

The Effects of Confinement on Kelvin-Helmholtz Billows

by

Anton Baglaenko

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

Waterloo, Ontario, Canada, 2016

© Anton Baglaenko 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

Turbulence is one of the most interesting and important, yet poorly understood, aspects of fluid mechanics. The transition from an unstable laminar flow to a fully turbulent flow can follow many paths depending on the background flow, stratification, and the boundary conditions. This study reconsiders a well-studied transition route, namely shear instability of a stratified fluid. The primary objectives of this study are the quantification of the effects of confinement by no slip walls on the evolution of Kelvin-Helmholtz (KH) billows, the application of metrics from turbulence theory to this flow, and the visualization of the dominant three-dimensional dynamics of the flow at a given time. The results are presented in three chapters.

First, the evolution of confined KH billows is examined as it is affected by the vertical domain in a tilted-tube type numerical experiment for a variety of vertical domains and stratifications. The results indicate that confinement dramatically affects the evolution of the KH instability and the induced regions of viscous dissipation. A novel flow structure is the development of anti-billows in the unstratified portion of the domain prior to spanwise destabilization of the flow.

Borrowing from turbulent literature the second and third invariants of the velocity gradient tensor (Q and R) are examined as a tool for the visualization and categorization of flow topology. These tools are used to analyze several canonical fluid flows (such as channel flow and Burgers vortex) from an analytical point of view; they are then applied to transitional flows (namely KH billow formation and evolution).

Simple extensions to the confined KH billows and turbulent metrics are presented in the final chapter with emphasis on verification of established techniques and an eye for future research directions. The role of no-slip boundaries is investigated and deemed critical to the results of the KH billow simulations. The turbulent metrics Q and R are applied to a different type of numerical simulation (large eddy simulation) and a variety of visualization techniques from the literature is compared with a goal of unambiguously identifying coherent structures.

Acknowledgements

Over the course of my studies at UW I've come into contact with so many people to thank that to list them all here would take days. First and foremost I'd like to thank my supervisor, Marek Stastna, for countless hours of excellent science, commiseration, and good rum. Thanks also go out to everyone in the fluids lab for camaraderie, encouragement, and willingness to talk through whatever problem had me stumped that week; oh and also the chess.

The numerical code used in this thesis was developed by Christopher Subich as his Ph.D. project, and I am deeply grateful to his hard work and programming know-how as well as everyone who took the time to run a simulation for me.

Of course I couldn't have survived the last four years without the support of family and friends; so special thanks go out to my parents, my brother (Yuriy), of course my cousin (Yuri - who is my rock), and Jenny for understanding and support.

Dedication

To my parents for inspiring a lifelong love of knowledge.

Table of Contents

| | |
|--|------------|
| Author's Declaration | ii |
| Abstract | iii |
| Acknowledgements | iv |
| Dedication | v |
| List of Figures | x |
| List of Tables | xxv |
| 1 Introduction | 1 |
| 1.1 Hydrodynamics: Qualitative Concepts | 1 |
| 1.1.1 Classical Turbulence Theory | 2 |
| 1.1.2 Hydrodynamic Instabilities | 3 |
| 1.2 Governing Equations | 3 |
| 1.2.1 Navier-Stokes Equations | 5 |
| 1.2.2 Boussinesq Approximation | 6 |
| 1.3 Dimensionless Parameters | 7 |
| 1.4 Shear Instability of Stratified Fluids | 9 |
| 1.4.1 The Taylor Goldstein Equation | 11 |

| | | |
|----------|--|-----------|
| 1.4.2 | Solutions of the T-G Equation for Linear Instability | 13 |
| 1.4.3 | Secondary Instabilities and Transition | 15 |
| 1.5 | Turbulent Metrics | 19 |
| 1.5.1 | Primary, Secondary, and Tertiary Variables | 19 |
| 1.5.2 | Q vs R Plots | 23 |
| 1.6 | Overview | 24 |
| 2 | Kelvin Helmholtz Billows: the classical picture | 26 |
| 2.1 | Characterizing the Flow | 28 |
| 2.2 | Numerical Methods | 29 |
| 2.3 | Two-Dimensional Simulations | 30 |
| 2.3.1 | 1cm Simulations | 32 |
| 2.3.2 | 1cm skewed case | 37 |
| 2.3.3 | 2 cm Simulations | 42 |
| 2.3.4 | 2 cm skewed domain | 44 |
| 2.3.5 | 0.75 cm skewed domain | 48 |
| 2.3.6 | Integrated Measures of Mixing and Dissipation | 53 |
| 2.4 | Three-Dimensional Simulations | 55 |
| 2.4.1 | Symmetric Two-Layer Flow | 56 |
| 2.4.2 | Asymmetric Two-Layer Flow | 71 |
| 2.4.3 | Linear Stratification | 82 |
| 2.4.4 | Discussion | 90 |
| 3 | Q and R: Tertiary Variable Analysis | 94 |
| 3.1 | Analysis Tools | 94 |
| 3.2 | Nonuniqueness | 95 |
| 3.3 | Canonical Flows | 96 |
| 3.3.1 | Two-Dimensional Analysis | 96 |

| | | |
|----------|--|------------|
| 3.3.2 | Irrotational Vortex | 98 |
| 3.3.3 | Solid Body Rotation | 99 |
| 3.3.4 | Internal waves in a Channel | 100 |
| 3.3.5 | Burgers Vortex | 102 |
| 3.3.6 | Taylor-Green Vortex | 105 |
| 3.4 | Visualization Tools | 109 |
| 3.4.1 | Multi-Region Scatter Plots | 112 |
| 3.4.2 | Summary Histograms | 113 |
| 3.4.3 | Augmented Primary Variable Plots | 115 |
| 3.5 | Case Studies | 116 |
| 3.6 | Kelvin-Helmholtz Billows | 116 |
| 3.6.1 | Symmetric Two-Layer Flow | 117 |
| 3.6.2 | Asymmetric Two-Layer Flow | 128 |
| 3.6.3 | Linear Stratification | 136 |
| 3.6.4 | Rayleigh-Taylor Instability | 141 |
| 3.7 | Discussion | 151 |
| 4 | Extensions | 154 |
| 4.1 | KH Extensions | 155 |
| 4.1.1 | Reynolds Stress | 158 |
| 4.2 | QR and Large Eddy Simulations | 163 |
| 4.2.1 | Simulated Filtering | 164 |
| 4.2.2 | 0.75 cm Symmetric Domain | 168 |
| 4.2.3 | 0.75 cm Skewed Domain | 172 |
| 4.3 | Visualization Tools | 176 |
| 4.3.1 | VISIT | 176 |
| 4.3.2 | Helicity | 180 |
| 4.3.3 | The λ_2 parameter | 184 |

| | | |
|-------|-------------------------------|------------|
| 4.4 | Discussion | 189 |
| 4.5 | Future Directions | 190 |
| 4.5.1 | KH | 190 |
| 4.5.2 | Turbulent Metrics | 191 |
| 4.5.3 | Visualizations | 192 |
| 4.6 | Concluding Comments | 192 |
| | References | 194 |

List of Figures

| | | |
|-----|---|----|
| 1.1 | This figure shows a typical background velocity profile (black) and background density profile (red dashed). This is the basic profile used both in the Taylor-Goldstein solver, as well as the induced current in the shear current experiments. | 11 |
| 1.2 | Growth rate versus wavenumber for a centered (red dashed line) and off-center at 75 % of depth (black solid line). Both setups show largest growth at low wavenumbers, indicating that it is the relatively large-scale instabilities which form first. The off-center pycnocline exhibits slightly faster growth at low wavenumbers. | 14 |
| 1.3 | A sample eigenfunction at wavenumber $k = \pi$. The solid line represents the real part of the eigenfunction while the dashed part represents the imaginary part. | 15 |
| 1.4 | This schematic (from [15]) shows the energy budget of a shear flow between kinetic \mathcal{K} , potential \mathcal{P} , and internal energy \mathcal{I} . Reversible changes are shown as solid arrows while irreversible changes (such as dissipation) are shown as dashed lines. | 16 |
| 1.5 | This schematic from [15] shows the formation of a primary instability, in this case a Kelvin-Helmholtz billow. Once the billow reaches small enough scales secondary instabilities begin to form in the spanwise direction, indicated in the diagram. | 17 |
| 1.6 | A schematic demonstrating the hierarchy of flow variables with primary variables such as velocity and density informing secondary variables such as vorticity and dissipation, which in turn can be combined to create tertiary diagnostic variables, such as Q , R , or the Okubo-Weiss parameter. | 20 |

| | | |
|-----|---|----|
| 1.7 | Schematic showing the dominant local flow topologies at various $Q - R$ regimes, reproduced from Ooi et al, [70], figure 1. The quadrants represent flow topologies (stable-focus/stretching (SF/S), unstable-focus/ contracting (UF/C), stable-node/saddle/saddle (SN/S/S) and unstable-node/saddle/saddle (UN/S/S)) that can exist in incompressible flows. The line separating vortex from shear-dominated flow corresponds to $D_A = 0$ | 22 |
| 2.1 | A schematic diagram showing the experimental setup in all simulations in this chapter. The lighter fluid rises to the top while the heavier fluid travels downward, creating shear across the density gradient. | 31 |
| 2.2 | Vertical profiles of density (left) and horizontal velocity (right) of the symmetric 1cm simulation at an early time - before billow evolution has begun. | 32 |
| 2.3 | Formation and evolution of Kelvin Helmholtz billows in a stratified shear flow between two (infinitely long in x and y) plates at 1.2 and 1.3 seconds. The initial evolution of the billows follows the classically prescribed route. The initial sinusoidal disturbance (top) grows until overturning occurs and billows form. | 34 |
| 2.4 | Evolution of Kelvin Helmholtz billows in a stratified shear flow between two (infinitely long) plates at 1.4 and 1.5 seconds. As the billows grow they begin to interact with and become deformed by the boundaries and each other. Notice the regions of significant overturning in the billow cores. | 35 |
| 2.5 | Dissipation ($2\mu e_{ij}e_{ij}$ scaled by 50% of max dissipation in simulation) of KH billows in a stratified shear flow at 1.2 and 1.3 seconds with a contour of density overlayed in white. Note that the initial dissipation occurs near the center of the domain, where the shear is strongest, but that for fully evolved billows the majority of grinding occurs near the boundaries and is entrained into the center of the domain and the billow cores, once again following the areas of highest relative shear. | 36 |
| 2.6 | Dissipation ($2\mu e_{ij}e_{ij}$ scaled by 50% of max dissipation in simulation) of KH billows in a stratified shear flow at 1.4 and 1.5 seconds with a contour of density overlayed in white. At these later simulation times the dissipation is focussed almost exclusively near the boundaries. | 37 |
| 2.7 | Formation and evolution of Kelvin Helmholtz billows in a skewed domain at 1.5 and 1.6 seconds. Note that the billows begin to form as described by classical theory, however quickly start to deform when their size forces them to interact with the boundaries. | 38 |

| | | |
|------|---|----|
| 2.8 | Density and dissipation (top and bottom panels respectively) of KH billows in a skewed domain at 1.7 seconds. Dissipation is scaled by 50% of max dissipation in simulation. The billows have grown to nearly the size of the domain and have deformed significantly. It is worth noting that the only significant locations of dissipation at this time occur at the far boundary. | 40 |
| 2.9 | Density and dissipation (top and bottom respectively) of KH billows in a skewed domain at 1.8 seconds. Dissipation is scaled by 50% of max dissipation in simulation. The fully mature billows begin to continue to be modified by the boundaries and begin to interact with each other. Note again that the significant locations of dissipation are near the boundaries and not the billow cores or braids. | 41 |
| 2.10 | Density contours of KH billows in a symmetric 2 cm domain at 1.4 and 1.5 seconds (top and bottom panels respectively). The initial billow formations follow the classically prescribed path, however before the billows can reach the boundaries the intense currents (and shear) induce significant billow-billow interaction. | 43 |
| 2.11 | Dissipation induced by KH billows in a symmetric 2 cm domain at 1.4 seconds scaled by 50% of max dissipation in simulation. The majority of dissipation occurs at the boundaries above and below the billow cores at the no-slip boundaries. | 44 |
| 2.12 | Density contours of Kelvin Helmholtz billows in a skewed 2 cm domain at 1.4 and 1.5 seconds (top and bottom panels respectively). The billows form and begin to interact with the top boundary, but are quickly deformed and in some cases (as with the billow on the far left) ejected away from the pycnocline location. | 45 |
| 2.13 | Dissipation contours of KH billows in a skewed 2 cm domain at 1.2 and 1.6 seconds (top and bottom panels respectively). We note that the dissipation for mature billows (bottom panel) is extremely weak and confined very strictly to the boundaries. | 47 |
| 2.14 | Log scaled dissipation during the pairing process between 1.3 and 1.6 seconds in 0.25 second increments. While the most significant dissipation occurs near the domain boundaries, we see significant secondary structure near the billow braids which is diffused during the pairing process. | 48 |

| | | |
|------|--|----|
| 2.15 | Density contours (top) and dissipation (bottom) of Kelvin Helmholtz billows in an asymmetric 0.75 cm domain at 1.5 seconds. The initial rollup occurs as classically prescribed (see Caulfield and Peltier and references therein [15]), however it is worth noting the large amount of dissipation occurs at the far boundary as before. | 49 |
| 2.16 | Density contours (top) and dissipation (bottom) of Kelvin Helmholtz billows in an asymmetric 0.75 cm domain at 1.75 seconds. The simulation displays mature billows interacting with the top boundary and exhibiting limited subharmonic pairing. | 51 |
| 2.17 | Density contours (top) and dissipation (bottom) of Kelvin Helmholtz billows in an asymmetric 0.75 cm domain at 2.0 seconds. Note the strong dissipation present at the far boundary. | 52 |
| 2.18 | Kinetic energy (left) and total mixing, M , (right) of all four one-dimensional simulations. Note that the energy of the symmetric simulations is monotonic, while the skewed simulations reach a peak just as the vertical velocities drive the billows away from the boundary. | 53 |
| 2.19 | Total Dissipation (solid lines) and dissipation in the interior - removing the top and bottom 20 grid points - (dashed lines) for the one-dimensional simulations. Notice that the vast majority of dissipation across all simulations occurs within 20 grid points of the boundary. | 55 |
| 2.20 | Isosurface of density showing $\pm 10\%$ of maximum density perturbation at 1.5 (top) and 1.75 (bottom) seconds. Notice that as predicted the initial rollup (top panel) occurs entirely in two dimensions with minimal activity (equivalent to initial noise) in the spanwise (y) direction until significant overturning has occurred in the billow cores (bottom panel). | 61 |
| 2.21 | Density (top panel showing $\pm 10\%$ of maximum density perturbation) and dissipation (bottom panel showing 10 and 50% of max dissipation) isosurfaces of the symmetric quasi-two layer flow at 2 seconds. At this time the billows have fully matured and begin to destabilize in the spanwise (y) direction. The majority of dissipation (denoted in red) occurs not near the billow cores, but the edge of the domain. | 62 |
| 2.22 | Spanwise standard deviation of kinetic energy of the symmetric quasi two-layer flow at 1.75 (top) and 2 (bottom) seconds saturated at 10% of the maximum kinetic energy. As predicted by classical theory, the spanwise (secondary) instabilities develop in the cores and braids of the KH billows where the majority of overturn and mixing occurs. | 63 |

