 km eastward.

[Schär et al., 2002] demonstrated that the terrain-following coordinate system produced significant errors, using their particular FD approaches, which have up to fourth order accuracy. [Schär et al., 2002] determined that the L_∞ -norm for their second order test case was 0.174. Also, for comparison, their fourth order centered scheme had a L_∞ -norm of 0.057 and the other methods they employed (upwind, MPDATA, MPDATA linearized) had L_∞ -norms of at least 0.269.

For their calculations, [Schär et al., 2002] used a grid resolution of 150×50 and their domain had a maximum height in the vertical of 25 km. However, they present the results for the bottom 15 km. Significant interaction with the top boundary was not observed in our calculations and so we considered a domain with maximum height of 15 km for comparison to their figures. This does change the problem, however the reduction in domain height poses a greater challenge in that the cells undergo a more rapid change. [Schär et al., 2002] used the leapfrog time-stepping scheme with the same Courant number of 0.25. The primary source of errors introduced in the simulations of [Schär et al., 2002] appeared to be the generation of oscillations produced by interaction with the topography. As well, the west half of the density anomaly was noticeably deformed after passing over the topography.

In the case of the leapfrog scheme with centered second order accuracy [Li et al., 2014] showed that a slight variant of the terrain-following coordinate system produced relative errors of over 200%. Their reference hybrid coordinate system is the same as the terrain-following coordinate system but their coordinates are defined such that all vertical cell

interfaces are parallel for a given column of cells. Therefore, there is no transition to a flat cell at the top of the domain. In that particular case, their relative errors were 2.05724 at the end of the integration.

We perform a series of simulations beginning with the reference grid with resolution of 150×30 and successively refine the grid by doubling the horizontal and vertical resolutions. For all the tests we use the WENO5 advection scheme coupled with the SPPRK3 time-stepping method. The FV advection equation is implemented. The results of these simulations are summarized in table 4.1. At the end of the integration, we see that in the case of the reference grid we have errors that are 29% lower than the second order method of [Schär et al., 2002], but 2.18 times larger than their fourth order method. Furthermore, we see that the FV formulation does yield a second order accurate method as was expected.

Grid resolution	Time step	L_∞ -norm	Order
150×30	25s	1.041e-1	
	25s	1.244e-1	
300×60	12.5s	1.953e-2	2.414
	12.5s	2.155e-2	2.529
600×120	6.25s	5.825e-3	2.080
	6.25s	6.356e-3	2.145

Table 4.1: Comparison of the L_∞ -norms for the Schär advection test using successively refined grids. For each resolution the top row represents the results at $t = 5000$ s and the bottom represents the results at $t = 10000$ s.

The WENO advection scheme, though theoretically lower order under the FV formulation, offers benefits over classical centered schemes. Even with the coarsest resolution, we see that there are relatively little oscillatory errors as shown in subfigure 4.2b. Also, as observed in subfigure 4.2a we see that the general shape of the density is also conserved. Considering a flat bottom simulation of the same resolution we observe the absolute errors shown in subfigure 4.2c. From this subfigure we see that the numerical error associated with the advection scheme alone is approximately two orders of magnitude less than the errors from the terrain-following transformation. Therefore, we conclude that the error in the variable topography case is due predominantly to the terrain-following transformation.

In the case of a simulation with a fine 600×120 resolution we see that both types of error are significantly reduced which can be observed in subfigures 4.3b and 4.3c. These figures show the absolute error at the midpoint and end of the integration for the Schär mountain and flat bottom test cases respectively. We see from 4.3a that with the finer resolution that the density anomaly maintains its original shape as it passes over the hill.

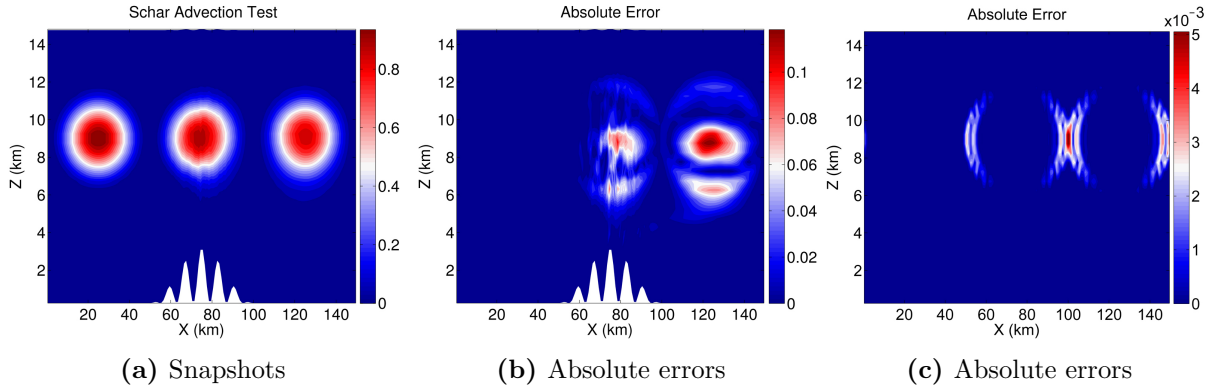


Figure 4.2: Schär advection test using a grid of 150×30 and Courant number of 0.25. Subfigure (a) shows the snapshots of the density anomaly as predicted by the WENO5 and SSPRK3 schemes. Subfigure (b) shows the absolute errors for the variable topography case and subfigure (c) shows the absolute error for the flat bottom case.

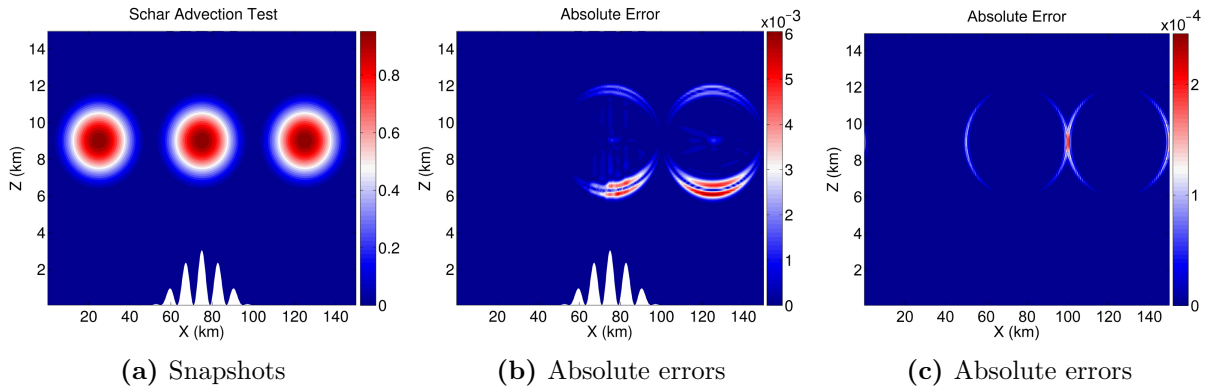


Figure 4.3: Schar advection test using a grid of 600×120 and Courant number of 0.25. Subfigure (a) shows the snapshots of the density anomaly as predicted by the WENO5 and SSPRK3 schemes. Subfigure (b) shows the absolute errors for the variable topography case and subfigure (c) shows the absolute error for the flat bottom case.

In both of the above cases, the WENO advection scheme suppresses oscillations, but it does not guarantee that the density anomaly will remain positive. We can apply the Zhang limiter to this in order to ensure positivity. In the absence of limiting the coarser grid above has a minimum density value of $-7.47e-4$. If we apply the Zhang limiter without reducing the Courant number then the Zhang limiter has a slight effect but still admits negative values. In this case the minimum is $-7.94e-6$. If we reduce the Courant number

to 0.01 respecting the strict CFL condition needed for the proper implementation of the Zhang limiter, then we obtain computationally zero values ($-2.24e-17$).

A benefit of the Zhang limiter is that away from areas that could potentially lead to negative values, the limiter respects the accuracy of the underlying advection scheme. If we consider the prediction from the Zhang limiter and look at the difference in the errors between the method with and without limiting we observe the error difference shown in figure 4.4. From above, the integration without limiting admitted errors of order 0.1. Therefore, we see that the errors introduced by the Zhang limiter are many orders of magnitude smaller than those introduced by the terrain-following transformation.

Therefore, this test has shown us that the WENO scheme used in the density equation of the Tempest model performs very well for this particular test case and that it acts to suppress oscillations that were caused by the topography. There are errors associated with the terrain-following transformation but we observed that the FV method still retains second order accuracy even with topographic features that have complex geometry. We also saw that the Zhang limiter, though it has a stricter CFL restriction, does prevent negative values while otherwise not affecting the results significantly.

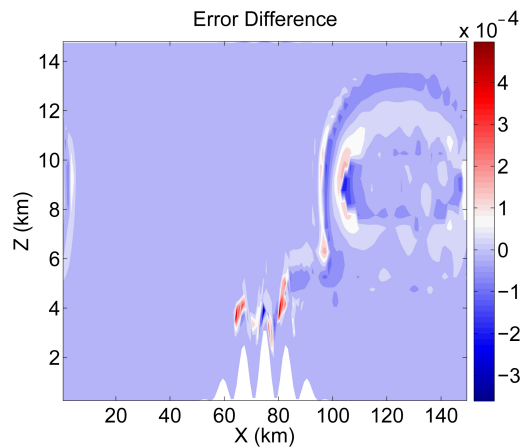


Figure 4.4: A plot of the error difference between the integration with and without limiting as predicted on the coarse 150×30 grid. The differences are computed at a time $t = 10000$ s associated with the end of the integration.

4.3 Passive Tracer Test

This next test is a coarse resolution simulation considering the advection of a passive tracer initially defined by non-negative values. This could represent the concentration of some biological or chemical tracer. The background flow is that of a barotropic jet, which is perturbed and generates instabilities, creating a complex velocity field. The WENO advection scheme is applied to this problem and though it acts to suppress oscillations in the solution it does not guarantee that the passive tracer will remain non-negative therefore leading to potentially nonphysical results. We apply the Zhang limiting technique to this problem to demonstrate how we can ensure positivity.

As we saw in the previous test, the Zhang limiter can be applied to linear advection methods to ensure positivity. However, to showcase the Zhang limiter in a practical test case we consider a barotropic jet which is perturbed and generates instabilities. The domain of this test problem is defined as

$$0 \text{ km} \leq x \leq 139.63 \text{ km}, \quad 0 \text{ km} \leq y \leq 100 \text{ km}, \quad 0 \text{ km} \leq z \leq 1 \text{ km}.$$

The background velocity associated with this barotropic jet takes the form:

$$u = U_0 \operatorname{sech}^2 \left(\frac{y - m_y}{y_d} \right),$$

where $U_0 = 1 \text{ m/s}$, m_y is the midpoint of the y extent, and $y_d = 10 \text{ km}$. On the initial time step, perturbations with a relative magnitude of $1\text{e-}3\%$ are added to this, exciting unstable modes.

A passive tracer τ is defined with a maximum value of 1 and minimum value of 0:

$$\tau = \exp \left[- \left(\frac{x - m_x}{s_x} \right)^2 - \left(\frac{y - m_y}{s_y} \right)^2 \right] \cdot \left(1 + \operatorname{erf} \left(\frac{z - H}{s_z} \right) \right),$$

where m_x and m_y are the midpoints of the zonal and meridional extents, $s_x = 13.96 \text{ km}$, $s_y = 10 \text{ km}$, and $s_z = 0.3 \text{ km}$. H represents the total domain height, 1 km . Therefore, this represents a Bi-Gaussian distribution that decays with depth.

A grid resolution of $64 \times 64 \times 64$ is used. The time step is set to 20 s and the equations are integrated for 25 days . The background stratification is set by $N = 1\text{e-}2 \text{ s}^{-1}$. The HY model is applied to this problem and tests with and without Zhang limiting were conducted. The run-time for this test without Zhang limiting was 55.6 minutes while the test with Zhang limiting required 65.9 minutes .

A cross-section of the passive tracer field at the end of the integration is shown in figure 4.5, where subfigures 4.5a, 4.5b, and 4.5c represent the integration without Zhang limiting, with Zhang limiting, and the relative difference between the two approximations respectively. This slice is taken near the surface at a relative height of 0.9. As in the previous section, it is observed that the Zhang limiter alters the approximation by a minimal amount. In this case the relative difference is on the order of $1e-3$, or 0.1%.

As figure 4.5 shows the results for one particular slice, the entire passive tracer field is also considered at the end of the integration. In the case when Zhang limiting is not applied the minimum value reaches -0.0015 , while the maximum is 0.1047 . On the other hand, when Zhang limiting is used the minimum value is $2.2e-11$ and the maximum is 0.1042 . Therefore, the relative difference of the maximum values is also on the order of $1e-3$ while positivity is strictly enforced. Also, the Courant number reaches a maximum of 0.0699 over the course of the simulation and, therefore, we see that even violating the strict CFL condition imposed by the Zhang limiter we still obtain positive preserving results for this particular case.




Figure 4.5: Cross-section of the passive tracer test using the WENO5 advection schemes as predicted by (a) the HY model without Zhang limiting and (b) with Limiting. The plots are shown for the end of the integration and (c) shows the relative difference between the two approximations.

4.4 Buoyancy Tests

Next, two convection based studies are considered. The first is a simple random perturbation experiment where dense fluid is introduced at the surface of an ocean basin. The dense fluid subsequently falls and produces plumes after passing through the upper, convective layer.

These plumes have a characteristic length scale. The second test is similar and introduces perturbations in a circular disc situated at the interior of a larger basin. For two days the perturbations are applied and then the forcing is turned off. The falling plumes having the same plume scale as the previous test and the dense fluid, after the forcing is turned off, form into conal structures with a characteristic cone scale. These tests showcase the NH model. The unstable stratification generated by the dense fluid randomly added to the surface leads to numerical instabilities in the case of solving the HY equations.

For the first test case, the configuration is equivalent to that used in the example configuration of the MITgcm [Adcroft et al., 2014]. In this experiment, the domain is defined as a flat bottom Cartesian system with the following definition:

$$0 \text{ m} \leq x \leq 3200 \text{ m}, \quad 0 \text{ m} \leq y \leq 3200 \text{ m}, \quad 0 \text{ m} \leq z \leq 1000 \text{ m}.$$

There is no background stratification and the initial density perturbation field is set to $\rho' = 4 \text{ kg/m}^3$. The Coriolis parameter f_0 is set to a value of $1\text{e-}4 \text{ s}^{-1}$ and the test is modelled on the f -plane.

Forcing is applied at the surface through a density source term. Consistent with the definition in MITgcm, we apply a heat flux with a mean value of 800 W/m^2 . We consider the heat flux Q to be defined as

$$Q = 800 (1 + 0.5l) \text{ W/m}^2$$

where l is normally distributed with zero mean and a standard deviation of one.

We use a specific heat coefficient for ocean water of $C_p = 3993 \text{ J}/(\text{kg}\cdot\text{K})$, as listed in [Kaye and Laby, 1995]. Therefore, the temperature flux across the top layer of the domain is defined as

$$\frac{dT}{dt} = \frac{Q}{\rho_0 C_p \Delta z} \text{ K/s}.$$

If we consider a linear equation of state without salinity and a thermal expansion coefficient of $\alpha = 2\text{e-}4 \text{ K}^{-1}$ then we get the following density flux term:

$$\frac{d\rho'}{dt} = -\alpha\rho_0 \frac{dT}{dt} = -\frac{2\text{e-}4 \cdot Q}{C_p \Delta z} \text{ kg}/(\text{m}^3\cdot\text{s}). \quad (4.1)$$

The random flux is stored initially and the same forcing is applied at each time step of the integration.

The other considerations are the viscosity, grid resolution, and time step. In this case, the viscosity is applied to the tracer and velocity components. For consistency with the

MITgcm test case viscosity coefficients of 0.1 m/s^2 are applied in both the horizontal and vertical. In this particular experiment, the grid resolution is $64 \times 64 \times 20$. The time step is chosen to be $\Delta t = 10 \text{ s}$ and the equations are integrated for one day.

During the initial stages the dense fluid undergoes convective motion. After passing through the purely convective layer, which is located near the surface, the convective motions form into plumes. These plumes form in response to rotational motion caused by the Coriolis forces, which play a role on time scales of f^{-1} and greater [Jones and Marshall, 1993]. The plumes generated have a characteristic horizontal length scale defined as

$$l_{\text{rot}} = \left(\frac{B_0}{f^3} \right)^{1/2}, \quad (4.2)$$

where B_0 is the magnitude of the Buoyancy flux. This length scale was observed in laboratory experiments by [Fernando et al., 1991]. Extensive numerical testing added further validation to this particular length scale [Legg and Marshall, 1993, Jones and Marshall, 1993].

For a heat flux averaging 800 W/m^2 we can compute the density flux using equation 4.1. Defined in terms of a general Δz the associated density flux is approximately $-(4\text{e-}5/\Delta z) \text{ kg}/(\text{m}^3 \cdot \text{s})$. If we multiply this quantity by $g\Delta z/\rho_0$ then we have that the buoyancy flux is $4\text{e-}7 \text{ m}^2/\text{s}^3$. For this value of B_0 , l_{rot} is approximately 600 m. In figure 4.6 we present the results of the experiment as predicted by the NH Tempest model. We see that our numerical solution gives a consistent result to that of the theory and to that presented by MITgcm. We see from subfigure 4.6a that the density flux initially undergoes a convective phase where there is little horizontal motion. Subsequently in subfigure 4.6b we see that plumes begin to develop with horizontal length scales comparable to those predicted by the theory.

Secondary Buoyancy Test

As a secondary convection test we reproduce a simulation comparable to the test outlined in [Jones and Marshall, 1993]. This test also considers a domain with rectangular geometry. The domain in this experiment is defined as

$$0 \text{ m} \leq x \leq 32000 \text{ m}, \quad 0 \text{ m} \leq y \leq 32000 \text{ m}, \quad 0 \text{ m} \leq z \leq 2000 \text{ m}.$$

The density flux is applied in the center of the domain by applying cooling to a disk with 8 km radius. The heat flux is set to an average of 800 W/m^2 , equivalent to that in the previous experiment, and the cooling is applied for the first two days of integration. The total integration is performed for four days. To limit the amount of viscosity we consider a viscosity coefficient of 0.1 m/s^2 as above.

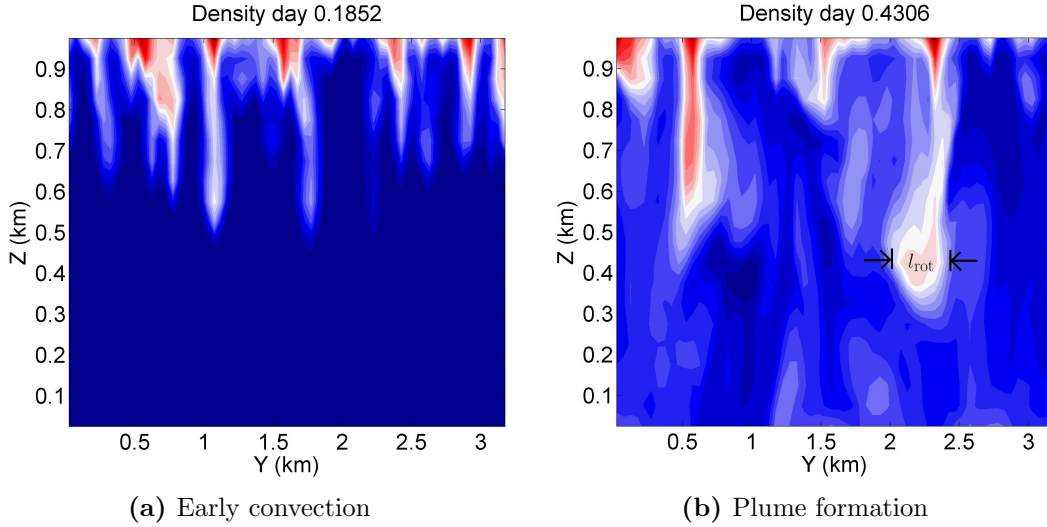


Figure 4.6: Simple buoyancy test. Random cooling at the surface first passes through the viscous boundary layer into the convective layer characterized by strong vertical motions (a). Upon passing through the convective layer plumes form with horizontal length scale on the order of l_{rot} (b).

In this particular simulation, there are a number of interesting physical features that evolve [Maxworthy and Narimousa, 1994]. The results at day two of integration are shown in figure 4.7 and we observe a number of the characteristic features we saw in the initial buoyancy test case above. After initially applying the flux we get convective motions. These convective motions next turn into plumes that evolve on the horizontal length scale l_{rot} as defined in equation 4.2. Between the dense falling plumes, there is upwelling of warmer fluid that reaches the surface and is subsequently cooled by the fluxes.

As the plumes reach the bottom of the domain they expand and over time cones of fluid affected by rotation are formed that trap the convected water. At the end of the convective period, the length scale of the geostrophically adjusted cones, as suggested by Maxworthy [Maxworthy and Narimousa, 1994], is set by the deformation radius:

$$l_{\text{cone}} \sim l_{\rho} = \frac{\sqrt{g'_{\text{rot}} H}}{f}, \quad (4.3)$$

where g'_{rot} is the buoyancy change, H is the domain height, and f is the Coriolis parameter. The buoyancy change is defined as $g'_{\text{rot}} = B_0/u_{\text{rot}}$ which is a product of B_0/h and h/u_{rot} ; in other words, the product of the forcing over the convective layer of thickness h and the time spent in the layer. At times on the order of f^{-1} and greater the rotationally dominated u

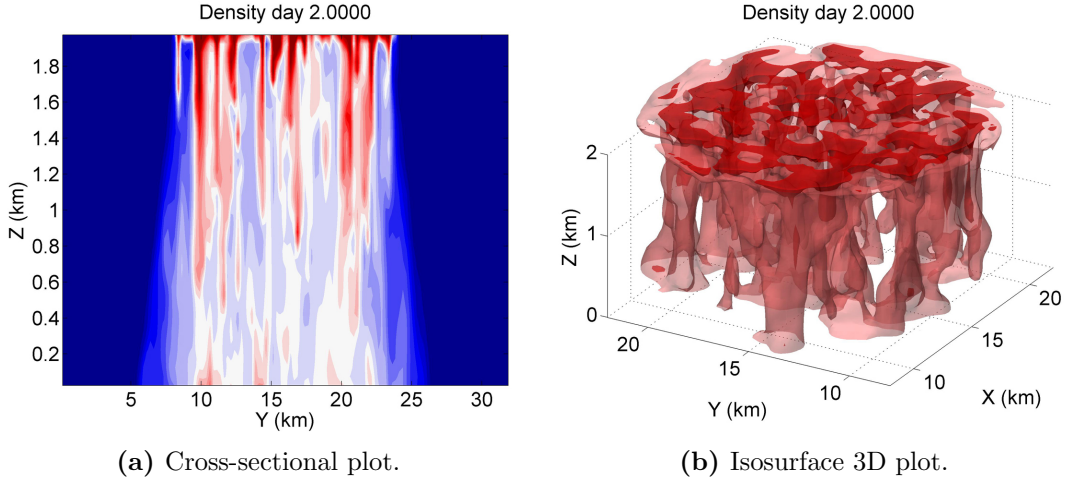


Figure 4.7: Buoyancy test of [Jones and Marshall, 1993], showing the cross-section (a) and 3D isosurface plot (b) for the dense fluid after two days of integration at which point a series of plumes have formed as a result of the applied density flux.

scales are defined as [Jones and Marshall, 1993]

$$u_{\text{rot}} \sim \sqrt{B_0 t} \sim \left(\frac{B_0}{f} \right)^{1/2}.$$

Therefore, upon substitution into equation 4.3 we determine the horizontal length scales of the cones to be [Jones and Marshall, 1993]

$$l_{\text{cone}} = \frac{B_0^{1/4} H^{1/2}}{f^{3/4}}.$$

For this particular test case, this is on the order of 1 km.

This is a reference scaling and for this particular configuration [Jones and Marshall, 1993] find that:

$$\frac{l_{\text{cone}}}{H} \sim (5.2 \pm 1.0) \times \sqrt{Ro^*}, \quad (4.4)$$

where

$$Ro^* = \frac{u_{\text{rot}}}{fH} = \frac{B_0^{1/2}}{f^{3/2}H} \sim 0.3,$$

which is consistent with the laboratory experiments of [Maxworthy and Narimousa, 1994].

After two days of integration the density flux is set to zero and the plumes are allowed to geostrophically adjust. It is at this point when the cones begin to fully form. After four days of integration the cones are well-established, shown in figure 4.8. The horizontal length scale of these plumes is observed to be between 5-7km consistent with the cone scaling, equation 4.4.

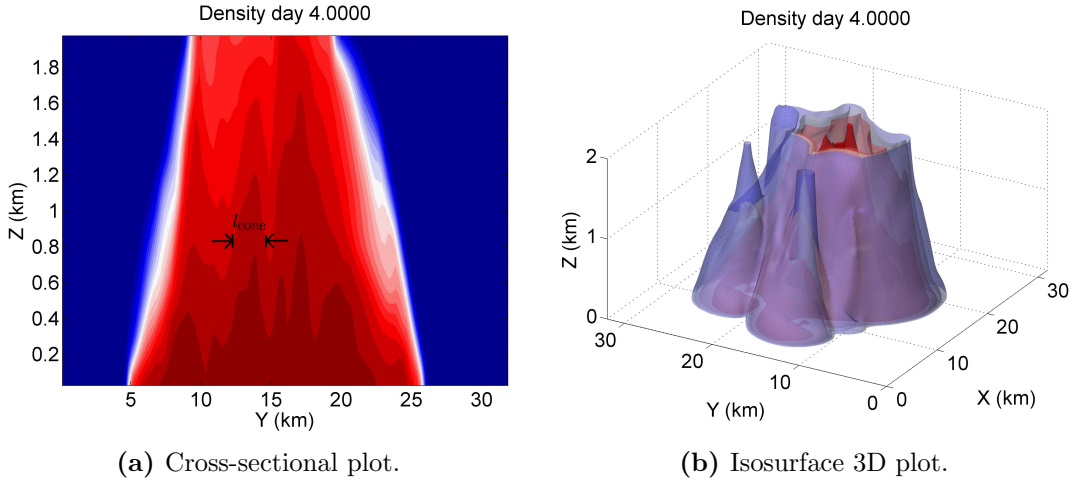


Figure 4.8: Buoyancy test of [Jones and Marshall, 1993], showing the cross-section (a) and 3D isosurface plot (b) for the dense fluid after day 4 of integration at which point the plumes have coalesced and formed into conal structures consistent with the findings of [Maxworthy and Narimousa, 1994].

In addition to confirming the plume and conal structures, there are a number of other interesting features to consider. In particular, there is well-defined structures in the surface vorticity field. The surface vorticity field is shown after a half day, two days, and four days of integration in subfigures 4.9a, 4.9b, and 4.9c respectively. In the earlier periods of the integration we see that there is a rim current around the applied density flux that is characterized by strong vorticity. This rim current is coherent early on and after two days of integration it begins to break up. By the end of the integration, the convective region has formed into areas of strong cyclonic vorticity with anticyclonic currents surrounding, consistent with [Jones and Marshall, 1993].

In the above simulations, we were able to test the NH equations of the Tempest ocean model. We observed that Tempest was able to capture a number of the physical features associated with these convective-driven flows. In particular, plumes were generated with the correct length scale and cones of dense fluid consistent with [Jones and Marshall, 1993]

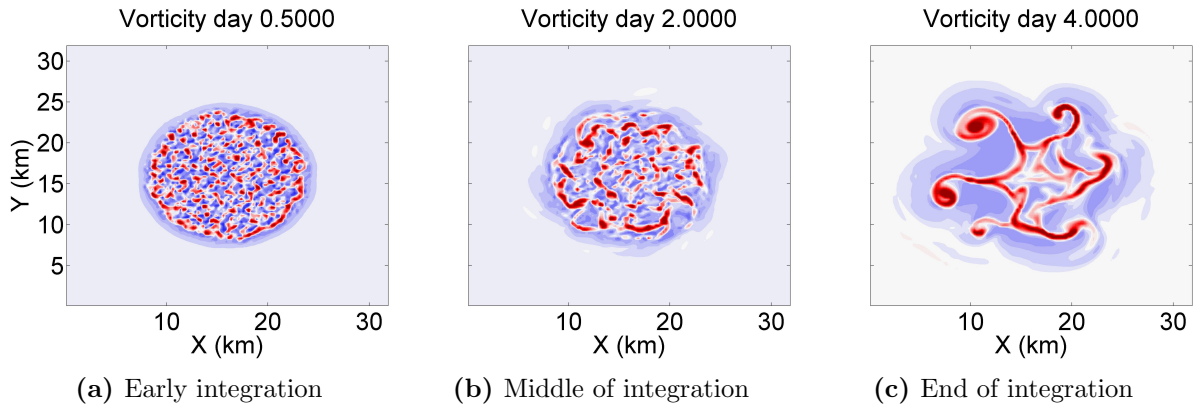


Figure 4.9: Buoyancy test of [Jones and Marshall, 1993], showing the surface vorticity at days (a) 0.5, (b) 2, and (c) 4 respectively.

were observed. The surface vorticity field also showed features very closely resembling those in the simulation of [Jones and Marshall, 1993]. During early periods of integration these test cases test the proper implementation of the buoyancy terms, NH pressure solve, and the buoyancy flux terms. As time progresses and the flow becomes affected by rotation these tests also allow us a means of assessing the Coriolis terms. The treatment of the nonlinear terms and the density equation are tested throughout the integration as they play a dynamic role in the evolution of the flow.

4.5 Flow Down Shelf Test

The next test case was used to assess the Tempest model at handling topography in a two-dimensional setting. This test models cold, dense fluid being added to the surface of the ocean above a continental shelf. The strongest perturbations were applied above the coastal, shallow region. The dense fluid falls and pools on the shelf. Under the effects of gravity, the dense fluid is then pulled down the shelf and, in the process, produces Kelvin-Helmholtz billows on the shelf. This test case is presented as an example in the MITgcm. It is an efficient simulation as it does not consider rotational effects and so can be modelled two-dimensionally [Adcroft et al., 2014].

The MITgcm test case has a varying grid in the horizontal, which Tempest does not allow for. Instead we consider an equally-spaced grid which has a resolution of 320×60

cells. The domain is defined as:

$$0 \text{ m} \leq x \leq 6400 \text{ m}, \quad 0 \text{ m} \leq z \leq 200 \text{ m}.$$

Therefore, the horizontal grid spacing is $\Delta x = 20 \text{ m}$ and vertical grid spacing is $\Delta z = 3.\bar{3} \text{ m}$.

The coastal topography is modeled by a hyperbolic tangent profile spanning up to a height of 80% of the domain depth. The topography profile takes the form:

$$h(x) = \frac{1}{2} A \left(1 - \tanh \left(\frac{x - m_x}{s_x} \right) \right),$$

where $m_x = 1700 \text{ m}$, $s_x = 533.\bar{3} \text{ m}$, and $A = 160 \text{ m}$.

The surface flux has a hyperbolic tangent profile as well, with the strongest cooling occurring at the Western boundary of the domain. The heat flux Q is defined as

$$Q(x) = \frac{1}{2} 200 \left(1 - \tanh \left(\frac{x - x_q}{s_q} \right) \right),$$

where $x_q = 700 \text{ m}$ and $s_q = 100 \text{ m}$. Initially, random perturbations are also introduced, as in the MITgcm setup. These act to promote the early onset of convection. These density perturbations have a mean magnitude of $2\text{e-}3 \text{ kg/m}^3$ consistent with the perturbations of $1\text{e-}2 \text{ K}$ introduced by MITgcm. We assume a linear equation of state and reference background density of 1000 kg/m^3 .

The density flux at the ocean surface is applied at each time step of the simulation. The equations are integrated for one day. The results at the end of this integration are shown in figure 4.10. The flux of buoyancy induces convection as in the open ocean test in the previous section, falling to the shelf and pooling. This pooled fluid forms a density current that begins to travel down the shelf accelerating to velocities comparable to $O(0.1) \text{ m/s}$ associated with a Courant number of order one, consistent with those found in the MITgcm.

The gravity current forms characteristic features having a noticeable gravity wave head [Firoozabadi et al., 2003] and a series of Kelvin-Helmholtz billows [Griffiths, 1986] which originate on the shelf. These features are observed in the density plot 4.10a. From the u velocity field 4.10b we see that there is a characteristic weak return flow above the strong positive u velocity associated with the gravity current. This return flow is characteristic of these types of gravity current experiments such as those of [Firoozabadi et al., 2003, Griffiths, 1986].

The MITgcm has the ability to consider variable horizontal grid spacing, which Tempest does not. For comparison to the above result we considered a finer resolution computation




Figure 4.10: Flow down a shelf problem of MITgcm [Adcroft et al., 2014] using a grid resolution of 320×60 .

by doubling the horizontal and vertical resolutions. This resolution allowed us to consider a grid resolution on-shelf comparable to the example setup in the MITgcm. The results at the end of the integration are shown in figure 4.11.

We observe that the finer resolution allows us to better resolve the small scale features on the shelf. In particular from the contour plot of the density 4.11a we see that the Kelvin-Helmholtz billows are much better resolved and they remain coherent as they are advected down the topography. Furthermore, the gravity wave head has a consistent structure to that of the coarser resolution simulation. With the finer resolution we observe, from the u velocity field 4.11b, that the velocity scales are comparable to that of the coarser resolution with slightly weaker velocities associated with the gravity current and a return flow comparable to that observed in the coarser resolution simulation.

These tests have given us the opportunity to assess the NH model of Tempest at handling steep topography. The Tempest model was able to accurately recreate a number of the characteristic physical features associated with these types of experiments. The Tempest model was able to generate the Kelvin-Helmholtz billows consistent with those observed in the example MITgcm test case. In the finer resolution run, it was possible to resolve the structure of the Kelvin-Helmholtz billows as they were advected down the topography. Lastly, the characteristic return flow was captured by Tempest and the horizontal velocities were consistent with the test case in the MITgcm.




Figure 4.11: Flow down a shelf problem of MITgcm [Adcroft et al., 2014] using a grid resolution of 640×120 .

4.6 Bowl Test

The next test considers two-dimensionally varying topography. This test defines dense fluid on a shelf in the Northern section of the region. To the South is a two-dimensionally varying bottom topography that resembles a bowl. The density gradient produces motion that acts to pull the dense fluid down into the interior of the bowl. This test case is similar to an example in the MOM open ocean test cases [Griffies, 2012]. We use a definition of the problem that closely follows that of [Winton et al., 1998]. The domain is defined as

$$0 \text{ km} \leq x \leq 1108.055 \text{ km}, \quad 0 \text{ km} \leq y \leq 1668.447 \text{ km}, \quad 0 \text{ km} \leq z \leq 3.5 \text{ km}.$$

The equations are defined on an f plane with $f = 2\Omega \sin(\pi/3)$. We consider the same grid resolution of $64 \times 64 \times 32$ as considered in [Winton et al., 1998].

The initial density configuration is shown in figure 4.12. The density is initially unstratified and is set to zero everywhere except on a shelf in the Northern one-third of the domain. In this region, we have that the density perturbation is set to $\rho' = 1 \text{ kg/m}^3$. In the absence of forcing this will induce a pressure gradient that causes the dense fluid to flow down the slope [Griffies, 2012].

The off-shelf topography is that of a bowl centered around $L_x/2, L_y/3$ where L_x and L_y are the spatial extents in the x and y directions respectively. This is formed by the product

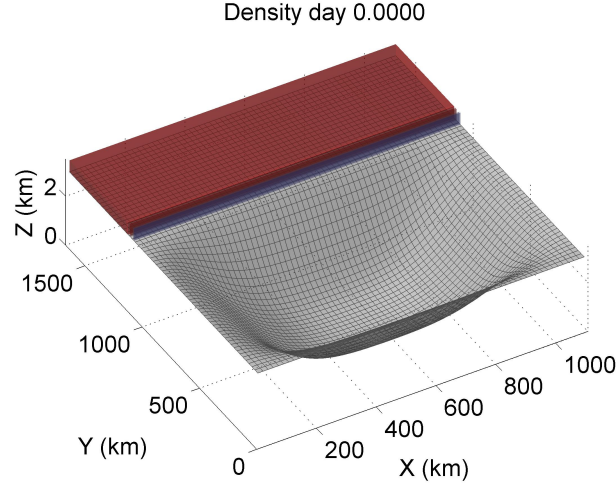


Figure 4.12: Initial configuration for the bowl test case. Dense fluid is located on a shelf and a lock-exchange experiment is conducted.

of transformed Gaussian functions:

$$h(x, y) = \begin{cases} \frac{6}{7}L_z & y > \frac{2}{3}L_y \\ \frac{6}{7}L_z (1 - G_x^1 G_x^2 G_y^1 G_y^2) & y \leq \frac{2}{3}L_y, \end{cases}$$

where L_z is the vertical extent of the domain and the transformed Gaussian functions are defined as

$$G_x^1 = 1 - \exp \left[- \left(\frac{x}{0.2L_x} \right)^2 \right], \quad G_x^2 = 1 - \exp \left[- \left(\frac{x - L_x}{0.2L_x} \right)^2 \right],$$

$$G_y^1 = 1 - \exp \left[- \left(\frac{y}{0.2L_y} \right)^2 \right], \quad G_y^2 = 1 - \exp \left[- \left(\frac{y - 2/3L_y}{0.2 \cdot 2/3L_y} \right)^2 \right].$$

In this test case we use a time step of one minute and calculate the solution for one year. [Winton et al., 1998] concluded that significant entrainment in a z -level model did not allow the dense water to flow to the bottom of the bowl. They compared the results to an isopycnal model and concluded that in the case of the accurate isopycnal model there was “very little dense fluid [...] above [a] 1-km depth” [Winton et al., 1998] and in the

case of the z -level model there was “very little cold water below this depth”. In the case of a terrain-following transformation this entrainment is not as severe, however pressure gradient errors can be a significant source of error [Griffies et al., 2000].

If we consider a y slice through the deepest part of the bowl after one year of integration we observe the density field as shown in figure 4.13. We see that Tempest is able to generate flow down the shelf consistent with that of the ρ model reference solution of [Winton et al., 1998]. We do however observe that in the Tempest model there is entrainment on the shelf which is not observed in the ρ model. That being said, the results are better than those observed in the z -level model used by [Winton et al., 1998] in that there is significant flow down the shelf.

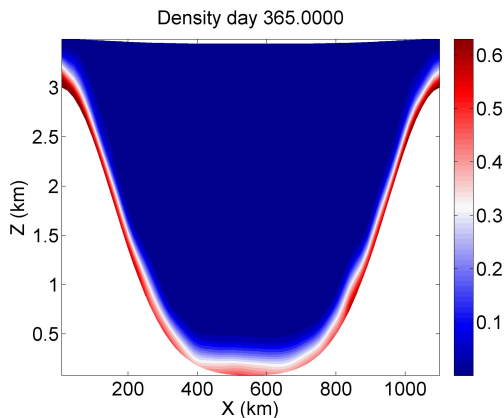


Figure 4.13: Cross-section of the density profile at $y = L_y/3$ using Tempest for the Bowl test case, as described in [Winton et al., 1998].

A potential explanation for the discrepancies between the Tempest results and those of [Winton et al., 1998] is that in this particular paper the authors apply biharmonic viscosity. As well, they use a flux-corrected transport scheme which has larger truncation errors and could potentially introduce greater numerical viscosity than the methods used in Tempest. For this study, the Tempest model was integrated using the inviscid equations. To reduce entrainment, one possible idea would be to introduce some additional dissipation into the model.

That being said, the downwelling flow observed in Tempest does generate a high pressure region near the center of the bowl. This is consistent with the ρ model of [Winton et al., 1998]. As can be seen from the vertically integrated vorticity field, shown in subfigure 4.14a, this high pressure region produces an anticyclonic vortex near the center

of the bowl. However, around this anticyclonic flow we observe a strong cyclonic circulation pattern around the exterior of the bowl. Comparing this to the surface vorticity field 4.14b we see minimal differences. The most notable differences are located on the shelf.

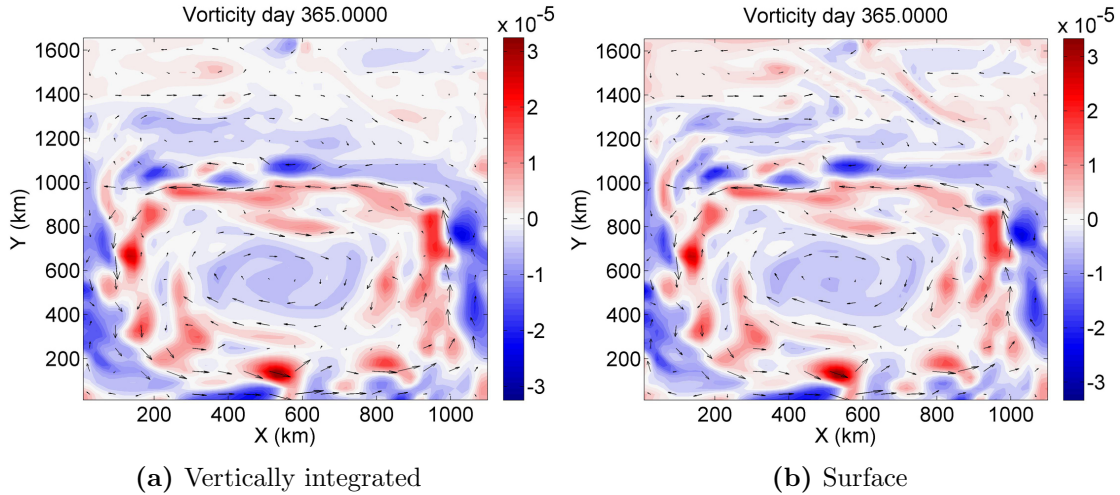


Figure 4.14: Vorticity field as predicted in the bowl test case at the end of the integration. Subfigure (a) shows the vertically integrated vorticity while subfigure (b) presents the surface vorticity.

Another very promising result is the ability to capture Kelvin-like waves that propagate in the Northern hemisphere with the coast to their right. However, these waves are not well resolved but are consistent with the results observed in [Winton et al., 1998]. From the midlatitude Rossby radii of deformation [North et al., 2014] we see that the decay scale is $O(10)$ km. This is consistent with a wavespeed on the order of 1 m/s:

$$L_D \sim \frac{u_K}{f} \quad \Rightarrow \quad u_K \sim 1\text{m/s.}$$

In figure 4.15, we see that Tempest does resolve these Kelvin wave-like structures. The approximate speed of these waves is 1.3 m/s consistent with the scaling given above.

Therefore, we have seen that Tempest is able to reproduce some of the downwelling of the dense fluid initially located on the shelf. In this regard, the Tempest model outperforms the z -level model of [Winton et al., 1998] but underperforms relative to the reference ρ model solution. Tempest however does reproduce a number of the dynamical features observed by [Winton et al., 1998], including the anticyclonic vortex located at the center of the bowl and Kelvin-like wave features located along the Western domain boundary.




Figure 4.15: Propagating Kelvin-like waves along the Western domain boundary as observed in the density field in the bowl test case. The cross-sections are taken through the middle of the domain.

4.7 Wind-Driven Gyre Tests

Lastly, the horizontal advection scheme and the viscosity terms are tested using single and double gyre problems. These tests were originally motivated by Western boundary current simulations in the MOM open ocean test cases [Griffies, 2012] and the MITgcm documentation [Adcroft et al., 2014]. The boundary conditions were set to free-slip conditions in all directions. The density profile has a stable stratification consistent with a buoyancy frequency of $N = 1e-2 s^{-1}$. The initial velocity fields are set to zero.

Dynamics are generated by applying a constant wind stress at each time step of the integration. This wind-stress is parameterized by a forcing applied to the u velocity, \mathcal{F}_u . If we are considering a problem with n gyres this wind stress takes the form:

$$\mathcal{F}_u = -\frac{1}{10 \rho_0 \Delta z} \cos\left(n \pi \frac{y - y_0}{L_y}\right) \text{ m/s}^2 \quad (4.5)$$

where y_0 is the minimum y value, L_y is the extent in the y direction, ρ_0 is the reference density, and Δz is the vertical grid spacing.

In terms of the other configuration settings we have that the domain is 1200×1200 km in the horizontal and the depth is chosen to be 2 km. The horizontal resolution is chosen as 256×256 . The Coriolis frequency is set to $f = 1e-4 s^{-1}$ and, therefore, $\beta \approx \frac{5}{3}e-11 \text{ (ms)}^{-1}$. The HY model is used for these test cases. We choose $\nu_h = 200 \text{ m}^2/\text{s}$ and, when we consider the cases with multiple vertical layers, $\nu_v = 0.02 \text{ m}^2/\text{s}$. The horizontal viscosity produces a

boundary current, often referred to as a Munk layer, which has the following characteristic length scale [Adcroft et al., 2014, Akuetevi and Wirth, 2015]:

$$M \sim \left(\frac{\nu_h}{\beta} \right)^{1/3}. \quad (4.6)$$

In all the test cases of this section we consider an integration spanning 500 days. In the single layer case, the only constraint on the time step is the CFL condition. In this case a time step of $dt = 1800$ s is sufficient for stability. In the multiple vertical level case, the time step must be chosen so as to resolve the inertial period which is $2\pi/N$ s. For $N = 1e-2$ this is consistent with a inertial period of ~ 600 s. Therefore, we consider $dt = 60$ s so that we have ten time steps per inertial period.

4.7.1 Single Gyre Test Cases

Single Layer

We begin by considering the single gyre, single layer simulation. As shown in subfigure 4.16a, after 65 days of integration we observe that the Western boundary current has formed. Using the scaling 4.6 we compute the scale of the Munk Layer thickness to be ~ 25 km. The observed result is consistent with this scaling. If we consider the vorticity plot 4.16b, we see that there are anticyclonic motions generated due to the interaction of the flow at the boundaries.

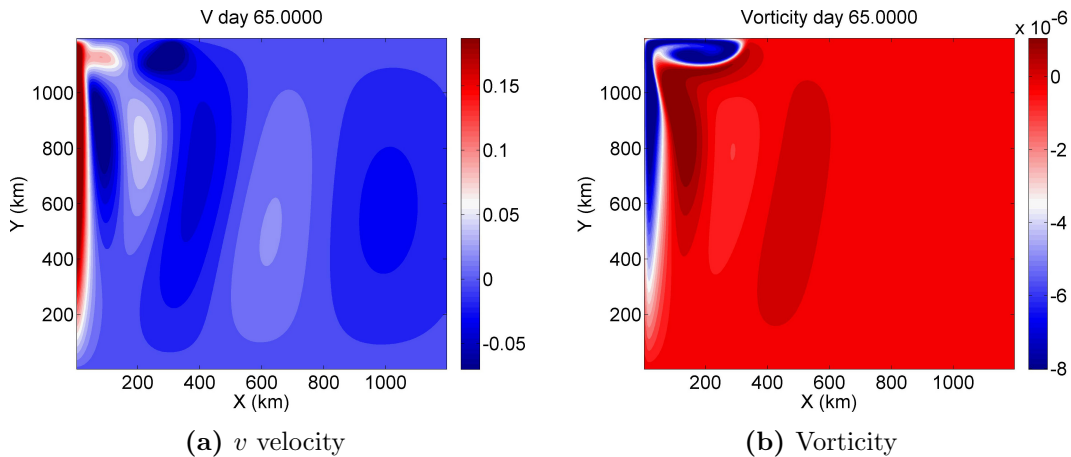


Figure 4.16: Single gyre, single vertical level test case.

Snapshots of the evolution of the vorticity are shown in figure 4.17. We see that the strongest zonal velocity is characterized by return flow along the Northern boundary, consistent with the findings in [Akueteve and Wirth, 2015]. However, the free-slip boundary conditions do not allow for the dipole formation as in [Akueteve and Wirth, 2015]. The generated current is similar in structure as that observed in [Pedlosky, 2010]. In addition, we observe cyclonic vortices formed in the interior of the flow that are generated as a consequence of the strong anticyclonic motions trapped along the Northern boundary. By the end of the integration, we observe a quasi-steady state.




Figure 4.17: 500 day integration of the single gyre, single vertical level test case showing snapshots of the evolution of the vorticity from day 50 to day 500 displayed in 50 day increments.

Multiple Layers

The next case we consider is the single gyre problem with sixteen vertical levels. We use the forcing defined in equation 4.5 consistent with that used by the MITgcm. Therefore, we see that the forcing will be sixteen times larger than in the case of the single vertical layer problem. Integrated over the depth of the top layer the same quantity of forcing is applied in both cases. We expect that this will generate strong motions near the surface.

If we consider snapshots of the evolution of the vorticity near the surface as shown in figure 4.18 we see that the evolution happens on a shorter timescale relative to the single vertical layer case. We also observe that there are a number of qualitative features that are equivalent between the single layer and multiple layer test cases, such as the width of the Western boundary current and the formation of the strong return flow along the Northern domain boundary. However, there does appear to be richer dynamics in the multiple layer case. In the multiple layer case the mesoscale vortices generated by the return flow appear to produce numerous internal gravity waves.




Figure 4.18: 500 day integration of the single gyre, 16 vertical level test case showing snapshots of the evolution of the vorticity from day 50 to day 500 displayed in 50 day increments.

The three-dimensional structure of the vorticity after 350 days of integration is shown in figure 4.19. We see that the magnitude of the barotropic structures 4.19a is comparable to the magnitude of the baroclinic structures 4.19b. We also observe from the baroclinic structure that a number of the small scale features in the interior of the domain are baroclinic in nature. The strong return flow along the Northern boundary is comprised of both barotropic and baroclinic structures with both, at the surface, contributing negative vorticity to the total vorticity field 4.19c.

At the end of the integration the vorticity has the three-dimensional structure shown

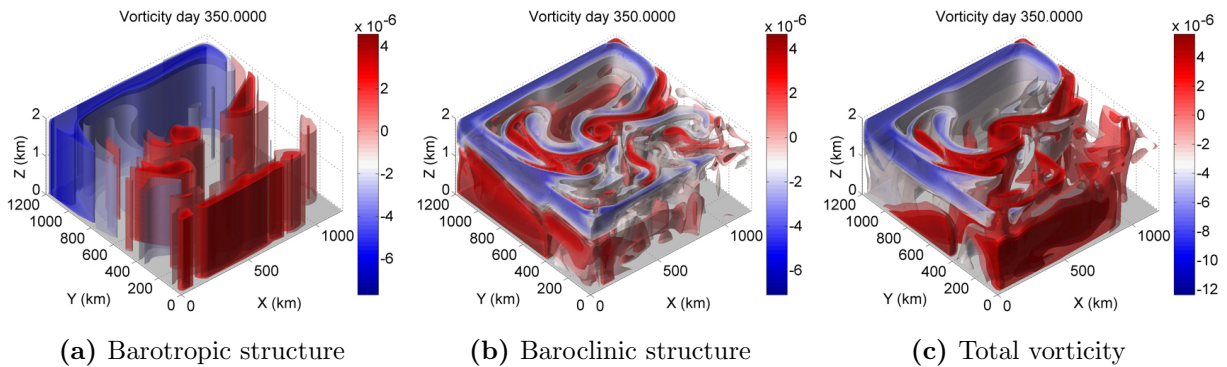


Figure 4.19: Three-dimensional structure of the vorticity in the single gyre, 16 vertical levels test case after 350 days of integration. Subfigure (a) shows the barotropic structure, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity field.

in figure 4.20. We observe from the barotropic structure 4.20a that the Northern return flow has a stronger barotropic structure at the end of the integration than that observed after day 350. Furthermore, we observe from the baroclinic structure 4.20b that the small scale features are again predominantly baroclinic in nature. Furthermore, there is a band of positive vorticity near the Northern boundary which is a surface-trapped baroclinic feature. The total vorticity field is shown in subfigure 4.20c.

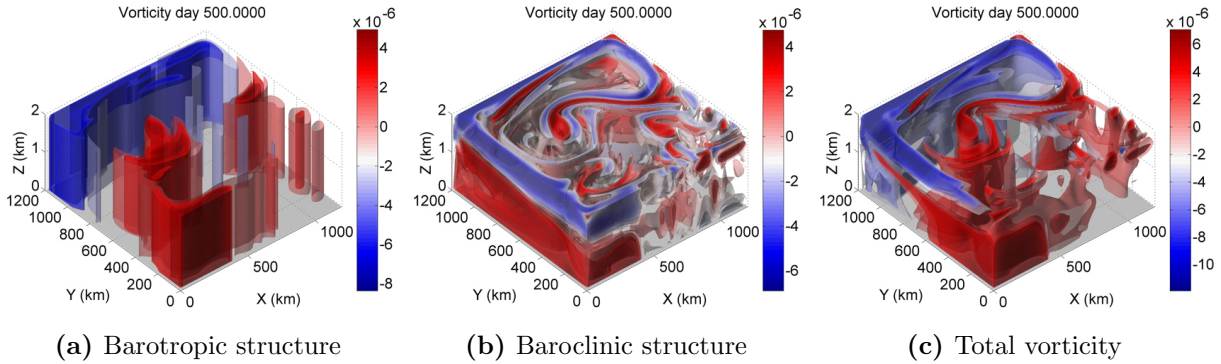


Figure 4.20: Three-dimensional structure of the vorticity in the single gyre, 16 vertical levels test case after 500 days of integration. Subfigure (a) shows the barotropic structure, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity field.

4.7.2 Double Gyre Test Cases

Single Layer

As a secondary set of tests we consider double gyre problems, which closely resemble those in the MOM open ocean test cases [Griffies, 2012]. In this case the domain is defined as

$$0 \text{ km} \leq x \leq 1074 \text{ km}, \quad 1668 \text{ km} \leq y \leq 3892 \text{ km}$$

in the horizontal. This is a classic configuration of a mid-latitude double gyre problem. The viscosity coefficients are set to $\nu_h = 200 \text{ m}^2/\text{s}$ and $\nu_v = 0.02 \text{ m}^2/\text{s}$ as in the case of the single gyre problem.

Snapshots of the evolution of the vorticity are shown in figure 4.21 for the single layer double gyre problem over the 500 days of integration. The boundary current has a width consistent with that of the theory. Furthermore, we observe qualitative features consistent with the analysis of [Fox-Kemper, 2005]. In particular, we see that in the case of the




Figure 4.21: 500 day integration of the double gyre, single vertical level test case showing snapshots of the evolution of the vorticity from day 50 to day 500 displayed in 50 day increments.

free-slip boundary conditions, which are imposed in this particular simulation, that there is a deviation from the mean y value in the intergyre dividing line. In particular, at day 250 we see that the Eastern extent of the dividing line deviates southward. Furthermore, consistent with [Fox-Kemper, 2005] we see that the initial jet crosses the full extent of the domain even with the formation of the mesoscale vortices.

Before the development of turbulent small scale features, we consider a time mean integration of the meridional velocity and the vorticity. This time mean is computed over the first 200 days and across the y level associated with the middle of the domain. The results are shown in figure 4.22. From subfigure 4.22a we see that there is a positive overshoot of velocity consistent with that observed in [Fox-Kemper, 2005]. Also, this overshoot produces sign-definite vorticity near the boundary as seen in subfigure 4.22b.

Multiple Layers

We next consider the case of sixteen vertical levels as in the single gyre problem. The snapshots of the evolution of the vorticity field near the surface are shown in figure 4.23. Using the same definition of the wind forcing 4.5 we have that the forcing is sixteen times larger at the surface than in the single layer simulation. It is therefore quite surprising that the evolution of the vorticity appears to have a similar timescale in the multiple layer case relative to the single layer case. One notable difference in the multiple layer case is that the gyre dividing line does not appear to deviate as significantly from the midpoint of the meridional extent.

For the double gyre problem we consider the three-dimensional structure of the vorticity after 300 days of integration as shown in figure 4.24. This particular time was chosen




Figure 4.22: Time mean of meridional velocity and vorticity at the midpoint of the zonal extent of the domain, integrated over the first 200 days.




Figure 4.23: 500 day integration of the double gyre, sixteen vertical levels test case showing snapshots of the evolution of the vorticity from day 50 to day 500 displayed in 50 day increments.

because of the coherent mesoscale vortices formed near the Eastern domain boundary. As seen from the barotropic and baroclinic structures, shown in subfigures 4.24a and 4.24b respectively, we see that in the double gyre case the barotropic instabilities are much larger than the baroclinic instabilities. This may be a potential explanation as to why the evolution in the multiple layer, double gyre problem occurs on a comparable timescale to the single layer problem. Similar to the single gyre problem, we see that a number of the smaller scale features are characterized by baroclinic instabilities.

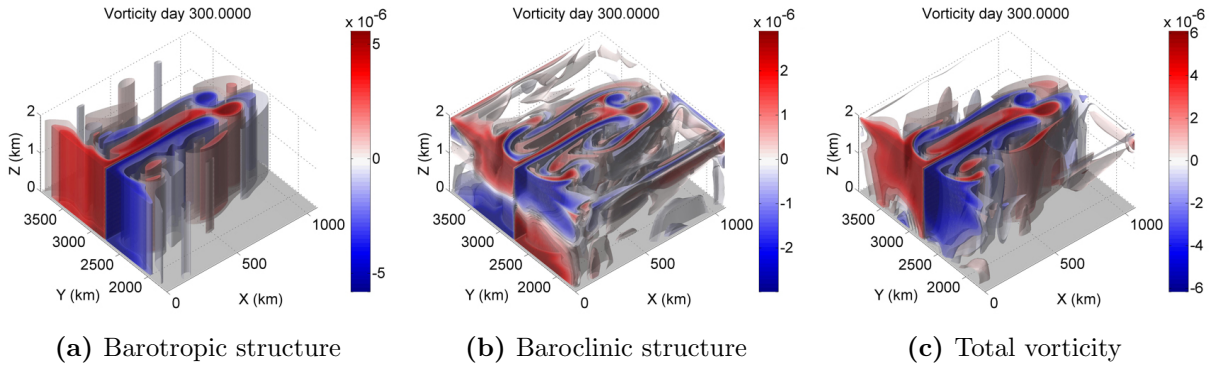


Figure 4.24: Three-dimensional structure of the vorticity in the double gyre, 16 vertical levels test case after 300 days of integration. Subfigure (a) shows the barotropic structure, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity field.

The three-dimensional structure of the vorticity at the end of the integration is shown in figure 4.25. We observe from the barotropic structure 4.25a and the baroclinic structure 4.25b that at this later time the baroclinic instabilities begin to contribute more to the total vorticity. In particular, we see that the anticyclonic vorticity associated with the barotropic and baroclinic instabilities appear to have nearly the same magnitude. On the other hand, the cyclonic vorticity appears to be dominated by barotropic motions. By the end of the integration there are a number of small scale features that are baroclinic and surface-trapped, consistent with that observed in the single gyre problem.

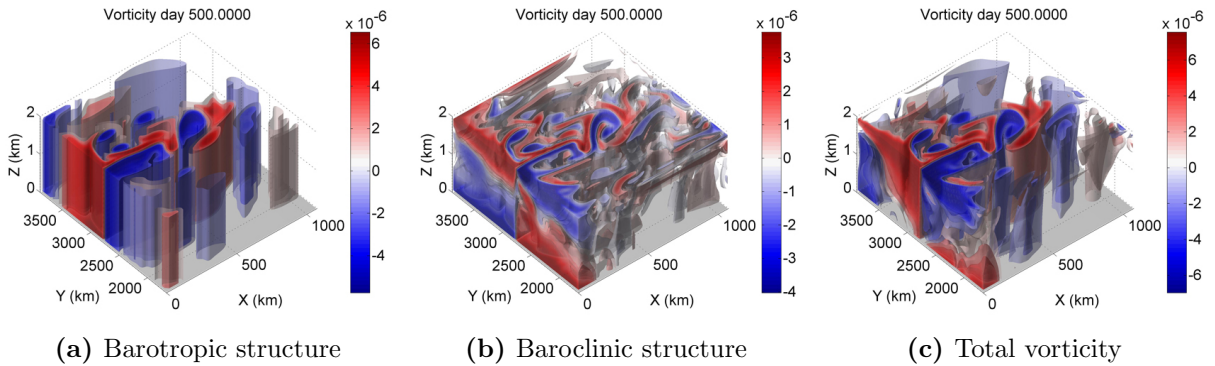


Figure 4.25: Three-dimensional structure of the vorticity in the double gyre, 16 vertical levels test case after 500 days of integration. Subfigure (a) shows the barotropic structure, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity field.

These tests have offered a means of assessing the viscosity terms in the Tempest ocean model. They also test the advection schemes, particularly in the horizontal, as well as the velocity surface forcing terms. The Tempest model reproduced a number of the classical features typical of these types of test cases. In particular, the model captured the correct length scale of the Western boundary current and also captured the mesoscale vortices produced from this Western boundary current. As well, single layer and multiple layer experiments were compared, allowing us a means of investigating the production of barotropic and baroclinic structures in the vorticity fields.

4.8 Conclusions

This chapter covered a set of validation tests that were used to assess various components of the Tempest ocean model, as well as the complete HY and NH models. Beginning with simple isolated tests of specific components it was seen in section 4.1 that the WENO5 method performed extremely well to suppress potential spurious oscillations. It outperformed the upwind fifth order method, which in turn outperformed the fourth order centered differencing method.

Next, the Schär advection test of section 4.2 was attempted in which a density anomaly was advected across a strongly varying mountain range. Though this was defined in an atmospheric setting it was simply used to test the FV density formulation. The WENO5 method was seen to suppress small scale oscillations that are prone to occur in this particular test [Schär et al., 2002]. Furthermore, the structure of the density anomaly remained coherent across the topographic features. The remaining validation tests considered either the full HY or NH Tempest model. The test cases and conclusions are summarized in table 4.2.

Test	Model	Comments
Passive Tracer	HY	<ul style="list-style-type: none"> – Difference with and without Zhang limiting was less than 1% – Zhang limiting ensured strictly positive results – Zhang limiting caused a slight increase in the computation time
Buoyancy	NH	<ul style="list-style-type: none"> – Correct plume and cone scales were calculated – Surface vorticity field compared to [Jones and Marshall, 1993] – HY model was unstable for unstably-stratified density
Shelf	NH	<ul style="list-style-type: none"> – Captured KH billows on shelf – Produced a classical return flow above flow – Tempest does not allow for variable grid spacing
Bowl	HY	<ul style="list-style-type: none"> – Remained stable in the presence of strong 2D topography – Produced significant flow into bowl interior – However, entrainment was observed on the shelf
Gyre	HY	<ul style="list-style-type: none"> – Correct Munk layer width was observed for boundary currents – Produced barotropic and baroclinic instabilities – Potentially some small errors causing spurious motions

Table 4.2: Summary of validation tests that were performed using the full NH or HY Tempest model.

Chapter 5

Linear Stability Analysis

Barotropic and baroclinic jets play a fundamental role in the dynamics of ocean coastal currents [Hristova et al., 2014]. Barotropic jets can be useful in the study of idealizations of coastal currents and generally make for a simple first step [Arbic and Flierl, 2004]. However, [Fratantoni and Pickart, 2003] concluded from observed data that baroclinic instability is the most likely cause of mesoscale variability and so the study of baroclinic jets will give valuable insights. [Chapman and Lentz, 1994] demonstrated how coastal density gradients can downwell to produce fronts that are parallel to the isobaths. The density gradient produces surface-intensified baroclinic jet structures which motivate the study of these particular baroclinic jet configurations.

There are a number of studies that have considered the case of barotropic jet instabilities, particularly through linear stability calculations. As the quasi-geostrophic two-layer model is able to capture many of the essential features of the instability of ocean currents they have been used extensively in previous studies [Arbic and Flierl, 2004, Irwin and Poulin, 2014]. These have given valuable insights into the growth rates of the most unstable modes however, the two-layer model is necessarily restricted in the number of modes that can be analyzed as well as the vertical structures that are produced. To get a better understanding of the modal structures and types of instabilities that can arise linear stability calculations based on primitive equation models can be studied [Xue and Mellor, 1993, Ménesguen et al., 2012, Lozier and Reed, 2005]. Though the primitive equation model can be used to study a range of unstable modes, many of the studies focus primarily on the barotropic and first baroclinic mode.

The effect of stratification has been studied in reference to the most unstable mode [Lozier and Reed, 2005], however it would be interesting to see how weak and strong stratifi-

cation affects the range of baroclinic modes which give rise to the variability in the jet structures. Topographical features have been shown to affect the stability characteristics of jets using a variety of models, including the RSW [Poulin and Flierl, 2005], QG [Ohshima, 1987, Li and McClimans, 2000, Esler, 2008], and PE [Barth, 1994, Lozier and Reed, 2005] models. However, it would be interesting to consider how various topographic heights affect the range of unstable modes.

By considering a range of unstable modes as well as calculations using both PE and continuously-stratified QG models, we can extend on previous studies to gain an understanding of how baroclinic modes affect the structure of the instabilities that arise from unstable barotropic and baroclinic jets. As well, by considering the affects on the barotropic and various baroclinic modes of instability we can get a better understanding of how background stratification and topography can affect the stability characteristics of these jets. This gives a more descriptive understanding of how the modal structures are altered. Extending on the results of [Li and McClimans, 2000] we can determine a series of symmetric (sinuous) and asymmetric (varicose) modes. Furthermore, by comparing the results from PE and QG calculations we can highlight specifically what non-QG effects arise in the PE calculations.

The remaining chapters focus on the study of barotropic and baroclinic coastal jets in which a Bickley jet (hyperbolic secant squared) profile is used. In particular, we investigate the instabilities that can arise from these jets through linear stability analyses. We can use the results from these calculations to help us determine interesting parameters to be studied in the context of nonlinear simulations.

5.1 LSA Governing Equations

In section 3.5 we derived the generalized eigenvalue problem that governs the linear stability calculations for these types of jets in a setting with potentially varying bottom topography in the context of the primitive equations. This generalized eigenvalue problem had the form $Ax = B\lambda x$ where

$$A = \begin{bmatrix} ikU & U_y - f & f^* + U_z & 0 & \frac{1}{\rho_0} ik \\ f & ikU & 0 & 0 & \frac{1}{\rho_0} (\partial_Y + Z_y \partial_Z) \\ -f^* & 0 & ikU & \frac{g}{\rho_0} & \frac{1}{\rho_0} (Z_z \partial_Z) \\ 0 & R_y & R_z - \frac{\rho_0}{g} N^2 & ikU & 0 \\ ik & \partial_Y + Z_y \partial_Z & Z_z \partial_Z & 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$x = [\hat{u}, \hat{v}, \hat{w}, \hat{\rho}, \hat{p}]^T$ and $\lambda = i\omega$. We saw that the eigenvalue with the maximum growth rate corresponded to the complex eigenvalue for which ω_i was largest. These equations are versatile in that they are valid for a wide range of physical parameters.

We also introduced the linear stability problem that is appropriate for the QG equations. This generalized eigenvalue problem was presented as $A\hat{\psi} = B\lambda\hat{\psi}$ where $\lambda = \frac{\omega}{k}$ and:

$$B = \left[(\partial_{yy} - k^2) + \frac{f_0^2}{N^2} \partial_{zz} \right], \quad A = \left[U (\partial_{yy} - k^2) + U \frac{f_0^2}{N^2} \partial_{zz} + Q_y \right] = [UB + Q_y].$$

This particular equation is much easier to solve than the PE problem, however these equations are only strictly valid for flows in which the QG assumption is valid, therefore flows with small Ro, small aspect ratio, and order one Bu. When solving this equation, the complex eigenvalue λ with the largest growth rate is associated with the λ that has the largest imaginary component.

5.2 Barotropic Jet: PE LSA Calculations

In the case of a barotropic Bickley jet, we consider a flow that is periodic in the zonal direction. In this case, we consider a domain defined as

$$x = [0, L_x], \quad y = [0, L_y], \quad z = [h(y), H],$$

where $h(y)$ represents our potentially varying bottom topography and $H = 1$ km. The domain has a meridional extent of 100 km and we consider a range of wavenumbers.

We consider a jet centered in the middle of the channel which has characteristic velocity and length scales defined as $U_{\text{jet}} = 1$ m/s and $L_{\text{jet}} = 10$ km respectively. We have a background profile that is independent of depth and varies only across channel. Therefore, the background flow has the following definition:

$$u = U(y), \quad v = w = 0, \quad \rho = R(y), \quad p = P(y),$$

where

$$U(y) = U_{\text{jet}} \operatorname{sech}^2 \left(\frac{y - m_y}{L_{\text{jet}}} \right).$$

From the quasi-hydrostatic and geostrophic balance we can solve for $R(y)$ and $P(y)$ from $U(y)$ as

$$R(y) = \frac{f^* \rho_0}{g} U(y), \quad P(y) = - \int \rho_0 f U(y) dy.$$

We use the following physical constants for all linear stability calculations:

$$g = 9.80616 \text{ m/s}^2, \quad f = f^* = 10^{-4} \text{ s}^{-1}.$$

Therefore, these flows have $\text{Ro} = 1$ and the QG approximation need not apply. However, by considering these flows we can discern how the QG and PE models differ and gain understanding about the limits of the QG model and how the model breaks down.

5.2.1 Convergence Analyses

Before a series of linear stability calculations were performed, various convergence analyses were considered to gauge how our SLEPc calculations depended on various factors. The first factor was the solver tolerance. The most unstable k value of 0.9 was considered and the calculations were performed on a grid with resolution 512×512 . Second and fourth order accurate FD approximations were considered and the solver tolerance was varied. The results from these calculations are summarized in table 5.1. We see that the tolerance appears to have no noticeable effect on the converged eigenvalues. For this reason, the default tolerance of 1e-8 has been used throughout the linear stability calculations.

tol.	order	iters.	converged eigenvalue	run-time
1e-8 (default)	2	4	$-1.606106\text{e-}5 + 4.053899\text{e-}5 i$	3m 45s
1e-10	2	4	$-1.606106\text{e-}5 + 4.053899\text{e-}5 i$	3m 56s
1e-30	2	5	$-1.606106\text{e-}5 + 4.053899\text{e-}5 i$	4m 1s
1e-8 (default)	4	4	$-1.606021\text{e-}5 + 4.053785\text{e-}5 i$	1m 28s
1e-10	4	4	$-1.606021\text{e-}5 + 4.053785\text{e-}5 i$	1m 29s
1e-30	4	5	$-1.606021\text{e-}5 + 4.053785\text{e-}5 i$	1m 39s

Table 5.1: Test of the dependence of convergence of solutions to SLEPc solver tolerances using fixed horizontal and vertical grid resolutions, $N_y = N_z = 512$

Next, the dependence on the order of accuracy of the finite difference approximations and grid resolution were assessed. For these calculations the number of grid cells varied but the same resolution was used in the horizontal and vertical for all calculations. As well, the default solver tolerance was used. The results are summarized in table 5.2. Due to the slower convergence, it is observed that a 400×400 grid using second order methods produced relative errors of 1e-2%, while a 200×200 grid using fourth order methods produced relative errors of 1e-4%. For this reason, the fourth order methods have been used throughout the calculations in the following sections.

N	order	iters.	converged eigenvalue	run-time	memory usage
50	2	4	$-1.597739e-5 + 4.041905e-5 i$	<1s	
100	2	4	$-1.603869e-5 + 4.050938e-5 i$	3s	
200	2	4	$-1.605547e-5 + 4.053157e-5 i$	13s	
300	2	4	$-1.605858e-5 + 4.053569e-5 i$	37s	
400	2	4	$-1.605967e-5 + 4.053713e-5 i$	1m25s	~2.5GB
500	2	4	$-1.606017e-5 + 4.053780e-5 i$	3m33s	~4.0GB
750	2	4	$-1.606066e-5 + 4.053846e-5 i$	13m51s	~8.0GB
50	4	3	$-1.607160e-5 + 4.052637e-5 i$	2s	
100	4	3	$-1.606141e-5 + 4.053835e-5 i$	8s	
200	4	4	$-1.606109e-5 + 4.053895e-5 i$	1m1s	
300	4	4	$-1.606107e-5 + 4.053898e-5 i$	2m50s	~3.5GB
400	4	4	$-1.606108e-5 + 4.053896e-5 i$	7m53s	~6.4GB

Table 5.2: Test of the dependence of convergence of solutions to grid resolution and order of the finite difference approximations. For all calculations $N_y = N_z$ and the default SLEPc solver tolerance of $1e-8$ was used.

As a final test, the dependence on the convergence of the solution to strictly the horizontal resolution was considered. In this test the vertical resolution was fixed to $N_z = 64$ and N_y was varied. The fourth order FD approximations and default SLEPc tolerance were used. The results of this test are summarized in table 5.3. We see that the $N_y \times N_z = 256 \times 64$ has the same growth rate as the 300×300 case from the previous table 5.2, and the phase is only off by $1e-4\%$. Therefore, $N_y \times N_z = 256 \times 64$ has been chosen as the reference resolution.

N_y	iters.	converged eigenvalue	total run-time
64	5	$-1.606314e-5 + 4.053584e-5 i$	3s
128	4	$-1.606119e-5 + 4.053875e-5 i$	8s
256	6	$-1.606107e-5 + 4.053897e-5 i$	22s

Table 5.3: Test of the dependence of convergence of solutions to horizontal grid resolution with a reduced vertical grid resolution of $N_z = 64$. The default SLEPc solver tolerance of $1e-8$ and fourth order finite difference approximations were used.

5.2.2 Reference State

As a reference test, the case of a flat bottom and stable stratification consistent with a buoyancy frequency $N = 1\text{e-}2\text{s}^{-1}$ was considered, while implies $\text{Bu} = 10$. Constant stratification is chosen for simplicity and to compare with the results of [Lozier and Reed, 2005]. The default SLEPc tolerance of $1\text{e-}8$ is used and a grid resolution of $N_y \times N_z = 256 \times 64$ is considered. The approximations are made using the fourth order FD formulas.

We consider wavenumbers in the range 0 to 2, where our k values are normalized by the characteristic jet width L_{jet} . Using the direct solver on a coarse grid, a range of 0 to 8 was initially considered. This showed two ranges of unstable modes. However, a calculation on an extremely fine grid of 8192×128 showed that the maximum growth rate of the second range of unstable k values was two orders of magnitude smaller than the maximum growth rate of the first range. Successive refinement of the grid led to successively decreasing growth rates. Therefore, these modes are believed to be spurious and regardless they are physically unimportant relative to the first range of unstable modes given their extremely small growth rates.

For this calculation 100 equally spaced wavenumbers in the range of $[0, 2]$ were considered and initial guesses were calculated using the coarse direct solver. We then use these guesses to initialize our SLEPc calculations. A maximum of 5000 iterations per wavenumber was considered and all but 2 wavenumbers converged. A comparison of the results as calculated by SLEPc and our coarse direct solver are shown in figure 5.1. Though the results are comparable, we see that the direct solver does inaccurately predict the structure of the range of unstable k values.

From our SLEPc calculation, we compute the maximum growth rate of the most unstable mode as $1.3877/\text{day}$ for a wavenumber of 0.9. This wavenumber coincides with a wavelength of approximately 69.81 km. The spatial structure of this mode is shown in figure 5.2 for this SLEPc calculation. The Matlab modal structure compares very closely and is not presented here. Subfigures (a)-(f) show the vertical structure of the u, v, w, ρ, p , and the vorticity while subfigures (g) and (h) show the horizontal structure of the v and vorticity field.

We observe that the vertical structure of this particular mode is barotropic with slight variations observed in the vertical velocity and density, which are both significantly smaller than the other fields. This is consistent with Squire's theorem which states that for every three-dimensional instability there corresponds a more unstable two-dimensional one or, in other words, the most unstable mode is two-dimensional [Vallis, 2006].

The structures in 5.2 have been normalized so that the maximum magnitude of the u velocity field is 1 m/s. Therefore, we also observe that the vertical velocity associated




Figure 5.1: Comparison of the growth rates of the most unstable mode for wavenumbers in the range $[0, 2]$. The results are shown for the coarse grid 10×10 matlab direct solver and the fine grid 256×64 SLEPc calculation.