| | | |
|------|--|----|
| 2.23 | Density (top panel showing $\pm 10\%$ of maximum density perturbation) and dissipation (bottom panel showing 10 and 50% of max dissipation) isosurfaces of the symmetric quasi-two layer flow at 2.25 seconds. Note the presence of large regions of high dissipations in the billow cores which did not exist in two-dimensional flows. | 64 |
| 2.24 | Enstrophy (top panel showing 10 and 50% of max enstrophy) and 3D enstrophy (bottom panel showing 10 and 50% of max 3D enstrophy) isosurfaces of the symmetric quasi-two layer flow at 2.25 seconds. Note the tube-like structure of enstrophy corresponding to billow cores, while the 3D enstrophy is relatively low due to the two-dimensionality of the simulation at this time. | 65 |
| 2.25 | Density (top panel showing $\pm 10\%$ of maximum density perturbation) and dissipation (bottom panel showing 10 and 50% of max dissipation) isosurfaces of the symmetric quasi-two layer flow at 3 seconds. The system has destabilized to the point that it now resembles isotropic turbulence. | 66 |
| 2.26 | Enstrophy (top panel showing 10 and 50% of max enstrophy) and 3D enstrophy (bottom panel showing 10 and 50% of max 3D enstrophy) isosurfaces of the symmetric quasi-two layer flow at 3 seconds. The total enstrophy is dominated by tube-like structures (corresponding to billow cores), while the 3D enstrophy identifies regions of spanwise instability. | 67 |
| 2.27 | Shear stress in $x - y$ plane at the top boundary $\mu\sqrt{u_z^2 + v_z^2}$ at 1.25 (top) and 1.5 (bottom) seconds scaled by the maximum local value (0.21 Pa for both outputs). At these times the flow is entirely two-dimensional, and as such the spanwise structure is uniform. | 68 |
| 2.28 | $x - y$ plane of shear stress near the top boundary $\mu\sqrt{u_z^2 + v_z^2}$ at 2.25 (top) and 3 (bottom) seconds scaled by the maximum local value (0.072 Pa and 0.038 Pa respectively). The spanwise destabilization appears as wave-like variations in top and bottom shear lines. | 69 |
| 2.29 | Summary plot showing total kinetic energy, total mixing, and dissipation (both internal and total) scaled by the maximum values achieved over the length of the simulation. Notice that as predicted the mixing and dissipation both peak during the initial, two-dimensional rollup event. As the system mixes and becomes nearly homogeneous gravitational adjustment stops and all motion ceases. | 70 |
| 2.30 | Density isosurface at 1.5 and 1.75 seconds showing $\pm 10\%$ of maximum density perturbation. The initial billow evolution is entirely two-dimensional, evolving as predicted by the 2D simulations. | 72 |

| | | |
|------|--|----|
| 2.31 | Density isosurface at 2.125 and 2.25 seconds showing $\pm 10\%$ of maximum density perturbation. | 73 |
| 2.32 | Standard deviation of kinetic energy in the spanwise direction scaled by the local maximum of kinetic energy (saturated at 10% of max). The majority of activity occurs near the center of the domain, where destabilization in the spanwise (y) direction is strongest. | 74 |
| 2.33 | Dissipation isosurface at 2.125 and 2.25 seconds showing $\pm 10\%$ of maximum density perturbation. | 75 |
| 2.34 | Enstrophy (left) and 3D enstrophy (right) isosurface at 2.125 and 2.25 seconds showing 10% (in green) and 50% (in red) of maximum enstrophy or 3D enstrophy, respectively. The highest regions of enstrophy are located at the far (bottom) boundary. We begin to see some regions of high enstrophy near the center of the domain. The structure of 3D enstrophy appears quite regular at the early times (due to lack of spanwise structure) becoming more interesting as the system destabilizes. | 76 |
| 2.35 | Vertical ($x - z$) slices through the spanwise center of the domain of density (left) and enstrophy (right) at 2.25, 2.375, and 2.5 seconds. As the billows destabilize and secondary braid and core instabilities form enstrophy is generated first near the boundaries then near the center. Despite the largest density fluctuations being confined to the top half of the domain, the bottom boundary experiences much greater dissipation and enstrophy. Additionally the generation of “anti-billows” near the bottom boundary provides a mechanism for interaction at a distance. | 79 |
| 2.36 | Summary plot showing total kinetic energy, total mixing, and dissipation (both internal and total) scaled by the maximum values achieved over the length of the simulation. Notice that unlike the symmetric case the asymmetric simulation experiences two distinct peaks in all the measured quantities, with mixing achieving its maximum during the initial (two-dimensional) event, while kinetic energy and dissipation peak during the second (more three-dimensional) event. | 80 |
| 2.37 | Spanwise-averaged density at later stages of the simulation. The average state transitions from a sharp pycnocline to a pseudo-linear regime. This figure provides a qualitative measure of the mixing which takes place in this simulation, whereby the sharp divide between light and heavy fluid is smoothed out. | 81 |

| | | |
|------|---|----|
| 2.38 | Density isosurfaces (left) showing $\pm 10\%$ of max perturbation (ρ') and spanwise mean (right) at 10 (top) and 13.5 (right) seconds. This early part of the simulation serves to demonstrate the drastically different timescale of the linear stratification as well as the consistency with classical theory in describing the initial billow rollup and formation event. | 82 |
| 2.39 | Dissipation isosurfaces (left) at 10 and (yellow) and 50 (red) percent of local max. Mean dissipation (right) scaled by one quarter of local maximum dissipation. Both figures are presented at 8.5 (top panels) and 10 (bottom panels) seconds. Early dissipation (top panels) - caused by the initial shear and broadening of the pycnocline - quickly gives way to boundary dominated dissipation once the billows have matured (bottom panels). | 83 |
| 2.40 | Enstrophy (left) and 3D enstrophy (right) isosurfaces at 10 (green) and 50 (red) percent of local max at 8.5 (top panels) and 10 (bottom panels) seconds. Mean dissipation (right) scaled by one quarter of local maximum dissipation. The enstrophy at these times is dominated by the billow cores and spanwise nearly constant, making the 3D enstrophy interesting to look at but not meaningful. | 84 |
| 2.41 | Enstrophy (left) and dissipation (right) isosurfaces (10% of local maximum in green and yellow, respectively, and 50% in red) at 7.5 seconds. The vast majority of activity is clearly confined to the top and bottom boundary layers. The induced dissipation and vorticity occur only at the top and bottom boundaries. | 85 |
| 2.42 | Density isosurfaces (top-left), spanwise-averaged density (top-right), dissipation isosurface (middle-left), standard deviation of kinetic energy saturated at 10% of max (middle-right), and enstrophy and 3D enstrophy isosurfaces (bottom left and right) at 12 seconds. Note that by this time the linear simulation has reached a quasi-mixed state, with no distinct billows or braid-core structures visible any longer. | 86 |
| 2.43 | Density isosurfaces (top-left), spanwise-averaged density (top-right), dissipation isosurface (middle-left), standard deviation of kinetic energy saturated at 10% of max (middle-right), and enstrophy and 3D enstrophy isosurfaces (bottom left and right) at 13.5 seconds. This corresponds well to the peak of the secondary instability which develops in the linearly stratified case. | 88 |

| | | |
|------|---|-----|
| 2.44 | Summary plot showing total kinetic energy, total mixing, and dissipation (both internal and total) scaled by the maximum values achieved over the length of the simulation (top) and a semi-log plot of the same variables (bottom). The majority of the interesting dynamics (dissipation, mixing, and the bulk of the enstrophy generation) occurs during the initial rollup event, around 3.5 seconds. Notice, however, the secondary event at 12-14 seconds which causes a significant rise in relative mixing as well as a small bump in dissipation and kinetic energy. | 89 |
| 3.1 | Line plot of Q (red) and R (blue) versus radius (ρ) for a Burgers vortex with $\Gamma = 15, \alpha = 2, \nu = 10^{-3}$. The region close to the origin (where the interesting dynamics in occurs) correspond to the first and third quadrants of the $Q - R$ plane. | 104 |
| 3.2 | A line plot showing Q (top-left) and R (top-right) versus height and Q vs R (bottom) in the Taylor-Green vortex for a vertical line at $x = \pi/4, y = \pi/4$. This relatively small sample demonstrates the presence of all regions of interest in the $Q - R$ plane. | 108 |
| 3.3 | A Q vs R scatter plot with an overlaid line (in black) at $D_A = 0$ showing a quasi-turbulent simulation. | 110 |
| 3.4 | A layered Q vs R plot showing the Q vs R scatter plot at fixed heights (z 's throughout a domain). Superimposed at each level is the line $D_A = 0$ (in black) ranged between the local minimum and maximum values of R . The points in the scatter plot are coloured by quadrant (corresponding to flow topology as seen in figure 1.7) | 111 |
| 3.5 | Q vs R slices at two simulation times. The topological regions (as seen in figure 1.7 are coloured blue for vortex stretching, red for vortex compression, green for planar shear, and yellow for axial strain. | 112 |
| 3.6 | Q vs R scatter plots at various levels - near bottom boundary (red), top boundary(green), and near the domain center (blue). | 113 |
| 3.7 | Histogram at 2 outputs (left and right) showing an overview (top panels) and zoomed in (bottom panels). | 114 |
| 3.8 | Slice of density at 3 seconds (both panels) into the symmetric two-layer simulation. Superimposed in the right panel are the countours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). The cyan points indicate regions of vortex stretching, while the white regions indicate high planar shear. | 115 |

| | | |
|------|---|-----|
| 3.9 | Slice of density at 1.25 seconds into the two-layer simulation. Superimposed are the countours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). The regions of high vorticity (cyan) correspond well to the billow cores and regions of high strain (white) correspond to the braids. | 118 |
| 3.10 | Slice of density at 1.5 seconds into the two-layer simulation. Superimposed are the countours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). Once the billows have grown to fill the domain the regions of extreme Q are in the billow cores, where vorticity clearly dominates, and near the boundary, where interaction with the no-slip boundaries induces significant shear - and dissipation. | 119 |
| 3.11 | Slice of density at 1.625 seconds into the two-layer simulation. Superimposed are the countours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). The simulation is at a state just before three-dimensionalization begins. The domain contains many regions of overturn and significant activity. The $Q - R$ markers have identified several regions of interest, such as the billow cores, where the overturn is most dramatic, braids, and regions where billows come into contact with the no-slip boundary layer near the top and bottom of the domain. | 120 |
| 3.12 | Q vs R scatter plot taken at each horizontal level ($z = z_i$) displayed in an (R, Q, z) space at 1.625 seconds. The black lines correspond to $D_A = 0$ between the minimum and maximum local (at fixed z) values of R . Note the high amount of activity in the first and second quadrants, corresponding to large enstrophy near the center of the domain. | 122 |
| 3.13 | Q vs R plot at three vertical levels, near the bottom boundary (red), center (blue, and bottom boundary (green) at 1.625 seconds with the black line corresponding to $D_A = 0$. Note that the majority of vorticity is generated near the interior of the domain, while the shear occurs primarily near the (no-slip) boundaries. | 123 |
| 3.14 | Vertically layered Q vs R plots between 1.875 and 2.25 seconds. This period shows the transition between two-dimensional rollup to fully-developed billows which begin to destabilize in the spanwise. | 124 |
| 3.15 | Q vs R plots by region (near boundary - red and green - and near center - blue) between 1.875 and 2.25 seconds. This period shows the transition between two-dimensional rollup to fully-developed billows which begin to destabilize in the spanwise. | 125 |

| | | |
|------|--|-----|
| 3.16 | Slice of density at 2.25 seconds into the two-layer simulation. Superimposed are the contours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). The cyan regions once again identify the original billow cores, while the white regions identify shear-dominated parts of this flow, such as the braids, near $x = 0.015$ and the billow-boundary interaction regions near the top and bottom of the domain. | 126 |
| 3.17 | Various $Q - R$ related plots at 4.5 seconds showing density isosurfaces, total Q vs R plots, layered Q vs R plots, and Q vs R plotted by region in the top left, top right, bottom left, and bottom right, respectively. The overall impression is that the state of the simulation corresponds well to classical turbulence, and indeed we can note the canonical “ice cream cone” shape begin to emerge. | 127 |
| 3.18 | A density slice at 1.625 seconds with regions of vortex stretching (cyan) and planar shear (white) as well as the line $D_A = 0$ (black) overlaid. We immediately observe that this early output follows the classically prescribed structure and mechanisms with vorticity dominating the billow cores and shear dominating in the braids. What is novel is the presence of shear where the billows come near the no-slip boundary. | 129 |
| 3.19 | Augmented density slices with extreme regions of $Q - R$ in the first (cyan) or third (white) quadrants representing vortex stretching or shear dominated flow, respectively between 1.75 and 2.125 seconds showing the evolution of confined billows in an asymmetric pycnocline and the formation of the “anti-billow” structures in the relatively density constant bottom half of the domain. | 130 |
| 3.20 | Vertically layered Q vs R plot at 1.875 seconds showing the Q vs R structure at constant heights (z) during the critical point of anti-billow formation. | 131 |
| 3.21 | Density isosurface (top-left), mean density (top-right), dissipation isosurface (bottom-left), and enstrophy isosurface (bottom-right) at 1.875 seconds. Note the obvious presence of dissipation in the anti-billow regions, but a total lack of enstrophy. | 132 |
| 3.22 | Density isosurface (top-left), mean density (top-right), dissipation isosurface (bottom-left), and enstrophy isosurface (bottom-right) at 2.125 seconds. The spanwise destabilization is well under way in both the billow and anti-billow regions. | 133 |

| | | |
|------|---|-----|
| 3.23 | Various $Q - R$ related plots at 1.875 seconds showing density isosurfaces, total Q vs R plots, layered Q vs R plots, and Q vs R plotted by region in the top left, top right, bottom left, and bottom right, respectively. The flow topology is clearly dominated by shear in the boundary and vorticity near the domain center. | 134 |
| 3.24 | Various $Q - R$ related plots at 2.25 seconds showing density isosurfaces, total Q vs R plots, layered Q vs R plots, and Q vs R plotted by region in the top left, top right, bottom left, and bottom right, respectively. Two important features of note is the fact that the most extreme vortex stretching happens near the domain center, and the region near the bottom (far) boundary is by far the very active, providing a unique superimposed profile. | 135 |
| 3.25 | Density isosurface (top-left), augmented density slice (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 9 seconds. The linear case evolves in much the same manner as the quasi two-layer case, albeit much more slowly. | 137 |
| 3.26 | Density isosurface (top-left), augmented density slice (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 12 seconds. The simulation at this time is still fundamentally two-dimensional with vortex stretching the dominant factor present. | 138 |
| 3.27 | Total Q vs R plots between 13 and 17 seconds showing the initial calming period as well as the rise of the second instability around 16 seconds. . . . | 139 |
| 3.28 | Total Q vs R plots between 13 and 17 seconds showing the initial calming period as well as the rise of the second instability around 16 seconds. Identical to figure 3.27 but with axes fixed for better comparison or the energy in the system. | 140 |
| 3.29 | Density isosurface (top-left), augmented density slice (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 17.5 seconds during the secondary rollup event. | 141 |
| 3.30 | The density evolution of the RT instability with axes set to scale. The evolution is shown from 0 to 1.4 seconds in 0.2 second increments. The total density difference shown is 2%. | 143 |
| 3.31 | Total Q vs R plots between 1.2 and 2.8 seconds showing the entire evolution of the Rayleigh-Taylor instability. | 144 |

| | | |
|------|---|-----|
| 3.32 | Total Q vs R plots between 1.2 and 2.8 seconds showing the entire evolution of the Rayleigh-Taylor instability. This figure is identical in content to figure 3.31 except with axes fixed for easier comparison of the energy in the system. | 145 |
| 3.33 | Augmented density plot (top-left), total Q vs R plot (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 1.2 seconds during the initial formation of the Rayleigh-Taylor instability. | 146 |
| 3.34 | Augmented density plot (top-left), total Q vs R plot (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 1.8 seconds. The most active region is still by far in the boundary layer and the overall shape at each level seems to correspond topologically to isotropic turbulence. | 147 |
| 3.35 | Augmented density plot (top-left), total Q vs R plot (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 2.4 seconds. The center has fully destabilized, causing increased velocity gradients and making the center of the domain dominate the Q vs R region plots. | 149 |
| 3.36 | Augmented density plot (top-left), total Q vs R plot (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 2.6 seconds. The plug has completely destabilized and is riddled with alternating regions of high shear and vorticity (in cyan and white on the augmented density plot). | 150 |
| 3.37 | Summary plot of Rayleigh-Taylor instability. As expected the kinetic energy and dissipation (both total and internal) rise monotonically as the plug accelerates. In contrast the mixing reaches a peak around 2.4 seconds. | 151 |
| 4.1 | Density plots of the asymmetric 0.75 cm simulation with free-slip boundaries at 1.2, 1.3, 1.4, and 1.5 seconds. The initial destabilization occurs slightly faster (1.2 for free slip compared to 1.3 seconds for no-slip boundaries). | 156 |
| 4.2 | Dissipation plots (with contours of density in white) of the asymmetric 0.75cm simulation with free-slip boundaries at 1 and 1.1 seconds. The only significant dissipation present is during the initial pycnocline broadening. Once the billow formation begins there is almost no dissipation present. | 156 |
| 4.3 | Log of dissipation plots of the asymmetric 0.75cm simulation with free-slip boundaries at 1.2 and 1.4 seconds. The structure of the dissipation follows fairly closely the shape of the billows, however the magnitude is so weak that the result is nearly negligible. | 157 |

| | | |
|------|---|-----|
| 4.4 | Density plots of the asymmetric 0.75cm simulation with free-slip boundaries at 1.6, 1.7, 1.8 and 2.0 seconds. Note the total lack of coherent structures in this flow with subharmonic pairing dominating nearly ever aspect of the billow evolution. | 158 |
| 4.5 | Mean velocities (left) and Reynolds stresses (right) at 2.1875s for the asymmetric simulation scaled by the maximum Reynolds stress at this time. The largest mean velocities (and the largest variations causing Reynolds stress) occur in the horizontal and vertical velocities. | 160 |
| 4.6 | Horizontally integrated mean velocities and Reynolds stresses for the asymmetric simulation at 2.1875 seconds.It is clear from the integrated values that the mean velocities are far more important than the reynolds stresses. | 161 |
| 4.7 | Horizontally integrated mean velocities and Reynolds stresses for the asymmetric simulation at 2.375 seconds. The mean and spanwise veocities are much higher at this time, and it is clear from the spanwise mean velocity (dotted line) and the $\langle uw \rangle$ component of the stress tensor that the spanwise instabilities are much more pronounced. | 162 |
| 4.8 | contour plots of horizontal velocity derivatives ($\frac{\partial u}{\partial x}$, $\frac{\partial u}{\partial y}$, and $\frac{\partial u}{\partial z}$ in top, middle, and bottom panels respectively. Unfiltered (DNS) variables are on the left and filtered (LES) variables are on the right. | 167 |
| 4.9 | Total Q vs R plot from a vertical slice of the 0.75 symmetric domain for unfiltered (top) and filtered (bottom) fields at 2.25 seconds. While the general shape is preserved, the filtered velocity field results in greatly reduced activity in the planar and axial shear regions of the Q vs R plot as well as an overall reduction of extreme values. | 169 |
| 4.10 | Total Q vs R plot from a vertical slice of the 0.75 symmetric domain for unfiltered (top) and filtered (bottom) fields at 2.75 seconds. The general shape is consistent, however we see dramatic reductions in the planar shear region as well as the first quadrant (near $x - axis$) activity. | 170 |
| 4.11 | Slice of density (unfiltered) at 2.75 seconds into the two-layer simulation. Superimposed are the countours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white) for the unfiltered (top) and filtered (bottom) variables. At this time in the simulation the flow is primarily two-dimensional. There are as yet no small-scale features and thus we are unsurprised to find the two plots almost identical. | 171 |

| | | |
|------|--|-----|
| 4.12 | Total Q vs R plot from a vertical slice of the 0.75 asymmetric domain for unfiltered (top) and filtered (bottom) fields at 2.25 seconds. As in the symmetric simulation we note the preservation of the overall structure (with enstrophy clearly dominating the flow topography) but a significant loss of outliers. | 173 |
| 4.13 | Slice of density (unfiltered) at 1.875 (left) and 2.25 (right) seconds into the asymmetric two-layer simulation. Superimposed are the contours of $D_A = 0$ (in black) as well as regions of high positive Q and R (cyan) and high negative Q and R (white) for the unfiltered (top) and filtered (bottom) variables. | 174 |
| 4.14 | Total Q vs R plot from a vertical slice of the 0.75 asymmetric domain for unfiltered (top) and filtered (bottom) fields at 2.75 seconds. At this relatively late, and quite turbulent, time the differences between the unfiltered and filtered variables become more significant. | 175 |
| 4.15 | Isosurfaces of density (2 levels in red and green) with velocity gradient tensor (in blue) overlaid at 1.875 seconds. While the isosurfaces are not capable of identifying the anti-billow regions, the tensor plot shows significant activity in that region. | 177 |
| 4.16 | Isosurfaces of density (2 levels in red and green) with velocity gradient tensor (in blue) overlaid at 2.3125 seconds. Note the regions of near-constant density with large tensor activity (the anti-billow regions) as well as regions of destabilization where the orientation of the tensor disk has changed to signify a new flow topology. | 178 |
| 4.17 | Isosurfaces of density (2 levels in red and green) with velocity gradient tensor (in blue) overlaid at 2.5 seconds. | 179 |
| 4.18 | Isosurfaces of helicity in the symmetric 0.75 cm domain, $H_n = 0.8$ at 2.25 seconds coloured by kinetic energy (left) and enstrophy (right) with the darkjet colorbar - blue representing low values and red representing high. Notice that the majority of regions identified by helicity have almost no enstrophy. | 181 |
| 4.19 | Isosurfaces of helicity in the asymmetric 0.75 cm domain, $H_n = 0.8$ at 2.25 seconds coloured by kinetic energy (left) and enstrophy (right) with the darkjet colorbar - blue representing low values and red representing high. The helicity has captured primarily the division between primary and anti-billow regions. | 182 |

| | | |
|------|--|-----|
| 4.20 | Isosurfaces of helicity in the asymmetric 0.75 cm domain, $H_n = 0.8$ at 2.5 seconds coloured by kinetic energy (left) and enstrophy (right) with the darkjet colorbar – blue representing low values and red representing high. The simulation at this time is very active and no longer contains any coherent structures, something the helicity reflects here. | 183 |
| 4.21 | Isosurfaces of λ_2 in the symmetric 0.75 cm domain at 10% of minimum value at 2.375 seconds coloured by vertical height. This plot successfully identifies the billow cores as well as the regions of intense dissipation near the boundaries – located directly above and below the billows. | 185 |
| 4.22 | Isosurfaces of λ_2 in the symmetric 0.75 cm domain at 10% of minimum value at 2.875 seconds coloured by vertical height. The coherent billow cores are still clearly visible, as is the interaction with the no slip boundaries at the top and bottom of the domain. | 186 |
| 4.23 | Isosurfaces of λ_2 in the symmetric 0.75 cm domain at 2.5% of highest minimum value achieved between 1.75 and 2.75 seconds in 0.125 second increments coloured by vertical height. The early times clearly identify both primary and anti-billow locations. Mid times (2.25–2.5 s) show the span-wise destabilization and start of the turbulent regime. Late times (2.5–2.75 s) show the onset and rise of turbulence and the destruction of the earlier coherent structures. | 187 |

List of Tables

| | | |
|-----|---|----|
| 2.1 | Relevant nondimensional simulation parameters. | 30 |
| 2.2 | Two-Dimensional simulations and their numerical parameters. | 32 |
| 2.3 | Three-Dimensional simulations and their numerical parameters. | 56 |

Chapter 1

Introduction

Few scientifically literate individuals would need to be convinced that the many naturally occurring motions of water are worth studying. These motions cover an amazing range of scales from droplets modified by surface tension, to ocean currents observable from space. A fundamental, and to this day poorly understood, aspect of many flows is the interaction between disparate scales and the role of this scale interaction in the transition between an initially laminar flow and a turbulent end state. In this work, we seek to better understand a well-described transition route - the formation and subsequent destabilization of Kelvin Helmholtz billows - in a confined setting, i.e. in the presence of boundaries. The interaction between large-scales (those associated with the background shear) and small-scales (such as the boundary layer produced by no-slip boundaries) is an aspect of a classical problem which has the potential to illuminate broader aspects of the transition to turbulence.

1.1 Hydrodynamics: Qualitative Concepts

Shear instability is a ubiquitous and important aspect of fluid flows at all scales - from the atmosphere, to the ocean and lakes, and all the way down to small-scale flows like pipe or tilted tube problems. The formation and growth of instabilities in a shear flow is the primary mechanism for converting the large background energy present in these systems into mixing, viscous dissipation, and other relevant dynamics. Shear instabilities also provide a well-studied ‘backbone’ of the transition to a turbulent state, a regime which is oft-described but still quite incompletely understood [25].

1.1.1 Classical Turbulence Theory

While historically, the motion of fluids has been described by a variety of approximate theories (e.g. potential flow, Stokes flow), turbulence is an essential aspect of real physics and engineering problems involving fluids. Since almost any real flow is turbulent to some extent, the nature of the transition to turbulence, and the subsequent evolution of the turbulent flow, is a widely applicable field of study that can help us better understand the nature of many flows.

As a brief motivation it is worth reiterating the classical picture of turbulence as described by Kolmogorov [53] as an early attempt to describe the behaviour of isotropic homogeneous turbulence. Interestingly, Kolmogorov's theory is not dependent on what actual partial differential equations govern fluid motion.

Kolmogorov, expanding on Richardson's idea of the spectral energy cascade - the process whereby energy is moved from the large (injection) scale down to the dissipation scale - hypothesized that the statistics of small scales are isotropic and depend only on viscosity, ν , and the rate of dissipation, ε . The outline of Kolmogorov's theory is that energy is transferred from larger to smaller eddies down the smallest scale. At this small scale is where the turbulent energy is dissipated. Assuming isotropy (turbulence has no preferred direction), similarity (statistics of small-scale motion are universal and determined by ε and ν), and the existence of a smallest scale, Kolmogorov argued that since at the smallest scales kinematic viscosity and dissipation are the most important terms, then it is the balance between these two quantities that provides the length scale of the smallest eddies.

Viscous dissipation in a fully turbulent flow is defined as the double contraction of the rate of strain tensor $\varepsilon = 2\nu e_{ij}e_{ij}$ from conservation laws. It is only at the small scales that the energy dissipation is effective. A critical part of the classical framework presented by Kolmogorov is that in turbulent flows large-scale mechanisms of vortex stretching create regions where small scale dissipation becomes important and dissipates the initially input energy. If we assume that ε does not depend on ν , but is instead determined by inviscid properties of large eddies (of size l and velocity scale u') we arrive at the scaling that energy of the eddy is proportional to $(u')^2$ and the timescale is of order $\frac{u'}{l}$. Thus the rate of energy dissipation in the fluid can be given by $\varepsilon \sim \frac{(u')^3}{l}$.

The dimensions of dissipation are $\frac{\text{length}^2}{\text{time}^3}$ and the units of kinematic viscosity are $\frac{\text{length}^2}{\text{time}}$. Based on dimensional analysis Kolmogorov concluded that the smallest scales are of size

$$\eta = \left(\frac{\nu^3}{\varepsilon} \right)^{\frac{1}{4}} .$$

This is known as the Kolmogorov length scale (see the discussions in [55] or [25] for more details)

Kolmogorov also postulated that between the injection scale and the smallest scale there exists an inertial subrange where the viscosity is not important and energy scales with dissipation and wavenumber (eddy size). This is the five thirds law which relates the energy to the wavenumber as $S \sim \varepsilon^{2/3} k^{-5/3}$. This result is in good agreement with experimental observations.

1.1.2 Hydrodynamic Instabilities

While Kolmogorov's scaling laws are a time-tested method of describing isotropic turbulence, they are a purely phenomenological theory and does not explicitly use the Navier-Stokes equations. In contrast to this hydrodynamic stability theory is a method of attempting to quantify the nature of the onset of instabilities of fluid flows which may develop into full turbulence. This method uses the Navier-Stokes equations - which are derived from conservation laws - to predict the rate of growth of certain modes for a perturbation (often assumed to be small) of the initial flow. This initial flow often takes a simple form that is assumed to be representative of a broader class of naturally occurring flows.

The transition between laminar and turbulent flow is a complex process. The exact nature of the transition depends on the type of flow and the driving forces, i.e. background shear flow for Kelvin Helmholtz (KH) billows, temperature gradients for convection problems, gravitational (or Rayleigh-Taylor) instability when more dense fluid overlies less dense fluid. Understanding the onset and evolution of these instabilities is the key to understanding transitional flows and in turn understanding the formation and evolution of any turbulent system. In classical (Kolmogorov) theory the energy spectrum in the inertial range scales with wave number as $k^{-5/3}$ up to a dissipation threshold. This leads to the estimate that the degrees of freedom in an isotropic turbulent flow scale as $Re^{9/4}$ [93], where $Re = \frac{UL}{\nu}$ is the Reynolds number. This implies that a high Reynolds number flow would require a great deal of computational power to resolve and hence simple flows provide a more tractable building blocks with which to understand at least some parts of turbulence.

1.2 Governing Equations

In this section background relating to the governing equations and approximations are discussed. In particular, we focus on the Boussinesq approximation.

A fluid system is described by the velocity field of a continuum of fluid particles. A fundamental aspect of continuum mechanics is the so called continuum hypothesis - the assumption that any matter within classical physics can be treated as a continuum when viewed for sufficiently large length scales when compared to the size of individual molecules. This assumption allows us to treat the displacement or velocity fields of a solid or fluid as a continuous field that is well defined at infinitely small points. Due to the large size of Avogadro's number (i.e. the large number of molecules in a mol of matter) the continuum hypothesis is a reasonable assumption which produces extremely accurate results. The continuum hypothesis leads naturally to the idea of the fluid parcel (sometimes called a fluid particle). The fluid parcel is an infinitely small amount of fluid (formally a finite volume taken in the limit to zero) which can be "tagged", whose motion can be considered as the fluid evolves and which carries well defined properties (temperature, density and so on).

The constantly evolving nature of fluid flows leads to a natural discussion of frames of reference. As an initially tagged fluid parcel evolves it changes not only position but also shape. This means that we must consider the path of the fluid parcel as it moves through space and time. The Lagrangian reference frame follows the position of an individual parcel. This can be likened to taking measurements on a boat as it is swept along by the current, or from a weather balloon as it is advected by the wind. The Eulerian reference frame assumes an external, stationary observer and considers the motion of the fluid within the region that observer's measurable region. This is often likened to watching a river from the bank, or measuring with a probe at a fixed location.

Since the field theories of classical physics, including fluid mechanics, are presented from the Eulerian point of view, one of the most fundamental concepts in fluid mechanics is the idea of the material derivative (also known as the substantial derivative, or particle derivative):

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + u_i \frac{\partial}{\partial x_i}$$

(following Einstein summation convention whereby repeating indices indicate a sum over that index). The material derivative represents how a variable at a fixed location changes with respect to time and due to the motion of the fluid particles. Thus the first term is a partial derivative with respect to time and the second is the inner product of the fluid velocity and the spatial change of the variable. The first part is known as the local derivative, and the second part as the advective derivative [55].