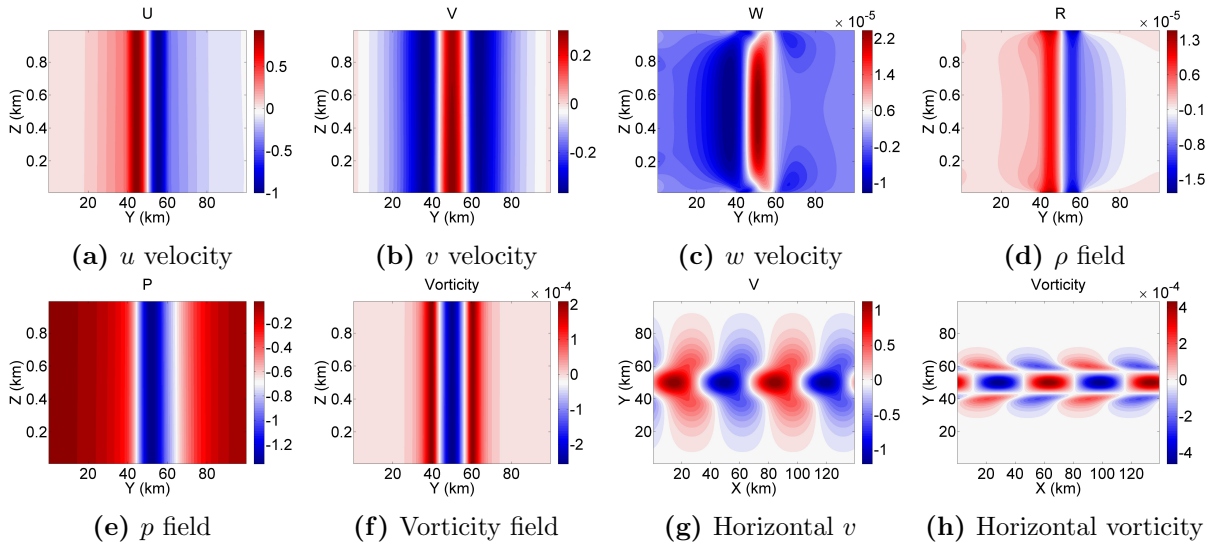


Figure 5.2: Subplots (a)-(f) present the vertical structure of the u, v, w, ρ, p and vorticity fields for the s_0 mode with the largest growth rate for the reference barotropic jet case. In subplot (g) and (h), the horizontal structure of the across channel velocity v and the vorticity are shown.

with this barotropic mode is many orders of magnitude smaller than the horizontal velocity components. This is typical of barotropic modes. This mode is classified as a sinuous mode since the across channel velocity and vorticity fields are symmetric about the flow. To represent this model as a sinuous mode with a mode zero (barotropic) vertical structure we

label it s_0 .

In the continuously stratified case the solution to the eigenvalue problem gives rise to a whole family of unstable modes. The above case considered only the most unstable mode and the results were extended to the five most unstable modes. In requesting more unstable modes, the computation time for the SLEPc calculation significantly increases. When only the most unstable mode is requested the computation time for this particular calculation was approximately 1.82 hours. On the other hand, when the five most unstable modes were requested the calculation required approximately 39.88 hours. When requesting multiple modes, the time for convergence can increase dramatically for those wavenumbers associated with weak growth rates. It was observed that this was particularly true for those modes associated with short wavelengths that were either stable or had very weak growth rates.

The growth rate curves for the five most unstable modes is shown in figure 5.3 (left) with a summary of the maximum growth rates and associated k values (right). We observe that the first and second modes have comparable growth rates and the wavenumber associated with the largest growth rate are equivalent in both cases. As the mode number increases the largest growth rate decreases and the tendency is for the largest growth rate instability to be associated with longer wavelengths (smaller wavenumbers).




Figure 5.3: Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the five most unstable modes in the reference barotropic case as predicted by the PE linear stability calculation.

The vertical structure of the across channel velocity v for these five modes are shown in the first row of figure 5.4. From these plots we observe that the vertical structure takes on the shape of the first five cosine functions, with the second mode having a mode one

structure, third mode having a mode two structure, etc. In order to discuss the symmetry of these modes we consider the horizontal structures as shown in the bottom row of figure 5.4. We observe that the first and second modes display near perfect symmetry about the flow. However, asymmetry begins to develop in the third mode onward. This asymmetry is not found in the QG model and is due to Ro being order one. Since all of these modes have v velocity maximums located at the center of the domain, we classify all of these modes as sinuous. The second through fifth modes represent the first four baroclinic sinuous modes, labeled s_1 to s_5 respectively, where the subscript represents the vertical structure. That is, s_1 represents a baroclinic sinuous mode with a mode 1 vertical structure. Note that all these modes are due to barotropic instability but the ones with a variable vertical structure are called baroclinic.




Figure 5.4: Vertical (top row) and horizontal (bottom row) structures of the across channel velocity v for the first five modes of the reference barotropic jet case, as predicted by the PE linear stability calculations.

As the vertical structure tends to smaller length scales the vertical velocities play a more significant role in the instabilities. However, it was observed that the vertical velocities remain at least two orders of magnitude smaller than the horizontal velocities for all the modes considered. The geostrophic scaling as described in section 2.4 indicated that $W \sim RoUH/L$. For this case $Ro \sim 1$ and $U \sim 1$ so we expect that the vertical velocity would have a maximum scale of that of the aspect ratio, which is $1/10$ using the domain height and characteristic horizontal length scale of the jet. What we observe is certainly within this range.

We typically observe a range of sinuous modes but over a smaller range of wavenumbers classical theory has shown that in the continuously stratified case we have a second set of unstable modes which are the varicose modes [Drazin and Reid, 2004, Irwin and Poulin, 2014].

These modes are characterized by an across channel velocity v that vanishes at the middle of the domain. Synonymous to the sinuous modes the first varicose mode represents the barotropic varicose mode and is labeled v_0 . The second varicose mode onward represents the family of baroclinic varicose modes and are labeled v_n where n indicates that the mode having a mode n vertical structure.

To gain a better understanding of the growth rates of the sinuous and varicose modes a reduced range of k values, $k \in [0, 1]$, was considered. In this calculation, the sixteen most unstable modes were requested. To account for the increasingly complex structure of the modes, a finer grid resolution of 512×128 was considered. The growth rate plots for these modes are shown in figure 5.5. The barotropic sinuous (s_0) and varicose (v_0) modes as well as the first baroclinic sinuous (s_1) and varicose (v_1) modes have been highlighted.




Figure 5.5: Plots of the growth rate curves for the sixteen most unstable modes in the reference barotropic case for a reduced range of k values, using the PE linear stability calculations. The sinuous and varicose barotropic and first baroclinic modes are highlighted.

The maximum growth rates of the barotropic and first baroclinic varicose modes are 0.3858/day and 0.2996/day associated with k values of 0.52 and 0.48 respectively. The vertical structure for the barotropic varicose mode is shown in figures 5.6 (a)-(f) and the horizontal structure is shown in figures 5.6 (g)-(h). We observe that, unlike the sinuous mode, that the varicose mode displays asymmetry about the flow. Similar to the sinuous mode though, it is observed that this barotropic mode has vertical velocities that are many orders of magnitude smaller than the horizontal velocity components. Also, as the varicose mode number increases the wavenumber of the maximum growth rate mode tends to decrease consistent with that observed for the sinuous modes.

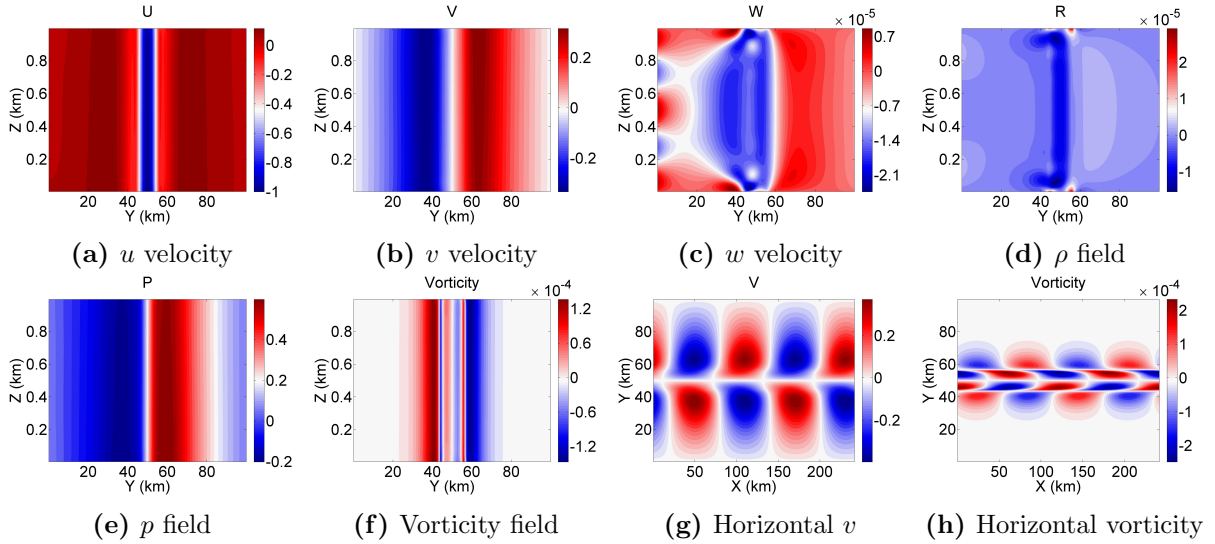


Figure 5.6: The structure of the v_0 mode in the reference barotropic jet calculations. Subplots (a)-(f) present the vertical structure of the u, v, w, ρ, p and vorticity fields and subplots (g) and (h) present the horizontal structure of v and the vorticity.

5.2.3 The Effects of Stratification

To investigate the effect that stratification has on the stability properties of the jet a series of tests were considered in which the strength of the background stratification was varied. The background stratification was first weakened by incrementally decreasing Bu by 1 from the reference value of $Bu = 10$ to a minimum of $Bu = 1$. To consider stronger stratification Bu was incrementally increased by 10 to a maximum of $Bu = 50$.

The growth rates of the most unstable mode over the full range of wavenumbers and stratifications are shown in subfigure 5.7a. The dashed line in this figure indicates the wavenumbers that give rise to the largest growth rates. Consistent with the reference calculation, over the full range of stratifications the wavenumbers are equal to 0.9. In the case of the second most unstable modes we obtain the results shown in figure 5.7b. For this mode, we see that there is a stronger dependence on the stratification. Furthermore, weaker stratification gives rise to much weaker growth rates.

The plots of the largest growth rates for the five most unstable modes as a function of Bu are shown in figure 5.8. It is observed that as the stratification increases the growth rates of multiple modes tend to the growth rate of the barotropic mode. Hence, the growth rates of the secondary modes increase in response to stronger stratification. On the other hand, the growth rate of the barotropic sinuous mode shows very little dependency on Bu .




Figure 5.7: Plots of the growth rates as a function of wavenumber, k , and the Burger number, Bu , for (a) the s_0 mode and (b) the s_1 mode of the barotropic jet calculations.




Figure 5.8: The plots of the largest growth rates for the five most unstable modes as determined from the PE linear stability calculations in the case of the barotropic jet.

Therefore, with increasing mode number, the stratification appears to have a greater effect.

The s_0 mode has the largest growth rate of 1.3886/day in the $Bu = 1$ case. This is a marginal increase over the reference stratification $Bu = 10$ case and could be due to numerical errors. As the stratification appears to have a strong effect on the growth rates the cases with weakest and strongest stratification have been considered in greater detail.

Weak Stratification Case

For the case of weak stratification the five most unstable modes were requested, however only the first three converged to physical solutions. Convergence analyses have indicated that the fourth and fifth modes appear to have stabilized as was also suggested in figure 5.8. The growth rate curves (left) and maximum growth rate summary (right) for this stratification are shown in figure 5.9.

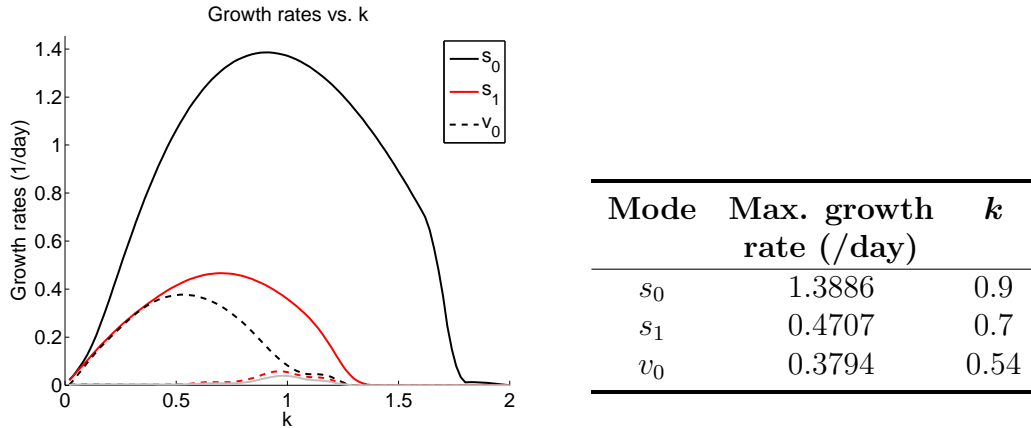


Figure 5.9: Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the three unstable modes in the weakly stratified $Bu = 1$ case as predicted by the PE linear stability calculations for the case of the barotropic jet.

In this case, the black and red solid lines represent modes s_0 and s_1 , respectively, and the black dashed line represents the v_0 mode. We see that the linear stability calculations predict that the first and second modes have maximum growth rates for different wavenumbers unlike in the reference case. However, as the mode number increases the wavelength of the most unstable mode also increases. The same tendency was observed in the reference case.

The top and bottom rows of figure 5.10 show the vertical and horizontal structures of the across channel velocity v for these three modes. It is observed that for this weakly stratified case mode s_0 has a comparable structure to that observed in the reference calculation. The meridional extent of the mode is comparable in both cases. As well, the zonal extent is equivalent since the most unstable mode occurs for the same wavenumber. The asymmetries observed in modes s_1 and v_0 are a result of non-QG effects caused by the order one Ro .




Figure 5.10: Vertical structure (top row) and horizontal structure (bottom row) of the v velocity for the k value associated with the maximum growth of that mode. The results are displayed for the first three modes for the case with $Bu = 1$ as predicted from the PE linear stability calculations.

Strong Stratification Case

In the case of the strongest stratification that we considered, the sixteen most unstable modes were calculated and the growth rate curves (left) and maximum growth rate summary (right) are shown in figure 5.11. The choice to consider a wide range of modes was motivated by figure 5.8 in which we saw the convergence of multiple modes as stratification increased. The calculation was performed over a reduced range of k values $k \in (0, 1]$ to facilitate convergence. These curves reinforce the fact that the background stratification plays a critical role in the growth and structure of the secondary modes. The increase in the growth rates of the secondary modes are readily observed and all sixteen modes are sinuous in contrast to the reference case, shown in figure 5.5, for which some of the sixteen most unstable modes were varicose.

From figure 5.11 it is observed that as the stratification increases the tendency is for a larger number of unstable modes to also share the same wavelength as well as comparable growth rates. By comparison, the sixteenth most unstable mode has a maximum growth rate of 1.0438/day for a wavenumber of 0.84 in this strongly stratified case. In addition to a dramatic increase in the growth rates, the strongly stratified case also gives rise to much weaker vertical velocities, which are four orders of magnitude smaller than the horizontal velocity components.




Figure 5.11: Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the most unstable modes in the strongly stratified $Bu = 50$ case as predicted by the PE linear stability calculations for the case of the barotropic jet.

The vertical structure of the first five modes in this case are shown in figure 5.12. From these figures, it is observed that these first five sinuous modes have near perfect symmetry about the middle of the channel. In contrast, the reference calculation of the previous section demonstrated that for the third mode and beyond, stronger asymmetries developed. Therefore, in terms of the modal structure, it is concluded that in the case of strong stratification that the non-QG effects do not play as significant a role as in the weaker stratification cases.




Figure 5.12: The vertical structures of the across channel velocity v for the first five modes of the strongly stratified barotropic jet case with $Bu = 50$, as predicted by the PE linear stability calculations.

5.2.4 The Effects of Topography

In the case with topography, the generalized eigenvalue problem in the PE linear stability calculations is determined by 3.18. To reduce the numerical error, the derivatives of

the background state are evaluated analytically. The perturbations are evaluated on the computational grid and the same interpolation formulas are used as in the flat bottom case.

We consider the same barotropic profile as above and a sloping topography modeled by a hyperbolic tangent function. We define a parameter, a , whose magnitude defines the percentage of the total height that the maximum of the topography reaches. That is, a value a of 0.5 represents topography that extends halfway up the domain. The sign of a represents the type of the topography with positive a representing prograde topography and negative a representing retrograde topography. The height of our topography $h(y)$ is therefore defined as

$$h(y) = -\operatorname{sgn}(a) \left(\frac{|a|H}{2 - |a|} \right) \tanh \left(\frac{y - m_y}{s_y} \right),$$

where s_y governs the width of the topography and H represents the mean depth of the domain.

In order to determine the impact prograde and retrograde topography has on the instabilities, we consider a range of topographic heights by varying the parameter a . For all cases the width of the topography was set by $s_y = 0.15 L_y$. A total of 31 different topographic heights was considered, with a ranging from -0.75 to 0.75 in increments of 0.05 . The five most unstable modes for each wavenumber were calculated. To isolate the effects of topography, we fix $\text{Bu} = 10$ as in the reference calculation. All calculations were performed on a 256×64 grid.

For the most unstable mode, the results for all k values and topography parameter a are shown in subfigure 5.13a, while the results for the wavenumber giving rise to the largest growth rates are shown in subfigure 5.13b. In subfigure 5.13a the dashed line indicates the wavenumber with the largest mean growth rates, while in subfigure 5.13b the dashed line indicates the flat bottom case.

From this figure, it is observed that the prograde topography acts to destabilize this particular jet, with larger topography having more of a destabilizing effect than shallower topography. In contrast, it is seen that the retrograde topography always has a stabilizing effect, with larger topography giving rise to greater stabilization. This trend is observed across the whole spectrum of wavenumbers and is consistent with the results found in [Poulin and Flierl, 2005], who studied the barotropic mode in a one-layer shallow water model with a free-surface. The largest growth rate is $1.3965/\text{day}$, for a k value of 0.9 and the strongest prograde topography profile. In general, the topographic slope appears to alter the growth rate of this s_0 mode by a small amount ($< 2\%$). However, from the convergence analysis, using the same grid resolution, it is expected that numerical errors would be on the order of $1e-4\%$ and so we expect this trend to be significant.




Figure 5.13: Plots of the growth rates as a function of wavenumber, k , and topography parameters, a , for the s_0 mode. The figure on the left shows the full range of k values and the figure on the right shows the k value that gives rise to the largest growth rates.

In the case of the second most unstable mode, the topography has a much greater effect. The results for all k values and topography parameter a are shown in subfigure 5.14a, while the results for the wavenumber giving rise to the largest growth rates are shown in subfigure 5.14b. In this case, shallow retrograde topography has a stabilizing effect while shallow prograde topography has a destabilizing effect. However, large prograde and retrograde




Figure 5.14: Plots of the growth rates as a function of wavenumber, k , and topography parameters, a , for the s_1 mode. The figure on the left shows the full range of k values and the figure on the right shows the k value that gives rise to the largest growth rates.

topographies both have a tendency of stabilizing this mode with the retrograde topography having a stronger stabilizing effect. Therefore, the largest growth rate occurs for a shallow prograde topography with $a = 0.15$ and is computed as 1.3020/day for a wavenumber of 0.9.

In the case of the third, fourth, and fifth modes, the plots of the growth rates as a function of a are shown in figure 5.15 (left). As the mode number increases we see that there is a shift from shallow prograde topography having a destabilizing effect to a stabilizing effect and, for higher mode number modes, the prograde topography is strictly stabilizing. In contrast, the shallow retrograde topography transitions from stabilizing to destabilizing. Furthermore, as the mode number increases the retrograde topography appears to have less of an effect on the growth rates. A summary of the maximum growth rates for the first five unstable modes as a function of k and a are shown in figure 5.15 (right).




Figure 5.15: Plots of the maximum growth rates (left) as a function of the topography parameters, a , for the s_2 , s_3 , and s_4 unstable mode of the barotropic jet case as predicted by the PE linear stability calculations. A summary of the maximum growth rates as a function of k and a are shown as well (right).

To study how the prograde and retrograde topography affects the structure of the instabilities, the intermediate cases for which $a = \pm 0.5$ were considered. The horizontal and vertical structures of the across channel v velocity of the most unstable mode are shown in subfigures 5.16a and 5.16b for the retrograde and prograde topography respectively. To consider the case when the topography has the strongest effect the horizontal profiles show the bottom vertical layer. Therefore, in the case of the s_0 mode it is observed that there are very little differences between the results from the prograde and retrograde calculations.

The horizontal and vertical structures of the across channel v velocity for the second mode are shown in subfigures 5.17a and 5.17b for the retrograde and prograde topographies




Figure 5.16: Plots of the vertical and horizontal structures for the across channel velocity v for the s_0 mode in the case of an intermediate $a = -0.5$ (a) retrograde topography. Subfigure (b) shows the same plots in the case of the intermediate $a = 0.5$ prograde topography.




Figure 5.17: Plots of the vertical and horizontal structures for the across channel velocity v for the s_1 mode in the case of an intermediate $a = -0.5$ retrograde topography (a). The subfigure (b) shows the same plots in the case of the intermediate $a = 0.5$ prograde topography.

respectively. Therefore, it is also observed that the s_1 mode in both the prograde and retrograde cases have comparable structure. However, in the case of the retrograde profile this mode appears to be more intensified to the South while in the prograde case this mode appears to be more intensified to the North, implying that intensification occurs in the direction of deep water. The same intensification pattern has been observed in the s_2 , s_3 , and s_4 modes.

5.3 Barotropic Jet: QG LSA Calculations

In the QG case we use the same definition of the barotropic Bickley jet as defined in section 5.2. This definition leads to the generalized eigenvalue problem 3.20 except with a background velocity and vorticity that only depends on the across channel coordinate. In the case of the QG model, the computations are significantly less demanding so the vertical resolution was increased. For these calculations a grid resolution of $N_y \times N_z = 256 \times 256$ was used.

5.3.1 Reference State

The configuration of the reference calculation in the QG case is equivalent to that used in the PE linear stability calculations of section 5.2.2. For this reference calculation $Bu = 10$. A set of 100 equally spaced wavenumbers was considered in the range $[0, 2]$. In section 5.2.2, it was indicated that a calculation requesting five unstable modes required approximately 39.88 hours using the primitive equations. By comparison, the QG calculation with the finer resolution was able to calculate the growth rates associated with the five most unstable modes in 0.89 hours. Clearly, when using the QG model there is a great increase in computational efficiency. The results of this QG calculation are summarized in figure 5.18.




Figure 5.18: Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the five most unstable modes in the reference barotropic as predicted from the QG linear stability calculations.

The QG calculations predict the maximum growth rate of the most unstable mode to be almost exactly equivalent to that predicted by the PE calculations, which were shown in figure 5.3. This is perhaps surprising since in that case we had $Ro = 1$ and therefore we should not expect QG to necessarily be valid. In the case of the second mode, the QG calculation predicts a 0.66% larger growth rate than that predicted in the PE calculations. Therefore, we see that the PE and QG linear stability calculations agree very well for the s_0 and s_1 modes.

By comparison, the QG model predicts the maximum growth rates of the third, fourth, and fifth modes are 4.06%, 23.96%, and 50.42% smaller than those predicted in the PE calculations. Therefore, we see that the discrepancies between the PE and QG results increase for increasing mode numbers. This indicates that the QG approximation starts to break down with a reduction in the vertical length scale of the unstable modes.




Figure 5.19: The horizontal structures of the across channel velocity v for the first five modes of the reference barotropic jet case as predicted by the QG linear stability calculations.

The horizontal structure of the across channel velocity v for the first five modes as predicted by the QG calculations is shown in figure 5.19. For these modes, we see that the maximum v velocity is located at the center of the channel. Therefore, these modes represent the first five sinuous modes. Comparing these to the first five sinuous modes as predicted by the PE calculations, shown in figure 5.4, it is observed that the first and second mode have nearly identical structures. The third, fourth, and fifth modes differ by increasing amounts. This is due to the asymmetries that develop in the PE calculations, whereas the QG calculations are always symmetric.

As was observed in the PE calculations there are two families of modes, the sinuous and varicose modes. The varicose modes occur over a reduced range of wavenumbers so a second calculation was performed over the range $k \in [0, 1]$. In this case the sixteen most unstable modes were requested, as in the PE calculations, however only the first eleven




Figure 5.20: Plots of the growth rate curves for the eleven most unstable modes in the reference barotropic case for a reduced range of k values, using the QG linear stability calculations. The sinuous and varicose barotropic and first baroclinic modes are highlighted.

modes converged. The growth rate curves for these eleven modes are shown in figure 5.20. In this plot $s_0, s_1, v_0,$ and v_1 have been highlighted. The maximum growth rate for the barotropic varicose mode is computed as 0.3860/day for a wavenumber of 0.53. This is nearly equivalent to the PE prediction of 0.3858/day for a wavenumber of 0.52. On the other hand, the maximum growth rate of the first baroclinic varicose mode as predicted by the QG calculation is 0.3208/day, an increase of 7.08% over that predicted by the PE calculation. The most unstable mode is also predicted to have a slightly shorter wavelength associated with a wavenumber of 0.51.

5.3.2 The Effects of Stratification

To consider the effects of stratification on the stability characteristics of the modes as predicted by the QG equations, the same series of tests were conducted as in section 5.2.3. The maximum growth rates associated with each Bu are shown in figure 5.21 for the five most unstable modes. Comparing these results to the PE calculations, shown in figure 5.8, it is observed that the stratification has a much stronger impact on the higher mode number modes in the QG calculations.

In the case of the strongest stratification the QG and PE calculations agree very well. However, as the stratification becomes weaker (Bu decreases) we see that the growth rates decrease much more rapidly in the QG case. In particular, for the weakest stratification it is observed that the third, fourth, and fifth modes have stabilized according to the predictions by the QG model. In contrast, the PE only predicted the fourth and fifth modes as stable.




Figure 5.21: The plots of the largest growth rates for the five most unstable modes, as a function of Bu, determined from the QG linear stability calculations in the case of the barotropic jet.

Therefore, it is observed that strong stratification leads to better agreement between the QG and PE models. Since the weakest and strongest stratification cases give very different levels of agreement between the QG and PE predictions these cases have been considered in greater detail.

Weak Stratification Case

For the weak stratification case the five most unstable modes were calculated however only the first two converged to unstable modes. The growth rate curves for this case are shown in subfigure 5.22. The black solid and dashed lines represent the s_0 and v_0 modes. The maximum growth rate of the most unstable s_0 mode is predicted as 1.3876/day representing a 0.07% difference between that observed in the PE calculations. The wavenumber associated with this growth rate is also 0.9. In the QG calculation the second unstable mode represents the v_0 mode. Therefore, in the case of weak stratification the QG equations were unable to resolve the first baroclinic sinuous mode captured by the PE calculations. The maximum growth rate of the barotropic varicose mode is 0.3858/day for a wavenumber of 0.52. This growth rate represents a 1.69% increase over the growth rate predicted by the PE model.




Figure 5.22: Growth rate curves in the weakly stratified $Bu = 1$ case as predicted by the QG linear stability calculations for the case of the barotropic jet. The two modes represent the s_0 and v_0 modes.

Strong Stratification Case

In the case of strong stratification, the sixteen most unstable modes were calculated. The predicted results are summarized 5.23. The calculated modes represent the first sixteen sinuous modes, s_0, s_1, \dots, s_{15} , as in the PE case however the maximum growth rates do decrease more rapidly in the QG case. The predicted maximum growth rates represent relative differences of 0%, 0.05%, 0.17%, 0.37%, and 0.56% when compared to the predictions from the PE calculations. By comparison, the growth rate of the sixteenth most unstable mode differs by 23.93% between the PE and QG results.

A plot of the maximum growth rates as predicted by the PE and QG calculations for all the modes are shown in subfigure 5.24a. We see that there is very strong agreement between the PE and QG calculations for the nine most unstable modes. These modes had very weak asymmetries in the PE calculations. Beyond the ninth mode, as the instabilities become more asymmetrical, it is observed that the predicted growth rates disagree by greater amounts. As more asymmetries form the QG model significantly underestimates the growth rates. By comparison, a calculation with a reduced Ro of 0.01 was performed and the predictions of the PE and QG calculations agreed for all sixteen modes as shown in subfigure 5.24b.

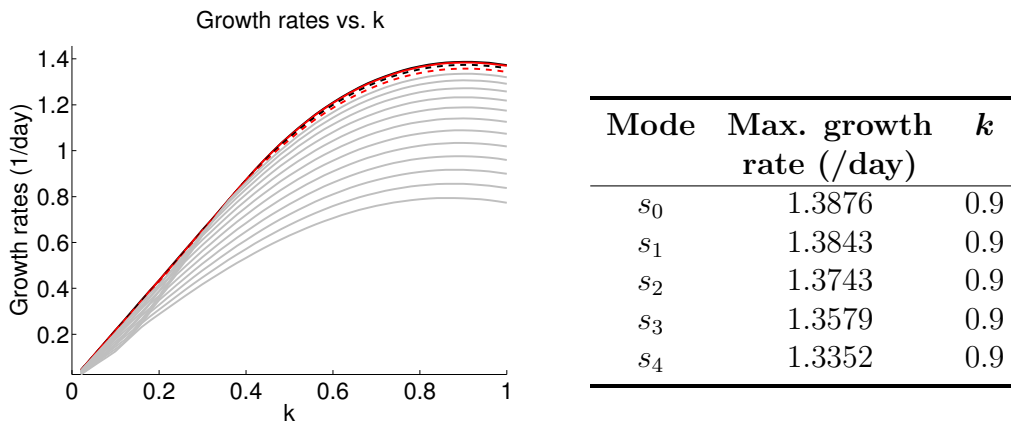


Figure 5.23: Plots of the growth rate curves (left) and a summary of the maximum growth rates and k (right) for the most unstable modes in the strongly stratified $Bu = 50$ case as predicted by the QG linear stability calculations for the case of the barotropic jet.

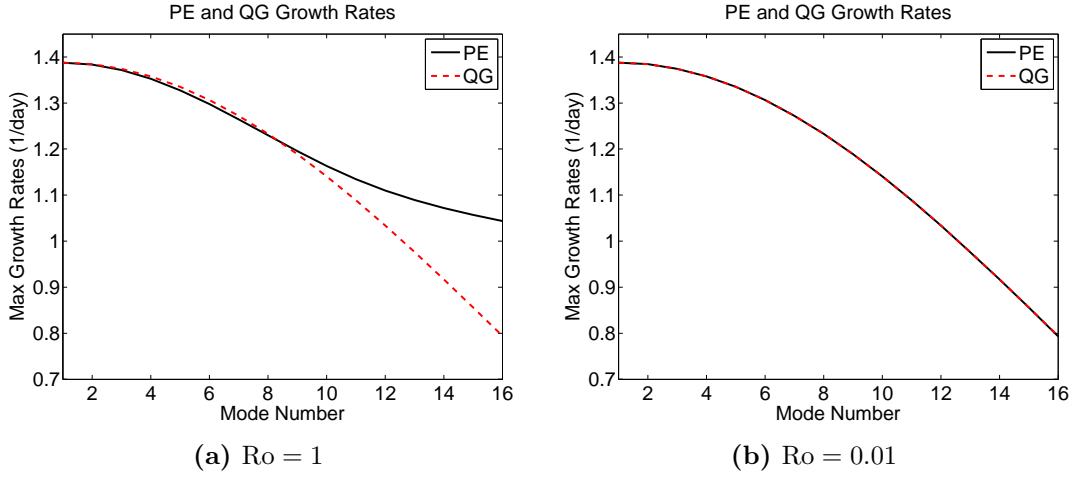


Figure 5.24: Plot of the maximum growth rates as a function of mode number as predicted by the PE and QG models, for the strongly stratified case of $Bu = 50$, when (a) $Ro = 1$ and (b) $Ro = 0.01$.

5.4 Baroclinic Jet: PE LSA Calculations

In the case of a baroclinic Bickley jet, we consider a flow that is periodic in the zonal direction with the same general domain definition as in section 5.2:

$$x = [0, L_x], \quad y = [0, L_y], \quad z = [h(y), H],$$

where $h(y)$ represents our potentially varying bottom topography and $H = 1$ km. The domain has a meridional extent of 100 km and we consider a range of wavenumbers.

The jet is centered in the middle of the channel and has characteristic velocity and length scales defined as $U_{\text{jet}} = 1$ m/s and $L_{\text{jet}} = 10$ km respectively. The background flow attributed to the jet is dependent on y and z and has the following definition:

$$u = U(y, z), \quad v = w = 0, \quad \rho = R(y, z), \quad p = P(y, z),$$

where

$$U(y, z) = U_{\text{jet}} \operatorname{sech}^2 \left(\frac{y - m_y}{L_{\text{jet}}} \right) \left(1 + \operatorname{erf} \left(\frac{z - H}{z_d} \right) \right)$$

and $z_d = 0.3 H = 300$ m. To compute the background pressure field we can use geostrophic balance to evaluate the pressure as

$$P(y, z) = - \int \rho_0 f U(y, z) dy. \tag{5.1}$$

Next, we can determine the density from quasi-hydrostatic balance as

$$R(y, z) = -\frac{1}{g} \frac{\partial}{\partial z} P(y, z) + \frac{f^* \rho_0}{g} U(y, z).$$

Upon substituting equation 5.1 for P we can rewrite this in the following form:

$$\begin{aligned} R(y, z) &= \frac{\rho_0 f}{g} \int \left(\frac{\partial}{\partial z} U(y, z) \right) dy + \frac{f^* \rho_0}{g} U(y, z) \\ \Rightarrow R(y, z) &= \frac{\rho_0}{g} \left[f^* U(y, z) + f \int \left(\frac{\partial}{\partial z} U(y, z) \right) dy \right]. \end{aligned}$$

We use the same set of physical constants as in the barotropic PE calculations:

$$g = 9.80616 \text{ m/s}^2, \quad f = f^* = 10^{-4} \text{ s}^{-1}.$$

Therefore, the $Ro = 1$ for these flows.

5.4.1 Convergence Analyses

Before performing a large set of SLEPc calculation, further convergence analyses were performed in specific application to the baroclinic jet case. From the convergence analyses of the barotropic jet cases covered in section 5.2.1 the default SLEPc solver tolerance of $1e-8$ and fourth order accurate FD approximations have been chosen. To assess the convergence of the SLEPc calculations to the grid resolution in the case of a baroclinic jet two convergence analyses were considered.

The first convergence study considered the case in which the number of grid cells varied but the horizontal and vertical resolutions were equal. The results are shown in table 5.4. It is observed that convergence does not occur as rapidly as in the barotropic case. However, the solution for $N = 200$ was accurate to four decimal places.

To consider the effect of vertical resolution on convergence we ran another test with fixed horizontal resolution of $N_y = 256$ and varied the vertical resolution N_z . We again used the default SLEPc solver tolerance. The results from this convergence study are shown in table 5.5. It is observed in this case that the convergence of the solutions show a strong dependence on the vertical resolution as well as the horizontal resolution. We also observe that a resolution $N_z = 128$ yields solutions that have converged to four decimal places. From these convergence studies, we have decided to consider a resolution of 256×256 for these baroclinic cases.

N	iters.	converged eigenvalue	total run-time
50	8	$6.117523\text{e-}6 + 1.552443\text{e-}5 i$	6s
100	6	$6.114726\text{e-}6 + 1.552534\text{e-}5 i$	17s
200	8	$6.114444\text{e-}6 + 1.552541\text{e-}5 i$	54s
300	8	$6.114423\text{e-}6 + 1.552541\text{e-}5 i$	3m1s
400	6	$6.114418\text{e-}6 + 1.552542\text{e-}5 i$	4m9s

Table 5.4: Test of the dependence of convergence of solutions to varying the grid resolution in both the horizontal and vertical $N_y = N_z = N$ for the baroclinic jet.

N_z	iters.	converged eigenvalue	total run-time
32	8	$6.122393\text{e-}6 + 1.552584\text{e-}5 i$	7s
64	8	$6.115128\text{e-}6 + 1.552537\text{e-}5 i$	19s
128	9	$6.114496\text{e-}6 + 1.552540\text{e-}5 i$	52s
256	6	$6.114428\text{e-}6 + 1.552541\text{e-}5 i$	1m48s

Table 5.5: Test of the dependence of convergence of solutions to varying the grid resolution in the vertical with fixed horizontal resolution $N_y = 256$ for the baroclinic jet.

5.4.2 Reference State

For our linear stability calculations a direct solver on a very coarse grid gives us initial guesses for our SLEPc solver. The SLEPc calculations are first performed on a coarser 128×128 grid. The solutions obtained from this coarser calculations are then used as the initial guesses to a finer 256×256 grid calculation. This technique was used to accelerate convergence. We begin with a reference calculation which has $Ro = 1$ and $Bu = 10$, synonymous to the parameters in the reference barotropic case 5.2.2. For these reference calculation 75 equally spaced wavenumbers in the range $[0, 1.5]$ were chosen.

The result for the coarser 128×128 SLEPc calculation versus those obtained from the coarse matlab direct solver is shown in figure 5.25. From this figure, it is observed that the direct solver gives much poorer results for the baroclinic case as compared to the barotropic case. As well, beyond a wavenumber of 1.2 the SLEPc calculation shows essentially no growth. This motivated the use of the 128×128 calculations as initial guesses to the finer 256×256 calculations.