The behaviour of the various fields of interest (such as velocity, density, temperature and pressure) is described by a system of partial differential equations known as the Navier-Stokes equations.

1.2.1 Navier-Stokes Equations

Once you add the discussion of Eulerian and Lagrangian variables above you can naturally define a material volume as the volume consisting of the fluid particles you “tag” at the initial time.

Conservation of mass means that the time derivative of the mass of a material volume, $V(t)$, is constant

$$\frac{dM(t)}{dt} = \int \int \int_{V(t)} \left(\frac{D\rho}{Dt} + \rho \frac{\partial u_i}{\partial x_i} \right) dV = 0.$$

Here the Reynolds Transport Theorem has been used to move the time derivative inside the integral. By the Dubois-Reymond Lemma the above expression implies that

$$\frac{D\rho}{Dt} + \rho \frac{\partial u_i}{\partial x_i} = 0.$$

This is known as the continuity equation.

From Newton’s laws (specifically the conservation of linear momentum, with an additional algebraic manipulation using the continuity equation) we know that

$$\rho \frac{Du_i}{Dt} = \rho g_i + \frac{\partial \tau_{ij}}{\partial x_j}.$$

This equation equates the evolution of the velocity field (the acceleration following a fluid particle) and the forces experienced by the fluid, namely the body force due to gravity, and the forces due to neighbours in the continuum as expressed (using Cauchy’s theorem) as the divergence of the stress tensor [55]. For a Newtonian fluid the stress can be written as

$$\tau_{ij} = -p\delta_{ij} + 2\mu e_{ij} + \lambda e_{mm}\delta_{ij}.$$

where $e_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)$ is called the rate of strain tensor. We can rewrite the trace of the rate of strain tensor as $\frac{\partial u_i}{\partial x_i}$. If the fluid is incompressible (as our system is) we instead of compressibility the condition that

$$\frac{\partial u_i}{\partial x_i} = 0.$$

Since the bulk viscosity, λ is multiplied by the divergence of velocity in the stress relation this term is often ignored. Note that the divergence of the velocity is not exactly zero but

rather very small. Since this approximation holds very well for water and specifically for the types of flows we are interested in this makes for an excellent approximation.

The equation of motion for a Newtonian fluid is given by

$$\rho \frac{Du_i}{Dt} = -\frac{\partial p}{\partial x_i} + \rho g_i + \frac{\partial}{\partial x_j} [2\mu e_{ij} - \frac{2}{3}\mu (\frac{\partial u_i}{\partial x_i}) \delta_{ij}], \quad (1.1)$$

where $\mathbf{u} = (u, v, w)$ is the velocity, ρ is the fluid density, g is the acceleration due to gravity, p is pressure, μ is the viscosity, and e_{ij} is the rate of strain tensor. The material derivative operator is (once again) defined as $\frac{D}{Dt} = \frac{\partial}{\partial t} + u_i \frac{\partial}{\partial x_i}$. This is the general form of the Navier-Stokes equation.

In the case where the temperature differences are small within the fluid and the viscosity can be taken as spatially constant the terms inside the square bracket can be rewritten as $\mu [\frac{\partial^2}{\partial x_k \partial x_k} u_i + \frac{1}{3} \frac{\partial}{\partial x_i} u_i]$. In the case of incompressible flow we have the additional restriction that $\frac{\partial u_i}{\partial x_i} = 0$, so the Navier-Stokes equations reduce to

$$\rho \frac{Du_i}{Dt} = -\frac{\partial p}{\partial x_i} + \rho g_i + \mu \frac{\partial^2}{\partial x_k \partial x_k} u_i \quad (1.2)$$

$$\frac{\partial u_i}{\partial x_i} = 0 \quad (1.3)$$

$$\frac{D\rho}{Dt} = 0. \quad (1.4)$$

1.2.2 Boussinesq Approximation

A fundamental cornerstone of modern fluid mechanics as it relates to water is the simplifying assumption of the Boussinesq approximation. Following the note by Spiegel and Veronis [85] and the derivation in Kundu's Fluid Mechanics textbook [55] we will present here a brief summary of how the full nonlinear equations governing a fluid can be simplified to produce the set of equations which we will be simulating throughout this work.

The two fundamental assumptions underlying the Boussinesq approximation are as follows [55, 85]: the fluctuations in density result primarily from thermal (and not pressure) effects; and in the equations for momentum and mass conservation the density fluctuations can be neglected except when they are coupled to the gravity force in the buoyancy term.

If the density field varies slowly throughout the domain we can decompose the density field into a mean reference density, ρ_0 , the stratification term, $\bar{\rho}(z)$, and the motion-induced disturbances ρ' as

$$\rho = \rho_0 + \bar{\rho}(z) + \rho'.$$

The magnitude of the stratification and motion-induced terms are both much smaller than the reference density for water flow.

Since the deviations in the density field are small (switching back to vector notation for ease of understanding) we replace the continuity equation $\frac{1}{\rho} \frac{D\mathbf{u}}{Dt} + \nabla \cdot \mathbf{u} = 0$ (conservation of mass) with

$$\nabla \cdot \mathbf{u} = 0$$

via the assumption that the term $\frac{1}{\rho} \frac{D\mathbf{u}}{Dt}$ is negligible in comparison to the velocity gradients.

Next we consider a reference state where the density is ρ_0 and the pressure is $p_0(z)$ such that $\nabla p_0 = \rho_0 \mathbf{g}$. After dividing by the reference density the momentum equation reduces to

$$\left(1 + \frac{\rho'}{\rho_0}\right) \frac{D\mathbf{u}}{Dt} = -\nabla p' + \rho' \mathbf{g} + \frac{\mu}{\rho_0} \nabla^2 \mathbf{u}.$$

Neglecting the variation in the inertia term due to the density fluctuation (which can be assumed to be small when compared to the background state) and including thermodynamic considerations of energy flux (full derivations in [55, 85]) we arrive at the set of equations:

$$\nabla \cdot \mathbf{u} = 0 \tag{1.5}$$

$$\frac{Du}{Dt} = -\frac{1}{\rho_0} \frac{\partial p}{\partial x} + \nu \nabla^2 u \tag{1.6}$$

$$\frac{Dv}{Dt} = -\frac{1}{\rho_0} \frac{\partial p}{\partial y} + \nu \nabla^2 v \tag{1.7}$$

$$\frac{Dw}{Dt} = -\frac{1}{\rho_0} \frac{\partial p}{\partial z} - \frac{\rho g}{\rho_0} + \nu \nabla^2 w \tag{1.8}$$

$$\frac{DT}{Dt} = \kappa \nabla^2 T \tag{1.9}$$

$$\rho = \rho_0 [1 - \alpha(T - T_0)]. \tag{1.10}$$

This approximation applies if the Mach number is small, so sound waves cannot be considered in this model. Additionally it assumes a perfect gas-type law and Fourier's law of heat conduction (that the heat is proportional to temperature flux).

1.3 Dimensionless Parameters

Before continuing to examine various methods for better understanding KH billow evolution we shall consider the nondimensional parameters that affect our initial setup, as well

as a brief discussion of methods of nondimensionalization for the results themselves.

First we must consider the nondimensional quantities which affect the initial evolution of the instability. We will begin with various definitions of Reynolds number in shear flow. The canonical [25, 55] definition of Reynolds number as a measure of the activity of a flow is

$$Re = \frac{UL}{\nu},$$

where ν is the kinematic viscosity of water (approximately $10^{-6} \text{ m}^2 \text{ s}^{-1}$) and U and L are ‘typical’ velocity and length. The way in which we define these typical values affects the resulting Reynolds number. For our simulations we used U as the maximum absolute value of velocity during the initial billowing event and L as the vertical domain length. It is worth noting that other work to date in shear flows has used several varying definitions. Smyth et al [84] use

$$Re = \frac{\Delta U \delta \rho}{\nu},$$

where ΔU is the maximum velocity difference across the shear layer and $\delta \rho$ is the initial pycnocline width. In contrast Caulfield and Peltier [15] use

$$Re = \frac{\frac{\Delta U}{2} \frac{\Delta H}{2}}{\nu},$$

where ΔU is again the maximum velocity difference and ΔH is the height of the shear layer.

While these three definitions appear quite different in practice they yield very similar estimates for the Reynolds number of a flow (in fact in our simulations all three metrics were within 5% of each other). It should be noted that for our simulations the tilted tube set up means that the ‘typical’ value of velocity is changing with time.

The next relevant parameter to the evolution is the bulk Richardson number. Commonly defined as

$$J = \frac{\Delta \rho g h_0}{\rho_0 \Delta U^2},$$

where h_0 is the vorticity interface thickness, $\Delta \rho$ is the density jump and ρ_0 is the reference density. The bulk Richardson number is often interpreted as setting the size of the billows. However in our set up the presence of the solid boundaries profoundly affects the size of the mature billows. To distinguish it from the gradient Richardson number we label the bulk Richardson number J_b .

The third relevant parameter is the gradient Richardson number, defined as

$$Ri = \frac{N^2 z}{U_z^2}.$$

Confusingly, it is also sometimes labelled J . This parameter is a standard part of the linear stability theory of stratified shear flows and will be discussed in detail in the next section.

The final relevant parameter is the scale ratio $R = \frac{h_0}{\delta_0}$, the ratio between the vorticity interface thickness and the density interface thickness. A key factor related to the scale ratio is the offset between the center of the vorticity profile and the density profile. When the two coincide (as they do in a tilted tube setup) the resulting instability takes the form of KH billows. If the two are not aligned, then we have the potential to create Holmboe waves.

Having completed an overview of the relevant nondimensional parameters we will now briefly discuss how our own simulations can be nondimensionalized. In stratified shear flow time is often nondimensionalized as [77, 76, 76]

$$t^* = t \frac{\Delta U}{h_0}.$$

This definition works well for experimental setups with a steady background shear flow profile. The problem with applying this system to our own simulation results is that our background velocity profile is determined completely by gravitational adjustment. As such the value of ΔU in a tilted tube simulation begins at 0 and grows throughout the experiment. The choice of appropriate ΔU is therefore not an obvious one, which makes nondimensionalization of time unintuitive. As a result, we will present our simulation results in terms of physical time, and the results are thus well geared for the experimentalist.

1.4 Shear Instability of Stratified Fluids

When a fluid undergoes shear, and specifically a density stratified fluid with a relatively sharp pycnocline, there are two primary instabilities which can arise: the Kelvin Helmholtz instability, and Holmboe waves. The shear flow takes the form $U(z) = \tanh(z)$ and the Brunt-Vaisala frequency (sometimes called the buoyancy frequency is the angular frequency at which a vertically displaced parcel will oscillate within a statically stable environment) is given by

$$N^2(z) = J \operatorname{sech}^2(Rz),$$

where R is the ratio of length scales over which the U and N^2 vary, and J is the gradient Richardson number at $z=0$ (the center of the pycnocline) defined as

$$J = \frac{N^2}{\left(\frac{dU}{dz}\right)^2}.$$

Here J acts as a critical value, where for values of $J < J_c$ the flow is more active and we result in Kelvin Helmholtz billows, while $J > J_c$ results in Homlboe waves, a similar but fundamentally different instability which consists of two opposite-propagating disturbances whose energy is concentrated either above or below the center of the shear layer. It is important to note that this critical value depends on R , the ratio between the vorticity and density layer thickenesses. This phenomenon as been studied analytically [38] in numerical simulations [83, 84, 81] as well as in lab experiments [26, 106]. While this instability contains many interesting aspects, we choose here to explore the older and better known instability in a new context.

The instability we have chosen to study is the classical shear instability of a density stratified fluid, often called Kelvin Helmholtz instability. Kelvin-Helmholtz billows are a well studied (see the literature review in [15] for a recent survey) mechanism for transition between laminar and turbulent flows. Since these billows, or waves, are frequently found in real world lake, ocean and atmospheric observations, they are worth studying at multiple scales. We have chosen to begin our investigations with small scale, confined billows.

One important goal of this study is to observe the formation, energy, and longevity of these waves and their interaction with solid walls. To our knowledge no such investigations (of confined Kelvin-Helmholtz billows) have been done, and as such we seek to expand the understanding of this transition to turbulence in the presence of no-slip boundaries. However, we will briefly present the canonical, inviscid stability regime first in order to provide a classical framework of our results.

The basic state that gives rise to KH billows is a shear background current. In two dimensions the background profile is given by $\mathbf{U} = (U(z), 0)$. $U(z)$ is shown in figure 1.1. Shear flow is present at a variety of scales from atmospheric (on the order of kilometres) down to lab scales (on the order of centimetres).

The background profile presented in figure 1.1 can result in many types of instabilities. The core of any turbulent transitional flow is that the initial driving force (in this case shear) first causes large-scale changes in the flow such as the KH billow. As the energy is distributed the features of the flow become finer, resulting in an eventual breakdown of all initial characteristics into finer features via secondary instabilities. When significant energy reaches the smallest scales, diffusion and viscosity become the most important aspect, and




Figure 1.1: This figure shows a typical background velocity profile (black) and background density profile (red dashed). This is the basic profile used both in the Taylor-Goldstein solver, as well as the induced current in the shear current experiments.

at these scales (sometimes referred to as the Kolmogorov length scale) energy is removed via viscous dissipation and the fluid is mixed irreversibly.

1.4.1 The Taylor Goldstein Equation

The fundamental concept of linear stability theory is that for small initial perturbation the equations can be linearized and as such the solutions can be written as a superposition of modes. These modes are assumed to be sinusoidal and written as $Ae^{ik(x-ct)}$, where A is a small initial amplitude, k is the wavenumber, and $c = c_r + ic_i$ is the complex phase speed. This formulation gives us an eigenvalue problem from the equation of motion for the phase speed, c .

Squire's theorem states that in a three-dimensional flow the first disturbances that begin are two-dimensional in nature [55]. By using the background profile from the previous section as well as the Boussinesq approximation we arrive at the following variables:

$$[U + u, 0, w], \quad P + p, \quad \bar{\rho} + \rho.$$

The Boussinesq equations become

$$\begin{aligned} \frac{\partial}{\partial t}(U_i + u_i) + (U_j + u_j) \frac{\partial}{\partial x_j}(U_j + u_j) &= \frac{-g}{\rho_0}(\bar{\rho} + \rho)\delta_{13} - \frac{1}{\rho_0} \frac{\partial}{\partial x_i}(P + p), \\ 0 &= -g \frac{\bar{\rho}}{\rho_0} \delta_{13} - \frac{1}{\rho_0} \frac{\partial P}{\partial x_i}. \end{aligned}$$

The nonlinear terms are only significant when the velocities and acceleration are large compared to the time derivative, and so we can simplify our equations for relatively quiescent flows by eliminating the nonlinear terms to the perturbation equations:

$$\frac{\partial u}{\partial t} + w \frac{\partial U}{\partial z} + U \frac{\partial u}{\partial x} = -\frac{1}{\rho_0} \frac{\partial p}{\partial x}, \quad (1.11)$$

$$\frac{\partial w}{\partial t} + U \frac{\partial u}{\partial x} = -\frac{g\rho}{\rho_0} - \frac{1}{\rho_0} \frac{\partial p}{\partial z}, \quad (1.12)$$

$$\frac{\partial \rho}{\partial t} + U \frac{\partial \rho}{\partial x} - \frac{\rho_0 N^2 w}{g} = 0, \quad (1.13)$$

$$\frac{\partial u}{\partial x} + \frac{\partial w}{\partial z} = 0. \quad (1.14)$$

Here the second-last equation (1.13) comes from the requirement that the density is conserved along the motion, and $N^2 = -\frac{g}{\rho_0} \frac{d\rho}{dz}$ is the buoyancy frequency. The last equation (1.14) comes from the incompressibility condition ($\nabla \cdot (u) = 0$).