The fine SLEPc calculation requesting the five most unstable modes required approximately 9.18 hours to complete using four processors. The calculation was able to solve for all the requested modes, however only the first two were physical. The growth rate curve

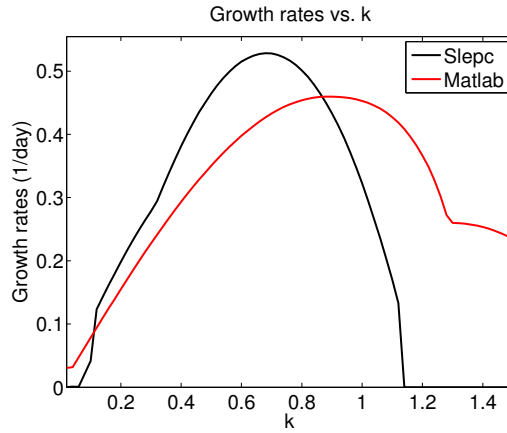


Figure 5.25: Comparison of growth rates using a 128×128 SLEPc computation versus the direct matlab solver on a very coarse 10×10 grid in the case of the reference baroclinic jet test.

showing the five most unstable modes is shown in figure 5.26. The two physical modes have maximum growth rates of 0.5283/day and 0.3859/day associated with wavenumbers 0.68 and 0.6 respectively. The third, fourth, and fifth modes had spurious modal structures.

In comparison to the reference barotropic calculation we see that the maximum growth rate of the most unstable mode is 2.04 times smaller than that predicted in the barotropic case. It is also observed that the wavelength of approximately 92.4 km is significantly larger than the 69.81 km wavelength predicted for the most unstable mode of the reference




Figure 5.26: Plots of the growth rate curves for the five most unstable modes in the reference baroclinic case as predicted by the PE linear stability calculations.

barotropic case. Comparable to the barotropic case, we observe a transition from the first mode to the second mode being most unstable at smaller wavenumbers.

The spatial structure of the most unstable mode, as predicted by the PE model, for the case of the maximum growth rate is shown in figure 5.27. The vertical structure of u, v, w, ρ, p and the vorticity fields are presented in subfigures 5.27a-5.27f, while the horizontal structure of the across channel velocity v and the vorticity are presented in subfigures 5.27g and 5.27h respectively. It is observed that this mode is symmetric about the flow, therefore it is characterized as a sinuous mode and we label this mode s_1 .

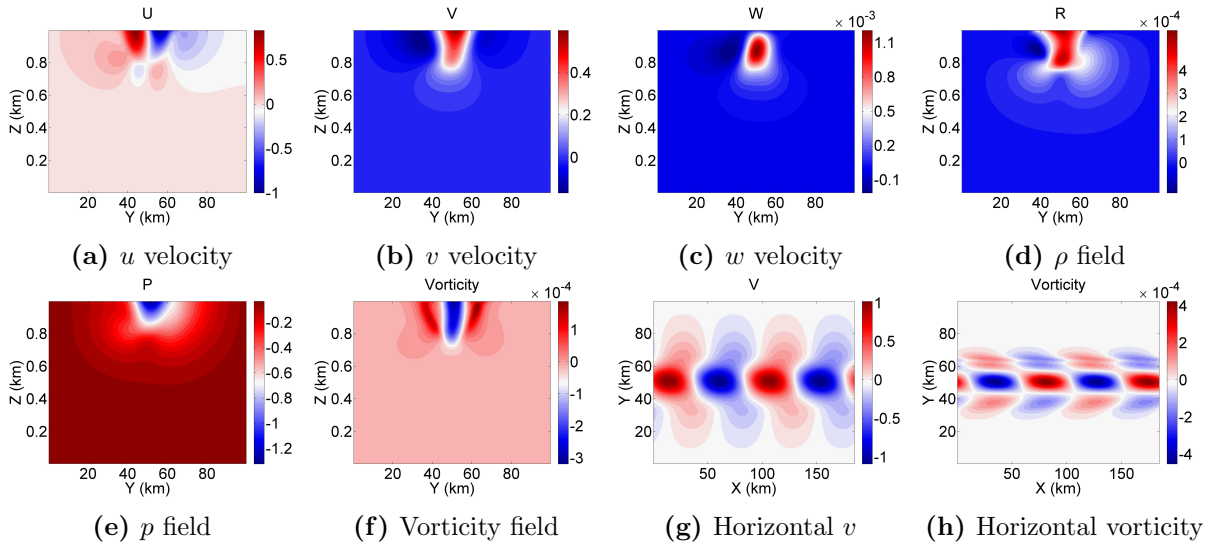


Figure 5.27: Subplots (a)-(f) present the vertical structure of the u, v, w, ρ, p and vorticity fields for the s_1 mode with the largest growth rate for the reference baroclinic jet calculation. In subplot (g) and (h), the horizontal structure of the across channel velocity v and the vorticity are shown.

There are fundamental differences between the most unstable mode in this reference baroclinic case as compared to that of the reference barotropic case. Specifically, it is observed that the mode structure includes non-QG effects. Furthermore, the relative vertical velocity is greater, which is a result of this jet having both strong horizontal and vertical shear [Lozier and Reed, 2005]. As non-QG effects are evident in the most unstable mode it is not expected that the QG calculations will agree with these PE calculations.

The spatial structure of the second mode associated with the largest growth rate is shown in figure 5.28 for this calculation. In comparison to the first mode, the second mode has much stronger asymmetries with the profile developing a Southward slant with depth.

In the horizontal vorticity plot, subfigure 5.28h, small scale features at the meridional extent of the instability are observed. These small scale features are more constrained in the North and have wider structures in the South.

It is difficult to classify this mode as sinuous or varicose, however from the growth rate curves 5.26 we saw that this unstable mode occurs over the same range of wavenumbers as the most unstable mode. Also, from the horizontal structures, subfigures 5.28g and 5.28h, we see that the maximum across channel velocity appears to occur near the center of the channel and the vorticity appears to be centered mid-channel. Therefore, this mode is classified as the sinuous s_2 mode.

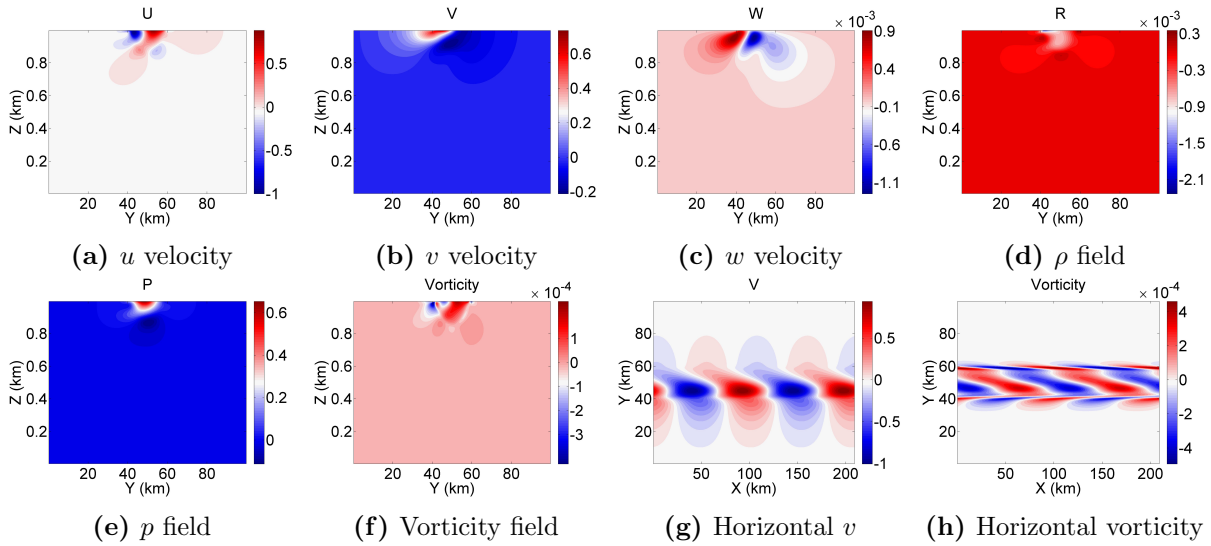


Figure 5.28: Subplots (a)-(f) present the vertical structure of the u, v, w, ρ, p and vorticity fields for the s_2 mode with the largest growth rate for the reference baroclinic case as predicted by the PE linear stability calculations. In subplot (g) and (h), the horizontal structure of the across channel velocity v and the vorticity are shown.

5.4.3 The Effects of Ro

To study the effects of Ro on the growth rates, calculations with decreased Ro were considered. As Ro decreases we tend more towards the QG regime. In the reference calculation strong non-QG effects were observed in both unstable modes. A reduction in Ro should predict more symmetrical modal structures. All other parameters were held constant, except the background velocity which was reduced by a factor of 10 and 100,

associated with $Ro = 0.1$ and $Ro = 0.01$ respectively. The predicted normalized growth rates for the first and second unstable modes are shown in subfigures 5.29a and 5.29b respectively.




Figure 5.29: Plots of the normalized growth rate curves for variable Ro as predicted by the PE linear stability calculations in the reference baroclinic case. Subfigure (a) and (b) show the results for the s_1 and s_2 modes.

It is observed that a reduction in Ro causes an increase in the growth rates for the most unstable mode. As well, the maximum growth rate mode has a shorter wavelength. On the other hand, the second mode has a growth rate that decreases and develops a spurious structure. Therefore, the second mode appears to have stabilized. The maximum normalized growth rates of the most unstable mode for the $Ro = 0.1$ and $Ro = 0.01$ cases are 0.6222/day and 0.6236/day respectively, both for a wavenumber of 0.82.

The spatial structures of the most unstable mode are presented in figure 5.30, where the top row represents the $Ro = 0.1$ case and the bottom row represents the $Ro = 0.01$ case. The columns from left to right represent the vertical across channel velocity, vertical vorticity, horizontal across channel velocity, and horizontal vorticity fields respectively. It is observed that a reduction in Ro does predict mode structures that are symmetric in y . Furthermore, both the $Ro = 0.1$ and $Ro = 0.01$ cases predict modes with comparable spatial structures (up to a correction in phase).




Figure 5.30: Plots of the vertical and horizontal across channel velocity and vorticity fields for the s_1 mode with the largest growth rate in the case of $Ro = 0.1$ (top row) and $Ro = 0.01$ (bottom row) for the baroclinic jet case as predicted by the PE linear stability calculations.

5.4.4 The Effects of Stratification

Next, the effects of stratification on the stability properties of the baroclinic jet were considered. In so doing, Ro was set to 1 and the background stratification was varied by considering a range of Bu , comparable to the barotropic case 5.2.3. First the background stratification was weakened by varying Bu from 10 to 5 in increments of 1. Decreasing Bu below a value of 5 produced spurious solutions. Subsequently, stronger stratification was considered by varying Bu from 10 to 100 in increments of 10. The growth rates for the most unstable mode over the range of wavenumbers and Bu are shown in figure 5.31




Figure 5.31: Plots of the growth rates as a function of wavenumber and Bu for the s_1 mode of the baroclinic jet in the case of $Ro = 1$ as predicted by the PE linear stability calculations.

It is observed from this figure that the tendency is for stronger stratification to destabilize the most unstable mode. Also, as stratification increases, the wavelength of the maximum growth rate mode tends to decrease. In the $Bu = 50$ case the maximum growth rate was calculated as 0.9610/day for a wavenumber of 0.82. This is a significant increase over the growth rates of 0.5283/day predicted in the reference baroclinic case. The spatial structure for this case is shown in figure 5.32. It is observed that, relative to the reference baroclinic case, that the mode is more surface intensified. This is caused by stronger stratification suppressing vertical shear and was also observed by [Lozier and Reed, 2005]. However, contrary to their findings, it has been observed that increased stratification acts to destabilize this mode.

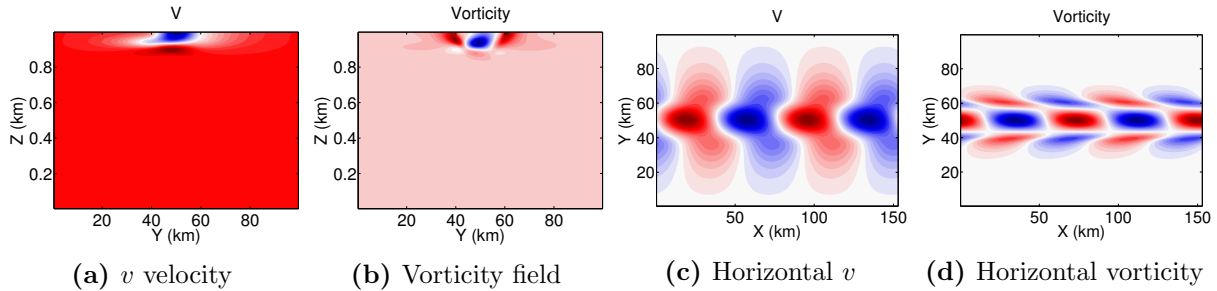


Figure 5.32: Plots of the vertical across channel velocity and vorticity fields, and the horizontal across channel velocity and vorticity fields in the baroclinic jet case as predicted by the PE linear stability calculations for the strongly stratified, $Bu = 50$, case.

5.4.5 The Effects of Topography

To consider the effect of topography, we consider the reference case and vary the topographic height by considering a range of a . This parameter defines the relative height of the topography and was defined in section 5.2.4. Positive a refers to prograde topography and negative a refers to retrograde topography. Since the background jet is surface-trapped, we expect that the topography would have less of an effect on the instabilities than in the case of the barotropic flow [Lozier and Reed, 2005, Barth, 1994].

We consider the same range of a as in the barotropic simulation, varying this parameter from -0.75 to 0.75 in increments of 0.05. The results for the most unstable mode are summarized in figure 5.33, with subfigure 5.33a showing the results for the full range of topographic heights and wavenumbers and subfigure 5.33b showing the maximum growth rates as a function of topographic height. We observe that for the majority of cases that the topography does have very little effect on the instabilities.




Figure 5.33: Plots of the growth rates as a function of topographic height and wavenumber for the s_1 mode of the baroclinic jet case as predicted by the PE linear stability calculations. The figure on the left shows the results for a full range of wavenumbers and the figure on the right shows the cases with largest growth rates.

In the most extreme cases, we see that prograde and retrograde topography has destabilizing effects on this mode. The maximum growth rates for the most extreme cases are computed as 0.5299/day and 0.5316/day for the prograde and retrograde profiles respectively, exhibiting a slight increase over the reference case which had a maximum growth rate of 0.5283/day. In these cases, the velocities at the bottom of the domain are approximately 4% of the maximum perturbation velocities. Therefore, the observed effect on the growth rate is likely due to the interaction of the background jet with the topography.

The vertical and horizontal spatial structures of the across channel velocity are shown for this mode in the extreme retrograde and prograde cases in subfigures 5.34a and 5.34b,




Figure 5.34: Plots of the vertical and horizontal structures for the across channel velocity for the s_1 mode in the case of (a) an extreme $a = -0.75$ retrograde topography. The subfigure (b) shows the same plots in the case of the extreme $a = 0.75$ prograde topography.

respectively. The spatial structure of these modes appear relatively unaffected by the topography, with only slightly more interactions with the topography observed in the retrograde case. It also appears that the vertical extent of the instability in the retrograde case is marginally larger than in the prograde case. The horizontal structures are virtually identical.

In the case of the second mode the results are comparable, as seen in figure 5.35. Subfigure 5.35a shows the results for the full range of topographic heights and wavenumbers and subfigure 5.35b shows the maximum growth rates as a function of topographic height. Over the range of shallow topographies we see that the growth rates of this mode are relatively unaffected. In the extreme cases, the retrograde topography has the effect of destabilizing this instability. However, the prograde topography acts to stabilize this mode. That being said, the effects are relatively small in that the maximum growth rates for the extreme retrograde (prograde) topographic profiles are 0.3866/day (0.3858/day), which is comparable to the reference case of 0.3859/day. The variations could be due to numerical errors or topography altering the definition of the background flow.




Figure 5.35: Plots of the growth rates as a function of topographic height and wavenumber for the s_2 mode of the baroclinic jet case as predicted by the PE linear stability calculations. The figure on the left shows the results for a full range of wavenumbers and the figure on the right shows the cases with largest growth rates.

The spatial structure for the second most unstable mode are shown in figure 5.36 with subfigures 5.36a and 5.36b showing the results for the extreme retrograde and prograde cases respectively. It is again observed that the prograde topography appears to cause the instability to be more surface intensified than the case of the retrograde topography. The horizontal spatial structure is consistent with that observed in the reference case 5.28.




Figure 5.36: Plots of the vertical and horizontal structures for the across channel velocity v for the s_2 mode in the case of an extreme $a = -0.75$ retrograde topography (a). The subfigure (b) shows the same plots in the case of the extreme $a = 0.75$ prograde topography.

5.5 Baroclinic Jet: QG LSA Calculations

In the case of the QG model, we consider the generalized eigenvalue problem 3.20. This is equivalent to that considered in the barotropic case 5.3 except now the background velocity and vorticity are dependent on the vertical coordinate z as well as the across channel direction y .

5.5.1 Reference State

The reference baroclinic case in for the QG calculations is consistent with that used in the PE calculation 5.4.2. Since $Ro = 1$ the QG equations are not expected to perform particularly well for this test case. As well, the changes in the isopycnals are large in the case of a surface-trapped front and so this could lead to further deviations between the QG and PE results. In the case of the QG calculation a grid resolution of 256×256 is considered, equivalent to that used in the PE calculations. Since only the first two modes were physical in the PE reference calculation, the QG calculation was performed with the request for the two most unstable modes. Computing the results for 75 equally spaced wavenumbers in the range $[0, 1.5]$ required only 0.51 hours.

The growth rate curves for these two most unstable modes are shown in figure 5.37. It is observed that the most unstable mode is the only physical mode predicted by the QG calculation and it has a growth rate curve that is similar to that predicted by the PE calculations. However, the most unstable mode as predicted by the QG calculation has a maximum growth rate of 0.5314/day for a wavenumber of 0.5. Therefore, the growth rate is 0.59% larger than that predicted by the PE calculation and the wavelength is 36% longer. This means that the QG model captures the growth rate quite well, but is significantly off




Figure 5.37: Plots of the growth rate curves for the s_1 and s_2 modes in the reference baroclinic case as predicted by the QG linear stability calculations.

on the wavelength that achieves this growth rate.

The vertical across channel velocity and vorticity structures are shown in subfigures 5.38a and 5.38b. Their associated horizontal structures are shown in subfigures 5.38c and 5.38d respectively. Since the QG model cannot capture non-QG effects it is observed that the structures predicted by the QG calculations are symmetric in y . Furthermore, the QG model predicts this mode to be more surface intensified compared to the PE prediction.

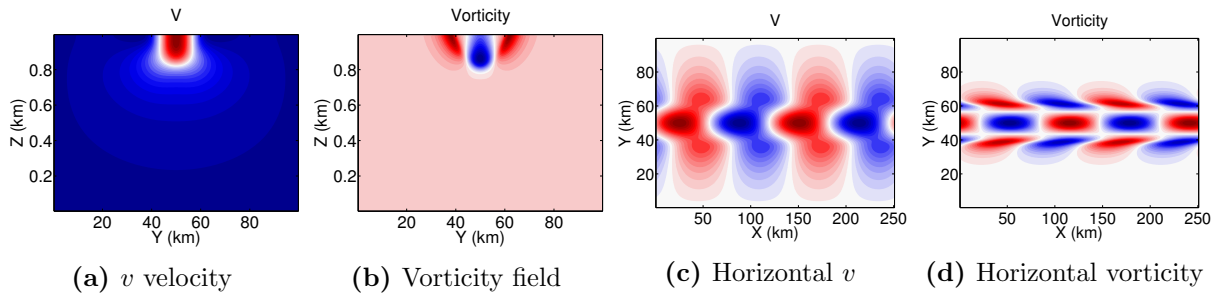


Figure 5.38: Plots of the vertical v (a) and vorticity (b) for the s_1 mode with the maximum growth rate as predicted by the QG case. Subplots (c) and (d) present the horizontal structures of v and vorticity respectively.

5.5.2 The Effects of Ro

The QG model is independent of Ro and so the calculations would lead to the same growth rate curves, scaled by Ro . Therefore, the maximum growth rates for the unstable mode in the QG calculation would have a normalized value of 0.5314/day for wavenumber 0.5. In addition, the spatial structure of the most unstable mode is identical to the case with $Ro = 1$. In the case of the PE calculations in section 5.4.3, the normalized growth rates for the $Ro = 0.1$ and $Ro = 0.01$ cases were predicted as 0.6222/day and 0.6236/day respectively for a wavenumber of 0.82.

Therefore, we conclude that a decrease in Ro does not necessarily mean better agreement between the QG and PE calculations in terms of the growth rates or wavelengths. The QG model is expected to have performed poorly for this particular calculation because of the large variation in the isopycnals. The spatial structures of the PE calculations with reduced Ro do, however, agree better with the QG predictions as seen in the reference case. The reduced Ro causes a decrease in the isopycnal tilt and it is proposed that this is the mechanism behind the increase in symmetry in the PE model predictions.

5.5.3 The Effects of Stratification

Since the decrease in Ro does not appear to yield much better agreement between the QG and PE calculations, we next consider the effects of stratification on the stability properties of the jet. Weaker and stronger background stratifications were parameterized by a range of Bu . The QG calculations yielded physically meaningful results up to the weakest stratification of $Bu = 1$, however the results are presented here up to $Bu = 5$ for comparison to the PE results 5.4.4. The growth rates for the most unstable mode over the full range of wavenumbers and Bu are shown in figure 5.39.

Plots of the maximum growth rates and associated wavenumbers as a function of Bu are shown in subfigures 5.40a and 5.40b respectively. To get a better understanding of how the stratification affects the agreement between the QG and PE results the PE calculations have been superposed. From these plots, it is observed that there is a very strong agreement in the predicted growth rates for stratification associated with $Bu \geq 10$. On the other hand, only in the cases of very strong stratification, $Bu \geq 50$, do the wavenumbers show a strong agreement.

Figure 5.40 suggests that the case of $Bu = 50$ shows strong agreement between the PE and QG calculations. Therefore, we consider this case in greater detail consistent with that analyzed in the PE calculations. For this stratification, the maximum growth rate as




Figure 5.39: Plots of the growth rates as a function of wavenumber and Bu for the s_1 mode of the baroclinic jet in the case of $Ro = 1$ as predicted by the QG linear stability calculations.

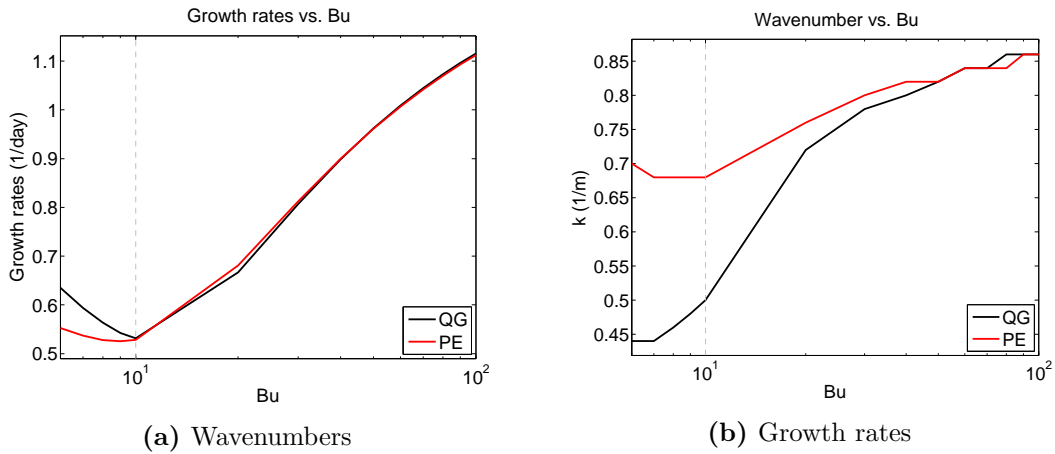


Figure 5.40: Plots of the growth rates and wavenumbers associated with the s_1 modes with maximum growth rates for the baroclinic jet case with varying stratification as predicted by the QG calculations. Subfigure (a) shows the maximum growth rates as a function of Bu and subfigure (b) shows the associated wavenumber at each stratification. Superposed is the results from the PE calculations for comparison.

predicted by the QG calculations is 0.9620/day for a wavenumber of 0.82. Therefore, similar to that observed in the barotropic jet case 5.3.2, the QG and PE predictions agree very well in the case of strong stratification. The calculated maximum growth rate only differs by 0.1% over the result predicted by the PE calculations and the QG and PE predictions share the same wavelength. The spatial structure of the vertical across channel velocity

and vorticity fields and the horizontal across channel velocity and vorticity fields are shown in figure 5.41. Comparing to the spatial structure predicted by the PE model as presented in section 5.4.4, we see that the spatial structures agree very well.

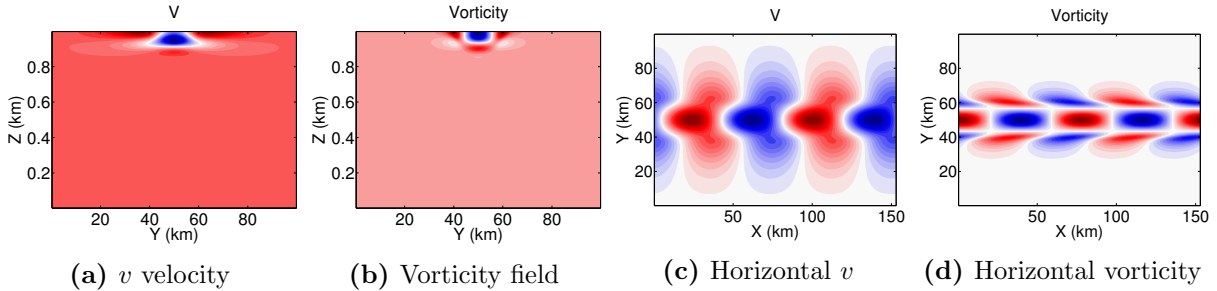


Figure 5.41: Plots of the vertical across channel velocity and vorticity fields and the horizontal across channel velocity and vorticity fields as predicted by the QG calculations for the strongly stratified case with $Bu = 50$.

5.6 Conclusions

In this chapter, linear stability calculations were used to perform a series of parameter tests in application to unstable barotropic and baroclinic jets. The benefit of linear stability calculations is that they can efficiently calculate the growth rates of the various unstable modes. Where nonlinear simulations can require many days to complete, linear stability calculations can be performed in less than an hour with significantly fewer processors allocated to the calculation. From these calculations, we can isolate particular problems of interest and perform in-depth nonlinear simulations to get a better understanding of the evolution of the instabilities that arise.

5.6.1 Barotropic Jet

In section 5.2, linear stability calculations were performed for a particular barotropic jet configuration. To determine the impact on the stability characteristics of this particular barotropic jet the background stratification and topographic heights were varied. Results of QG linear stability calculations were later presented in section 5.3 and comparisons were made to the PE results.

Reference Case

Initially, a reference calculation was performed which considered a jet in a domain that was

100 km wide and 1 km deep. The jet had a characteristic width of 10 km and a background velocity of 1 m/s. The stratification for this reference calculation was set by $Bu = 10$ and the Rossby number was set to $Ro = 1$. Though the QG approximation need not apply, it was found that the PE and QG calculations showed very strong agreement particularly for the first two modes of instability. The most unstable mode was predicted to have a growth rate of 1.3877/day and 1.3876/day, for the PE and QG calculations respectively, both for a wavelength of 69.81 km. This mode was a barotropic sinuous mode.

Both the PE and QG calculations predicted a range of unstable modes, with varying degrees of vertical structure. The second most unstable modes were predicted to have growth rates of 1.2979/day and 1.3065/day by the PE and QG calculations both again for a wavelength of 69.81 km. This mode was characterized by a mode 1 vertical structure. As we considered higher mode number modes the PE and QG predictions differed by increasing amounts. The PE model predicted non-QG effects in the form of asymmetries in the modal structures. The growth rate of the next mode was predicted to be 1.1349/day and 1.0888/day by the PE and QG calculations.

Varying Stratification

Varying the stratification had very little effect on the barotropic sinuous mode in both the PE and QG calculations. However, as stratification was increased, it was observed that a range of modes converged to comparable growth rates in both the PE and QG calculations. In particular, for the strongly stratified $Bu = 50$ case, the PE and QG predictions agreed very well up to the ninth mode. From the ninth mode onward the PE model predicted growth rates that were greater than those predicted by the QG model, similar to the trend observed in the reference case. On the other hand, weak stratification of $Bu = 1$ led to a dramatic decrease in the growth rates of secondary modes in both the PE and QG calculations. However, in this case the PE model predicted 3 unstable modes while the QG predicted only 2.

Varying Topographic Height

To consider a range of topographic heights, predictions using the PE calculations were performed. It was found that the most unstable s_0 mode was only affected by a small amount ($< 2\%$) consistent with the findings in [Lozier and Reed, 2005]. However, retrograde (prograde) topography was always found to have a stabilizing (destabilizing) effect on the s_0 mode. By comparison, for all other modes considered, both strong retrograde and prograde topography showed a stabilizing effect. Strong prograde topography appeared to have a strong stabilizing effect on all the baroclinic modes considered, while strong retrograde topography had a greater stabilizing effect on those modes with weaker vertical structure.

Summary

Therefore, the linear stability calculations in these sections validated and expanded on previous studies [Lozier and Reed, 2005, Irwin and Poulin, 2014, Poulin and Flierl, 2005]. By considering both PE and QG calculations we were able to study how the calculations agreed in the case of the barotropic jet and at what point the agreement began to break down. From the spatial structures of the modes, we were able to determine how non-QG effects presented themselves in the PE predictions. Furthermore, as many modes were considered we were able to get an in-depth look into how the stratification and topography affected the various modes of instability adding to previous knowledge.

5.6.2 Baroclinic Jet

In section 5.4 we computed the stability characteristics for a particular baroclinic jet configuration. This particular jet had a background flow of 1 m/s and characteristic depth of 300 m in a domain that was 1 km deep. Similar to what was done in the barotropic calculations, the jet had characteristic width of 10 km and was defined in a channel that was 100 km across. The reference jet was characterized by $Ro = 1$ and $Bu = 10$.

In sections 5.4 and BCQGLSA, linear stability calculations were performed using the PE and QG models respectively and it was found that there was fairly strong disagreements between the two models. The QG model performed really well in the barotropic jet cases even for the parameter studies that were outside the QG regime, however the same parameters gave strong disagreement in the baroclinic case. This was attributed to the strong variations in the isopycnals of the background jet.

Reference Case

In the reference case, the PE calculations captured two modes of instability having maximum growth rates of 0.5283/day and 0.3859/day respectively. The wavelengths associated with these growths were 92.4 and 104.72 km respectively. By comparison, the QG model only captured one mode of instability in this reference case. The QG model predicted the growth rate as 0.5314/day which showed strong agreement with the PE growth rate, but the wavelength of 125.67 km was significantly longer. Reducing Ro did not promote greater agreement between the results from the two models, though it did give indication that the secondary mode of the PE model could not be obtained in the QG model.

Varying Stratification

In the case of strong stratification for which $Bu = 50$ it was found that there was a significant increase in the growth rates as predicted by both models. Furthermore, in cases of strong stratification the PE and QG models agreed very well. In the specific case for $Bu = 50$ the PE and QG models predicted growth rates of 0.9610/day and 0.9620/day both for a

wavelength of 76.62 km. In both the PE and QG model the modal structure showed a surface intensification.

Varying Topographic Height

To test the impact of topography a range of topographic heights were considered using an idealized hyperbolic tangent coastal configuration. The most unstable mode was unaffected by the topography in all cases but the most extreme. In these cases, it was hypothesized that the observed differences were predominantly caused by the alteration of the definition of the background jet profile than actual interactions with the topography. The secondary mode showed similar behavior.

Summary

Therefore, the linear stability calculations in these sections expanded on the results of [Irwin and Poulin, 2014, Lozier and Reed, 2005] by considering both first and second modes of instability, as well as a comparison between PE and QG calculations. We were able to determine that the second unstable mode was not able to be captured by the QG model. By considering a continuously stratified domain we were able to obtain information about the horizontal and vertical spatial structure of these modes. As well, we observed that stronger stratification appeared to increase the growth rates of the most unstable mode, contrary to findings of [Lozier and Reed, 2005]. Topography had very little affect on the modes of instability, which agreed with the results of [Lozier and Reed, 2005].

Chapter 6

Nonlinear Simulations

In this chapter we focus on nonlinear simulations for a specific set of barotropic and baroclinic jet test cases. The choice of which tests to consider were guided by the results of the previous chapter. The benefit of combining nonlinear simulations with linear stability calculations is that it offers us a means to validate the two sets of codes. Also, it sheds some insight into how easily these modes form and grow in a nonlinear context. Furthermore, nonlinear simulations allow us a way to analyze the behavior for later times, when linear theory need not apply.

6.1 Diagnostics

In the discussion of the nonlinear simulations, the growth rates as predicted by the nonlinear simulations will be compared to those from the linear stability analysis. As well, the incremental changes in the potential and kinetic energy and the relative energy loss are shown for each simulation. The change in mass is not discussed as the FV method ensures conservation of mass. There is however rounding error but these were considered negligible. In all the nonlinear simulations presented in this chapter the mass was conserved to within $1e-9\%$.

6.1.1 Growth Rate Curves

To define the growth rate curves, we plot the L1-norm of the horizontal velocity perturbations and the density perturbations as a function of time at each time step. From the linear

theory, we expect that the perturbations are governed by exponential growth and so we can fit an exponential curve to the perturbations to determine the growth rate. That is, we expect that the perturbations have the following form:

$$[u', v', \rho'] = [u'_0, v'_0, \rho'_0] \exp(\gamma t + \mathcal{C}),$$

where γ is the growth rate and \mathcal{C} is some real-valued constant.

To fit the curve we take the natural logarithm of the above equation and fit the linear function $\gamma t + \mathcal{C}$ to the data. To determine the maximum growth rate, a numerical approach is taken. A particular span of n days is set and the linear function is fit to all possible periods of n days. This yields a series of γ values and the maximum growth is determined by the maximum of these values.

6.1.2 Mechanical Energy Equation

An equation governing the evolution of the mechanical energy can be derived from the Boussinesq equations. In the absence of viscosity, it is found that the sum of the kinetic energy and the potential energy are conserved. To begin, the equation for the kinetic energy density is found by taking the dot product of the momentum equation with the velocity. Taking the sum of the resulting equations we obtain the following:

$$\rho_0 \frac{D}{Dt} \left\{ \frac{u^2 + v^2 + w^2}{2} \right\} = - \left(u \frac{\partial p'}{\partial x} + v \frac{\partial p'}{\partial y} + w \frac{\partial p'}{\partial z} \right) - g\rho'w.$$

The first set of terms on the right hand side represents the convergence of energy flux due to the work done by pressure [Vallis, 2006] and the second term is the conversion to potential energy. The potential energy term is combined into the material derivative on the left hand side by rewriting w as Dz/Dt . We then integrate over the entire domain V to obtain:

$$\frac{D}{Dt} \int_{\Omega} \left\{ \rho_0 \frac{\mathbf{u} \cdot \mathbf{u}}{2} + g\rho'z \right\} dV = - \int_{\partial\Omega} p' \mathbf{u} \cdot \mathbf{n} dS.$$

where we have made use of the continuity equation and the divergence theorem to rewrite the pressure term. In the above dS and dV represent differential area and volume elements. In terms of a discretized grid we can therefore determine the mechanical energy of the system as the sum of the mechanical energy density times the element volumes.

In the absence of sources or sinks, the pressure term is identically zero and we have that

$$\frac{D}{Dt} \int_{\Omega} \left\{ \rho_0 \frac{\mathbf{u} \cdot \mathbf{u}}{2} + g\rho'z \right\} dV = 0.$$

This can be written in terms of a mechanical energy density:

$$E = \rho_0 \frac{\mathbf{u} \cdot \mathbf{u}}{2} + g\rho'z.$$

E is conserved exactly when there is no numerical and physical dissipation. However, in the case of numerical approximations it is expected that there will be a net loss of energy due to numerical errors. As well, there is physical dissipation that occurs predominantly at small scales through the conversion of kinetic energy to internal energy.

In the case of stratified flows we can rewrite the above conservation equation using the following alternative. By multiplying the density equation by ρ' we obtained the following:

$$\frac{D}{Dt} \left(\frac{\rho'^2}{2} \right) = \frac{\rho_0 N^2}{g} \rho' w.$$

Therefore, in cases when N is considered constant, we conclude that the equation for the evolution of the perturbation potential energy density is

$$\frac{D}{Dt} \left(\frac{g^2}{2\rho_0 N^2} \rho'^2 \right) = g\rho' w,$$

which is written in terms of the square of the perturbation density. Combining this result with the equation for the evolution of kinetic energy above we determine the equation that evolves the mechanical energy density in the Boussinesq system:

$$\frac{D}{Dt} \left\{ \rho_0 \frac{\mathbf{u} \cdot \mathbf{u}}{2} + \frac{g^2}{2\rho_0 N^2} \rho'^2 \right\} = 0. \quad (6.1)$$

To summarize, the kinetic and potential energy can be obtained from the following discrete equations:

$$KE = \int_V \rho_0 \frac{\mathbf{u} \cdot \mathbf{u}}{2} dV, \quad PE_1 = \int_V g\rho'z dV, \quad PE_2 = \int_V \frac{g^2}{2\rho_0 N^2} \rho'^2 dV.$$

All the simulations in this chapter are defined with background stratification so the following definition of total energy is used throughout:

$$TE = KE + PE_2.$$

However, as the derivative of PE_1 and PE_2 must agree, we see that these two definitions of PE can differ by at most a constant.

From the above conservation equation 6.1, we expect that TE will be conserved. This implies that a loss (gain) of PE_2 will coincide with a gain (loss) of KE . Therefore, when discussing the evolution of the energy we will consider the incremental changes in the kinetic and potential energy ($\Delta KE, \Delta PE$) and the relative total energy loss defined as:

$$rel. TE = \frac{TE_n}{TE_0} - 1,$$

where TE_n refers to the total energy at the n th time step.