This system of equations can be further simplified by adopting the streamfunction defined as

$$u = \frac{\partial \psi}{\partial z}, \quad w = -\frac{\partial \psi}{\partial x}.$$

Furthermore, adopting the ansatz

$$[\rho, p, \psi] = [\hat{\rho}(z), \hat{p}(z), \hat{\psi}(z)] \exp ik(x - ct),$$

produces the set of equations

$$(U - c)\hat{\psi}_z - U_z \hat{\psi} = -\frac{1}{\rho_0} \hat{p} \quad (1.15)$$

$$k^2(U - c)\hat{\psi} = -\frac{g\hat{\rho}}{\rho_0} - \frac{1}{\rho_0} \hat{p}_z \quad (1.16)$$

$$(U - c)\hat{\rho} + \frac{\rho_0 N^2}{g} \hat{\psi} = 0. \quad (1.17)$$

Finally simplifying this set of equations to remove pressure and density we arrive at the single equation

$$(U - c) \left(\frac{d^2}{dt^2} - k^2 \right) \hat{\psi} - \frac{\partial^2 U}{\partial z^2} \hat{\psi} + \frac{N^2}{U - c} \hat{\psi} = 0. \quad (1.18)$$

Which is known as the Taylor-Goldstein (henceforth T-G) equation. Equation 1.18 describes the behaviour of perturbations in a stratified shear flow. It is important to note that the eigenvalue of this equation, c is complex so that any nonzero imaginary part of c indicates the presence of an unstably growing mode.

1.4.2 Solutions of the T-G Equation for Linear Instability

Solving the Taylor-Goldstein equation for a continuous profile of shear velocity and density is best done numerically (see Hazel [41] for one of the earliest examples). However, by introducing a new variable ($\phi = (U - c)^{-\frac{1}{2}} \hat{\psi}$), multiplying the new equation by the complex conjugate of ϕ and integrating in the vertical direction with no-flux boundary conditions we arrive at the requirement

$$c_i \int \frac{N^2 - \frac{1}{4} U_z^2}{|U - c|^2} |\phi|^2 dz = -c_i \int [|\phi_z|^2 + k^2 |\phi|^2] dz.$$

Since the right hand side of this equation is strictly positive, this relation can only be satisfied for a nonzero c_i if $\frac{N^2}{U_z^2} < \frac{1}{4}$. This combination of variables is known as the Richardson number as $Ri = \frac{N^2}{U_z^2}$, and thus we arrive at the restriction that the flow is always linearly stable if $Ri > \frac{1}{4}$.

In practice it is worth noting that Richardson number below one quarter is a necessary but not sufficient condition for instability. That is to say that a flow with a Richardson number below the critical value is not guaranteed to be unstable, but a flow with a Richardson number above the critical value **is** guaranteed to be linearly stable. We can solve the Taylor-Goldstein equation by specifying a profile (density, velocity, etc.) and solving the eigenvalue problem for the phase speed c by specifying the wavelength, k . Interestingly while equation 1.18 is most often used for analytical results it is actually equations 1.15 - 1.17 which are easiest to solve numerically since the single equation is actually a quadratic eigenvalue problem, while the separate equations can be rewritten as a standard eigenvalue problem.

Solving the eigenvalue problem numerically we can determine the growth rates for a specific profile given some initial shear. By solving such an eigenvalue problem for a range of wavenumbers, $k = \frac{2\pi}{\lambda}$ where λ is wavelength, we can identify the wavelength for the perturbation which has the greatest growthrate. Figure 1.2 shows the growth rates for a profile with the pycnocline located in the center (red dashed) and 75 percent down (black solid). In both cases the height of the domain is 1, the width of the pycnocline is five percent of the total depth and the average velocity is 0.1. Notice that the off-center pycnocline exhibits faster growth at small wavenumbers. This is a feature which will appear in two and three-dimensional simulations of shear flow. Figure 1.3 shows the real (solid) and imaginary (dashed) parts of a sample eigenfunction at wavelength $\lambda = 1$.




Figure 1.2: Growth rate versus wavenumber for a centered (red dashed line) and off-center at 75 % of depth (black solid line). Both setups show largest growth at low wavenumbers, indicating that it is the relatively large-scale instabilities which form first. The off-center pycnocline exhibits slightly faster growth at low wavenumbers.

Hazel’s paper “Numerical studies of the stability of inviscid stratified shear flows” in the Journal of Fluid Mechanics [41] is one of the best early examples of numerically solved Taylor-Goldstein equations for an inviscid shear flow in two dimensions. The paper considers cases where the Richardson number is strictly positive. The eigenvalue problem is then written in terms of four parameters and solved using two numerical methods depending on whether Hazel is solving for the phase speed, c , or the profile shape. The shooting method is used for solving the phase speed, c . The paper examines both shear flows and jets. Hazel then computes the growth rates and curves of maximum growth for various profiles and

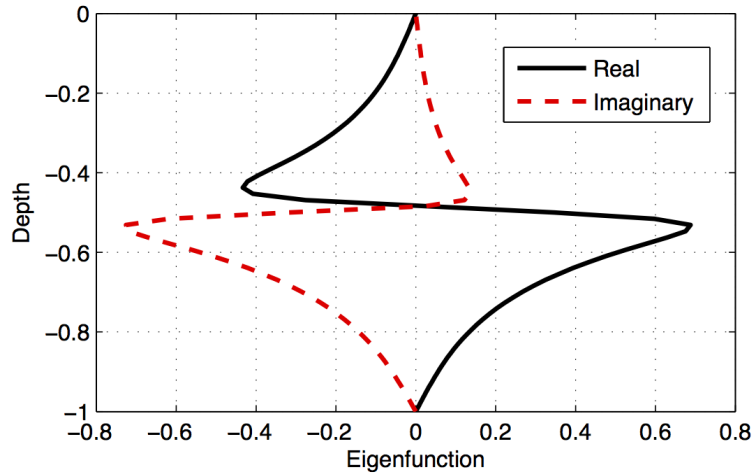


Figure 1.3: A sample eigenfunction at wavenumber $k = \pi$. The solid line represents the real part of the eigenfunction while the dashed part represents the imaginary part.

initial conditions. These results are compared with experimental data for various shear flows, and this shows good agreement. Hazel does briefly discuss the role of confinement and the effect on boundaries on the growth rate of instabilities. However, the parameters used in his paper do not correspond to physical flows, making our investigations unique in that respect.

1.4.3 Secondary Instabilities and Transition

In a review of the instability of inviscid flows [73] Pierrehumbert concludes that two-dimensional vortices occurring in an incompressible inviscid flow are unstable to two-dimensional perturbations. This review shows that short-wave instability is compatible with the spectral cascade. The energy transfer into a given scale from the set of larger scales is influenced by the nearby length scales. Thus the short-wave instability is a mechanism of cascade.

Klaassen and Peltier discussed the onset of turbulence in KH billows in a 1985 JFM paper [50]. They conclude that the collapse of KH billows involves a transition from two-dimensional to a three-dimensional flow. By assuming a mode decomposition of the flow instability and building a matrix operator Klaassen and Peltier were able to perform an eigenvalue mode analysis to show that for a stratified fluid the unstable modes in a KH

billow arise not in the braids, but in the cores, where the overturning is strongest.

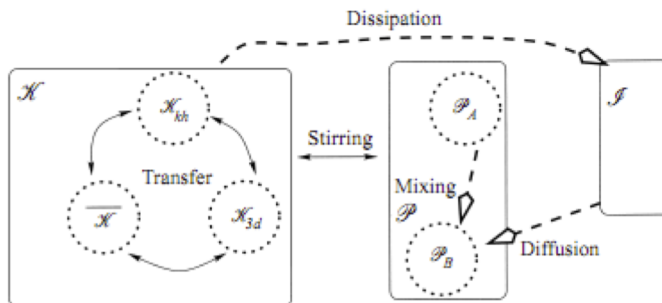


Figure 1.4: This schematic (from [15]) shows the energy budget of a shear flow between kinetic \mathcal{K} , potential \mathcal{P} , and internal energy \mathcal{I} . Reversible changes are shown as solid arrows while irreversible changes (such as dissipation) are shown as dashed lines.

An important aspect of understanding turbulent transitional flows is the classification of various mechanisms of energy exchange. The energy budget (figure 1.4 from Caulfield and Peltier [15]) shows how kinetic energy, potential energy, and internal energy of the flow. Dissipation converts kinetic energy to internal energy, which in turn (through dissipation) changes the potential energy state of the system. Similarly stirring can convert between kinetic and potential energy, although unlike dissipation this process is reversible.

With the advent of more powerful computing capability (both hardware and software) fluid dynamics is able to move away from two-dimensional simulations and towards fully-realized three-dimensional simulations of real flows. A key result in three dimensional simulations is the development, evolution, and structure of spanwise (or secondary) instabilities.

Squire's theorem holds well for shear flows in Newtonian fluids, so the initial instability response is almost entirely two-dimensional, however, as the flow develops and becomes more turbulent and complex three-dimensional effects begin to form. These effects possess their own unique characteristics worth exploring in greater detail. Many papers (Orszag and Patera, 1983 [71], Herbert, 1983 [42], Metcalfe et al, 1987 [63], Bayly and Orszag, 1988 [3], Butler and Farrel, 1992 [9]) have considered the effects of three-dimensional perturbations on various types of shear flows (such as plane and pipe Poiseuille flow, plane Couette flow and flat-plate boundary layers). It is often the case that three-dimensional waves transfer energy from the mean flow to three-dimensional instabilities but are not directly responsible for the instability, i.e. the waves feed the perturbations which grow,




Figure 1.5: This schematic from [15] shows the formation of a primary instability, in this case a Kelvin-Helmholtz billow. Once the billow reaches small enough scales secondary instabilities begin to form in the spanwise direction, indicated in the diagram.

while the three-dimensional waves do not [71]. Small three-dimensional (spanwise to the primary instability) perturbations which are considered stable by linear theory have been shown to extract energy from the mean flow and grow on a convective timescale [63]. This growth can lead to almost isotropic turbulent states.

In a later study Klaassen and Peltier, in a JFM paper [51], analyzed the stability of horizontally periodic, two-dimensional finite amplitude KH billows with respect to small three-dimensional perturbations. An energy budget established that perturbation growth was relatively insensitive to stratification strength, which suggests that the basic underlying mechanism is dynamical shear instability. This paper concludes that billow core overturning acts as a site for secondary instabilities with the maximum energy.

The paper by Caulfield and Peltier (“The anatomy of the mixing transition in homogeneous and stratified free shear layers” [15]) provides a more recent review of the mixing transition through which turbulence develops in shear and homogeneous flows. This paper discusses in depth the formation and structure of KH billows in three dimensions. By performing three-dimensional direct numerical simulations (DNS) the authors are able to test predictions of secondary stability analysis. The first idea tested that in a homogeneous fluid the vorticity braids (regions connecting billow cores) are the source of instabilities. The second idea is that in a strongly stratified fluid it is convective destabilization of the unstable sublayers (i.e. the cores themselves) which is the main mechanism of instability.

In regards to general wave steepening and breaking for internal waves in a stratified shear flow, the Staquet and Sommeria [86] review provides an excellent discussion of instability. Staquet presents a review of a relatively complete picture from the steepening of unstable plane waves and reflecting waves on a sloping boundary to the breakdown to small-scale turbulence. The review concludes that it is not clear whether the classical k^{-3} spectrum from Garrett and Munk [34] for the energy dissipation rate of waves can be explained by weakly interacting waves alone. Staquet also points out that classical grid turbulence collapses in a stratified medium and is “always quite different from the ‘wave turbulence’ sustained by internal gravity waves” [86]

Mixing efficiency is defined as the ratio of irreversible increases in potential energy of the density distribution to the loss of kinetic energy. Peltier and Caulfield, in a review letter [72] consider recent work on the transitions in stratified shear flows which examine the mixing efficiency of three-dimensional flows. The review concludes that for KH-type instabilities the process begins with a two-dimensional inviscid-like KH billows which, as the billows saturate, form secondary instabilities. These secondary instabilities are convective and focused in regions of overturning. The efficiency of mixing varies dramatically throughout the evolution of these instabilities, starting with a relatively high mixing efficiency (0.7) and declining until finally arriving at the classical estimate (on the order of 0.2). The interaction of turbulent mixing with internal gravity waves leads to both temporal and spatial variability. The review concludes that internal waves provide a mechanism by which energy and momentum can be transported away from generation sites and as such understanding the interplay between the onset of eddy-like instabilities and internal wave dynamics is an important part of the energy and momentum budgets of turbulent flows.

Smyth et al. in a 2005 paper [82] numerically compares the turbulent diffusivities of heat and salt during the formation and destabilization of KH billows. The results, closely following Klaassen and Peltier’s 1991 paper on secondary instabilities [51], determine that the two are roughly equivalent for Reynolds numbers greater than $O(100)$. Smyth et al. compare ocean, lab, and numerical simulations and find that the ratio between the cumulative turbulent diffusivity of salt to that of temperature approaches unity with increasing Reynolds number.

A recent review in a series of papers by Mashayek and Peltier [60, 61] considers the various secondary instabilities known to co-exist in a stratified shear flow which undergoes Kelvin-Helmholtz instability, specifically focussing on whether it is the braid or the core which exhibits primary destabilization. Mashayek and Peltier identify that for weakly stratified domains or low Reynolds numbers it is the braid which is most susceptible to secondary convective instabilities, while strongly stratified high Reynolds number flows maintain the importance of the destabilization in the billow cores. Mashayek and Peltier

also discuss billow pairing - a process whereby billows combine and destabilize and occurs only for high enough shear regimes.

Our own simulations to date have supported these results. The spanwise instabilities form after the initial instability, namely the billows, form and even mature. As the scale of the two-dimensional instabilities decreases the three-dimensional effects become more important until they two instabilities become comparable. Our simulations also consider the effects of confined (or trapped) billows near no-slip boundaries. The effects of close boundaries on three-dimensionalization is one of the key goals of this project.

1.5 Turbulent Metrics

With the rise of computing power and the growing availability and reliability of numerical software ranging from direct numerical simulations, to large eddy simulations, to $k - \varepsilon$ models, there are more tools available now to help study the evolution of the rise and development of turbulent flows. Added to this the refinement of tools such as laser Doppler velocimetry and particle tracking velocimetry has resulted in myriad ways to analyze a flow setup (whether numerical or experimental) with access to the full velocity and acceleration fields.

With access to the full velocity field we now concern ourselves with the best method of visualizing a flow. To this end we will consider one such metric, namely the invariants of the velocity gradient tensor, $A = \nabla \mathbf{u}$.

1.5.1 Primary, Secondary, and Tertiary Variables

The first fact to note, is that there are several types of variables which provide information about the flow. Primary variables (such as $\mathbf{u} = (u, v, w)$, ρ , T , or p) are direct descriptions of physical flow variables. However at times, these variables are of limited usefulness in analyzing characteristics of the flow, especially for vorticity-dominated and turbulent flows. Due to the complex nature of such flows it would take several primary variables in tandem to coherently describe the relevant dynamics.