6.2 Barotropic Jet

To study mesoscale instabilities in the oceans there are two broad classes of instabilities that can arise, those being barotropic and baroclinic instabilities [McWilliams, 2006]. Barotropic instabilities extract kinetic energy from the basic state. This is in contrast to baroclinic instabilities that extract both potential and kinetic energy from the basic state. We begin the analysis of these types of instabilities by considering an idealized barotropic jet. This type of jet is characterized by horizontal shear, does not vary with depth, and as observed in the linear stability calculations the primary mode of instability tends to be barotropic. These types of jets give rise to barotropic instability and this can manifest itself as barotropic or baroclinic modes. As observed in the LSA chapter, the baroclinic modes can have growth rates comparable to the sinuous barotropic mode.

Barotropic jets are an intuitive first step in the simulation of coastal jets because they are much simpler to solve numerically when compared to baroclinic jets. One example where these jets may be applicable is in the case of the Antarctic Circumpolar Current (ACC) [Killworth and Hughes, 2002]. From the thermal wind balance, it is necessary that the jet does not alter the background stratification and so the dynamics are easier to capture than in the case of a strongly baroclinic jet. However, of particular interest is how the modes of instability evolve both in the linear and nonlinear regimes. Particular attention is given to determining what factors contribute to the decrease in vertical length scales that can arise.

From the previous chapter, we expect that the PE and QG simulations will show strong agreement in the reference case as the most unstable mode was found to have a comparable growth rate for both models. As the background stratification is increased we would expect greater vertical structure in both the PE and QG simulations. On the other hand, the weak stratification cases should show very little vertical structure. In the cases with topography it is expected that the PE model will predict growth rates comparable to those observed in

the reference calculations, with the prograde topography potentially having a greater effect on the structure through the secondary modes.

6.2.1 Hydrostatic Reference Simulation

The same definition of the jet is used as that of the linear stability calculations in section 5.2. The jet is centered in the middle of the channel and has characteristic velocity and length scales defined as $U_{\text{jet}} = 1 \text{ m/s}$ and $L_{\text{jet}} = 10 \text{ km}$ respectively. We have a background profile that is independent of depth and varies only across channel. The background flow is defined as

$$u = U(y), \quad v = w = 0,$$

where

$$U(y) = U_{\text{jet}} \operatorname{sech}^2 \left(\frac{y - m_y}{L_{\text{jet}}} \right).$$

As the traditional approximation is made, the jet does not yield variation in the density field. The following physical constants are used throughout the barotropic simulations:

$$g = 9.80616 \text{ m/s}^2, \quad f = 10^{-4} \text{ s}^{-1}.$$

Using the Tempest model, we compute the evolution of the instabilities using the HY equations. As the most unstable mode is barotropic the HY equations are expected to capture the dynamics of the instabilities. It will later be shown that the NH and HY models agree very well in the case of the reference baroclinic jet simulation for which NH effects are more likely to appear due to the density variations. As well, motivation to use the HY model lies in their computational simplicity and efficiency.

To stimulate the onset of instabilities, a perturbation of random noise with relative magnitude 0.1% of the background flow is applied to the u velocity. The default solver tolerances of 1e-8 are used for the calculations and a grid resolution of $256 \times 256 \times 256$ is considered. A time step of $dt = 50 \text{ s}$ is used to ensure that we have at least 10 time steps per inertial period $2\pi/N \sim 600 \text{ s}$. This inertial period poses the strictest constraint on the time step. The value of $N = 1\text{e-}2 \text{ s}^{-1}$ implies a top to bottom density difference of 10 kg/m^3 over the 1 km domain depth.

On 128 cores, the computation required 53.59 hours for a total of 43200 time steps, representing 25 days of integration. The viscosity coefficients for this particular simulation are set to $\nu_h = 0 \text{ m}^2/\text{s}$ and $\nu_v = 1\text{e-}3 \text{ m}^2/\text{s}$, used to suppress some small scale oscillations in

the vertical produced by the random perturbations on or near the grid scale. The domain is defined as

$$0 \text{ km} \leq x \leq 139.62 \text{ km}, \quad 0 \text{ km} \leq y \leq 100 \text{ km}, \quad 0 \text{ km} \leq z \leq 1 \text{ km},$$

where the zonal extent has been set by double the wavelength of the most unstable mode as predicted by the linear theory. This is a choice based upon the linear calculations and will emit modes with wavelengths that are integer divisors of this zonal extent.

For this particular simulation, the growth rate plot is shown in subfigure 6.1a and the changes in the energy are shown in subfigure 6.1b. For the growth rate curve, we fit an exponential function to the L1-norms of the u , v , and ρ perturbation fields as explained in subsection 6.1.1. The fitted curves are shown as the solid lines in subfigure 6.1a. For the energy plot, we consider the incremental changes in the kinetic and potential energy (ΔKE , ΔPE) and the relative total energy loss as defined in 6.1.2. For this particular simulation, the total energy loss over the 25 days of integration is 2.0228%. Note that during the phase of exponential growth, first 10 days, the energy is better conserved. After the phase of nonlinear equilibration, the energy decreases the most. Therefore, the vortical motions cause the decrease in total energy.

From the fitted exponential functions, we compute the growth rate to be 1.2776/day which is approximately 7.93% lower than that predicted by the linear theory. If we consider

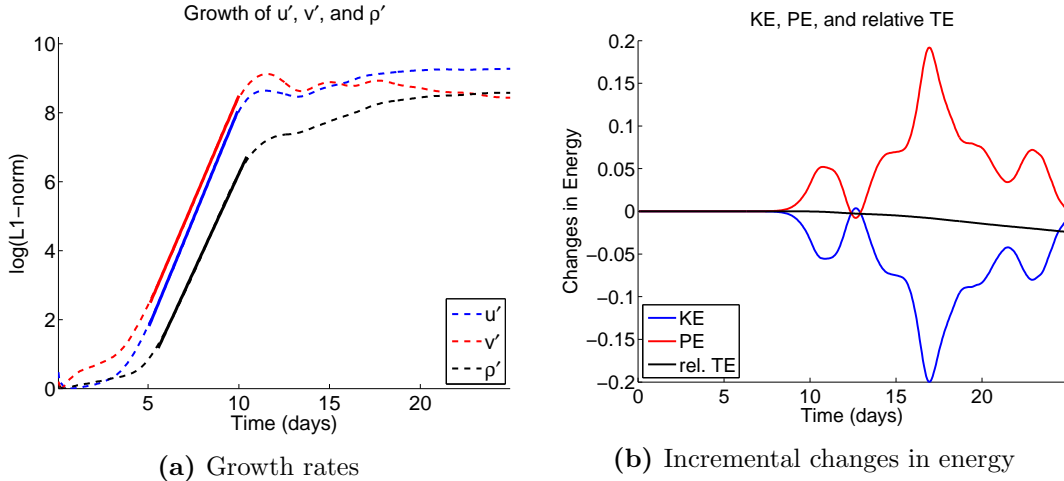


Figure 6.1: Plots of (a) the growth rates and (b) incremental changes in the energy and relative energy loss for the reference barotropic simulation with $Ro = 1$ and $Bu = 10$ as predicted by the Tempest HY model.

a particular snapshot of the perturbation vorticity fields during the period of maximum growth we observe the structures shown in figure 6.2 at day 8. It is clear from subfigure 6.2a that the instability shows a mode 1 vertical structure. From the linear calculations, it was predicted that this s_1 mode had a comparable growth rate to the s_0 mode and both are presumed to be captured by the nonlinear simulations. The excited wavelength as demonstrated from the horizontal structure of the instability, shown in figure 6.2b is consistent with that predicted in the linear theory. However, the wavelengths of the two unstable modes show slight differences.

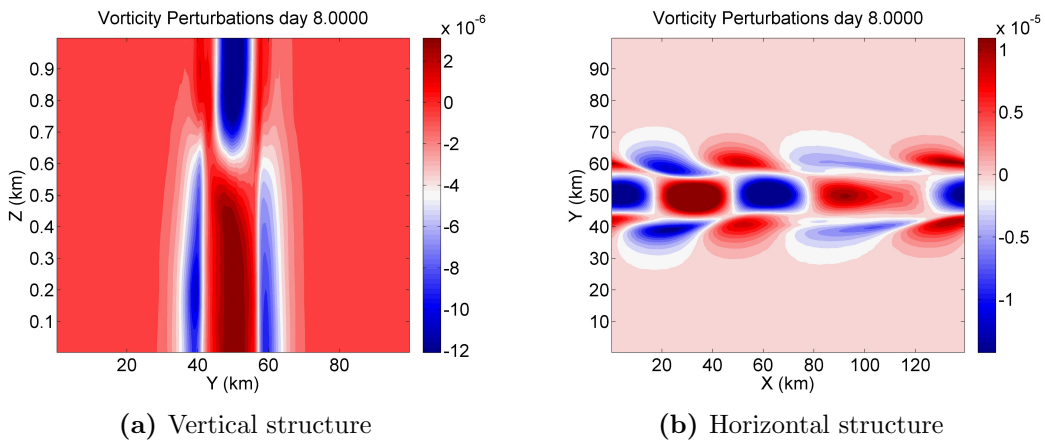


Figure 6.2: The structure of the perturbation vorticity field at a snapshot during the period of maximum growth. Subfigure (a) shows the vertical structure of the instability and subfigure (b) shows the horizontal structure of the instability.

Snapshots of the evolution of the vorticity field taken near the surface (relative height of 0.9) are shown for the 25 days of integration in figure 6.3. As this flow is for $Ro = 1$ slight asymmetries develop in the solution. The results are displayed in 5 day increments. As is predicted by the growth rate curve, by day 5 the perturbations have not grown enough to significantly alter the vorticity field. However, by day 10 the perturbations have undergone significant exponential growth and mesoscale vortices have begun to form. These vortices are advected along and subsequently merge, as observed by day 20. By the end of the integration the merged vortices are interacting with the boundary.

To consider the vertical structure of the vorticity field we plot isosurface plots for day 10 of integration, near when the mesoscale vortices are first starting to form, as shown in figure 6.4. Subfigures 6.4a, 6.4b, and 6.4c show the barotropic, baroclinic, and total vorticity fields. It is observed that the barotropic mode is slightly larger in magnitude than the baroclinic mode, consistent with the stronger growth rate of the barotropic mode. As

well, we see that the baroclinic mode has both complex horizontal and vertical structures, but there does appear to be a mode one vertical structure that dominates this field, as we anticipated.




Figure 6.3: Snapshots of the evolution of the vorticity field in the reference barotropic case as predicted by the Tempest HY model. The results are displayed in 5 day increments.

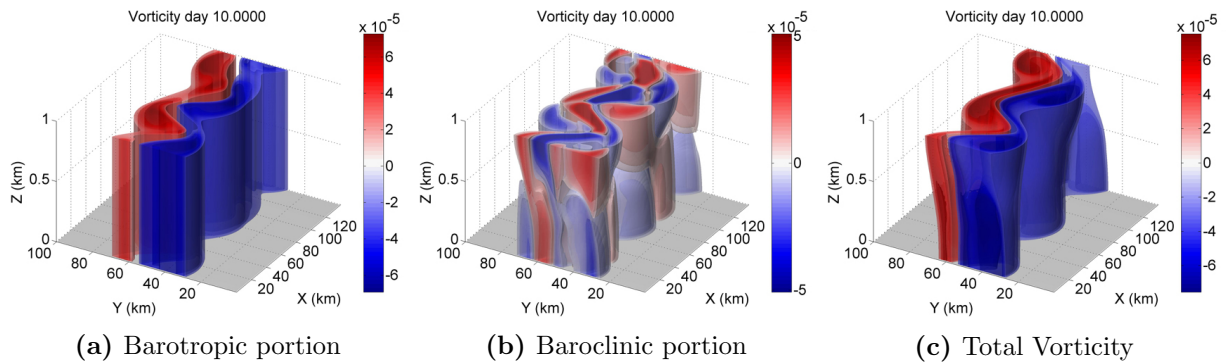


Figure 6.4: The three-dimensional structure of the vorticity field in the reference barotropic simulation as predicted by the Tempest HY model at day 10. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.

6.2.2 Quasi-Geostrophic Reference Simulation

For comparison, we consider a simulation of the QG model. For this simulation a grid of $256 \times 256 \times 128$ was used. As the QG model does not admit inertial gravity waves, there is not the same limitation on the time step as that of the HY model. In this case, a time step of $dt = 100$ s was considered, imposed by the CFL condition, and the model was integrated for 25 days. Though less complex than the NH equation set, the QG model does require solving a three-dimensional linear system. On 64 processors, and a total of 21600 time steps, the computation required 64.4 hours. To accelerate the computation and for stability the solver tolerance was increased to $1e-6$. Using the default tolerance of $1e-8$ caused the solution to diverge after a period of approximately 15 days. As well, HYPRE was used as other solvers did not yield stable results. To ensure that the perturbations were large enough to be captured by the reduced tolerance, random perturbations with a relative magnitude of 1% were applied to the vorticity field.

To determine the growth rate in the case of the QG model, exponential curves were fit to the u, v , and vorticity fields. The growth rate curves are shown in subfigure 6.5a. In this case, the computed maximum growth rate is 1.2500/day which is very close to the predicted value of the HY model of 1.2776/day from the previous subsection. As well, the wavelength has been verified to be consistent with that predicted in the HY simulation and linear stability calculations. The incremental changes in energy and relative total energy

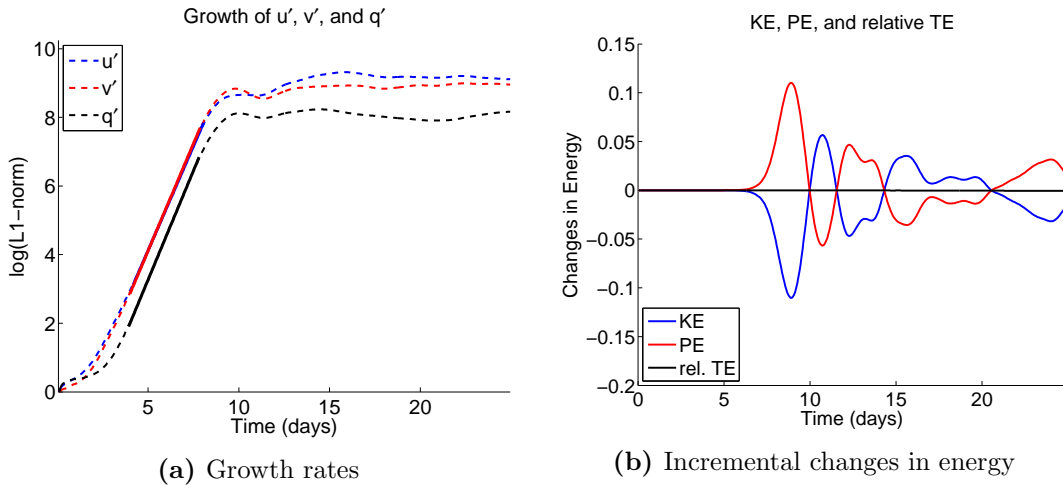


Figure 6.5: Plots of (a) the growth rates and (b) incremental changes in the energy and relative energy loss for the reference barotropic simulation with $Ro = 1$ and $Bu = 10$ as predicted by the Tempest QG model.

loss is shown in subfigure 6.5b. The relative energy loss is 0.056%, which is significantly smaller than that calculated in the HY reference simulation.

Snapshots of the evolution of the vorticity field as predicted by the QG model are shown in figure 6.6. We see that the general features in the evolution are equivalent to that of the HY model; that of mesoscale vortices arising from the exponential growth of the perturbations, followed by vortex merging. Though the characteristic length scale of the mesoscale instabilities appear to agree in the two simulations, there are a number of small, filament-like motions in the QG simulation that are not observed in the HY simulation. This demonstrates a qualitative difference between the QG and HY results in this particular test case.




Figure 6.6: Snapshots of the evolution of the vorticity field in the reference barotropic case as predicted by the Tempest QG model. The results are displayed in 5 day increments.

To make a consistent comparison of the vertical structure with that of the HY model, we consider isoplots of the vorticity field at a time associated with the end of the period of exponential growth. We consider three-dimensional plots at day 8 as shown in figure 6.7 where subfigures 6.7a, 6.7b, and 6.7c show the barotropic, baroclinic, and total vorticity fields. It is observed that the QG model predicts that the baroclinic modes are stronger than in the case of the HY model. Even though the third and higher mode number modes were predicted to have a weaker growth rate in the case of the QG model over the PE model the mode 1 baroclinic sinuous mode did have a slightly larger growth rate in the case of the QG linear stability calculations, 1.3065/day compared to 1.2979/day, which could be

a potential explanation for the stronger baroclinic motions. The barotropic instabilities 6.7a do show smaller scale features in the QG simulation as compared to the HY model.

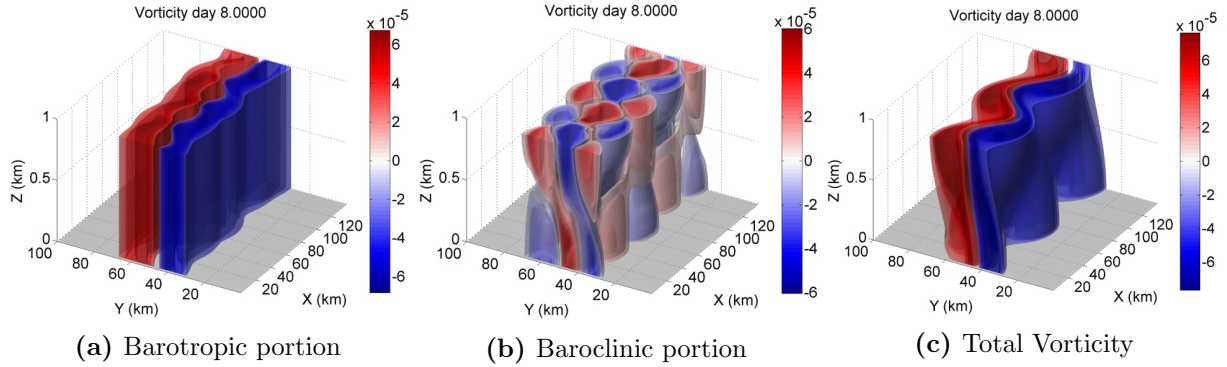


Figure 6.7: The three-dimensional structure of the vorticity field in the reference barotropic simulation as predicted by the Tempest QG model at day 8. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.

6.2.3 Hydrostatic Simulations: Effects of Stratification

From the linear stability calculations, it appeared stratification had very little effect on the growth rate of the barotropic sinuous mode. That being said, stronger stratification did result in stronger vertical variations and better agreement between the PE and QG models for more of the baroclinic sinuous modes. Therefore, two nonlinear simulations were considered; that of weak stratification for which $Bu = 1$ and that of strong stratification for $Bu = 50$. The computed top to bottom density differences are approximately 0.1 kg/m^3 and 250 kg/m^3 for the weak and strong stratification cases respectively. The extreme, strong stratification case certainly pushes the limits of the Boussinesq approximation however the dynamics of this case are of particular interest. In both cases a grid resolution of $256 \times 256 \times 128$ was considered, motivated by the reduction in time step required for stability in the strong stratification case. In the stronger stratification case, the time step set by the inertial period was 10s. Therefore, the integration was performed for 216,000 time steps and required 226.6 hours on 64 cores. In the weakly stratified case, the time step was set to 50s as in the reference case and the computation required 45.5 hours on 64 cores.

The growth rate curves for these weak and strong stratification cases as predicted by the HY model are shown in subfigures 6.8a and 6.8b, from which the predicted growth rates are 1.2931/day and 1.2742/day respectively. These growth rates are very close to the 1.2776/day as predicted in the reference case. Furthermore, in both cases, the wavelength

of the most unstable mode of the nonlinear calculation agrees with that predicted by the linear theory. Consistent with the linear stability calculations, the weak stratification does show a minor increase in growth rates while the strong stratification yields a minor decrease over that of the reference case.

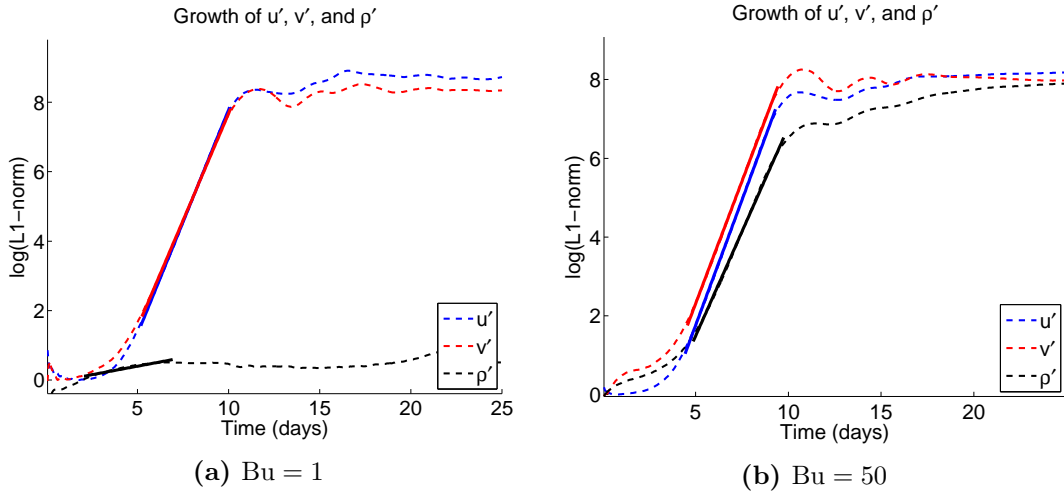


Figure 6.8: Growth rate curves for the barotropic jet simulations as predicted by the Tempest HY model equations for the (a) weak $Bu = 1$ stratification case and the (b) strong $Bu = 50$ stratification case.

However, it is observed that the growth rates as predicted in the nonlinear simulations are less than those predicted by the linear theory. This was also observed in the reference calculation of section 6.2.1. In the cases of weak and strong stratification, the nonlinear simulation growth rates are 6.88% and 7.91% weaker than those of the linear stability calculations. Therefore, it is possible that in the nonlinear simulations the s_0 mode interacts with the secondary modes causing a reduction in growth rates. In the case of the weakest stratification for which secondary modes are the weakest, we see the best agreement between the nonlinear predictions and the linear calculations. It is also possible that the weaker growth rates in the nonlinear simulations could be caused by numerical and physical dissipation.

The evolution of the energy in the weak and strong stratification cases are shown in subfigures 6.9a and 6.9b respectively. The weak stratification case creates very little potential energy as very few coherent vertical structures form in the density field. The relative energy loss for this weak stratification case is only 0.3795%. On the other hand, dissipation at the smaller scale is more prevalent in the case of strong stratification. The relative energy loss for this simulation is 4.3392%. In addition, the strongly stratified case

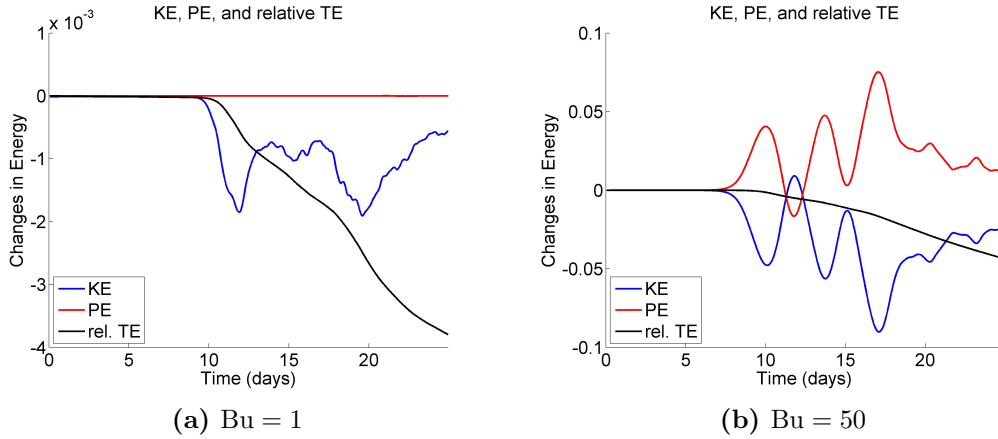


Figure 6.9: Incremental changes in the energy for the barotropic jet simulations as predicted by the Tempest HY model equations for the (a) weak $Bu = 1$ stratification case and the (b) strong $Bu = 50$ stratification case.

produces vortices with strongly varying vertical structure. This structure creates gradients of density giving rise to an increase in potential energy.

The evolution of the instabilities is qualitatively similar to that of the reference barotropic case. However, if we consider a horizontal cross-section of the density perturbation field at an intermediate point in the integration we clearly observe differences between the three cases. A cross-section taken near the surface (relative height 0.9) is shown in subfigures 6.10a, 6.10b, and 6.10c for the weak, reference, and strong stratification cases respectively.

In the case of weak stratification it is observed that there are a number of small scale features that persist as they are able to easily transfer between vertical layers. Consistent with the Rossby radius of deformation, L_d , these features have length scales on the order of 10 km. In a stratified flow the Rossby radius of deformation is defined as [Vallis, 2006]

$$L_d = \frac{NH}{f} \quad \Rightarrow \quad L_d = Bu \cdot L,$$

where in the context of these jet studies $L \sim L_{jet}$ and H is the vertical scale of the motions, which for the most unstable mode is 1 km (the total domain depth).

It is also observed that an increase in stratification causes an increase in the magnitude of the density perturbations as observed from figure 6.10. The weak density perturbations in the case of $Bu = 1$ is consistent with the negligible potential energy observed in this case. On the other hand, the moderate (reference) and the strong stratification cases show smooth density perturbation profiles with the strongest density observed in the strongly

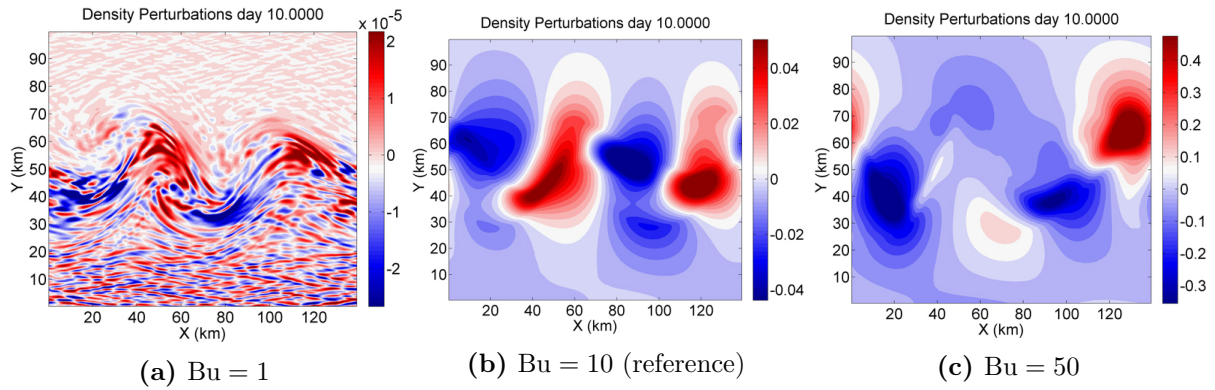


Figure 6.10: Horizontal cross-section of the density perturbation field near the surface (relative height 0.9) after 10 days of integration for the (a) weakly stratified $Bu = 1$, (b) reference $Bu = 10$, and (c) strongly stratified $Bu = 50$ cases in the barotropic jet simulations as predicted by the Tempest HY model.

stratified case. This is characterized by the increase in potential energy observed in the plot of the incremental changes in energy.

We consider isosurface plots of the weakly and strongly stratified cases, as shown in subfigures 6.11a and 6.11b respectively. In the weakly stratified $Bu = 1$ case, there is minimal variations in the vertical and the dynamics are essentially two-dimensional in nature.

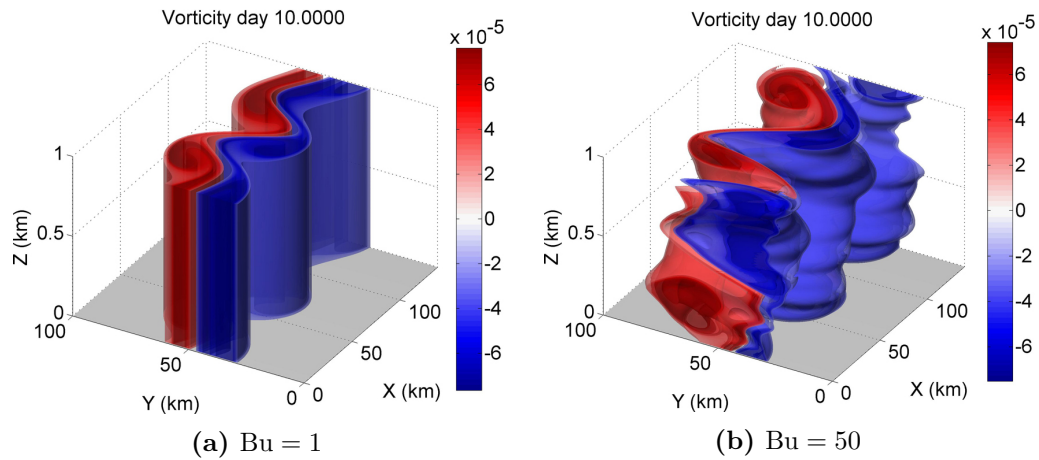


Figure 6.11: Three-dimensional isosurface plots of the vorticity after 10 days of integration for the (a) weakly stratified $Bu = 1$ and (b) strongly stratified $Bu = 50$ cases in the barotropic jet simulations as predicted by the Tempest HY model.

At the opposite end of the spectrum, the strongly stratified case for which $Bu = 50$ gives rise to significant variations in the vertical, which agrees with the theoretical predictions from our LSA.

If we further consider barotropic and baroclinic portions of the vorticity after 10 days of integration in the strongly stratified case, we obtain the isoplots shown in subfigures 6.12a and 6.12b respectively. This is particularly interesting as in the strongly stratified case, it is observed that the sum of the baroclinic modes contributes more to the total vorticity than the barotropic s_0 mode; the baroclinic portion has double the magnitude of the barotropic portion.

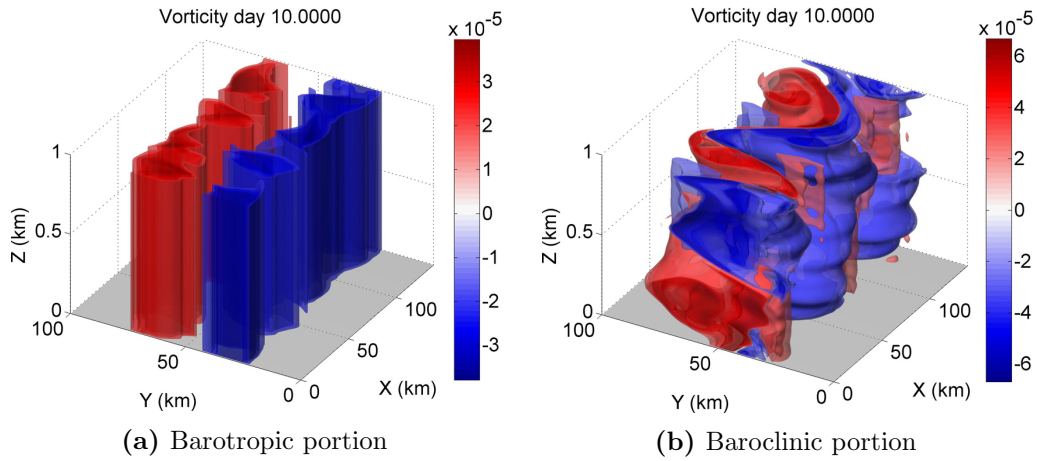


Figure 6.12: Three-dimensional isosurface plots of the barotropic and baroclinic portions of the vorticity after 10 days of integration in the strongly stratified $Bu = 50$ cases in the barotropic jet simulations as predicted by the Tempest HY model.

6.2.4 Quasi-Geostrophic Simulations: Effects of Stratification

Through our LSA calculations, it was observed that as the stratification increased the QG and PE models predicted similar results. In particular, for the strongly stratified case the QG model was able to capture more of the higher mode number instabilities. A nonlinear simulation of the strongly stratified case with $Bu = 50$ on a grid with resolution $256 \times 256 \times 128$ was conducted using the QG model. As the QG model does not depend on the stratification, in particular the buoyancy frequency N , there is no constraint on the time step other than that imposed by the CFL condition. Therefore, the same time step of 100s was used as in the reference case. The computation required 82.5 hours on 64 cores.

The growth rate plot is shown in figure 6.13. From the fitted exponential function, the growth of the most unstable mode is determined to be 1.2905/day. This is a slight increase over the 1.2500/day growth rate in the reference case. In the reference case, the difference between the growth rates as predicted by the HY and QG model was 2.16%. In this strongly stratified case, the 1.2905/day growth rate is a 1.26% difference from the 1.2742/day growth rate of the HY model. Therefore, the HY and QG model do demonstrate better agreement in the strongly stratified case as predicted in the linear stability calculations.




Figure 6.13: Growth rate curves for the barotropic jet simulations as predicted by the Tempest QG model strongly stratified $Bu = 50$ case.

The three-dimensional structure of the vorticity is shown in figure 6.14 after 8 days of integration, coinciding with a time after the period of strongest exponential growth. Subfigure 6.14a, 6.14b, and 6.14c show the barotropic, baroclinic, and total vorticity fields. The general structure of the most unstable mode as predicted by the QG simulation is comparable in structure to that of the HY model prediction. In particular, the baroclinic portion of the instability is stronger than the barotropic portion and there is greater vertical variations.

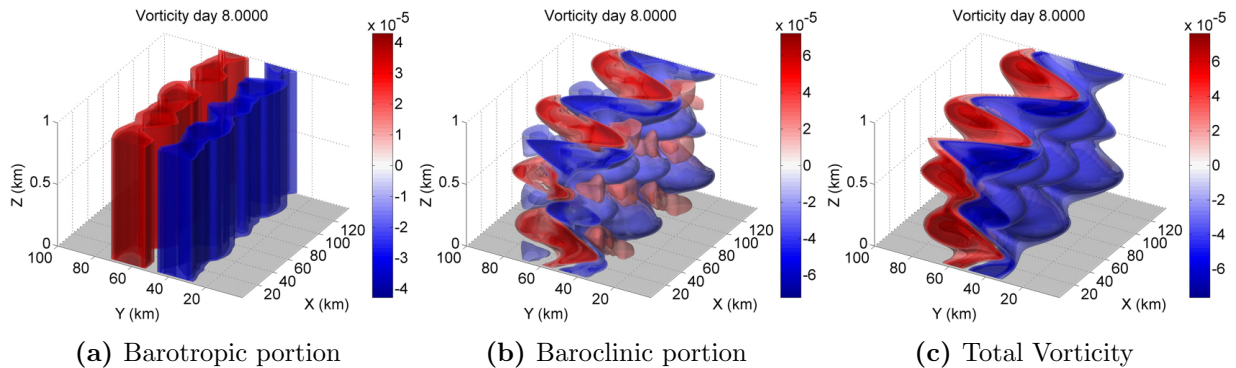


Figure 6.14: The three-dimensional structure of the vorticity field in the strongly stratified $Bu = 50$ barotropic simulation as predicted by the Tempest QG model at day 8. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.

6.2.5 The Effects of Topography

To assess the impact of underlying topography, nonlinear simulations for a specific prograde and retrograde configuration are considered. An intermediate topographic height spanning 50% of the domain depth is considered. This height was chosen as it represents a gently sloping topography for which the terrain-following coordinate system is well-suited. A fine resolution simulation of $256 \times 256 \times 256$ was used for both simulations, as small computational features were being produced near the bottom boundary in cases with coarser vertical resolutions. The most probable cause for these small scale oscillations was the terrain-following coordinate system as they originated at the location of the steepest slope in the topography.

However, regardless of grid resolution, these small scale features could be removed with vertical viscosity. A finer resolution was chosen though to allow us to consider a vertical viscosity of $\nu_v = 1e-3 \text{ m}^2/\text{s}$ consistent with that used in the reference case. No horizontal viscosity was used in these simulations. Using a time step of $dt = 50 \text{ s}$ and integrating for 25 days (43200 time steps), the simulation required 50.2 hours and 49.6 hours for the prograde and retrograde simulations respectively. These simulations were run on 128 processors.

The plots of the growth rates are shown in subfigures 6.15a and 6.15b for the retrograde and prograde simulations. The growth rate in the retrograde case is computed as 1.2574/day and in the prograde case as 1.3029/day. Therefore, relative to the 1.2776/day growth rate predicted in the case of the reference, flat-bottom simulation we see that retrograde (prograde) topography has a stabilizing (destabilizing) effect on the most unstable mode.

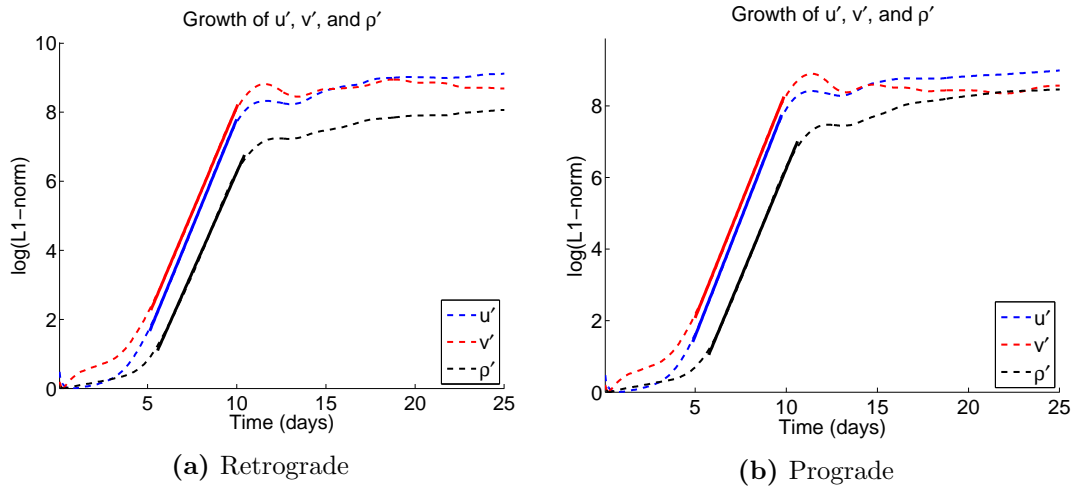


Figure 6.15: The growth rate curves in the case of the barotropic jet simulation over (a) retrograde and (b) prograde topographic profiles as predicted by the Tempest HY model.