Secondary variables are manipulations of the primary variables, such as vorticity ($\vec{\omega} = \nabla \times \mathbf{u}$), or the rate of strain ($e_{ij} = \frac{1}{2}(u_{i,j} + u_{j,i})$). These variables describe particular physical characteristics of the flow better than primary variables. For example the double contraction $2\mu e_{ij}e_{ij}$ is a scalar that demonstrates the amount of viscous dissipation in the

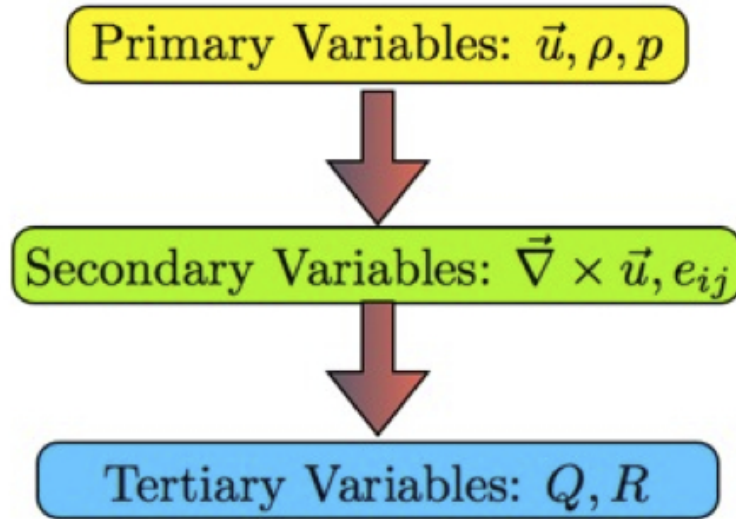


Figure 1.6: A schematic demonstrating the hierarchy of flow variables with primary variables such as velocity and density informing secondary variables such as vorticity and dissipation, which in turn can be combined to create tertiary diagnostic variables, such as Q , R , or the Okubo-Weiss parameter.

flow, or the irreversible conversion of mechanical energy to internal energy. While the quadratic, positive semi-definite quantity, enstrophy $\frac{1}{2}\omega_i\omega_i$, is the analog of kinetic energy for the vorticity field.

In contrast, tertiary variables are variables derived from combinations of primary and secondary variables. Q and R are variables (from the turbulence literature) defined as

$$\begin{aligned}
 P &= \lambda_1 + \lambda_2 + \lambda_3 \\
 Q &= \frac{1}{2}(\lambda_1^2 + \lambda_2^2 + \lambda_3^2) \\
 R &= \frac{1}{3}(\lambda_1\lambda_2\lambda_3)
 \end{aligned}$$

Here λ_i are the eigenvalues of the Velocity Gradient, $A_{ij} = \frac{\partial u_i}{\partial x_j}$ - sometimes called the velocity gradient tensor (even though it does not satisfy the tensor transformation properties). P , Q , and R , are the first, second, and third invariants of this quantity, respectively. This measure was first developed by in a seminal paper by Chong and Perry

[16] as a means of classifying three-dimensional flow-fields. The idea was further refined (and articulated) by Ooi et al. in a JFM review [70]

The values of Q and R represent a balance between vorticity and shear, and axial and bi-axial dominance in the flow, respectively. Q large and positive implies regions of intense enstrophy, while Q large and negative implies regions of strong strain rate. R relies on Q but gives a general measure of axial (large positive) versus planar (large negative) stretching. See figure 1.7 for a visualization of the types of flow represented here. Note that Q and R are both strictly real variables although the eigenvalues are complex (this is a consequence of some algebraic restrictions on the tertiary variables). Details and variations of these variables are available in Davidsons turbulence book [25], which references the landmark review paper by Ooi et al. [70] in which these variables are used for the analysis of forced homogeneous isotropic turbulence [70].




Figure 1.7: Schematic showing the dominant local flow topologies at various $Q - R$ regimes, reproduced from Ooi et al, [70], figure 1. The quadrants represent flow topologies (stable-focus/stretching (SF/S), unstable-focus/ contracting (UF/C), stable-node/saddle/saddle (SN/S/S) and unstable-node/saddle/ saddle (UN/S/S)) that can exist in incompressible flows. The line separating vortex from shear-dominated flow corresponds to $D_A = 0$.

1.5.2 Q vs R Plots

By considering these “tertiary” variables we can characterize the flow into regions of local flow topologies (see figure 1.7). The curve separating these regions is given by the relation $D_A = 0$, where D_A is the discriminant of the velocity gradient matrix and is given by

$$D_A = \frac{27}{4}R^4 + Q^3$$

Mathematically this curve is the division between the areas where the eigenvalues of A_{ij} are real or complex (in the region below the curve the velocity gradient matrix has strictly real eigenvalues). These diagnostic tools (See [70], for details and variations) are commonly used in analyzing direct numerical simulations of homogeneous isotropic turbulence. They are also widely used by the forced turbulence community (see references in Ooi et al. for uses in numerical methods and Luthi et al. [59] for uses in experimental results). Davidson in his turbulence text [25] derives these methods as a general tool, and as such we feel they will prove useful in characterizing transitional flows, though we do note that their most widely accepted use is as diagnostic for fairly low Reynolds number, homogeneous, forced isotropic turbulence.

Experimentally tertiary variables appear, at first glance, to be of limited use, as they require the full velocity gradient tensor (and the eigenvalues thereof). Consequently the majority of work with these tensor invariants involve numerical methods (sepecifically DNS) in some way.

Q and R have been used fairly widely in conjunction with DNS. One of the first notable examples of invariant analysis is Yeung and Pope in 1989 [102] performed a study of Lagrangian statistics of velocity, acceleration, and other quantities in isotropic turbulence using high-resolution (at the time) 128^3 DNS simulations. Using invariants of the velocity gradient tensor (though not precisely Q and R) they derived the classical tear-drop profile for characterizing isotropic turbulence. More recent work by da Silva and Pereira in a 2008 paper [22] analyzed the turbulent/nonturbulent interface in jets streams. The use of Q and R (as well as the invariants of the rate of rotation tensor, allowed them to conclude that the primary mechanism for almost all points at the interface is irrotational dissipation. This paper also noted a lack of coherent structures at the interface, something $Q - R$ analysis can be used to identify. Pirozzolia and Grasso [74] analyzed statistical properties of compressible isotropic turbulence and used Q and R in order to identify potential similarities with incompressible turbulence. They determined that Q vs R has a universal structure between compressible and incompressible turbulence. For a recent overview Meneveau has an excellent review in Annual Review Fluid Mechanics [62] of

several direct numerical simulations of turbulent flows using Q and R demonstrating how these variables can be used to highlight the fine-scale structure of turbulence.

While there is no theoretical reason why the tensor invariants cannot be applied to LES-generated flow fields it appears that there has been almost no work done on the subject. We will address this issue later (extensions chapter) in order to help understand why this technique is not more widely applicable.

While it is more difficult to apply $Q - R$ metrics to experimental results improvements in measurement techniques and laboratory-scale experiments in recent years have made the velocity gradient tensor accessible under certain conditions. A hallmark study by Luthi et al. [59] used three-dimensional particle tracking velocimetry (PTV) to measure homogeneous isotropic turbulence using neutrally buoyant particles, and demonstrated the utility of Q and R in an experimental setting. This study is on a similar (lab) scale as our simulations - specifically a cube 10 cm to a side. Luthi et al. were able to apply these turbulent metrics to demonstrate the validity of the classical “ice-cream cone” (sometimes called teardrop) profile for isotropic turbulence. Other experiments using PTV to access the full velocity field in isotropic turbulence such as Worth et al. [99] who use tomographic PTV to compare isotropic turbulent simulations to experimental results using (in part) Q and R metrics; and Guala et al. [37] who use QR analysis to investigate Lagrangian correlations in small-scale turbulence.

The use of Q and R metrics in the analysis of turbulence is well established, as we have seen. Armed with an understanding of the origins of these turbulent metrics our investigations will consider the applicability of Q and R (and specifically various types of Q vs R plots) to active, but not fully turbulent simulations, namely the evolution and destabilization of KH billows.

1.6 Overview

The first objective of this thesis is to use direct numerical simulations (DNS) to explore the behaviour - namely the formation, evolution, and eventual destabilization - of Kelvin-Helmholtz billows in a confined domain. Our goal is to determine how the presence of boundaries and the confinement caused by a small vertical domain will impact the evolution of the KH instability: namely we expect that the proximity of no-slip boundaries will lead to significantly higher dissipation and thus may alter the way in which KH billows grow and destabilize. We also investigate how pycnocline placement (whether at the center, or at 25% of the domain depth) alters the evolution of the billows and the development of secondary

instabilities therein. We begin our investigations in two dimensions in order to observe the initial rollup and determine what parameter regime (within those we have selected to vary, namely domain size and pycnocline placement) are worth further consideration. We then move on to fully three-dimensional simulations.

The next section of the thesis deals with the turbulent metrics Q and R as they apply to complex, near-turbulent flows. We first apply these metrics to some canonical flows in order to determine the applicability of these metrics and justify their use in more complex domains. Next we consider the visualization strategy for complex, three-dimensional flows such as those presented in the Kelvin-Helmholtz instability chapter. We apply the turbulent metrics Q and R to this problem in order to facilitate the process of identifying coherent structures and characterizing the flow topology at a given time. Finally we apply these metrics to a totally different problem in order to determine the applicability of our methodology to a problem which we knew almost nothing about *a priori*.

Finally we present some simple extensions to the ideas put forth in the first two chapters. First we present a confirmation of the importance of our choice to simulate the confined domain with no-slip boundaries. Next we discuss the application of Q and R to other types of numerical methods, namely large eddy simulations (LES). And finally we outline the usefulness of Q and R as a visualization strategy by comparing it to some established techniques such as the Visit toolbox, Helicity, and the λ_2 parameter.

Chapter 2

Kelvin Helmholtz Billows: the classical picture

Kelvin Helmholtz (henceforth KH) billows form when a fluid undergoes some form of shear, for example in a closed tube filled with stably stratified fluid that is suddenly tilted to a fixed angle from the horizontal. In a continuously stratified fluid, billows and internal waves can coexist and terms such as breaking and overturning are often used interchangeably. With a fixed stratification, the greater the shear, the more rapid and intense the formation of the billows. KH billows occur over a wide range of scales, from flow over small scale topography, to flow due to internal waves in the ocean, in lakes or the atmosphere ([4, 8, 26, 96, 95, 49, 50]). As was demonstrated in the literature overview, KH billows are a well-studied problem. The uniqueness of our investigation arises from the very small scale of the (numerical) experiments being performed and the presence of no-slip boundaries, which drastically affect the formation and evolution of KH billows.

While stratified shear instability, often in the form of Kelvin-Helmholtz billows, is known to occur in almost any density stratified fluid which undergoes shear, the majority of research on the subject has been done by those interested in the solution in the absence of boundaries. Whether motivated by atmospheric and oceanic phenomena [8, 95, 96, 50, 72] or numerical computation of the transition to turbulence [14, 15, 32], the result is that there is almost no work discussing the effects of confinement on the formation and evolution of KH billows in the fully nonlinear regime, one notable exception being Hazel in his 1972 paper [41] who discusses inviscid walls and their effect on stability. Yet, in physical laboratory experiments, typified by the seminal work on the formation of KH billows by Thorpe [92] and his tilted tube experiments, the experimental setup involves a stratified fluid in a long, rectangular tube. The tube is then tilted to some angle (45 degrees for our simulations)

and gravitational adjustment induces shear across the pycnocline, all in the presence of boundaries at which no slip boundary conditions hold.

The classical evolution of KH billows (in the absence of boundaries) follows a well-described path [15, 49, 72]: an initial disturbance forms across the shear region, with an along tube length scale given by the fastest growing mode of linear instability theory. This mode grows in magnitude entirely in two dimensions, forming discrete billows connected by thin regions of high shear, known as braids. Following the initial billow formation, the flow begins to destabilize in the spanwise direction, collapsing into turbulence and becoming well mixed. Recent extensions of this classical regime into high-Reynolds number flow by Mashayek et al. [60, 61] highlight new mechanisms for mixing and the transition to turbulence. In this Chapter we consider the qualitatively and quantitatively different regime of instability in a confined, small-scale setting with a typical Reynolds number of around 100.

As mentioned above, the classical evolution of KH billows (in the absence of boundaries) follows a well-described path. Once formed, the billow can broadly be broken up into two categories: the cores, which form the majority of the fluid that has ‘rolled up’, and braids, which are thin regions of fluid connecting cores together. While the billow cores contain the majority of mixing and overturning, and are thus most interesting from a dynamical standpoint, it is important to remember that the braids exhibit the largest shear and thus have potential for the largest dissipation in the flow [60, 61]. The braid regions produce intense regions of dissipation which are prone to secondary braid instabilities as well as disruptions due to subharmonic pairing - an aspect of confined billows that we will explore in the relevant simulations.

We expect the confinement of this shear flow and the presence of no-slip boundaries to change the way that the billows evolve. Indeed, with a small enough gap between the two walls, we expect laminarization and a complete absence of instability. However, for a slightly larger gap the specific mechanism of interaction is not as simple to predict. This chapter conducts a parameter space exploration with several goals, a notable one being to determine relative timing of the onset of the two-dimensional instability, strong interaction with the walls, and three-dimensionalization. We expect that in some cases the dissipation and mixing will be dominated by the early two-dimensional stage of the instability. The strong interaction of billows with the walls will be referred to as ‘grinding’ in the following text.

2.1 Characterizing the Flow

In our analysis of the evolution of confined shear flows there are several important characteristics of the flow we wish to capture, and these are best discussed in terms of different variables. The first variable of interest is naturally the density profile. For the two-dimensional simulations we present the density perturbation (i.e. the ρ' part of $\rho_0 + \rho'$ in the Boussinesq equations). For three-dimensional simulations, we present isosurfaces of $\pm 10\%$ of $\max \rho'$. This gives an accurate representation of the deformation of the pycnocline (or linear stratification) as well as allowing us to visualize the areas of interest, namely the cores and braids of the KH billows.

A secondary characteristic of the flow that we are interested in is the viscous dissipation. Following standard practice (see [55] for a detailed derivation and justification), we define the viscous dissipation in the flow as

$$\varepsilon = 2\mu e_{ij}e_{ij}.$$

This is a valuable metric because it allows us to visualize the locations of most significant dissipation, for example those related to the no-slip boundaries, and contrast them with the billow cores and the braids. It also allows us to numerically calculate the total dissipation induced by the various initial conditions (via the integral of ε over the entire domain). In order to meaningfully observe and examine the qualitative structure, the dissipation figures in this chapter are scaled by one quarter of the maximum dissipation for the current output; on top of this plot we overlay a single contour of density (in white) in order to contrast the locations of highest dissipation with the structure of the billows.

Another fundamental aspect of the flow we wish to consider is the amount of mixing that takes place in the simulation. A simple way to calculate how much possible mixing will occur (i.e. when the most amount of mixing will take place) is to calculate the gradient of the density: $\nabla\rho$. In order to calculate the total amount of mixing at any given time we can make this into a positive definite quantity, and sum over the domain giving us:

$$M = \int_{Domain} |\nabla\rho|^2 dV.$$

This total quantity when summed is only loosely related to potential energy. When domain-averaged it gives a measure of how sharp the gradients of density are and thus how rapidly diffusion will take effect. However, it does not locally represent the amount of mixing that will occur at a given point and as such the only way we will consider this measure is in the domain-integrated sense.