This was also found to be the case in the linear stability calculations.

The plots of the vertical and horizontal cross-sections of the vorticity are shown in subfigures 6.16a and 6.16b for the retrograde and prograde cases respectively. In both the prograde and retrograde cases the wavelength of the excited mode is in agreement with the LSA. As well, it is observed that the topography suppresses the secondary modes from growing, consistent with that observed in the linear theory. Both of the unstable modes appear to develop a lean in the same direction. This is also observed in the flat bottom case




Figure 6.16: Plots of vertical and horizontal cross-sections of the vorticity after 10 days of integration for the barotropic jet simulation over (a) retrograde and (b) prograde topographic profiles spanning a maximum height of 50% of the domain height as predicted by the Tempest HY model

(not shown) and is a result of the vertical structure introduced by the baroclinic modes.

The incremental changes in the energy and relative energy loss for the retrograde and prograde simulations are shown in subfigures 6.17a and 6.17b respectively. We observe that the retrograde topography causes a greater loss of energy relative to the prograde topography; a relative loss of 7.2419% as opposed to 3.2394%. This is consistent with the observed decrease in the growth rate. From the linear calculations it was observed that strong retrograde topography did have a greater stabilizing effect on the first baroclinic mode relative to the strong prograde topography. It is possible that these interactions are causing a potential sink in energy.

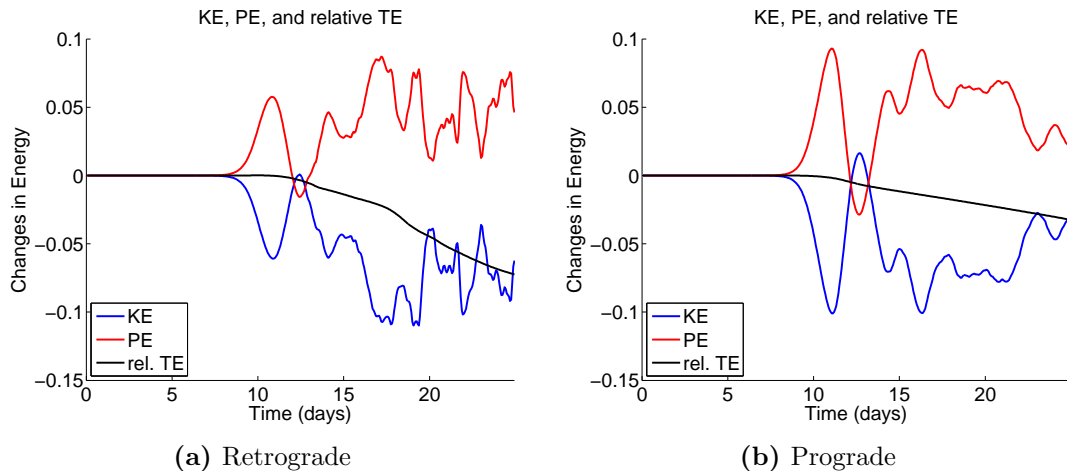


Figure 6.17: Plots of the incremental changes in the energy and relative energy loss in the case of the barotropic jet simulation over (a) retrograde and (b) prograde topographic profiles as predicted by the Tempest HY model.

To get a better understanding of the evolution of the instabilities in the two cases, we first plot horizontal cross-sections of the vorticity as shown in figure 6.18 where the top and bottom rows show the cases with the retrograde and prograde topographic profiles respectively. These cross-sections are considered near the surface at a relative height of 0.9. The results are displayed in 5 day increments. We observe that the evolution appears consistent in both cases until day 15 (fourth column). At this time the submesoscale vortex in the Northern half of the domain is predicted as a monopole in the retrograde case and a dipole in the prograde case. Small variations early on lead to significant difference in the two cases as time progresses. In particular, at later times, we observe that instabilities over prograde topography show a tendency to move offshelf, while retrograde topography appears to have very little effect on the location of the mesoscale vortices.




Figure 6.18: The evolution of the vorticity field in the case of the barotropic jet simulation over retrograde (top row) and prograde (bottom row) topographic profiles respectively. The outputs are printed in 5 day increments.

To consider the three-dimensional, vertical structure of the vorticity after the instabilities have undergone the period of maximum exponential growth isoplots are displayed at day 10 of integration. First, the retrograde case is presented in figure 6.19 where subfigures 6.19a, 6.19b, and 6.19c represent the barotropic, baroclinic, and total vorticity fields. The general barotropic and baroclinic structure at this day are similar to those observed in the reference case. However, the topography does have a strong impact on the baroclinic structure. This explains the strong stabilizing effects that the topography had on the baroclinic modes of the linear stability calculations. On the other hand, the structure of the barotropic mode is relatively unaffected by the topography.

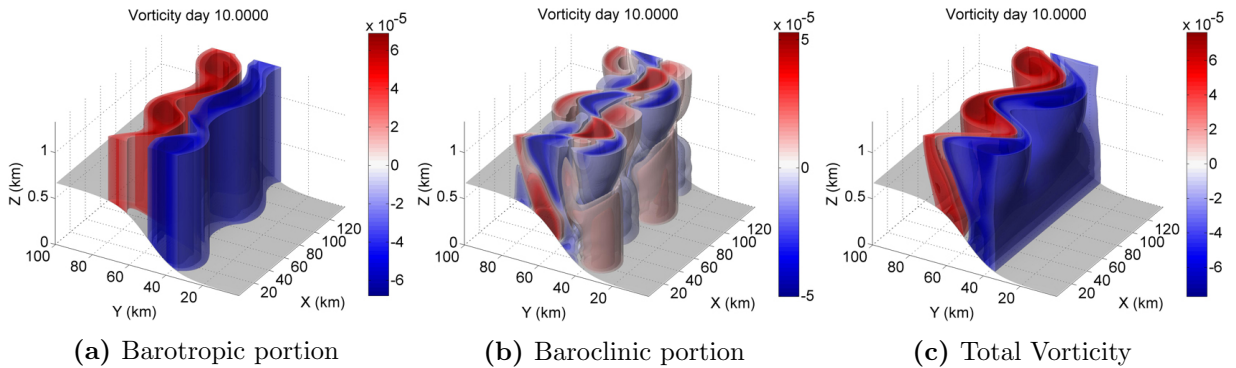


Figure 6.19: The three-dimensional structure of the vorticity field in the retrograde barotropic simulation as predicted by the Tempest HY model at day 10. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.

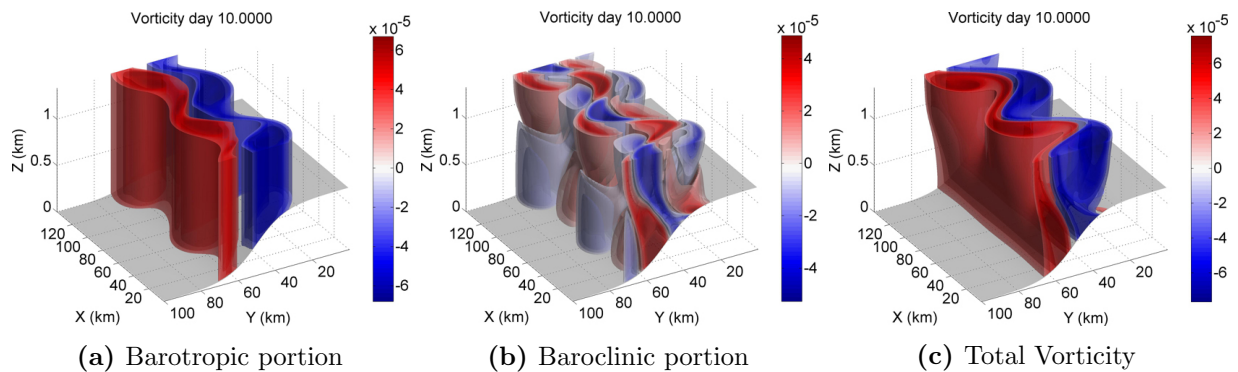


Figure 6.20: The three-dimensional structure of the vorticity field in the prograde barotropic simulation as predicted by the Tempest HY model at day 10. Subfigure (a) shows the barotropic portion, subfigure (b) shows the baroclinic structure, and subfigure (c) shows the total vorticity.

By way of comparison the three-dimensional, vertical structure of the prograde case is presented in figure 6.20 where subfigures 6.20a, 6.20b, and 6.20c represent the barotropic, baroclinic, and total vorticity fields. At this time in the integration, the qualitative effect of the prograde topography is similar to the retrograde case. The topography interacts with the baroclinic modes of instability.

The three-dimensional structures are displayed at the end of the integration in figure 6.21, where subfigure 6.21a and 6.21b display the results for the retrograde and prograde cases

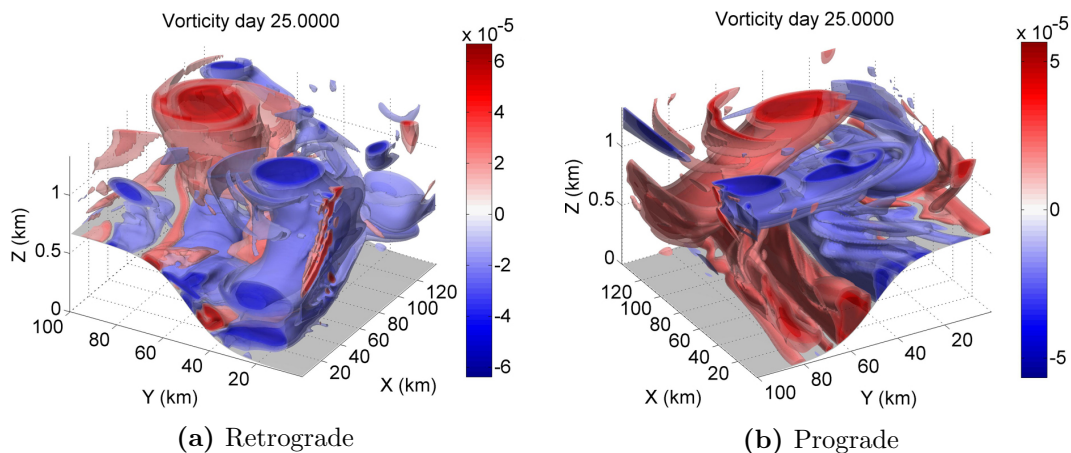


Figure 6.21: The three-dimensional structure of the vorticity field in the (a) retrograde and (b) prograde barotropic simulations as predicted by the Tempest HY model at the end of the integration.

respectively. At this time, there are strong formations present along the bottom boundary, particularly noticeable in the prograde case. These features have a large horizontal length scale and are therefore assumed to be physical in nature. Though not presented, the baroclinic modes are stronger than their barotropic counterparts in both the prograde and retrograde cases emphasizing the strong vertical variations that have developed.

6.3 Baroclinic Jet

To continue the analysis of shear instabilities that arise in the oceans, we consider next an idealized baroclinic jet. This type of jet is characterized by both horizontal and vertical shear. As the velocity profile varies with depth, through the thermal wind balance we have that the density profile must vary with depth as well. These types of jets are more difficult to model than the barotropic jets due to the balance of the background state and because of the complex instabilities that can arise.

From the previous chapter, we expect that the PE and QG simulations should have comparable growth rates in the reference case however there may be disagreement in the wavelengths. As the background stratification is increased we would expect better agreement between the results of the PE and QG simulations. In the cases with topography it is expected that the PE model will predict growth rates similar to the reference case as the unstable modes in the case of the surface-trapped baroclinic jet appeared relatively unaffected by the underlying topography.

6.3.1 A Hydrostatic/Non-Hydrostatic Comparison

In the case of the baroclinic jet, NH and HY simulations were considered using a grid with resolution $128 \times 128 \times 64$ in order to determine the effectiveness of the HY model at predicting potentially NH motions. Setting all other parameters equal, the numerical integration of the HY model required 12.7 hours and the NH model required 42.8 hours. Therefore, the NH model required over triple the amount of computation time for this particular case.

As a first point of comparison, we consider the predicted growth rates using the two models. The growth rate curves are shown in figure 6.22 where subfigure 6.22a and 6.22b show the results for the HY and NH models respectively. It is observed that early on there appears to be some sort of adjustment process occurring as seen in the u' and ρ' curves. Furthermore, the fitted curves appear to have slightly different slopes which may suggest




Figure 6.22: Plots of the growth rate curves in the case of the reference baroclinic test case for a grid resolution of $128 \times 128 \times 64$ using the (a) HY and (b) NH models.

an effect of stratification on the evolution of the density perturbations. We will later see that in the strongly stratified case that the growth of ρ' weakens. In the case of the NH model the maximum growth rate is $0.50534/\text{day}$ while for the HY model the maximum growth rate is $0.50535/\text{day}$, an almost identical result. The predicted wavelengths and the modal structures are also equivalent qualitatively in both cases.

Next, plots of the incremental changes in energy and the relative energy loss were compared. These are shown in figure 6.23, where subfigures 6.23a and 6.23b show the results as predicted by the HY and NH models respectively. In both cases, there is initially a slight increase in potential energy followed by a very strong increase in kinetic energy. The energy plots in the two cases are equivalent and by the end of the integration the relative energy loss in both cases is 3.2922% .

Additionally, a qualitative comparison of a horizontal cross-section of the vorticity at the final day of integration shows very little difference between the NH and HY results. These results are displayed in figures 6.24a for the HY model and 6.24b for the NH model. It is also evident from these plots that there are very little submesoscale motions at this level of resolution and, as reiterated from [Mahadevan, 2006], it may be with increased resolution that NH effects would become more important particularly when the horizontal grid spacing is on the order of the vertical instabilities. As this instability is surface trapped predominantly to the top 300 m of the domain we would expect NH effects to occur on the grid scale for resolutions with at least 300 cells across channel. In addition, it is possible

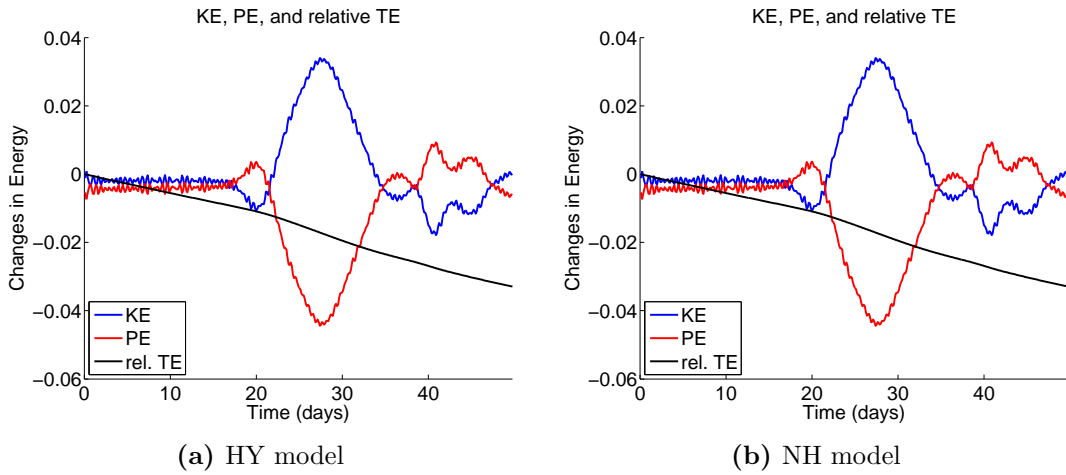


Figure 6.23: Plots of the incremental changes in energy and the relative energy loss in the case of the reference baroclinic test case for a grid resolution of $128 \times 128 \times 64$ using the (a) HY and (b) NH models.




Figure 6.24: A horizontal cross-section of the vorticity field at the final day of integration in the case of the baroclinic jet as predicted by the HY model (a) and the NH model (b)

that topographic variations could lead to NH motions.

However, from the above calculations we determined, at least for this particular configuration, that the HY model performed very well in describing the onset of the linear instability and the nonlinear equilibration. We also saw in the parallel efficiency tests that the HY model tends to scale better than the NH model. Therefore, in an effort to resolve

smaller scale features, we have chosen to apply the HY model exclusively for the baroclinic jet tests using an increased resolution.

6.3.2 Hydrostatic Reference Simulation

The same definition of the baroclinic jet is used as that of the linear stability calculations in section 5.4. The jet is centered in the middle of the channel and has characteristic velocity, horizontal length, and vertical length scales defined as $U_{\text{jet}} = 1 \text{ m/s}$, $L_{\text{jet}} = 10 \text{ km}$, and $H_{\text{jet}} = 300 \text{ m}$ respectively. We have a background profile that depends both on y and z . The background flow is defined as

$$u = U(y, z), \quad v = w = 0,$$

where

$$U(y, z) = U_{\text{jet}} \operatorname{sech}^2 \left(\frac{y - m_y}{L_{\text{jet}}} \right) \left(1 + \operatorname{erf} \left(\frac{z - H}{H_{\text{jet}}} \right) \right).$$

Next, we can determine the density from hydrostatic balance as

$$R(y, z) = -\frac{1}{g} \frac{\partial}{\partial z} P(y, z).$$

To evaluate the pressure we make use of geostrophic balance and upon substitution into the equation for $R(y, z)$ we obtain:

$$R(y, z) = \frac{\rho_0 f}{g} \int \left(\frac{\partial}{\partial z} U(y, z) \right) dy.$$

The following physical constants are used throughout the baroclinic simulations:

$$g = 9.80616 \text{ m/s}^2, \quad f = 10^{-4} \text{ s}^{-1}.$$

In the nonlinear simulations we perturb the zonal background velocity by adding on random noise with a relative magnitude of 0.1%. The evolution of the instabilities in the baroclinic case have much smaller growth rates than those predicted in the case of the barotropic jet. Therefore, we integrate the nonlinear model for a total of 50 days. For the baroclinic cases, the domain is defined as

$$0 \text{ km} \leq x \leq 184.8 \text{ km}, \quad 0 \text{ km} \leq y \leq 100 \text{ km}, \quad 0 \text{ km} \leq z \leq 1 \text{ km},$$

with the zonal extent set by double the wavelength of the most unstable mode as predicted by the linear theory.

For this reference test case we considered a HY model simulation with grid resolution $256 \times 256 \times 128$. The stratification is set so that $Bu = 10$ and the Rossby number for this simulation is $Ro = 1$. The horizontal and vertical viscosity coefficients are $\nu_h = 0 \text{ m}^2/\text{s}$ and $\nu_v = 1\text{e-}3 \text{ m}^2/\text{s}$. The time step was set to $dt = 50 \text{ s}$, set by the inertial period, and the equations were integrated for 86,400 time steps. On 128 processors, the numerical integration required 56.32 hours to complete.

The growth rate curve for this particular reference simulation is shown in figure 6.25a. From the fitted exponential function, the growth is determined to be 0.5026/day, showing a very strong agreement to the predicted value from the linear stability calculations with a relative difference of 4.86%. Also, the excited mode has a wavelength consistent with that predicted by the linear theory. However, as was also observed in the barotropic test cases, the nonlinear simulation does predict a growth rate that is less than the predicted growth rate from the linear stability calculations. This is to be expected because of the dissipation that is present in the nonlinear simulations.

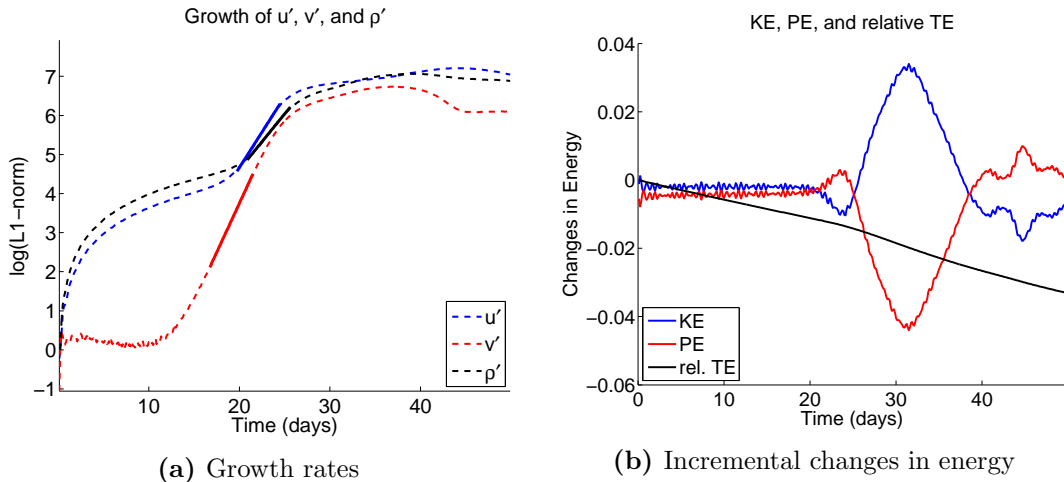


Figure 6.25: Plots of (a) the growth rates and (b) incremental changes in the energy and relative energy loss for the reference baroclinic simulation with $Ro = 1$ and $Bu = 10$ as predicted by the Tempest HY model.

The incremental changes in the energy and the relative energy loss are shown in figure 6.25b. By the end of the integration, the relative energy loss is 3.3025%. Given the slight change from the coarser resolution run, it is believed that the energy loss is due

predominantly to numerical and physical viscosity. As the initial instability forms we get a slight increase in potential energy followed by a rapid decrease, which is accompanied by an increase in kinetic energy.

To consider the modal structure as predicted by the nonlinear HY simulation, a horizontal and a vertical cross-section of the vorticity is shown at day 20 of integration. This represents a point in the integration characterized by the strongest exponential growth as determined from the growth rate plot 6.25a. The horizontal cross-section is shown in figure 6.26a and the vertical cross-section, taken at a relative height of 0.9, is shown in figure 6.26b.

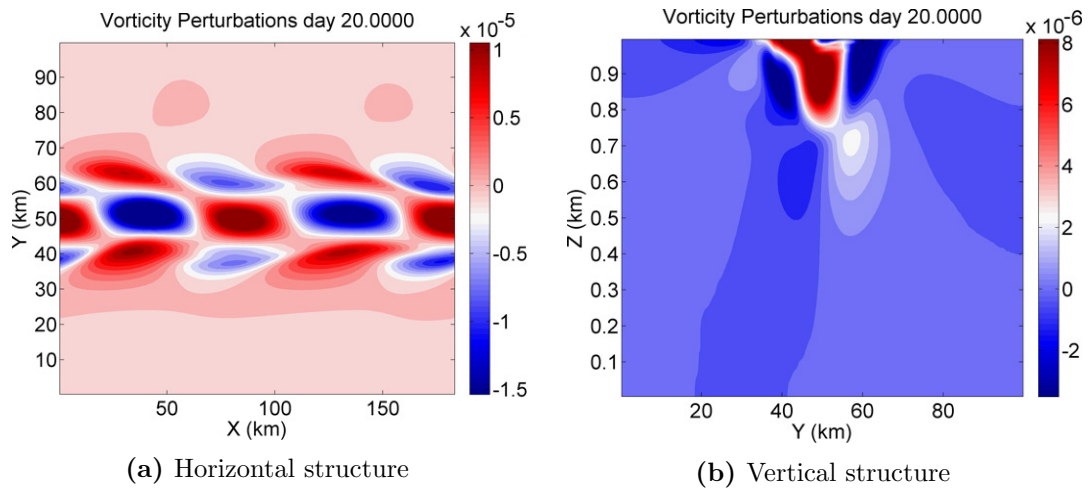


Figure 6.26: Cross-sections of the vorticity plot in the case of the reference baroclinic jet as predicted by the Tempest HY model at day 20 of integration which represents a time during the period of strongest exponential growth. Subfigure (a) shows the horizontal cross-section and subfigure (b) shows the vertical cross-section at a relative height of 0.9.

As the baroclinic linear stability calculations predicted only two modes of instability we can readily compare these structures to the modes predicted by the linear theory. The structures of the two unstable modes as predicted by the linear stability calculations are shown in figure 6.27 for reference. Subfigure 6.27a shows the horizontal and vertical structure of the vorticity of the first unstable mode and subfigure 6.27b shows the horizontal and vertical structure of the vorticity of the second unstable mode. It appears that the modal structure predicted by the nonlinear calculations 6.26 is a combination of the two mode structures predicted by the linear calculations. In particular, looking at the vertical structure 6.26b we see that the structure is representative of the first unstable mode 6.27a, however there are small features near the surface indicative of the second unstable mode 6.27b.




Figure 6.27: Horizontal and vertical structures of the vorticity for the (a) s_1 mode and (b) s_2 mode of the reference baroclinic jet case as predicted by the linear stability calculations.

In an effort to understand the evolution of the instabilities, snapshots of the vorticity are shown in figure 6.28, displayed in 10 day increments. By day 20 the mesoscale vortices are just forming and by day 30 there are coherent structures. They propagate along and do not appear to begin merging until much later (by the end of the integration). In contrast, in the barotropic jet simulations the vortices merged almost immediately. At day 40, there are small scale feature formed from interactions with the boundary, however they do not remain coherent and could be a result of numerical errors.




Figure 6.28: Snapshots of the evolution of the vorticity in the case of the reference baroclinic jet as predicted by the Tempest HY model. The cross-sections are taken near the surface at a relative height of 0.9. The results are displayed in 10 day increments.

To gain information about the three-dimensional structure of the jet isoplots of the vorticity field are shown in figure 6.29, where subfigures 6.29a and 6.29b represent day 30 and 50 of integration respectively. It is clear that the instabilities remain surface-trapped, consistent with the predictions in the linear stability calculations. Furthermore, the cyclonic motions appear to remain more surface intensified with a wider horizontal scale than the anticyclonic motions. This is physically intuitive as anticyclones represent high pressure areas and, therefore, would produce a stronger downward force that would generate greater vertical length scales. At the end of the integration a region of anticyclonic motion is observed near the Northern boundary, which pinched off from the stronger cyclonic vortices in the South. This was observed in the snapshots of the evolution shown in figure 6.28.

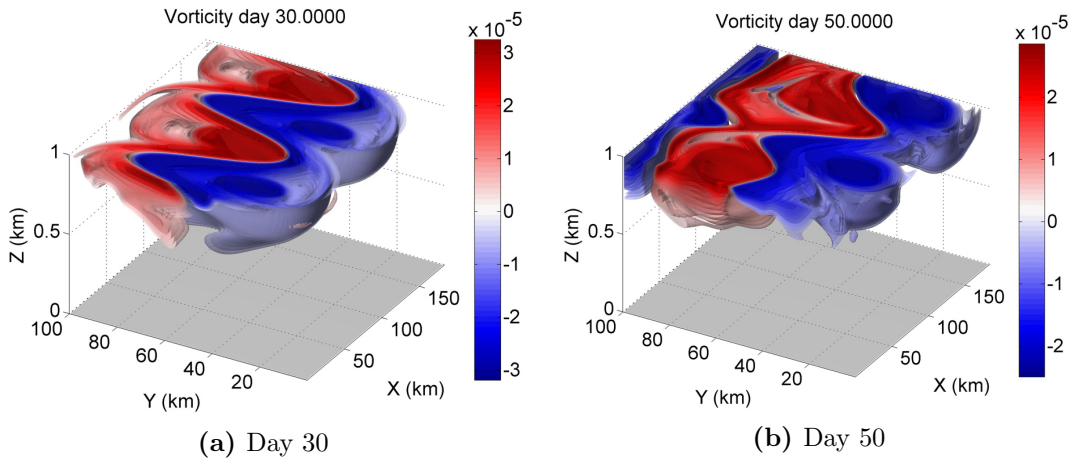


Figure 6.29: Three-dimensional structure of the vorticity field of the reference baroclinic case as predicted by the Tempest HY model. Subfigure (a) shows the structure at day 30 and subfigure (b) shows the structure at the end of the integration.

6.3.3 Quasi-Geostrophic Reference Simulation

The above reference calculation was compared to a QG simulation with the same configuration. The time step was set to $dt = 100$ s and integration was performed for 50 days, requiring 142.5 hours of computation time using 64 cores. The growth rate curve for this simulation is shown in subfigure 6.30a. The maximum growth rate as determined from the fitted exponential function is 0.6344/day. Compared to the predicted growth rate of 0.5314/day from the linear stability calculations this represents a significant overestimate. However, interestingly enough, this growth rate is comparable to the $Ro \ll 1$ linear stability

calculations of the PE model, for which the growths were 0.6222/day and 0.6236/day for Ro of 0.1 and 0.01 respectively.

The incremental changes in the energy and the relative total energy loss are shown in subfigure 6.30b. The relative total energy loss is approximately 10.58% which is significantly greater than the 3.30% in the reference HY case. The relative total energy loss curve is not plotted so that the changes in KE and PE can be observed more clearly. In the HY case we saw a more balanced trade-off between PE and KE with a constant energy loss over the course of the simulation. In subfigure 6.30b, we see that energy in the QG model is lost from the lack of balance between the KE and PE.

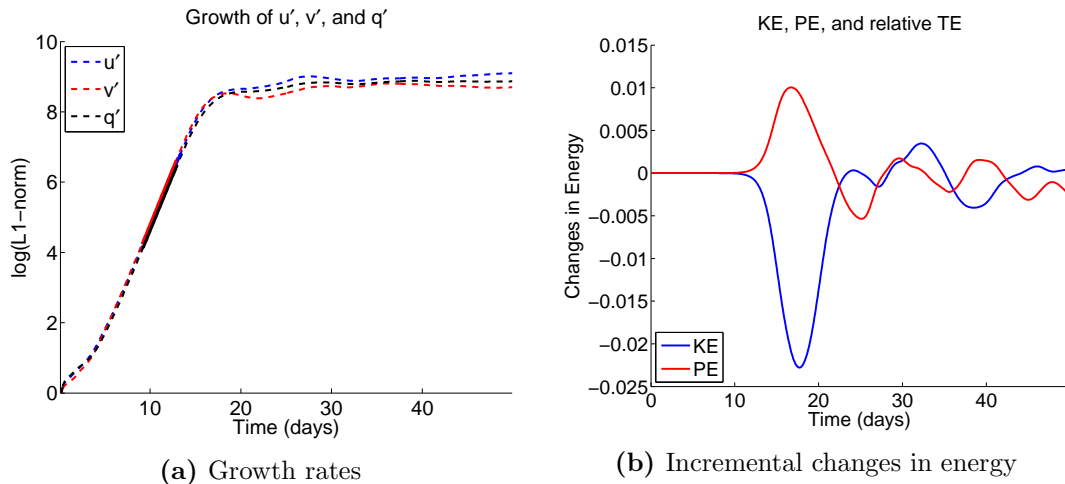


Figure 6.30: Plots of (a) the growth rates and (b) incremental changes in the energy and relative energy loss for the reference baroclinic simulation with $Ro = 1$ and $Bu = 10$ as predicted by the Tempest QG model.

Snapshots of the evolution of the vorticity are shown in figure 6.31, displayed in 10 day increments. From these plots, we observe a significant disagreement between the HY and QG simulations. In the HY simulation the wavelength of the most unstable mode was 92.4 km, however in the case of the QG model the most unstable mode has a significantly shorter wavelength of approximately 60 km. This is in contrast to the linear stability calculations for which the QG model predicted the most unstable mode had a wavelength longer than that of the PE calculations.

However, following the initial mesoscale vortex formation, the mesoscale vortices merge as observed by day 30. Two pairs of mesoscale vortices persist by the end of day 40 and are advected along. Comparing to the spatial structures in the HY case 6.28, we see that




Figure 6.31: Snapshots of the evolution of the vorticity in the case of the reference baroclinic jet as predicted by the Tempest QG model. The cross-sections are taken near the surface at a relative height of 0.9. The results are displayed in 10 day increments.

there are a number of small scale features predicted in the QG case that are not present in the HY case. This could be due in part to the lack of viscosity in the QG model. It is also observed that in the case of the QG model the magnitude of the vorticity remains constant, whereas in the HY case the vorticity appeared to decrease in magnitude as time passed further indicating potential dissipation in the HY model.

Isoplots of the vorticity field are shown for day 25 and day 50 in subfigures 6.32a and 6.32b respectively. In the case of the QG model, there appears to be numerical errors associated with the boundary conditions. This could be due to the natural boundary condition approximation that we have assumed. It could also be due to the strong vertical variation that is difficult to capture by the QG model. Regardless, we conclude that the QG model predicts significantly different dynamics than the PE model for this particular test case.

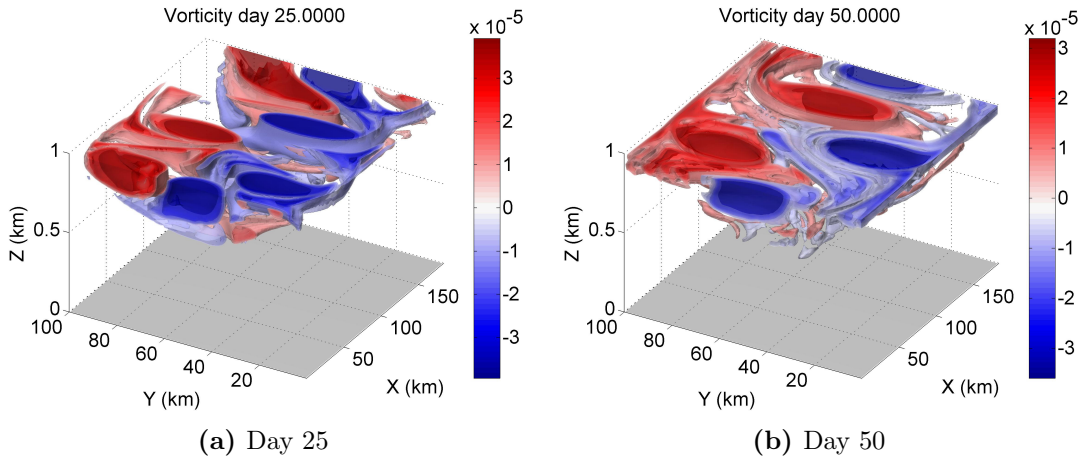


Figure 6.32: Three-dimensional structure of the vorticity field of the reference baroclinic case as predicted by the Tempest QG model. Subfigure (a) shows the structure at day 25 and subfigure (b) shows the structure at the end of the integration.

6.3.4 Hydrostatic Simulations: Effects of Stratification

To consider the effects stratification has on the nonlinear evolution of the baroclinic jet we considered the strongly stratified case of $Bu = 50$. In this case, there was much stronger agreement between the PE and QG linear stability calculations and it would be interesting to see if the nonlinear calculations yield better agreement. In response to the shorter inertial period the time step was reduced to 10s and a coarser resolution of $128 \times 128 \times 128$ was used. The equations were integrated for 25 days, instead of 50 days, as the linear stability calculations predicted a much greater growth rate in this strongly stratified case.

The growth rate curve is shown in figure 6.33a and the incremental changes in energy are shown in figure 6.33b. The maximum growth rate is predicted as 0.8551/day associated with a wavelength of 76.62 km consistent with the linear stability calculations. This growth rate is a significant increase over the 0.5026/day growth rate computed in the reference case, however it is not as large as the 0.9610/day predicted by the linear stability calculations. However, from the energy plot it is clear that there is a significant sink of energy which could be an explanation for the differences. In this case the relative energy loss is not shown as by the end of the integration the relative loss is 23.6530%. The increase in energy loss with increasing stratification was also observed in the barotropic jet simulations of the previous section.

Figure 6.34 shows snapshots of the evolution of the vorticity at day 5 6.34a, day 15 6.34b, and day 25 6.34c. In the case of strong stratification, the baroclinic jet instabilities

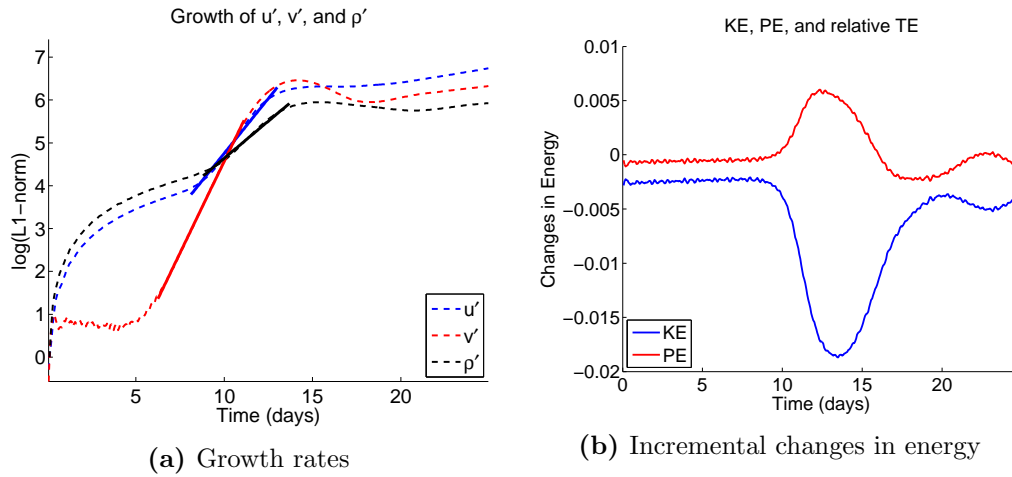


Figure 6.33: Plots of (a) the growth rates and (b) incremental changes in the energy for the strongly stratified $Bu = 50$ baroclinic simulation as predicted by the Tempest HY model.

appear to evolve almost as a series of two-dimensional layers. It requires significantly more energy for motions to interact vertically and it appears that each vertical level evolves on their own timescales. In particular, it is observed that by the end of the integration the bottom of the initial jet has not changed. On the other hand, the vorticity at the surface has evolved on a much shorter time scale. Over the course of the simulation, the magnitude of the vorticity field decreases quite dramatically consistent with the kinetic energy loss observed in the energy plot.




Figure 6.34: Isoplots of the vorticity field in the strongly stratified $Bu = 50$ baroclinic jet case as predicted by the Tempest HY model. Subfigures (a), (b), and (c) show the snapshots at days 5, 15, and 25 respectively.