Since three-dimensionalization is a fundamental aspect of the flow which we cannot overlook, we can also consider (for fully three-dimensional simulations) the total and spanwise enstrophy. Defining (local) enstrophy as

$$\mathcal{E} = \frac{1}{2}(\omega_x^2 + \omega_y^2 + \omega_z^2),$$

where $\omega = (\omega_x, \omega_y, \omega_z)$ is the vorticity defined as

$$\begin{aligned}\omega_x &= \frac{\partial v}{\partial z} - \frac{\partial w}{\partial y} \\ \omega_y &= \frac{\partial w}{\partial x} - \frac{\partial u}{\partial z} \\ \omega_z &= \frac{\partial u}{\partial y} - \frac{\partial v}{\partial x}.\end{aligned}$$

This allows us to capture the total amount of vorticity present in the simulation. And if we consider a new variable

$$\mathcal{E}_{3D} = \frac{1}{2}(\omega_x^2 + \omega_z^2),$$

we can contrast some measure of spanwise vorticity with total vorticity. Note that both of the above metrics are positive definite which allow for easy visualizations of intensity.

Further analysis of the evolution of the flow can be performed using turbulent metrics, Q and R ; which are discussed in detail, both in terms of theory and as case studies, in the $Q - R$ chapter which follows the present one.

2.2 Numerical Methods

All simulations have been performed using the scalable pseudo spectral solver SPINS [88], which solves the incompressible, stratified Navier-Stokes equations. This code uses the Boussinesq approximation, and so is only capable of handling initial conditions with relatively small density fluctuations. Since our simulations require a density difference of less than five percent, this restriction does not pose a problem.

The code uses pseudospectral methods, implemented in C++ using the open source Blitz library for array handling. Parallelization is handled using MPI, with the spectral and pseudo spectral transforms parallelized to ensure scalability. In particular, since our simulations will deal with a regularly spaced grid in the along channel and spanwise directions, thus we can use the FFT (through the FFTW package) and achieve both high

Table 2.1: Relevant nondimensional simulation parameters.

| | |
|----------------------|-----|
| Reynolds Number (Re) | 100 |
| Schmidt Number (Sc) | 10 |
| Prandtl Number (Pr) | 10 |

accuracy, and scalability with the number of processors. A Chebyshev grid is employed in the vertical direction, allowing for a high concentration of points around the boundaries and thus high accuracy in the boundary layer. This is especially important as we wish to see how confined billows differ from the standard evolution paradigm and thus will be paying close attention to the boundaries.

The relevant nondimensional parameters of the simulation are given in table 2.1. The typical velocity chosen for the calculation of these numbers is the mean downstream velocity. The molecular and viscous diffusion rates are set to physically relevant values for a salt or temperature stratified fluid. Note that typical Reynolds Number values are within classically prescribed laminar regimes, however the flow exhibits typical turbulent features, such as mixing and an energy cascade to small scales.

2.3 Two-Dimensional Simulations

Simulations presented herein can be divided into three categories, those with centered pycnoclines, those with off-center pycnoclines, and linearly stratified flow. In the non-centered (asymmetric) pycnocline simulations the pycnocline center is placed at 25 percent of the total depth from the top boundary. The goal is to observe the evolution of KH billows in the presence of these boundaries both with intact and broken symmetry and to study the effects of no-slip walls on the dissipation, mixing, and evolution of the billows themselves.

The system is initialized by a stably stratified fluid with a total density jump of five percent and a pycnocline width of 10 percent of total depth. The tube is then tilted 45 degrees and allowed to equilibrate. The simulation is done on a periodic grid in the x (length) direction, which is equivalent to observing the center portion of a long capped tube (see Thorpes benchmark paper [92] on the observation of KH billows in a tilted tube for experimental justification of this study).




Figure 2.1: A schematic diagram showing the experimental setup in all simulations in this chapter. The lighter fluid rises to the top while the heavier fluid travels downward, creating shear across the density gradient.

Since the initial onset of the instability is two-dimensional (as predicted by Squires theorem and verified by three-dimensional simulations) the initial investigations and parameter space explorations were performed in two dimensions. The most promising results were then repeated in full 3D DNS and analyzed further using turbulent metrics - namely Q and R . The horizontal domain size was kept constant at 10 cm while the vertical domain was either 1, 2, or 0.75 cm tall. The initial investigations were run with 8 processors on a 768 by 128 (or 192 for the 2 cm case) grid in x and z respectively (see table 2.2). The Prandtl numbers explored vary between 1, 2, and 10, but were kept at 10 for the presented results.

Table 2.2: Two-Dimensional simulations and their numerical parameters.

| Simulation | Domain Size | | Pycnocline Depth | Grid Size | |
|----------------|-----------------------|---------------|------------------|------------------|--------------------|
| | Max Grid Length | 1cm symmetric | | 0.5cm | 768×128 |
| 1cm skewed | 10 cm \times 1cm | | 0.75cm | 768×128 | 1.24×10^6 |
| 2 cm symmetric | 10 cm \times 2cm | | 1cm | 768×192 | 1.42×10^6 |
| 2 cm skewed | 10 cm \times 2cm | | 1.5cm | 768×192 | 1.42×10^6 |
| 0.75 cm skewed | 10 cm \times 0.75cm | | 1.5cm | 768×128 | 9.28×10^5 |




Figure 2.2: Vertical profiles of density (left) and horizontal velocity (right) of the symmetric 1cm simulation at an early time - before billow evolution has begun.

2.3.1 1cm Simulations

The first set of simulations we discuss here is the symmetric 10 cm by 1cm domain. Figures 2.3 and 2.4 show the evolution of the density field at four output times corresponding to various stages in the development and evolution of the billows. The evolution of centered-

pycnocline billows can be classified into distinct stages. The first is the formation and onset of the instability. This stage evolves exactly as predicted by classical and deep-domain theories ([15, 72]). The pycnocline deforms first in a sinusoidal shape and then forms the classical cores and braids described in the KH instability literature. This stage of the simulation is not impacted by the presence of boundaries and as such we have chosen only to present a snapshot demonstrating that the formation of our confined billows follows the classically prescribed path.

The next stage of the simulation concerns the growth of billows to their mature (domain-dependent) size. At 1.4 seconds, figure 2.4 shows the mature billows which have reached the size of the domain. At this point it is clear that although the system is still extremely active the billows cannot grow any further without either interacting with the boundaries or each other. It is at this point that we will want to consider the effects of confinement on billow-billow interaction, as well as the induced dissipation ($2\mu e_{ij}e_{ij}$).

The bottom panel of figure 2.4 shows the details of the interaction between the mature billows and the no-slip boundaries. Note that at 1.5 seconds the deformation of the billow shape near the boundaries as the effects of the walls begin to manifest. The billows maintain their basic shape - that of cores connected by braids - however there is clear deformation near the boundaries, which has caused secondary turnover in the regions between cores. By 1.6 seconds this boundary-induced effect has created a region which is likely very active and we expect can be classified as transitioning to a fully turbulent flow - a feature which our two-dimensional simulations are not sufficient to explore.

At this point in the simulation it is clear that if 3D effects (spanwise instabilities) exist in this regime they must take effect during or after this time, and so no further progress can be made in two dimensions.

Now that we have formed a coherent picture of the evolution of confined billows we can explore the effects that dissipation and mixing play in the evolution of this system. For the same simulation we consider the locations and magnitude (qualitative and quantitative) behaviour of dissipation in figures 2.5 and 2.6. All values of dissipation presented herein are scaled by some reference value - determined on a per-simulation basis as 50 percent of max dissipation achieved over the course of the simulation.

The initial phase of the numerical experiment shows that the viscosity-induced broadening of the shear flow across the pycnocline generates a significant amount of dissipation, essentially due to the large relative velocities of the fluid. Looking further ahead in time we notice a distinct pattern emerging: regions of highest dissipation are generated near the boundaries due to the presence of no-slip walls, and as the flow becomes more active progresses further into the center of the domain, entrained by the velocities induced by the




Figure 2.3: Formation and evolution of Kelvin Helmholtz billows in a stratified shear flow between two (infinitely long in x and y) plates at 1.2 and 1.3 seconds. The initial evolution of the billows follows the classically prescribed route. The initial sinusoidal disturbance (top) grows until overturning occurs and billows form.

billow-boundary interaction.

The initial locations of induced dissipation are (as expected) strongest where the nearly horizontal billow velocities interact with the no-slip boundary. Secondary locations of noteworthy dissipation include the braids connecting billows, as well as the centers of




Figure 2.4: Evolution of Kelvin Helmholtz billows in a stratified shear flow between two (infinitely long) plates at 1.4 and 1.5 seconds. As the billows grow they begin to interact with and become deformed by the boundaries and each other. Notice the regions of significant overturning in the billow cores.

the billows (to a lesser extent). It is worth noting that these secondary locations are significantly weaker than the near-wall locations.

At later times (figure 2.6) we can clearly see that despite a very active center of the domain, the majority of dissipation occurs only very near the boundaries. It is also worth




Figure 2.5: Dissipation ($2\mu e_{ij}e_{ij}$ scaled by 50% of max dissipation in simulation) of KH billows in a stratified shear flow at 1.2 and 1.3 seconds with a contour of density overlayed in white. Note that the initial dissipation occurs near the center of the domain, where the shear is strongest, but that for fully evolved billows the majority of grinding occurs near the boundaries and is entrained into the center of the domain and the billow cores, once again following the areas of highest relative shear.

noting that the total dissipation is significantly lower than the earlier times.




Figure 2.6: Dissipation ($2\mu e_{ij}e_{ij}$ scaled by 50% of max dissipation in simulation) of KH billows in a stratified shear flow at 1.4 and 1.5 seconds with a contour of density overlaid in white. At these later simulation times the dissipation is focussed almost exclusively near the boundaries.

2.3.2 1cm skewed case

Since it is clear that boundaries have a profound effect upon both the evolution of billows and the dissipation induced by them we now consider the effects of breaking symmetry.

This section presents the results of an asymmetric (skewed) simulation with the pycnocline placed at one quarter of the total depth. Figure 2.7 shows the skewed 10 by 1cm domain at 1.5 and 1.6 seconds (top and bottom panels respectively). It is worth noting that due to the symmetric nature of the system placing the pycnocline 25% from the top or the bottom of the domain should maintain the same qualitative behaviour.




Figure 2.7: Formation and evolution of Kelvin Helmholtz billows in a skewed domain at 1.5 and 1.6 seconds. Note that the billows begin to form as described by classical theory, however quickly start to deform when their size forces them to interact with the boundaries.

The first observation we can make about this simulation is that the billows develop much slower than the centered pycnocline case, in this case mature billows do not fully form until 1-2 outputs (about 0.2 seconds) later. This can be at least partially explained by the fact that the gravity-induced shear is slightly smaller than for the centered pycnocline case. Another important thing to note is that the flow does not become nearly as complex as the centered pycnocline case, at least before three-dimensionalization takes place.

In the asymmetric simulations the billows form in much the same way as before. They are smaller, as they are now restricted to effectively half the domain before symmetry is broken. At 1.7 seconds figure 2.8 shows the state of the billows before interaction with the boundary significantly affects the evolution of the flow. Following along to the next output time (figure 2.9 we see that the presence of boundaries has significantly impacted the evolution of the flow, causing deformation of the classical billow shape.

The top panel of figure 2.8 shows the highly asymmetrical nature of this simulation as the billows have collided with the wall and deflected towards the farther boundary. An interesting feature of this simulation is the preservation of the internal structure of the core while the billow shape is manipulated by the boundary. We can continue to clearly observe the classic core and braid structure even as the billows are deformed (stretched or sheared) by the interaction with the boundary. It is also worth noting that the skewed case demonstrates a significant amount of dissipation at a distance. Namely we see that the rightmost and leftmost billows in figure 2.8 do not appear to contact the bottom boundary, however we see significant dissipation at the far boundary.

Another fundamental feature clearly present in the asymmetric simulation is subharmonic pairing, or as we shall sometimes refer to it, billow-billow interaction. In the symmetric quasi-two layer simulation we clearly observed the individual billows form, grow to the size of the domain and then become deformed, as is clearly visible in the density and dissipation plots. The overturning here was violent and induced significant mixing and dissipation, both effects that we are greatly interested in. The end result was a quasi-turbulent state which we were quite convinced would result in three-dimensionalization. What we did *not* observe in any capacity was an intermediate stage where the billows grew large enough to self-interact. This was likely due to the small domain size and the greater importance of boundaries than adjacent billows.

In the asymmetric case we are privy to a very different growth regime. The billows, placed far from the bottom boundary, grow to roughly half the size of the domain before coming into contact with the top no-slip boundary. At this point the induced vertical velocities cause several things to happen. First (as we can see from the density plots in figure 2.7, 2.8, and 2.9) the billows are ejected away from the nearest boundary. Additionally we




Figure 2.8: Density and dissipation (top and bottom panels respectively) of KH billows in a skewed domain at 1.7 seconds. Dissipation is scaled by 50% of max dissipation in simulation. The billows have grown to nearly the size of the domain and have deformed significantly. It is worth noting that the only significant locations of dissipation at this time occur at the far boundary.

witness the beginnings of sub-harmonic pairing - the newly ejected billows are now in a regime where adjacent cores provide the most significant sources of resistance and as such begin to mutually interact.




Figure 2.9: Density and dissipation (top and bottom respectively) of KH billows in a skewed domain at 1.8 seconds. Dissipation is scaled by 50% of max dissipation in simulation. The fully mature billows begin to continue to be modified by the boundaries and begin to interact with each other. Note again that the significant locations of dissipation are near the boundaries and not the billow cores or braids.

One somewhat strange (at least at first) feature of this simulation is the seemingly empty space between billows. The fact that the billows don't grow to fill the domain in the horizontal is a result of there being significantly more dense fluid present in the domain.

However it does raise the question of what kind of motion is present in this region and why mixing is not significantly present. The answer is that this strange behaviour is an indicator of a kind of coherent structure which we will explore in more detail in the 0.75 cm simulations; however, it is worth noting here that this behaviour is present in some way in larger simulations.

While we can clearly observe billows pairing and amalgamating (especially around $x = 0.025$), the amount of pairing activity is severely limited by the strict confinement of the flow. This phenomenon will be worth revisiting in a slightly larger domain.

If we consider the dissipation of the same simulation (see figures 2.8 and 2.9) we note that as before the initial broadening of the pycnocline causes a large amount of dissipation, however the later times show something somewhat unexpected. While the dissipation in the center of the domain is again very weak, the amount of dissipation generated at the far (bottom) boundary is nontrivial, a surprising result considering the billows did not grow large enough to reach the far boundary until 1.7 seconds. This implies that these billows are still very energetic, causing significant grinding and dissipation whenever they come in contact with the farther boundary (or each other). Additionally it clearly demonstrates that the induced velocities reach throughout this (relatively small) domain despite the asymmetric nature of the initial setup.

2.3.3 2 cm Simulations

Before beginning to explore the 3D effects associated with spanwise instabilities and overturning found in the core it is useful to consider the effects of confinement, and specifically domain size, on the billows development. As such, the next set of simulations we consider is the 10 cm by 2 cm domain (2cm symmetric).

Figure 2.10 shows the density results of the 2 cm tall simulation. We see that the initial development and growth of the billows follows the classically prescribed path [15]. Notice also that the timing of the 2 cm symmetric domain corresponds well to that of the 1cm symmetric domain - the simulations exhibit similar timings for the initial formation of the sinusoidal disturbance (around 1 second) as well as the formation of mature billows (around 1.4 seconds).