6.3.5 Quasi-Geostrophic Simulations: Effects of Stratification

In the case of the QG model the time step is not dependent on the stratification so in the strongly stratified $Bu = 50$ case we can consider the same grid resolution of $256 \times 256 \times 128$ as in the reference case. For this particular case the growth rate curve is shown in figure 6.35 and the growth rate has been computed as 1.035/day. This growth rate is 21.04% greater than the growth rate in the HY simulation. However, this is a slightly better agreement than the reference case for which the difference between the QG and HY model growth rates was 26.22%. As well, the unstable mode has the correct wavelength and a comparable growth rate to the 0.962/day predicted by the QG linear stability calculations.




Figure 6.35: Growth rate for the strongly stratification $Bu = 50$ baroclinic simulation as predicted by the Tempest QG model.

To compare the evolution to that of the HY simulation, isoplots of the vorticity are presented in figure 6.36. These snapshots represent day 5 6.36a, day 15 6.36b, and day 25 6.36c of the integration. In comparing the results with the HY model, the most notable difference is that the QG model does not predict a significant decrease in the magnitude of the vorticity. As well, the QG model predicts a number of small scale features that are not captured by the HY model. This was also observed in the reference calculation. That being said, the results do show a strong agreement to the HY model. By day 15 there are coherent mesoscale vortices and by the end of the integration they have formed into surface-intensified mesoscale vortices with a very large length scales.

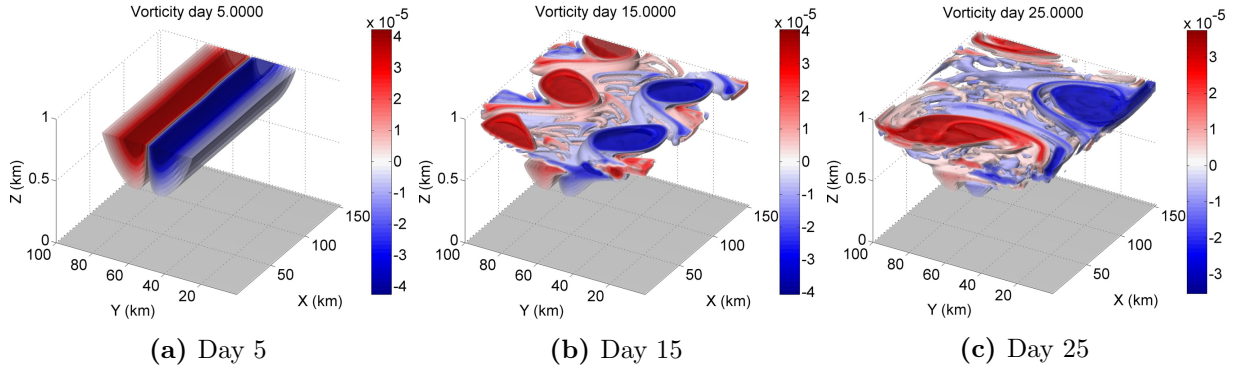


Figure 6.36: Isoplots of the vorticity field in the strongly stratified $Bu = 50$ baroclinic jet case as predicted by the Tempest QG model. Subfigures (a), (b), and (c) show the snapshots at days 5, 15, and 25 respectively.

6.3.6 The Effects of Topography

To study the effects of topography on the baroclinic jet, nonlinear simulations for a specific prograde and retrograde configuration are considered. In each case, the topographic height spans 50% of the domain depth. From the linear theory we predict that the growth rate will be comparable in both cases and similar to the growth in the reference case. In order to resolve some numerical errors observed in the beginning of the simulation, we consider a grid with resolution $128 \times 128 \times 256$. The results were computed on 64 processors and required 48.5 and 52.2 hours for the prograde and retrograde cases respectively. The integration was performed for 50 days and required 86400 time steps. The vertical viscosity coefficient was set to $\nu_v = 1e-3 \text{ m}^2/\text{s}$ and no horizontal viscosity was included.

The plots of the growth rates are shown in subfigures 6.37a and 6.37b for the retrograde and prograde simulations. The growth rate in the retrograde case is computed as 0.5040/day and in the prograde case as 0.5059/day. This is very close to the growth rate of 0.5026/day in the reference HY case. Though the trend was not significant in the linear stability calculations, it is interesting to note that the topography has a destabilizing effect in the nonlinear simulations as well. Also, the most unstable mode, as predicted by the linear theory, is consistent with the most unstable mode in these nonlinear simulations.

Though the unstable modes were predicted to be surface-trapped by the linear stability calculations, it is clear from the nonlinear simulations that there are a number of interaction with the topography. Figure 6.38 shows a vertical cross-section of the vorticity perturbations at day 0.5 of integration. There are approximately ten grid cells across the features in these figures, which leads us to believe that they are fairly well-resolved. It is observed

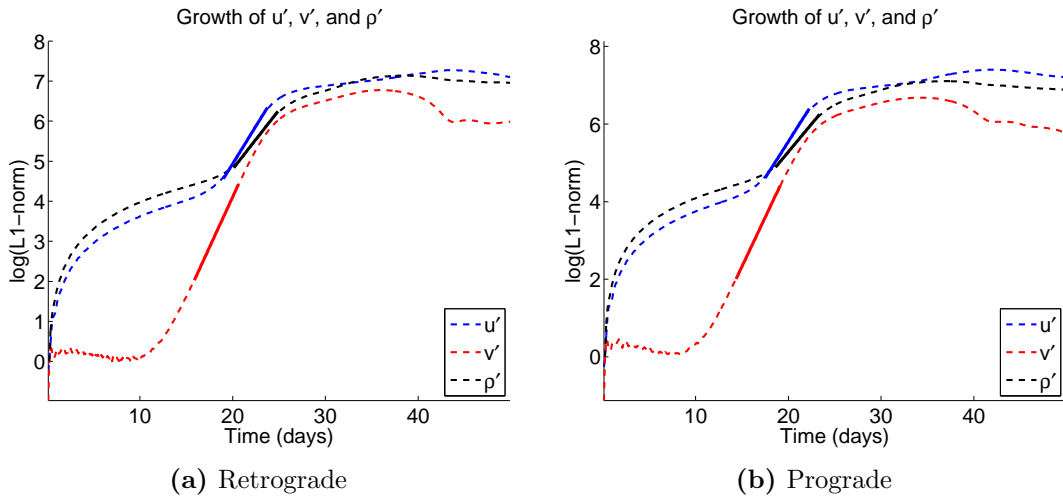


Figure 6.37: The growth rate curves in the case of the baroclinic jet simulation over (a) retrograde and (b) prograde topographic profiles as predicted by the Tempest HY model.

from the reference simulation that the perturbations do extend downwards and in the case of the prograde and retrograde topographic profiles we observe interactions with the topography. These interactions lead to reflections and wave scattering in the shallow regions. Consistent with the Rossby radius of deformation, this produces smaller scale features in the shallow regions. Though it is not believed to affect the overall evolution, it is an interesting difference not captured by the linear calculations.

The incremental changes in the energy and relative energy loss for the retrograde and




Figure 6.38: Plots of the vorticity perturbations at day 0.5 of integration for the retrograde (left), flat-bottom (middle), and prograde (right) topographies in the case of the baroclinic jet simulations as predicted by the Tempest HY model.

prograde simulations are shown in subfigures 6.39a and 6.39b respectively. The relative energy losses for the retrograde and prograde cases were 3.4274% and 3.4833% respectively. Although the growth rates marginally increase in this case, the losses in energy show a marginal increase over the 3.3025% loss reported for the reference case.

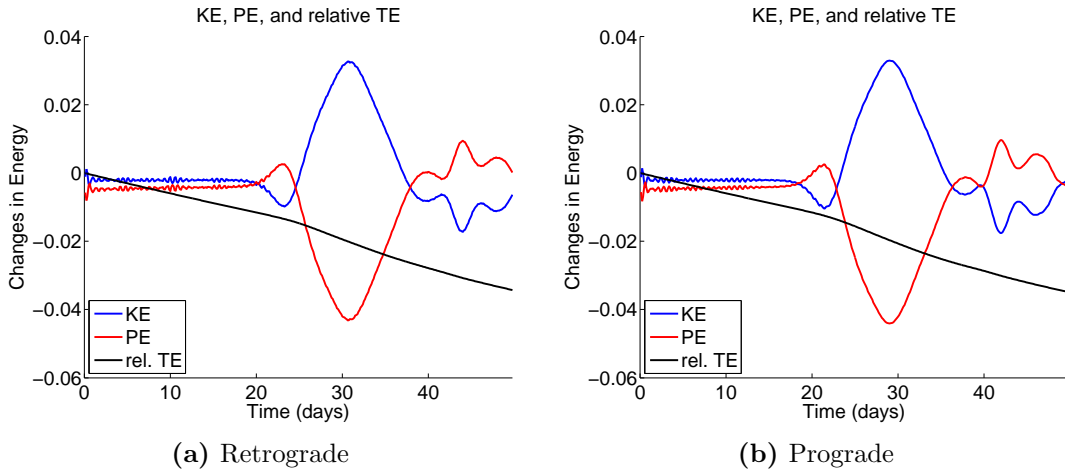


Figure 6.39: Plots of the incremental changes in the energy and relative energy loss in the case of the baroclinic jet simulation over (a) retrograde and (b) prograde topographic profiles as predicted by the Tempest HY model.

Snapshots of the evolution of the vorticity field are shown in figure 6.40. The top row shows the evolution in the case of the retrograde topography and the bottom row shows the evolution in the case of the prograde topography. This cross-section is taken at a relative height of 0.9. The topography has little effect on the overall evolution of the mesoscale vortices at this level height, consistent with that observed in the barotropic simulations. We see that by day 25 mesoscale vortices have formed having the same wavelength in both prograde and retrograde topographic cases. This wavelength is consistent with that predicted in the linear theory.

There are however qualitative differences between the two cases. In particular, the retrograde topography shows some smaller scale features in the shallow region, consistent with the reduction in the Rossby radius of deformation. As well, at the end of the integration we see that the anticyclones in the prograde case have a slightly smaller length scale than the anticyclones in the retrograde case coinciding with the shallower region. These are not dramatic differences and the overall evolution of the instabilities appears to be in close agreement with the reference case. In particular, comparing to figure 6.28 we have the same mesoscale vortex structure with the cyclonic motions having a longer horizontal length




Figure 6.40: Snapshots of the evolution of the vorticity in the case of the retrograde (top row) and prograde (bottom row) topographic profiles. Plots are displayed in 12.5 day increments starting from day 25.

scales than the anticyclonic motions. As well, comparing to the figures in the reference we see a similar reduction in the vorticity by the end of the integration.

The evolution of the instabilities are comparable to those observed in the reference case. Isoplots are instead shown in figure 6.41 for the end of the integration with subfigures 6.41a and 6.41b representing the retrograde and prograde topographic cases respectively. As time progresses, we see noticeable features develop at the bottom boundary in the prograde simulation. In coarser resolution calculations these features persisted and were of the same length scale implying that these features are indeed physical. The fact that these features only develop in the prograde simulation is consistent with the orientation of the instabilities observed in the linear calculations particularly in the case of the secondary mode.

6.4 Conclusions

In this chapter, we presented nonlinear simulations of barotropic and baroclinic jets using the HY and QG Tempest model. It was observed that vertical variations in the instabilities that arose from the perturbation of these jets were dependent on the background stratification as was concluded in the linear stability calculations. As well, we were able to gain a better

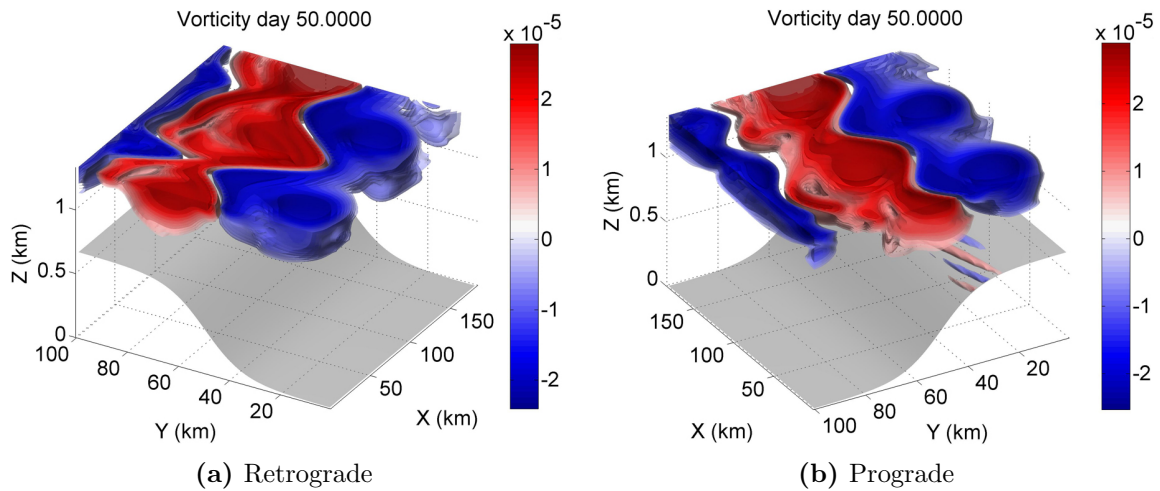


Figure 6.41: The three-dimensional structure of the vorticity field at the end of the integration in the case of (a) retrograde and (b) prograde topographic profiles for the baroclinic jet test case as predicted by the Tempest HY model.

understand of the evolution of the linear and nonlinear instabilities that were generated.

In the barotropic jet cases, we found very strong agreement between the HY and QG models. In the comparison of the energy plots in the reference simulations the QG model was observed to have significantly better energy conservation versus the HY model. In the strongly stratified case, there were qualitative similarities between the two models however the HY model appeared to capture more of the unstable baroclinic modes as observed by the more complex vertical structure than in the case of the QG model. For variable topography, it was observed that the growth rates of instability were comparable to the reference case.

In the baroclinic jet cases, it was found that for the particular resolution chosen that the HY model performed very well. It was proposed that increasing the resolution could potentially lead to NH effects [Mahadevan, 2006]. In the reference case, the simulations from the QG and HY models differed greatly. The HY model agreed with the linear theory, whereas the QG model overestimated the growth rate and significantly underestimated the wavelength. Furthermore, a much greater energy loss was observed in the QG case.

Increasing the stratification did promote better agreement between the QG and HY model results, however it did significantly increase the energy loss in the HY model. In contrast, the QG model performed very well for this particular test. The growth rate and wavelength of the most unstable mode matched fairly closely the linear theory. Therefore, in both the barotropic and baroclinic cases it was observed that the QG model performed

well for strongly stratified cases. The growth rates in the case of topography were nearly equivalent to the reference case and there was no significant change in the energy loss.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

In this thesis, a new fluid dynamics model named Tempest was introduced. The Tempest model was presented with the QG, HY, and NH options. A set of test cases were used to validate various methods and components used in the model. In particular, we showed that the WENO advection scheme performed very well for the Schär advection test, even though terrain-following coordinate systems conventionally perform poorly for this test.

The Tempest model captured the physical features associated with the convection tests as proposed by the MITgcm for the validation of NH ocean models; that of flow down a shelf and a randomly forced open ocean. Further testing was done to compare to [Jones and Marshall, 1993] and Tempest reproduced plume and cone structures associated with randomly forced convection at the surface. The model performed well in the Bowl test case whose configuration had two-dimensionally varying bottom topography [Winton et al., 1998]. Small amounts of entrainment did occur however the results were fairly accurate relative to the isopycnal model presented by [Winton et al., 1998]. The Tempest model was able to capture the Kelvin wave-like features indicated by [Winton et al., 1998], though they were not well-resolved. Plots of the vorticity showed a circulation pattern very closely resembling that seen in the article.

To consider the ability of the Tempest model to capture horizontal motions the model was tested against classical Western boundary current tests and the features as presented by [Fox-Kemper, 2005] were captured. As well, the growth of the Western boundary current was consistent with the Munk layer theory [Adcroft et al., 2014, Akuetevi and Wirth, 2015].

Linear Stability Analyses

The Tempest model was then applied to barotropic and baroclinic jet studies. These studies began with thorough parameter testing using LSA techniques. The studies conducted are similar in nature to those of [Lozier and Reed, 2005] though there are differences in the configurations and conclusions. In particular, their jet had an exponentially decaying structure in the vertical while the jet used in our studies decayed like the error function. This was consistent with our assumed homogenous Neumann boundary conditions on density. Furthermore, the linear stability calculations performed extended on previous studies which looked at a limited number of unstable mode [Xue and Mellor, 1993, Ménesguen et al., 2012, Lozier and Reed, 2005].

In the study by [Lozier and Reed, 2005], it was found that in the case of a weakly baroclinic jet prograde topography leads to a greater growth rate than no topography which, in turn, has a greater growth rate than retrograde topography. They find that in the case of strong stratification all these jets show comparable growth rates. In the case of a strongly baroclinic jet, they find that topography has very little effect on the growth rates while increasing stratification leads to stabilization. Furthermore, they find that the observed changes in stratification are on the same order as the changes introduced by a significant change in Ro . This would explain why, for strong stratification, PE and QG results compare in an otherwise non-QG regime.

We began with a reference barotropic jet calculation for which $Ro = 1$ and $Bu = 10$. It was found that there was very strong agreement between the PE and QG calculations for the first and second sinuous (s_0 and s_1) modes. By varying the background stratification, it was found that weak stratification stabilized a number of modes. In the weakest stratification case of $Bu = 1$ the QG model was able to capture the barotropic sinuous s_0 mode but failed to capture the s_1 mode observed in the PE calculation. In contrast, the strongly stratified case demonstrated the convergence of growth rates to that of the barotropic sinuous mode and a strong level of agreement in the PE and QG cases. The convergence of modes is comparable to the results of [Lozier and Reed, 2005].

From the results, it is concluded that the QG model performs very well in capturing barotropic instabilities. This is particularly interesting as the parameter regime was outside the standard QG regime. If a study is conducted with strong stratification, then it is proposed that the QG calculations could perform as well as the PE calculations. This is beneficial as the QG calculations were significantly more efficient allowing for much finer calculations.

The effects of varying topographic height were studied using a particular hyperbolic tangent idealization of a coastal topography. The results expanded on previous studies

using the RSW [Poulin and Flierl, 2005], QG [Ohshima, 1987, Li and McClimans, 2000, Esler, 2008], and PE [Barth, 1994, Lozier and Reed, 2005] models. We considered the effects that various topographic heights had on a range of unstable modes. In both prograde and retrograde topographic profile the growth rate of the barotropic mode was relatively unaffected, showing $< 2\%$ variability. However, prograde topography led to an increase in growth rates over the flat bottom case which, in turn, had a larger growth rate than the retrograde cases. This was consistent with the findings of [Lozier and Reed, 2005]. For the other modes considered, the first four baroclinic sinuous modes, the extreme prograde and retrograde topographies showed stabilizing effects with the extreme retrograde topography having a greater stabilizing effect on the first baroclinic mode. For modes $s_2, s_3,$ and s_4 the retrograde topography had a decreasing effect while the prograde topography showed relatively strong stabilization in all cases. The topography had very little effect on the most unstable mode and a fairly weak effect on the secondary modes. It is postulated that the QG model could potentially capture the instabilities very well in this case.

In the case of the baroclinic jet the PE and QG models differed substantially. The PE calculations predicted two unstable sinuous modes, s_1 and s_2 . On the other hand, the QG model captured only the s_1 mode and the wavelengths predicted by the PE and QG models differed significantly. However, when the stratification was increased the most unstable mode showed nearly perfect agreement between the PE and QG calculations. It was also concluded that strong stratification caused a significant increase in the growth rate for both cases. This was in contrast to what was observed by [Lozier and Reed, 2005] who showed that an increase in stratification caused a decrease in the growth rates for their strongly baroclinic jet. However, in one of their particular studies [Lozier and Reed, 2005] varied the background stratification by varying the definition of their jet.

Therefore, even outside the QG regime, instabilities in strongly stratified cases can be captured very well by QG LSA calculations. In these cases, it was observed that the QG model can be applied to jets characterized by both barotropic and baroclinic instabilities. As the modal structure was slightly different in both cases it would be beneficial to calculate the structure using PE and QG calculations for at least one wavenumber for comparison. In particular, QG calculations can be used to determine the most unstable mode and a small range around this wavenumber could then be analyzed using the PE calculations.

In the case of topography the baroclinic linear stability calculations showed very little difference between that of the flat bottom. It was only in the case of extremely strong topography that the growth rate of the first or second mode was affected. This was also consistent with [Lozier and Reed, 2005]. In these extreme cases, the prograde and retrograde topography had a slight stabilizing effect. It was assumed that this trend was caused by the interactions of the jet with the topography.

Nonlinear Simulations

Based on the results of the linear stability calculations specific nonlinear simulations were performed. We began with reference calculations in both the barotropic and baroclinic cases consistent with the definition in the linear stability calculations. These simulations were characterized by $Ro = 1$ and $Bu = 10$. In the case of the barotropic jet, it was found that there was very strong agreement between the HY and the QG models and both agreed closely with the linear calculations. Furthermore, the QG model performed exceptionally well with only a 0.056% decrease in energy over the course of the integration.

Weakening the stratification gave rise to nearly vertical invariant motions, while the stronger stratification gave rise to an increase in modal structures in the vertical. Similar sinusoidal structures in the vertical were observed by [Potylitsin and Peltier, 1998, Potylitsin and Peltier, 2003] in their study of columnar vortices. The results of these simulations were consistent with the linear theory. Also, the energy loss appeared to be a function of stratification with the smallest energy loss observed in the weakest stratification case.

The introduction of topography caused an increase in energy loss, with retrograde topography having a more dramatic increase in energy loss over prograde topography with the same topographic height. The growth rates were comparable to the reference case in both the variable stratification cases and the topography cases.

Comparable to the linear stability calculations, we observe that the QG model performs exceptionally well for those studies characterized by barotropic instability. However, for nonlinear simulations the QG model does not appear to be as efficient as the HY model.

In the case of the baroclinic jet, we began our nonlinear calculations by considering a comparison between the HY and NH models. Even in the case of a predominantly hydrostatic regime, it was found that the NH model took over triple the amount of computation time. The results were qualitatively equivalent in both cases and the growth rates were equivalent.

Therefore, the HY model can be used in a multitude of application and offers great computational efficiency over the NH model. This was also observed by [Mahadevan, 2006]. It is possible that with finer resolutions, or for studies characterized by strong vertical motions (in the case when the aspect ratio tends to one), that the NH model may have more applicability.

In the reference baroclinic case, the QG and HY models gave very different results. The HY simulations predicted both a growth rate and wavelength consistent with that of the linear theory. On the other hand, the QG simulations overestimated the growth rate and significantly underestimated the wavelength. Furthermore, the energy was not conserved as

well by the QG model in this case showing an approximate 10.6% energy loss relative to the 3.30% in the case of the HY model.

In an attempt to find better agreement between the QG and HY nonlinear simulations the case of strong $Bu = 50$ stratification was considered. For the baroclinic jet, an increase in stratification caused an increase in the growth rate, demonstrated with both the linear stability calculations and the nonlinear calculations. However, the Tempest HY model did not perform well for this test case, causing a significant increase in energy loss over 25 days of integration. This was also observed to impact the vorticity, which displayed a significant reduction in magnitude.

In contrast, the QG model appeared to perform very well for this particular test. It did show a number of small scale features, some of which may have been computational, but the magnitude of the vorticity was comparable at the beginning and end of the simulation indicating little dissipation. Furthermore, the growth rate and wavelength of the most unstable mode matched fairly closely the linear theory. Therefore, consistent with that observed in the linear theory the QG model performs very well for instabilities characterized by barotropic and baroclinic instabilities when the stratification is sufficiently strong. Additionally, the increase in stratification does not pose a greater time stepping constraint as it does in the HY model since the QG model does not admit inertial gravity waves.

In the case of topography, it was found that small scale features did interact with the bottom boundary causing wave scattering. However, these small scale features did not appear to alter the overall linear growth of the most unstable mode and the results matched very closely the linear theory. The growth rates in the case of topography were nearly equivalent to the reference case and there was no significant change in the energy loss.

All-in-all, the HY model was able to perform exceptionally well in almost all cases aside from the baroclinic case with strong stratification. In comparison, the QG model also did a great job considering the regime of study was not in the classical QG regime. In particular, the QG model did exceptionally well for the barotropic case and appeared to perform very well in the case of the strongly stratified baroclinic jet. Linear stability analysis was used religiously throughout and offered a powerful tool to study a wide range of parameters efficiently and effectively.

7.2 Future Work

As a first step, it would be great to update the QG LSA code to allow for variable topography. It is clear that the QG model performs well even for those cases not strictly within the

QG regime, particularly in those cases characterized by barotropic instability and those cases with strong stratification. The QG LSA calculations are significantly more efficient than the PE LSA calculations and would allow us to look at more topographic profiles and heights efficiently.

In terms of the NH model, it would be interesting to test an FD pressure solve as this could potentially lead to higher order accuracy overall in the momentum equations. Updating the FV density equation could potentially be done using the methods of [Buchmüller and Helzel, 2014]. Testing this against the current second order approach could be valuable. Performing these two updates could allow for a higher global accuracy when vertical velocities are small.

It would be beneficial to attempt finer resolution nonlinear simulations, particularly for the baroclinic jet cases for which NH effects are expected. Comparing HY and NH calculations for the cases with variable topography may shed some light on potential development of NH effects. Steepening the topography by decreasing the width could also generate smaller scale motions, though this may not be handled well by the terrain-following coordinate system. The development of a subsidiary code based upon the immersed boundary methods [Marshall et al., 1997a] could be more adept at handling these types of problems.

As unstable stratification led to instabilities in the HY model when considering convection tests, it would be beneficial to test the NH model in application to the strongly stratified test cases. In these cases, the density perturbations were largest and overturning could possibly lead to NH effects. These cases do pose a stronger restriction on the time step so it would be valuable to first consider nearly two-dimensional tests by reducing the number of grid cells in the along channel direction.

It would also be interesting to extend some of the results of the preliminary test section, particularly the analysis of Western boundary currents. These could be cast into smaller process studies that could emphasize the formation of mesoscale vortices from the Western boundary current. Increasing the grid resolution or isolating a smaller portion of the domain would allow us to capture a lot more of the smaller scale features that were not able to be captured by the coarse grid simulations.

By considering more realistic forcing and comparison to oceanographic studies, we could extend on the insights gained and how the instabilities affect particular coastal regions. [Fratantoni and Pickart, 2003] presented data on how mesoscale variability was affected by baroclinic instability in coastal jets. As well, the studies of [Chapman and Lentz, 1994] and [Gawarkiewicz et al., 1990] demonstrate how dense influx at coastal regions can cause downwelling and produce surface-intensified baroclinic jets. It would be interesting to try to recre-

ate these types of formations through numerical simulations. Furthermore, saline intrusions were observed to be a valuable component of the jet variability [Gawarkiewicz et al., 1990] and so including a more realistic equation of state with an explicit equation for the evolution of saline concentrations could also add further insights to the results.

Either of the barotropic or baroclinic test cases could be extended to a family of test cases considering two-dimensionally varying bottom topographies. If the same hyperbolic tangent profile for topography was to be used, then studies of the impact of topographic heights on the results of nonlinear simulations could be tested by considering a series of different heights. The baroclinic profile in particular could be studied with an increased height, though the sharp topography would be difficult to resolve with the terrain-following coordinate system. With an increase in resolution, computational instabilities could be tamed. As well, the baroclinic test could be extended to consider meddies or topographic boundary currents. There is nearly a limitless number of tests and model updates that can be included in the future.

APPENDICES

Appendix A

Maple Code for FD Coefficients

In this appendix we present the `Coeffs` Maple procedure which has been used throughout the code development to determine the FD coefficients for interpolation and differentiation. The `Coeffs` procedure is based off the `Tay` procedure, which is shown below.

```
restart; with(ArrayTools):with(LinearAlgebra):
Tay := proc(order, dxMult,dyMult) local i,j,k, func,deriv, derivhold,derivhold2,ans;
# Produces the taylor polynomial of f(x + dxMult·Δx, y+dyMult·Δy)
# up to the requested order about the point x=x0,y=y0
ans := 0;
func := (x,y)→f(x,y);
for i from 0 to order do
  deriv := 0;
  if i > 0 then
    for j from 0 to i do
      derivhold := func(x,y);
      for k from 0 to i-1 do
        if j > k then
          derivhold := diff(derivhold,y)*dyMult*Δy;
        else
          derivhold := diff(derivhold,x)*dxMult*Δx;
        end if
      od;
      derivhold := binomial(i,j)*derivhold;
      deriv := deriv+derivhold;
    od;
  else
    deriv := func(x,y);
  end if;
ans := ans+deriv/(factorial(i));
```

```

od;
RETURN(ans);
end:

```

Making use of the Tay procedure the Coeff procedure is defined as follows:

```

Coeffs := proc (order, m, deriv) local func, div, numb, eqns, coeffs, i, j, const,
    err, val;
# Uses Tay procedure to compute the approximation to the derivative defined by deriv
# where 1 implies first derivative, 2 implies second, etc.
# Computes the approximation as: a*f(x+m[1]*Δx)+b*f(x+m[2]*Δx)+\ldots
numb := Size(m, 2);
func := 0;
const := [aa, bb, cc, dd, ee, ff, gg, hh, ii, jj, kk, ll, mm];
for i from 1 to numb do
    func := func+const[i]*Tay(order, m[i], 0);
od;
func := expand(func);
for i from 1 to numb do
    if i = 1 then
        eqns := [coeff(coeff(func, Δx, 0), Δy, 0)];
    else
        eqns := [op(1 .. nops(eqns), eqns), coeff(coeff(func, Δx, i-1), Δy, 0)];
    end if;
    div := f(x, y);
    if 1 < i then
        for j from 1 to i-1 do
            div := diff(div, x);
        od;
    end if;
    if i-1 = deriv then
        eqns[i] := simplify(eqns[i]/div - 1/Δxi-1);
    else
        eqns[i] := simplify(eqns[i]/div);
    end if;
od;
coeffs := op(1 .. numb, op(1, solve(eqns, const[1 .. numb])));
func := 0;
for i from 1 to numb do
    func := func+op(2, coeffs[i])*f(x[0]+m[i]*Δx, y[0]);
od;
func := simplify(func);
for i from 1 to numb do
    func := subs(f(x[0]+m[i]*Δx, y[0]) = f,i',m[i], func);
od;

```

```

err := 0;
for i from 1 to numb do
  err := err+op(2, coeffs[i])*Tay(order, m[i], 0);
od;
err := expand(err);
err := subs(f(x, y) = f(x[0], y[0]), err);
print(Expanded Approximations = err);
val := 0;
for i from 1 to order-deriv do
  val := coeff(coeff(err, Δx, i), Δy, 0);
  val := val*Δxi;
  if val ≠ 0 then
    break;
  end if;
od;
err := val;
err := subs(f(x, y) = f(x[0], y[0]), err);
print(Approximation = func);
if err = 0 then
  printf("Error could not be determined: Try increasing the order.\n");
else
  print(Truncation Error = err);
end if;
end:

```

This procedure is limited to stencils containing no more than 13 points, the number of values in the `const` array. It would be beneficial to extend this to the general case but 13 points has never posed a limitation as the majority of the stencils contain no more than 5 points.

As input, the `Coeffs` procedure requires an order k , which determines the order of the Taylor expansion for each value of the function about $x = x_0$, a set of values that indicate the stencil, and a non-negative integer value indicating the derivative of which we wish to approximate. A value of 0 implies that we wish to approximate the value of the function. The set of function location values are listed as rational values that indicate the number of Δx steps from the approximation location $x = x_0$. The `Coeffs` procedure will output the FD approximation of highest order to the value at the indicated location.

An example calling sequence for the standard second order FD second derivative would be:

`Coeffs(6, [-1,0,1], 2)`

where 6 represents our order, the $[-1, 0, 1]$ represents the function values at $f(x_0 - \Delta x)$, $f(x_0)$, $f(x_0 + \Delta x)$, and the 2 indicates we are considering an approximation to the second derivative. The order argument is passed to the **Tay** procedure above. The **Coeffs** procedure will output the expanded approximation, which is the approximation plus the error terms up to the sixth order terms in the Taylor expansion. As well, it will present the FD approximation and the truncation error.

Appendix B

Nonlinear Case Study

In this particular case study we will go through the barotropic test case. It is assumed that PETSc is installed and the base module of the Tempest package has been compiled.

B.1 Defining the `TestCase#.cpp` File

To initialize a test case a particular `TestCase#.cpp` file within the Tempest model must be defined. These files are called by the preference files found in the `Prefs` folder of the Tempest model directory. Each of the test case files is given a numeric value, which is used as a reference in the `TestCase` option of the appropriate preference file. The majority of the options for a particular test case are set within the test case files. It is here that we set the treatment of the Coriolis terms, the boundary conditions, the advection methods, perturbations, initial state variables, topography, and external forcing terms.

The first step in creating a test is to define this test case file. Only those test case files with an integer attached should ever be modified. For this particular case, most of the settings are already predefined in the `TestCase1.cpp` file, however each option will be explained. In the `Initialize(...)` function, we set the treatment of the Coriolis terms either to traditional or non-traditional, and to f or β plane. For this case, we set these to traditional and f plane respectively. Next, we choose the boundary conditions. In this particular case we consider a jet profile periodic in x and slip conditions in y and z . This is reflected in the following settings:

```
param.m_eBoundaryConditionsX = SimulationParameters::BoundaryConditions_Periodic;  
param.m_eBoundaryConditionsY = SimulationParameters::BoundaryConditions_Slip;
```



```
param.m_eBoundaryConditionsZ = SimulationParameters::BoundaryConditions_Slip;
```

Then the nonlinear advection scheme and the advection in the tracer update equations are set. In this case we will consider a WENO5 scheme for both. Therefore, we choose the settings:

```
param.m_eNonlinearAdvection = SimulationParameters::AdvectionMethod_Weno5;  
param.m_eDensityAdvection = SimulationParameters::AdvectionMethod_Weno5;
```

The background profile of our jet is assumed to be in geostrophic balance so we set `param.m_fGeostrophicBalance = true;` and we do not consider any external forcing, via surface winds or density fluxes. Therefore, we set `param.m_fWindForced = false;` and `param.m_fDensityFlux = false;`. We only apply perturbations to the u velocity field since this velocity is nonzero in the background configuration. We apply a random perturbation with a magnitude of 0.1% relative to the magnitude of our background profile. Therefore, we set `param.m_dPerturbU = 1.0e-3;` while the other perturbation are set to 0.

The remainder of the initialization routine simply reads in preferences from the `barotropic.pref` preference file, defined in the next section, and defines the associated terms in the model. For example, the maximum y value is read from the `YDomainEnd` preference and set to the simulation parameter `m_dYDomainEnd` of the `SimulationParameters` class. This is done using the following command.

```
param.m_dYDomainEnd = aPrefs.GetPreferenceAsDouble("YDomainEnd");
```

The `TestCase1.cpp` file also contains functions to set the value of the u, v, w , and ρ background fields. In our case we are considering a barotropic jet profile of the following form:

$$U(y) = U_0 \operatorname{sech}^2 \left(\frac{y - m_y}{y_d} \right),$$

where $U_0 = 1$ m/s, m_y is the midpoint of the y extent, and $y_d = 0.1 L_y$. This is defined in the `EvaluatePointwiseUVelocity(...)` function as:

```
// Get the system state  
const SystemState & state = block.GetSystemState();  
  
// Get the simulation parameters  
const SimulationParameters & param = state.GetSimulationParameters();
```

```

// Physical y value
double y = dY;

// Initialized background state
double U0 = 1.0;
double y0 = param.m_dYDomainBegin + 0.5 * m_Ly;
double a = (y-y0)/yd;
double sech = 1.0 / cosh(a);

// Barotropic jet
return U0 * sech * sech;

```

The stratification is accounted for in the background density profile, which is set in the EvaluatePointwiseState function:

```

// Get the system state
const SystemState & state = block.GetSystemState();

// Get the simulation parameters
const SimulationParameters & param = state.GetSimulationParameters();

// Constant background density
double dRho0 = phys.GetWaterRho();

// Background buoyancy density (N^2)
double dRhoZ = -pow(m_dN,2.0)*phys.GetWaterRho()*(dZ-m_Lz) / phys.GetG();
dTracer[0] = dRhoZ;

```

This function is called by the EvaluateElementwiseState of the parent TestCase.cpp file. Within the EvaluateElementwiseState Gaussian quadrature is used to define the cell averaged density values.

The topographic profile is set in the GetTerrainHeight function. In this flat bottom case we could set the contents of this function simply to return 0;. However, in general we consider variable topography with a hyperbolic tangent profile. This takes the following form:

```

double A = m_dTerrainAmplitude;

double my = m_dMidY;
double sy = m_dSigmaY;

if (fabs(A) < 1.0e-14) return 0.0;

```

```

double a = (y-my)/sy;
if (m_fPrograde) {
    return 0.5 * A * ( 1.0 - tanh(a));
} else {
    return 0.5 * A * ( 1.0 + tanh(a));
}

return 0.0;

```

Therefore, in the flat bottom case when $A = 0$ we just have that this function will return a value of zero. The derivatives of this terrain profile are set in the `GetTerrainDerivX`, `GetTerrainDerivXX`, `GetTerrainDerivY`, and `GetTerrainDerivYY` functions.

Lastly, we can set the external forcing via the surface wind forcing functions `GetSurfaceU`, `GetSurfaceV`, or the forcing from a density flux via the `GetDensityFlux` function. The seed argument is an integer defined in terms of the global cell index and is set by the system state. It is used to seed the C++ random number generator and the call `srand(seed)`; should be made inside the function for which a random value is requested. This will prevent the parallel processing from reseeding with the same values on different processors.

B.2 Defining the Preference File

The next step is to define a preference file which will call this particular test case. As we are considering a barotropic jet we can use the preferences file `barotropic.pref` located within the `./Prefs` folder. The switch in the preference file `Restart` defines whether the simulation is to be restarted (1) or not (0). The `InputIndex` determines the particular output file that will be inputted upon a restarted simulation. Since we are starting a fresh simulation we will set `Restart = 0`. The `InputIndex` is therefore ignored. The `OutputDir` defines where the grid and output files will be located. The `OutputDir` preference is set as `data/Run_Outputs/folder_name` and in this case we will use the folder name `out-barotropic`. These are all the I/O options we need to set and it will tell the Tempest model where to output the data fields at each output step.

The next set of preferences sets the model and topography configuration. The `Model` preference is set to `HY` or `NH` to indicate which equation set to use. In this case we set the equations to `HY`. The `TestCase` preference is set to 1 to indicate that we will be referencing the `TestCase1.cpp` file for this particular test. To consider the simplest profile we will set the height of the topography relative to the mean domain depth, `Terrain`, to 0. The `TopoProfile`, `TerrainMx`, `TerrainMy`, `TerrainSx`, `TerrainSy` are irrelevant in

the case of the flat bottom but are read into the `TestCase1.cpp` file and so should be left at their default values when `Terrain = 0`. In the jet test case we have the option of setting our jet profile to either barotropic or baroclinic, where here we set the preference `Profile = barotropic`.

Next, we establish some grid information. In this particular case, let us consider a $256 \times 256 \times 128$ grid which we do by setting the `XElements`, `YElements`, and `ZElements` preferences to 256, 256, and 128 respectively. The `HaloElements` preference defines the width of our halo region and should be left at the default value of 4. The MPI splitting is set by `XProcs` and `YProcs` indicating the number of processors to split the data in the x and y directions respectively. It is important that `XProcs` divides `XElements` and `YProcs` divides `YElements` since Tempest does not allow for variable amounts of data per processor. One-dimensional splitting occurs when one of these values is set to 1. The total number of processors used is `XProcs * YProcs`. The beginning and end of the zonal and meridional extents are set as well as the `ZDomainEnd` which is the mean height of the domain. In the case of our hyperbolic tangent profile the mean depth of the domain is set by this preference.

After configuring the grid we set some information about the time-stepping. We set the time step as `DeltaT`, the total integration time to `EndTime`, and the increments at which to print the fields as the `OutputTime`. All of these are set in seconds. In this case we set these settings to 50, 2160000, 5400, to reflect a 50 s time step, an end time of 25 days, and 400 outputs over the course of the simulation. Following this, stratification is set by the preference `N`, the Brunt-Väisällä frequency.

Then we set the solver tolerances for the elliptic problem for pressure and the vertical velocity inversion operator. The relative and absolute tolerances are referenced as `RTol` and `ATol`. In this case, we set all these settings to $1e-8$, the default values. These values can be used for almost every test case. To accelerate convergence they can be increased but it is then important to monitor iterative solver to ensure that it is converging to an updated solution. If more accuracy is required, then these can be reduced. Lastly, we set the viscosity coefficients used in the momentum equations, `Water`, and the density equation, `Tracers`. We will consider only weak dissipation for this test case and set `nuHWater = 0`, `nuVWater = 1e-4` and set the `Tracers` coefficients to 0.

Other preferences can always be defined and retrieved in the `TestCase#.cpp` file. For example, in the case of the barotropic vortex there is a `VortexWidth` preference, which is accessed via the command `m.VortexWidth = aPrefs.GetPreferenceAsDouble("VortexWidth")`. This value is then passed to the simulation parameters class, which can be accessed by the time stepping scheme.

B.3 Running the Test

Between the preference file and the `TestCase#.cpp` file all settings are set for the particular test case. More elaborate test cases may require the modification of other features of the code but this is not considered here. The next step is to run the code. If modifications were done to the `TestCase#.cpp` file or the code has never been compiled then the following command should be called within the TEMPEST model directory.

```
make clean all -s
```

Once the compilation is complete, we can run the code using the following command:

```
mpirun -n N ./all ./Prefs/barotropic.pref
```

where `N` is the total number of processors, `XProcs * YProcs` in the preference file.

At each output step there is text printed to `stdout` and it gives us a means of tracking our progress. An example output is the following:

```
OUTPUT (1)...
Current time:          5.400000000000000e+03
Global courant number: 0.092081652177759474
Total Mass:           7.29730081211883e+13
Total Energy:         6.83333930965094e+01
Runtime 108:          (TOTAL) 00:08:12
                    (AVERAGE) 00:00:04
```

This indicates the total elapsed time and the average time required per time step. After the simulation completes we will get displayed text similar to the following:

```
Total Run Time: 44hr 35min 38sec
-- END SIMULATION -----
-- RESULTS -----
-- END RESULTS -----
```

In this particular case, we see that the simulation required just under 2 days to complete using 64 processors.

B.4 Analyzing the Results

B.4.1 Viewing Slices

One of the easiest ways to view the results of the simulation is to consider slices through the three-dimensional domain. The matrices of the simulation are stored in the `OutputDir` specified by the preference file. If the simulation grid is coarse enough that the data is stored on a local computer then the slices of the data can be viewed through the matlab script `testread.m` in the `./data/MatlabFiles` folder.

Within this file the only settings that need to be configured are the `simname` which specifies the *folder_name* of the `OutputDir`. Either an x , y , or z slice can be chosen by setting the appropriate value of `slice`. The particular field to view can be specified in the `Field` variable. A plot of quivers can be chosen to be displayed by setting `Pquiver` to 1 or 0, where 1 implies quivers are to be displayed. The `nrange` variable indicates which outputs to consider where the outputs are numbered starting at 0. In the above barotropic case we considered 400 outputs so to view all outputs we would consider `nrange = 0:400;`. This `nrange` can be larger than the number of outputs from the simulation without errors being produced. Lastly the `rslice` can be set optionally. This specifies the relative location of the slice where 0 would indicate the first slice and 1 would indicate the last. As default values we consider x and y slices through the middle of the domain and z slices 10% from the top. This then outputs figures to a sub-folder of the `./data/Run_Outputs/0_Figures` folder. This sub-folder has the name *simname_Field_slice*, which allows for a unique identifier for each particular case.

If the simulation is too large to be computed on a local machine and results are saved on a computational server, then a similar octave script has been created which will produce the same figures. This file is called `fnc_testread.m` and is located in the `./data/OctaveFiles` folder. This can be called from a linux terminal and takes in 3 arguments, the `simname`, `slice`, and `Field` as described in the matlab file. Other options can be defined within the `fnc_testread.m` file exactly how they were defined in the matlab script.

On some computational servers the rendering of the figures may produce errors in the plots. If this is the case a `fnc_savedata.m` file has been created in the same octave folder. This takes the same arguments as the `fnc_testread.m` and outputs the data used in the plotting of the figures. Since this data is a particular slice of one three-dimensional field it does not represent much data per output. These data files are stored in the `./data/Run_Outputs/0_Data` folder in a sub-folder of the same name as the figures sub-folder above. This can be read into matlab using the `./data/MatlabFiles/datatofig.m`

script, which will print the data as a figure and output it to a `./data/Run_Outputs/1_Figures` folder in a sub-folder of the same name.

B.4.2 Viewing Three-Dimensional Isosurface Plots

To view isosurface plots particular files of the simulation need to be transferred or, if the simulation was run locally, stored in a `./data/Run_Outputs/folder_name` folder. Assuming the simulation is completed on a computational server the `outGrid.mat`, `nElements.dat`, and `out#####.mat` file associated with the particular output we wish to view need to be transferred. For example, the `out000100.mat` file implies that we are considering the 100th output step.

In the `./data/Matlab_Files` folder we can then use the `isoplot.m` script to view the results. In this script we must set `folder` to `folder_name` above. As well we must set `i` to the numeric value associated with the `out#####.mat`, which for `out000100.mat` implies `i = 100`. A `Field` can be specified similar to above with `Vort` being the most commonly used. If the data is bimodal, such as in the case of the vorticity, then `datatype` should be set to `bimodal`. On the other hand, if it is unimodal than `unimodal` should be used. As well, the mean depth of the domain `H` should be set, which is equivalent to the value `ZDomainEnd` in the preference file. Lastly, `dOT` should be set, equivalent to the value `OutputTime` in the preference file. The displayed isoplot figure is then saved to the `./data/Run_Outputs/0_Figures/simname_Field` folder.

The isocaps and isosurfaces in this script can take a substantial amount of time to render depending on the graphical processing unit of the host computer. If the simulation resolution is too fine than other scripts have been developed to handle this situation. In the `./data/Octave_Files` folder a `fnc_compress_output.m` script will be found. All the variables must be set locally in this file but are all equivalent to those set in the `testread.m` scripts described above. The only exception is the variable `inc` which is a compression factor. A compression factor of `n` implies that we sample every `n`th element in each of the coordinate directions. That is, a value of `inc = 2` compresses a three-dimensional field by a factor of 8.

These compressed files are stored in a sub-folder local to `./data/Octave_Files` with a name `c-simname`. This folder can be transferred to `./data/Run_Outputs` and the script `isoplot_compressed.m` of the `./data/Matlab_Files` can be used to reference these compressed data fields. The displayed isoplot figure is then saved to the `./data/Run_Outputs/0_Figures/c_simname_Field` folder.

B.4.3 Energy and Growth Plots

At each time step of the simulation diagnostics are outputted to the output folder for the simulation. All of these diagnostics are stored in a `stats.dat` file. This is the only file needed to calculate the energy and growth rate plots. Using the `readstats.m` script in the `./data/Matlab_Files` folder we can print both the growth rate and energy plots.

In this script there are only a couple of variables that need to be set. The variable `nsteps` defines the number of equally spaced points we should consider from the `stats.dat` file. For example, the above barotropic case has 43200 time steps and we only want to plot a subset. The other variables that need to be set are the `span`, `mint`, and `maxt`. The `span` determines the period in days over which we should fit the growth rate curve. Reducing this value will generally increase the growth rate marginally. Therefore, it is important to keep this value consistent across all simulations for a given problem type. For example, for all barotropic simulations we set this value to 5 days. The other variables determine the search range, `[mint, maxt]`. Generally `mint` is more important than `maxt` and a value of 5 days for `mint` works well in general.

The figure of the energy presents incremental changes in the potential energy and kinetic energy as well as the relative loss in the total energy. The growth rate curve presents the L1-norm of the perturbations of the u , v , and ρ variables. It then fits a curve to the period of exponential growth of the perturbations, which determines the growth rate of the perturbations. The w velocity curve is not presented since, in almost all applications that we consider, the vertical velocity is significantly smaller than the horizontal velocity components.

In the case of the QG model, we can also consider a simulation by specifying a unique preference file and `TestCase#.cpp` file. In the definition of the QG model we only require the vorticity to be defined in the associated test case file. On the first step, the velocity components can be determined by solving for the streamfunction and then taking the derivative to determine the geostrophic velocities. Currently all fields are being outputted at every time step where the tracer field is now associated with the QGPV. The same files can be used to analyze the results as was presented above. The growth rate is now determined from curves that are fitted to plots of the L1-norm of the u , v , and vorticity perturbation fields.

Appendix C

LSA Configuration

When running the SLEPc LSA code, each test is set by a preference file similar to the nonlinear simulations in appendix B. The `OutputDir` preference is determined from the `folder` variable of the initial guess direct solver. The other options in the preference file are comparable to those in the nonlinear Tempest model, aside from `NEigenvalues`, which determines the number of eigenvalues to compute for each k in `krange`. The solver tolerances are not used by the SLEPc script; instead the tolerance can be set when making the function call. The `TestCase` option is used by the `TestCase.cpp` file and calls one of the particular cases. Currently there are two cases, a barotropic and a baroclinic jet, set by setting `TestCase` to 0 and 1 respectively. In addition, the `Ujet` and `JetWidth` options set the background velocity and size of the jet.

In the `TestCase.cpp` file the background state is set for these particular profiles. Any time a new profile wishes to be tested these values must be changed accordingly and the code recompiled. The important note is to ensure that these profiles are consistent with those used in the initial guess matlab file. The `Analysis.cpp` file is responsible for computing the eigenvalues and eigenvectors over the whole wavenumber spectrum. It first reads in the wavenumbers and initial guesses and then broadcasts them to each processor.

For each of the k values the solver routine `EPSSolve` is then seeded with these initial guesses. This solver routine uses the `EPS` environment of SLEPc which, like the `KSP` environment of PETSc, houses the information about the problem being solved. In the general eigenvalue problem the `EPS` solver holds the matrices A and B of the generalized eigenvalue problem $Ax = B\lambda x$. The calling sequence for a particular eigenvalue problem is then:

```
mpirun -n 4 ./all ./Prefs/Jet.pref -st_type sinvert -st_pc_factor_mat_solver
```

```
_package mumps -eps_smallest_real
```

Other run-time options are:

- `-eps_max_it n` , where n represents the maximum number of iterations before the solver moves on to the next wavenumber.
- `-eps_tol t` , where t represents the required tolerance to determine if an eigenvalue is a valid solution (the default is $1e-8$).
- `-mat_mumps_icntl_14 a` , where a represents the % increase in allocation space used for mumps over the default allotted amount. In the case when the mumps solver is failing it is generally worthwhile to try $a = 200 \dots 1000$.

In the case of the QG LSA code, all the initial setup is the same in terms of defining the initial guesses and configuring the preference file. The fields defining the background state are different in the `TestCase.cpp` and need to be changed according to the test case being considered. The model equations are different but the SLEPc configuration are similar. The calling sequence for this case changes only slightly to:

```
mpirun -n 4 ./all ./Prefs/Jet.pref -st_type sinvert -st_pc_factor_mat_solver  
_package mumps -eps_largest_imaginary
```

This reflects the fact that in the QG model the eigenvalue associated with the maximum growth rate is that with the largest imaginary part.

In both the PE and QG LSA cases the results from the simulations can be analyzed in matlab using the `B_ReadResults.m` scripts in the `./data/Matlab_Files` folder. This file presents the growth rate curve over the entire range of k values and the spatial distribution of the fields for the k value giving rise to the maximum growth. The modenumber can be changed within this file by setting the value of `mode` to something other than 1. If only the growth rate curve wants to be considered we can turn off the printing of the spatial distributions by setting `printvec` to 0.

References

- [Adcroft et al., 2014] Adcroft, A., Campin, J., Dutkiewicz, S., Evangelinos, C., Ferreira, D., Forget, G., Fox-Kemper, B., Heimbach, P., Hill, C., Hill, E., et al. (2014). Mitgcm user manual.
- [Ahrens, 2008] Ahrens, C. (2008). *Essentials of meteorology : an invitation to the atmosphere*. Thomson Brooks/Cole, Belmont, CA.
- [Akuetevi and Wirth, 2015] Akuetevi, C. and Wirth, A. (2015). Dynamics of turbulent western-boundary currents at low latitude in a shallow-water model. *Ocean Science*, 11(3):471–481.
- [Allen et al., 2003] Allen, S., Dinniman, M., Klinck, J., Gorby, D., Hewett, A., and Hickey, B. (2003). On vertical advection truncation errors in terrain-following numerical models: Comparison to a laboratory model for upwelling over submarine canyons. *Journal of Geophysical Research: Oceans*, 108(C1).
- [Amestoy et al., 2000] Amestoy, P., Duff, I., and L’Excellent, J. (2000). Multifrontal parallel distributed symmetric and unsymmetric solvers. *Computer methods in applied mechanics and engineering*, 184(2):501–520.
- [Amestoy et al., 2001] Amestoy, P., Duff, I., L’Excellent, J., and Koster, J. (2001). A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal on Matrix Analysis and Applications*, 23(1):15–41.
- [Arakawa and Lamb, 1977] Arakawa, A. and Lamb, V. (1977). Methods of computational physics. *Academic Press*, 17:174–265.
- [Arbic and Flierl, 2004] Arbic, B. and Flierl, G. (2004). Baroclinically unstable geostrophic turbulence in the limits of strong and weak bottom ekman friction: Application to midocean eddies. *Journal of Physical Oceanography*, 34(10):2257–2273.

- [Armfield and Street, 2005] Armfield, S. and Street, R. (2005). A comparison of staggered and non-staggered grid navier–stokes solutions for the 8: 1 cavity natural convection flow. *ANZIAM Journal*, 46:918–934.
- [Bai et al., 2000] Bai, Z., Demmel, J., Dongarra, J., Ruhe, A., and van der Vorst, H. (2000). *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*. Software, Environments and Tools. Society for Industrial and Applied Mathematics.
- [Balay et al., 2015] Balay, S., Abhyankar, S., Adams, M., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Eijkhout, V., Gropp, W., Karpeyev, D., Kaushik, D., Knepley, M., et al. (2015). *Petsc users manual*. Technical Report ANL-95/11 - Revision 3.6, Argonne National Laboratory (ANL). <http://www.mcs.anl.gov/petsc>.
- [Balay et al., 1997] Balay, S., Gropp, W., McInnes, L., and Smith, B. (1997). Efficient management of parallelism in object oriented numerical software libraries. In Arge, E., Bruaset, A., and Langtangen, H., editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhauser Press.
- [Barth, 1994] Barth, J. (1994). Short-wavelength instabilities on coastal jets and fronts. *Journal of Geophysical Research: Oceans*, 99(C8):16095–16115.
- [Barth et al., 2000] Barth, J., Pierce, S., and Smith, R. (2000). A separating coastal upwelling jet at cape blanco, oregon and its connection to the california current system. *Deep Sea Research Part II: Topical Studies in Oceanography*, 47(5):783–810.
- [Beckmann and Haidvogel, 1993] Beckmann, A. and Haidvogel, D. (1993). Numerical simulation of flow around a tall isolated seamount. part i: Problem formulation and model accuracy. *Journal of Physical Oceanography*, 23(8):1736–1753.
- [Borges et al., 2008] Borges, R., Carmona, M., Costa, B., and Don, W. (2008). An improved weighted essentially non-oscillatory scheme for hyperbolic conservation laws. *Journal of Computational Physics*, 227(6):3191–3211.
- [Boyd, 2000] Boyd, J. (2000). *Chebyshev and Fourier spectral methods*. Courier Corporation.
- [Buchmüller and Helzel, 2014] Buchmüller, P. and Helzel, C. (2014). Improved accuracy of high-order WENO finite volume methods on cartesian grids. *Journal of Scientific Computing*, 61(2):343–368.
- [Chanona et al., 2015] Chanona, M., Poulin, F., and Yawney, J. (2015). The stability of oceanic fronts in a shallow water model. *Journal of Fluid Mechanics*, 785:462–485.

- [Chapman and Lentz, 1994] Chapman, D. and Lentz, S. (1994). Trapping of a coastal density front by the bottom boundary layer. *Journal of Physical Oceanography*, 24(7):1464–1479.
- [Charney, 1948] Charney, J. (1948). On the scale of atmospheric motions. *Geophys. Publ.*, pages 1–17.
- [Charney and Stern, 1962] Charney, J. and Stern, M. (1962). On the stability of internal baroclinic jets in a rotating atmosphere. *Journal of the Atmospheric Sciences*, 19(2):159–172.
- [Chorin, 1968] Chorin, A. (1968). Numerical solution of the navier-stokes equations. *Mathematics of computation*, 22(104):745–762.
- [Cohen and Kundu, 2008] Cohen, I. and Kundu, P. (2008). *Fluid mechanics*. Academic Press, Amsterdam Boston.
- [Colella and Woodward, 1984] Colella, P. and Woodward, P. (1984). The piecewise parabolic method (ppm) for gas-dynamical simulations. *Journal of Computational Physics*, 54(1):174 – 201.
- [Cotter and Shipton, 2012] Cotter, C. and Shipton, J. (2012). Mixed finite elements for numerical weather prediction. *Journal of Computational Physics*, 231(21):7076–7091.
- [Crone and Munday, 2014] Crone, J. and Munday, L. (2014). Parallel performance of linear solvers and preconditioners. Technical report, DTIC Document.
- [Danaila et al., 2006] Danaila, I., Joly, P., Kaber, S., and Postel, M. (2006). *An Introduction to Scientific Computing: Twelve Computational Projects Solved with MATLAB*. Texts in applied mathematics. Springer New York.
- [Drazin and Reid, 2004] Drazin, P. and Reid, W. (2004). *Hydrodynamic Stability*. Cambridge Mathematical Library. Cambridge University Press.
- [Du et al., 2015] Du, J., Shu, C., and Zhang, M. (2015). A simple weighted essentially non-oscillatory limiter for the correction procedure via reconstruction (cpr) framework. *Applied Numerical Mathematics*, 95:173–198.
- [Durrant, 2010] Durrant, D. (2010). *Numerical methods for fluid dynamics : with applications to geophysics*. Springer, New York.

- [Esler, 2008] Esler, J. (2008). The turbulent equilibration of an unstable baroclinic jet. *Journal of Fluid Mechanics*, 599:241–268.
- [Ezer, 2006] Ezer, T. (2006). Topographic influence on overflow dynamics: Idealized numerical simulations and the faroe bank channel overflow. *Journal of Geophysical Research: Oceans*, 111(C2).
- [Fernando et al., 1991] Fernando, H., Chen, R., and Boyer, D. (1991). Effects of rotation on convective turbulence. *Journal of Fluid Mechanics*, 228(5):13–547.
- [Firoozabadi et al., 2003] Firoozabadi, B., Farhanieh, B., and Rad, M. (2003). Hydrodynamics of two-dimensional, laminar turbid density currents. *Journal of Hydraulic Research*, 41(6):623–630.
- [Fox-Kemper, 2005] Fox-Kemper, B. (2005). Reevaluating the roles of eddies in multiple barotropic wind-driven gyres. *Journal of physical oceanography*, 35(7):1263–1278.
- [Fratantoni and Pickart, 2003] Fratantoni, P. and Pickart, R. (2003). Variability of the shelf break jet in the middle atlantic bight: Internally or externally forced? *Journal of Geophysical Research: Oceans*, 108(C5).
- [Gawarkiewicz, 1991] Gawarkiewicz, G. (1991). Linear stability models of shelfbreak fronts. *Journal of Physical Oceanography*, 21(4):471–488.
- [Gawarkiewicz et al., 1990] Gawarkiewicz, G., McCarthy, R., Masse, A., Church, T., and Barton, K. (1990). A gulf stream-derived pycnocline intrusion on the middle atlantic bight shelf. *Journal of Geophysical Research*, 95(22,305):22–313.
- [Grandy, 1997] Grandy, J. (1997). Efficient computation of volume of hexahedral cells. Technical report, Lawrence Livermore National Lab., CA (United States).
- [Griffies, 2012] Griffies, S. (2012). A suite of ocean-only test cases. Technical report, NOAA/GFDL, Princeton, USA.
- [Griffies et al., 2000] Griffies, S., Böning, C., Bryan, F., Chassignet, E., Gerdes, R., Hasumi, H., Hirst, A., Treguier, A., and Webb, D. (2000). Developments in ocean climate modelling. *Ocean Modelling*, 2(3):123–192.
- [Griffiths, 1986] Griffiths, R. (1986). Gravity currents in rotating systems. *Annual Review of Fluid Mechanics*, 18(1):59–89.

- [Guy and Fogelson, 2005] Guy, R. and Fogelson, A. (2005). Stability of approximate projection methods on cell-centered grids. *Journal of Computational Physics*, 203(2):517–538.
- [Haidvogel and Beckmann, 1998] Haidvogel, D. B. and Beckmann, A. (1998). Numerical models of the coastal ocean. *The Global Coastal Ocean, Processes and Methods*, pages 457–482.
- [Harwood, 2005] Harwood, R. (2005). Atmospheric dynamics. University Lecture.
- [Henrick et al., 2005] Henrick, A., Aslam, T., and Powers, J. (2005). Mapped weighted essentially non-oscillatory schemes: achieving optimal order near critical points. *Journal of Computational Physics*, 207(2):542–567.
- [Hereman, 2011] Hereman, W. (2011). Shallow water waves and solitary waves. In *Mathematics of Complexity and Dynamical Systems*, pages 1520–1532. Springer.
- [Hernandez et al., 2005] Hernandez, V., Roman, J., and Vidal, V. (2005). SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans. Math. Software*, 31(3):351–362.
- [Hernandez et al., 2003] Hernandez, V., Roman, J. E., and Vidal, V. (2003). SLEPc: Scalable Library for Eigenvalue Problem Computations. *Lecture Notes in Computer Science*, 2565:377–391.
- [Hoefer et al., 2007] Hoefer, T., Gottschling, P., Lumsdaine, A., and Rehm, W. (2007). Optimizing a conjugate gradient solver with non-blocking collective operations. *Parallel Computing*, 33(9):624–633.
- [Hristova et al., 2014] Hristova, H., Kessler, W., McWilliams, J., and Molemaker, M. (2014). Mesoscale variability and its seasonality in the solomon and coral seas. *Journal of Geophysical Research: Oceans*, 119(7):4669–4687.
- [Irwin and Poulin, 2014] Irwin, R. and Poulin, F. (2014). The influence of stratification on the instabilities in an idealized two-layer ocean model. *Journal of Physical Oceanography*, 44(10):2718–2738.
- [Iskandarani et al., 1995] Iskandarani, M., Haidvogel, D., and Boyd, J. (1995). A staggered spectral element model with application to the oceanic shallow water equations. *International Journal for Numerical Methods in Fluids*, 20(5):393–414.

- [Jiang and Shu, 1996] Jiang, G. and Shu, C. (1996). Efficient implementation of weighted ENO schemes. *Journal of Computational Physics*, 126:202–228.
- [Jones and Marshall, 1993] Jones, H. and Marshall, J. (1993). Convection with rotation in a neutral ocean: A study of open-ocean deep convection. *Journal of Physical Oceanography*, 23(6):1009–1039.
- [Kaye and Laby, 1995] Kaye, G. and Laby, T. (1995). *Tables of physical and chemical constants: and some mathematical functions, (16th ed.)*. Longmans, Green and Company.
- [Kent et al., 2012] Kent, J., Jablonowski, C., Whitehead, J., and Rood, R. (2012). Down-scale cascades in tracer transport test cases: an intercomparison of the dynamical cores in the community atmosphere model cam5. *Geoscientific Model Development*, 5(6):1517–1530.
- [Killworth and Hughes, 2002] Killworth, P. and Hughes, C. (2002). The antarctic circumpolar current as a free equivalent-barotropic jet. *Journal of marine research*, 60(1):19–45.
- [Legg and Marshall, 1993] Legg, S. and Marshall, J. (1993). A heton model of the spreading phase of open-ocean deep convection. *Journal of physical oceanography*, 23(6):1040–1056.
- [Leonard, 1979] Leonard, B. (1979). A stable and accurate convective modelling procedure based on quadratic upstream interpolation. *Computer Methods in Applied Mechanics and Engineering*, 19(1):59 – 98.
- [LeVeque, 1998] LeVeque, R. (1998). Balancing source terms and flux gradients in high-resolution godunov methods: the quasi-steady wave-propagation algorithm. *Journal of computational physics*, 146(1):346–365.
- [Li and McClimans, 2000] Li, S. and McClimans, T. (2000). On the stability of barotropic prograde and retrograde jets along a bottom slope. *Journal of Geophysical Research: Oceans*, 105(C4):8847–8855.
- [Li et al., 2014] Li, Y., Wang, B., Wang, D., Li, J., and Dong, L. (2014). An orthogonal terrain-following coordinate and its preliminary tests using 2-d idealized advection experiments. *Geoscientific Model Development*, 7(4):1767–1778.
- [Liu and Osher, 1996] Liu, X. and Osher, S. (1996). Nonoscillatory high order accurate self-similar maximum principle satisfying shock capturing schemes i. *SIAM Journal on Numerical Analysis*, 33(2):760–779.

- [Liu et al., 1994] Liu, X., Osher, S., and Chan, T. (1994). Weighted essentially non-oscillatory schemes. *Journal of computational physics*, 115(1):200–212.
- [Lozier and Reed, 2005] Lozier, M. and Reed, M. (2005). The influence of topography on the stability of shelfbreak fronts. *Journal of physical oceanography*, 35(6):1023–1036.
- [Mahadevan, 2006] Mahadevan, A. (2006). Modeling vertical motion at ocean fronts: Are nonhydrostatic effects relevant at submesoscales? *Ocean Modelling*, 14(3):222–240.
- [Marshall et al., 1997a] Marshall, J., Adcroft, A., Hill, C., Perelman, L., and Heisey, C. (1997a). A finite-volume, incompressible navier stokes model for studies of the ocean on parallel computers. *Journal of Geophysical Research: Oceans*, 102:5753–5766.
- [Marshall et al., 1997b] Marshall, J., Hill, C., Perelman, L., and Adcroft, A. (1997b). Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling. *Journal of Geophysical Research: Oceans*, 102(C3):5733–5752.
- [Maxworthy and Narimousa, 1994] Maxworthy, T. and Narimousa, S. (1994). Unsteady, turbulent convection into a homogeneous, rotating fluid, with oceanographic applications. *Journal of Physical Oceanography*, 24(5):865–887.
- [McWilliams, 2006] McWilliams, F. (2006). *Fundamentals of geophysical fluid dynamics*. Cambridge University Press, Cambridge.
- [Mellor et al., 1998] Mellor, G., Oey, L., and Ezer, T. (1998). Sigma coordinate pressure gradient errors and the seamount problem. *Journal of Atmospheric and Oceanic Technology*, 15(5):1122–1131.
- [Ménèsquen et al., 2012] Ménèsquen, C., McWilliams, J., and Molemaker, M. (2012). Ageostrophic instability in a rotating stratified interior jet. *Journal of fluid mechanics*, 711:599–619.
- [Miller, 2007] Miller, R. (2007). *Numerical modeling of ocean circulation*. Cambridge University Press, Cambridge.
- [Mueller, 2015] Mueller, J. (2015). *Essentials of Computational Fluid Dynamics*. CRC Press.
- [NASA, 2005] NASA (2005). Wind driven surface currents: Western boundary currents.
- [North et al., 2014] North, G., Pyle, J., and Zhang, F. (2014). *Encyclopedia of atmospheric sciences*, volume 2. Elsevier.

- [Ohshima, 1987] Ohshima, K. (1987). Stability of a barotropic jet on a sloping bottom. *Journal of the Oceanographical Society of Japan*, 43(1):49–60.
- [Pacanowski, 1995] Pacanowski, R. C. (1995). Mom2 documentation, users guide and reference manual, version 2. *GFDL Ocean Group Tech. Rep*, 3(3):232pp.
- [Pedlosky, 1992] Pedlosky, J. (1992). *Geophysical Fluid Dynamics*. Second edition. Springer-Verlag New York.
- [Pedlosky, 2010] Pedlosky, J. (2010). *Ocean Circulation Theory*. Springer Berlin Heidelberg.
- [Phillips, 1951] Phillips, N. (1951). A simple three-dimensional model for the study of large-scale extratropical flow patterns. *Journal of Meteorology*, 8(6):381–394.
- [Potylitsin and Peltier, 1998] Potylitsin, P. and Peltier, W. (1998). Stratification effects on the stability of columnar vortices on the f-plane. *Journal of Fluid Mechanics*, 355:45–79.
- [Potylitsin and Peltier, 2003] Potylitsin, P. and Peltier, W. (2003). On the nonlinear evolution of columnar vortices in a rotating environment. *Geophysical & Astrophysical Fluid Dynamics*, 97(5):365–391.
- [Poulin and Flierl, 2003] Poulin, F. and Flierl, G. (2003). The nonlinear evolution of barotropically unstable jets. *Journal of physical oceanography*, 33(10):2173–2192.
- [Poulin and Flierl, 2005] Poulin, F. and Flierl, G. (2005). The influence of topography on the stability of jets. *Journal of physical oceanography*, 35(5):811–825.
- [Qiu and Shu, 2011] Qiu, J. and Shu, C. (2011). Conservative semi-lagrangian finite difference WENO formulations with applications to the vlasov equation. *Communications in Computational Physics*, 10(4):979.
- [Rabenseifner, 2010] Rabenseifner, R. (2010). Optimization of MPI applications. *High Performance Computing Center Stuttgart*.
- [Roman et al., 2015] Roman, J., Campos, C., Romero, E., and Tomas, A. (2015). SLEPc users manual. Technical Report DSIC-II/24/02 - Revision 3.6, D. Sistemes Informàtics i Computació, Universitat Politècnica de València.
- [Sadourny, 1975] Sadourny, R. (1975). The dynamics of finite-difference models of the shallow-water equations. *Journal of the Atmospheric Sciences*, 32(4):680–689.

- [Schär et al., 2002] Schär, C., Leuenberger, D., Fuhrer, O., Lüthi, D., and Girard, C. (2002). A new terrain-following vertical coordinate formulation for atmospheric prediction models. *Monthly Weather Review*, 130(10):2459–2480.
- [Sekora and Colella, 2009] Sekora, M. and Colella, P. (2009). Extremum-Preserving Limiters for MUSCL and PPM. *arXiv:0903.4200*.
- [Shchepetkin and McWilliams, 2003] Shchepetkin, A. and McWilliams, J. (2003). A method for computing horizontal pressure-gradient force in an oceanic model with a nonaligned vertical coordinate. *Journal of Geophysical Research: Oceans*, 108(C3).
- [Shchepetkin and McWilliams, 2005] Shchepetkin, A. and McWilliams, J. (2005). The regional ocean modeling system: A split-explicit, free-surface, topography-following-coordinate ocean model. *Ocean Modelling*, 9:347–404.
- [Sheu and Lin, 2003] Sheu, T. and Lin, R. (2003). An incompressible navier-stokes model implemented on nonstaggered grids. *Numerical Heat Transfer: Part B: Fundamentals*, 44(3):277–394.
- [Shi et al., 2002] Shi, J., Hu, C., and Shu, C. (2002). A technique of treating negative weights in WENO schemes. *Journal of Computational Physics*, 175(1):108–127.
- [Shu, 1991] Shu, C. (1991). High order ENO and WENO schemes. In *High-Order Methods for Computational Physics*, volume 9, chapter 5, pages 439–582. Springer Science & Business Media.
- [Shu, 1998] Shu, C. (1998). Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws. In Quarteroni, A., editor, *Advanced Numerical Approximation of Nonlinear Hyperbolic Equations*, volume 1697 of *Lecture Notes in Mathematics*, pages 325–432. Springer Berlin Heidelberg.
- [Shu, 2009] Shu, C. (2009). High order weighted essentially nonoscillatory schemes for convection dominated problems. *SIAM review*, 51(1):82–126.
- [Shu and Osher, 1988] Shu, C. and Osher, S. (1988). Efficient implementation of essentially non-oscillatory shock-capturing schemes. *Journal of Computational Physics*, 77(2):439–471.
- [Simpson, 1982] Simpson, J. (1982). Gravity currents in the laboratory, atmosphere, and ocean. *Annual Review of Fluid Mechanics*, 14(1):213–234.

- [Sowa, 1990] Sowa, J. (1990). Stability of a runge-kutta method for the navier-stokes equation. *BIT Numerical Mathematics*, 30(3):542–560.
- [Subich et al., 2013] Subich, C., Lamb, K., and Stastna, M. (2013). Simulation of the navier-stokes equations in three dimensions with a spectral collocation method. *International Journal for Numerical Methods in Fluids*, 73(2):103–129.
- [Suresh and Huynh, 1997] Suresh, A. and Huynh, H. (1997). Accurate monotonicity-preserving schemes with runge-kutta time stepping. *Journal of Computational Physics*, 136(1):83–99.
- [Thomas et al., 2008] Thomas, L., Tandon, A., and Mahadevan, A. (2008). Submesoscale processes and dynamics. *Ocean Modeling in an Eddy Regime*, pages 17–38.
- [Ullman and Cornillon, 1999] Ullman, D. and Cornillon, P. (1999). Satellite-derived sea surface temperature fronts on the continental shelf off the northeast us coast. *Journal of Geophysical Research: Oceans (1978–2012)*, 104(C10):23459–23478.
- [Ullrich and Jablonowski, 2012] Ullrich, P. and Jablonowski, C. (2012). Operator-split runge-kutta-rosenbrock methods for nonhydrostatic atmospheric models. *Monthly Weather Review*, 140(4):1257–1284.
- [Vallis, 2006] Vallis, G. (2006). *Atmospheric and oceanic fluid dynamics : fundamentals and large-scale circulation*. Cambridge University Press, Cambridge.
- [Wang, 2002] Wang, Z. (2002). Spectral (finite) volume method for conservation laws on unstructured grids. *Journal of Computational Physics*, 178:210–251.
- [Wang et al., 2013] Wang, Z., Fidkowski, K., Abgrall, R., Bassi, F., Caraeni, D., Cary, A., Deconinck, H., Hartmann, R., Hillewaert, K., Huynh, H., et al. (2013). High-order cfd methods: current status and perspective. *International Journal for Numerical Methods in Fluids*, 72(8):811–845.
- [Wicker and Skamarock, 2002] Wicker, L. and Skamarock, W. (2002). Time-splitting methods for elastic models using forward time schemes. *Monthly Weather Review*, 130(8):2088–2097.
- [Winton et al., 1998] Winton, M., Hallberg, R., and Gnanadesikan, A. (1998). Simulation of density-driven frictional downslope flow in z-coordinate ocean models. *Journal of Physical Oceanography*, 28(11):2163–2174.

- [WWF-Deutschland, 2014] WWF-Deutschland, U. (2014). *Living Planet Report 2014*. Umweltstiftung WWF - Deutschland.
- [Xue and Mellor, 1993] Xue, H. and Mellor, G. (1993). Instability of the gulf stream front in the south atlantic bight. *Journal of physical oceanography*, 23(11):2326–2350.
- [Zalesak, 1979] Zalesak, S. (1979). Fully multidimensional flux-corrected transport algorithms for fluids. *Journal of Computational Physics*, 31(3):335 – 362.
- [Zhang et al., 2011] Zhang, S., Jiang, S., and Shu, C. (2011). Improvement of convergence to steady state solutions of euler equations with the WENO schemes. *Journal of Scientific Computing*, 47(2):216–238.
- [Zhang and Shu, 2010] Zhang, X. and Shu, C. (2010). On maximum-principle-satisfying high order schemes for scalar conservation laws. *Journal of Computational Physics*, 229(9):3091–3120.
- [Zhang and Shu, 2011] Zhang, X. and Shu, C. (2011). Maximum-principle-satisfying and positivity-preserving high-order schemes for conservation laws: survey and new developments. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, page rspa20110153. The Royal Society.