However, an important difference from the 1cm domain is that once the billows have grown large enough to interact with the boundary, they instead begin to form billow-billow interactions - sometimes referred to as subharmonic pairing (see [32, 60, 61] and references therein). These interactions dramatically deform the classical shape of the cores and braids, however this is all possible with a strong enough shear current even in the




Figure 2.10: Density contours of KH billows in a symmetric 2 cm domain at 1.4 and 1.5 seconds (top and bottom panels respectively). The initial billow formations follows the classically prescribed path, however before the billows can reach the boundaries the intense currents (and shear) induce significant billow-billow interaction.

absence of boundaries [60, 61]. From this we conclude that 1cm or smaller is the most appropriate vertical domain size in order to study the effect of confinement on individual billows with this stratification and a 45 degree tilt.

Figure 2.10 also demonstrates that once the billows have self-interacted, they begin to

be affected by the presence of no-slip boundaries. This is most visible in the dissipation plot (figure 2.11).

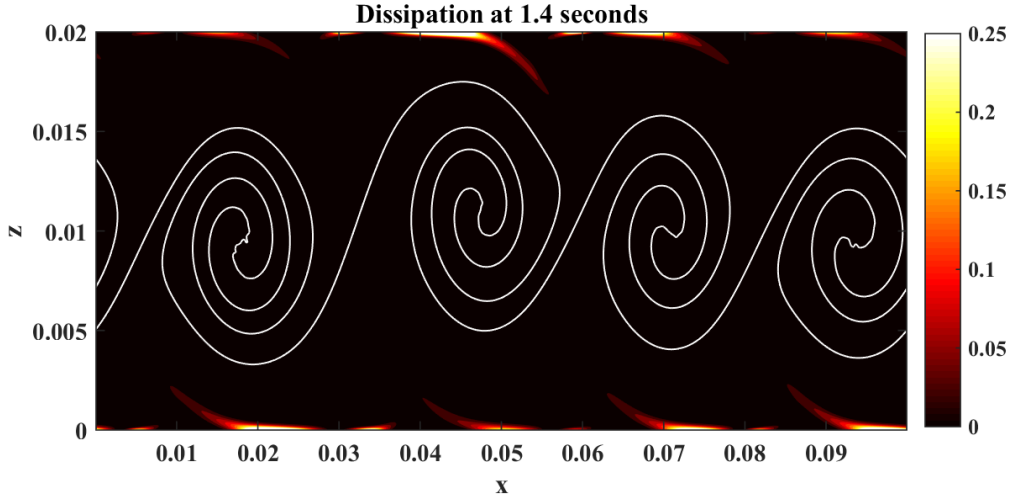


Figure 2.11: Dissipation induced by KH billows in a symmetric 2 cm domain at 1.4 seconds scaled by 50% of max dissipation in simulation. The majority of dissipation occurs at the boundaries above and below the billow cores at the no-slip boundaries.

Figure 2.11 shows the dissipation ($2\mu e_{ij}e_{ij}$) of the 2 cm symmetric domain again scaled by a percentage of the maximum value achieved throughout the simulation. Only one output time is presented because the dissipation present in this run is minimal. We notice immediately that the dissipation looks very similar to the smaller vertical extent case, and that the vast majority of late-time dissipation occurs near the boundaries. It is worth noting that even during the billow-billow interactions, there is enough velocity generated near the boundaries (away from the primary interaction) to induce significant dissipation even away from the billow cores.

2.3.4 2 cm skewed domain

Considering the 2cm domain with an off-center (asymmetric) pycnocline is interesting for two reasons. First we wish to consider the effects of broken symmetry on this stratified shear flow in the presence of boundaries - in a setting similar to the 1cm skewed case. Second, placing the pycnocline at one quarter of the total depth (0.5 cm from the top

boundary) is effectively the same as the very first simulation we ran (the 1cm centered pycnocline case) with the bottom boundary moved an extra 1cm away from the billow cores. Thus we expect this simulation to proceed very similarly to the 1cm centered case with more dissipation in the closer boundary and much less activity in the bottom half of the domain.




Figure 2.12: Density contours of Kelvin Helmholtz billows in a skewed 2 cm domain at 1.4 and 1.5 seconds (top and bottom panels respectively). The billows form and begin to interact with the top boundary, but are quickly deformed and in some cases (as with the billow on the far left) ejected away from the pycnocline location.

Figure 2.12 shows that proximity to a boundary dramatically alters the evolution of KH billows. While the very early stages are similar to previous simulations (and have thus been omitted) it is clear that once the billows have matured (1.6 seconds and onwards) rather than remaining in place the billows are deflected away from the boundary and ‘ejected’ into the body of the fluid. At 1.5 seconds we see the formation of a billow core ‘dipole’ as it is ejected from the pycnocline, trapping some light fluid and advecting it towards the far (bottom) boundary.

By the later times (4 seconds and onwards, not shown) the density field has lost all resemblance to classical shear flow theory and we can be certain that three-dimensional effects would begin to dominate the dynamics. This will be confirmed below when we examine the onset of spanwise secondary instabilities in various regimes and the transition to turbulence.

Considering the dissipation of the 2cm simulation with the off-center pycnocline (figure 2.13) we see that the structure of the dissipation is very similar to previous simulations. Interestingly, despite the boundary being much further away from the pycnocline than previous simulations we still see significant dissipation induced there even from early times (top panel of figure 2.13), despite no billows having reached that far into the fluid until almost 0.4 second later.

A key aspect of this simulation at later times is the effects of subharmonic pairing, similarly to that observed in the 1cm skewed case. Like in the 1cm skewed domain billow growth becomes limited by the size of the top region and adjacent billows become the prevailing feature - rather than the presence of boundaries. As such, at 1.4 seconds we see billows pair and combine into larger, less regular, structures. A log plot of dissipation seen in figure 2.14 shows several interesting features. First, as predicted by Mashayek and Peltier ([60, 61]) the pairing process is accompanied by a spreading of dissipation from the braid. Unlike in the simulations presented in Mashayek’s landmark paper the presence of boundaries means that the most dramatic dissipation occurs near the boundaries, rather than in the braid. However the log-scaled plot clearly shows that the braid experiences significant dissipation and comparable spatial spreading of dissipation.

This process demonstrates that we have chosen a domain size (namely 0 – 2cm) that captures two types of shear flow that we are interested in - near-wall flow and far-wall behaviour. In our simulations boundary dominated behaviour is characterized by billow growth being limited once interaction with the boundary begins in earnest (i.e. when the grinding against the walls becomes dominant). This type of flow has not been thoroughly examined before, and so we will be more interested in the effects of three-dimensionalization (and spanwise instability) in this regime. Additionally, the interaction with the boundaries

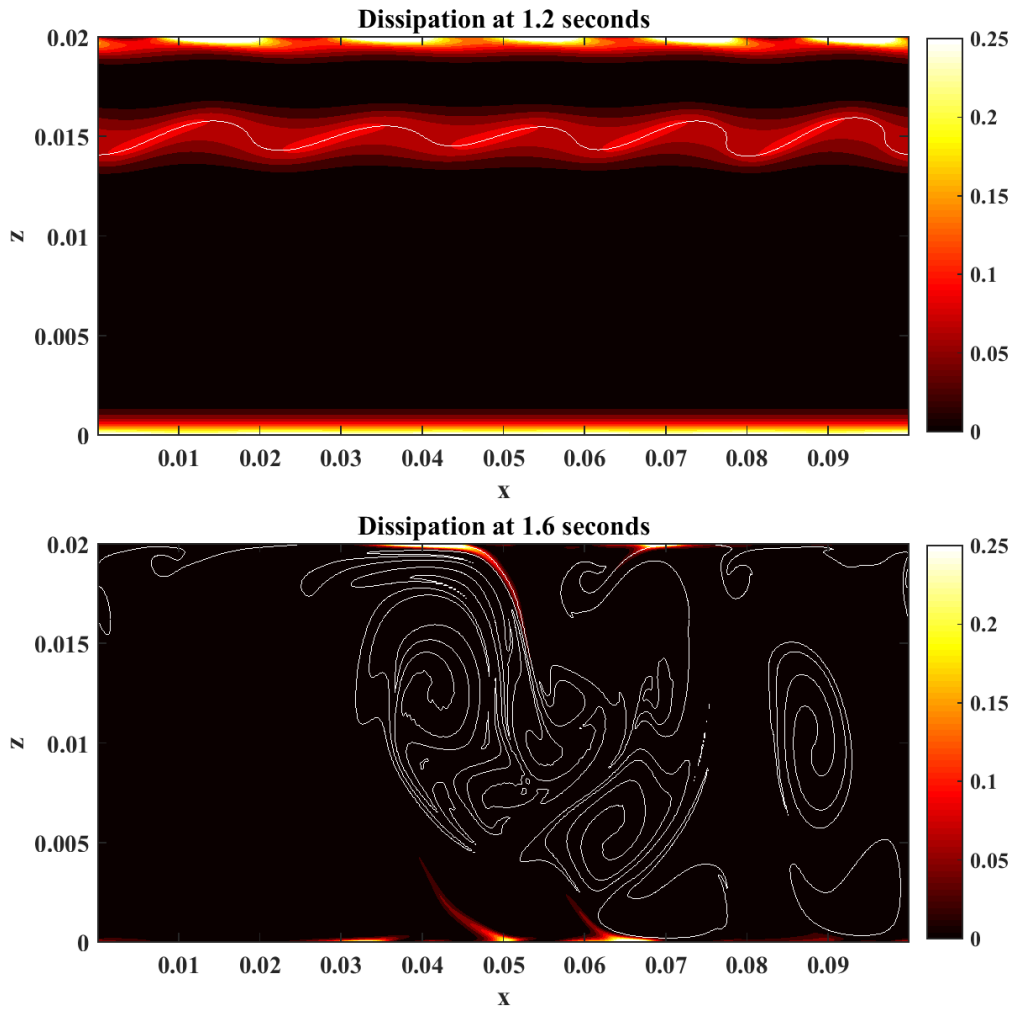


Figure 2.13: Dissipation contours of KH billows in a skewed 2 cm domain at 1.2 and 1.6 seconds (top and bottom panels respectively). We note that the dissipation for mature billows (bottom panel) is extremely weak and confined very strictly to the boundaries.

provides much greater dissipation in both cases, and it is this interaction which is of interest at this scale.

In contrast, the 2cm simulation exhibits subharmonic pairing, an indication that the shear is large enough so that this is effectively a far-wall-like regime. While this is an interesting result, it has been extensively studied before ([14, 15, 49, 60, 61, 72]) and thus

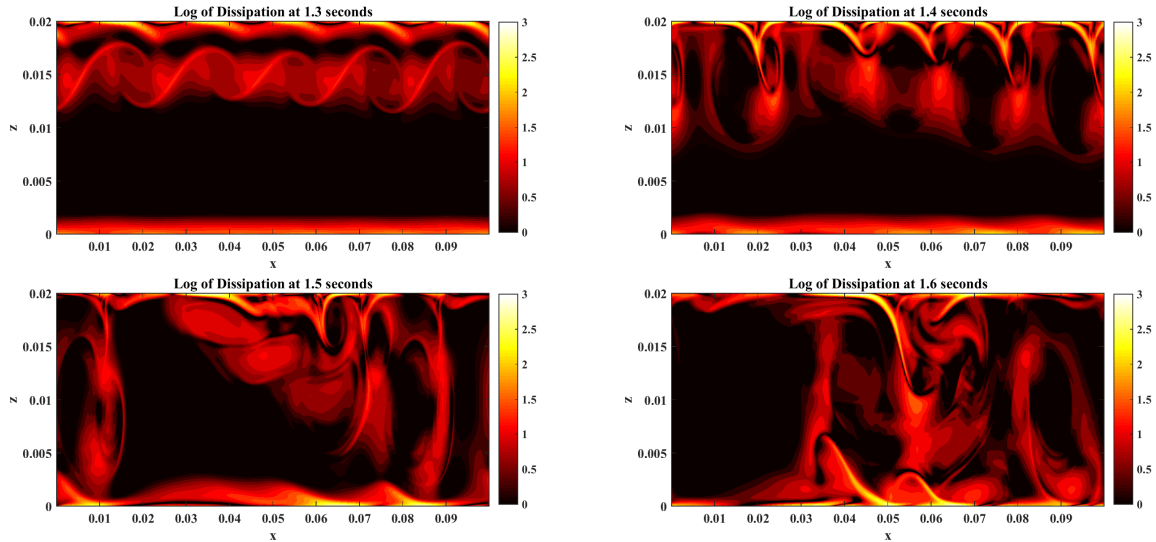


Figure 2.14: Log scaled dissipation during the pairing process between 1.3 and 1.6 seconds in 0.25 second increments. While the most significant dissipation occurs near the domain boundaries, we see significant secondary structure near the billow braids which is diffused during the pairing process.

we will limit our further investigations to smaller domains. While it is possible that a regime exists where the subharmonic pairing results in billows being ejected towards a far boundary that, in colliding, results in large grinding; however, this is well beyond the scope of this work.

2.3.5 0.75 cm skewed domain

We have investigated both symmetric and skewed domains in 1 and 2 cm heights. While the symmetric domains exhibit very similar behaviour, there is a significant difference between 1 and 2 cm skewed domain. In order to investigate this further we will consider a slightly smaller domain (0.75 cm) in order to see if we can force a regime change with a relatively minor domain alteration.

The initial evolution of the system is well described in previous sections and once again follows the classical evolution [14, 15] and is omitted here for the sake of brevity. The first output we will examine is a very early time, order to confirm




Figure 2.15: Density contours (top) and dissipation (bottom) of Kelvin Helmholtz billows in an asymmetric 0.75 cm domain at 1.5 seconds. The initial rollup occurs as classically prescribed (see Caulfield and Peltier and references therein [15]), however it is worth noting the large amount of dissipation occurs at the far boundary as before.

Figure 2.15 shows density and dissipation (top and bottom panels respectively) the 0.75 cm tall asymmetric domain at 1.5 seconds. This rollup occurs exactly as classically prescribed [14, 72], and the billows are well confined to the top half of the boundary, as expected from previous simulations and our intuition. One feature of great interest at this

time is that although there is no density disturbance in the bottom half of the domain, it is here that we witness the largest locations of dissipation. The bottom boundary at this early time is host to large regions of dissipation directly below the primary billows. This effect was present (albeit to a lesser degree) in the 1cm skewed simulation, and much less so in the 2cm skewed domain. It is a worthwhile observation (although perhaps obvious in retrospect) that this ‘action at a distance’ occurs more dramatically in the smaller domains. This will provide an interesting point of study in the three-dimensional simulations and is worth keeping an eye on at later times.

Turning to the next output, figure 2.16, we see a snapshot of the simulation at 1.75 seconds with density and dissipation (top and bottom panels) displayed. The billows at this time are clearly quite mature, exhibiting strong interactions with the top boundary - evidenced here by the flattening/broadening of the billow cores near the top and the triangular elongation near the center of the domain. Additionally we see clearly the effects of limited subharmonic pairing - what we called billow-billow interactions, as evidenced by the braids becoming absorbed into adjacent billow cores as well as significant core deformation. Altogether, this output is clearly very active and while we are still reasonably confident that a fully three-dimensional simulation would remain mostly 2D, it is easy to convince oneself that instabilities in the (quite deformed) braids or cores may begin to manifest here.

Regardless of our doubts as to the total validity of this two-dimensional simulation it is worth noting that the dissipation once again displays significant action at the far boundary, where the light density fluid is just now beginning to reach. The other noteworthy fact here is that the bottom dissipation region appears to be getting entrained into the center of the domain. The direct mechanism for this dispersal of dissipation is not clear, however we can be reasonably confident that it originated at the bottom boundary.

Turning to the next output (at 2.0 seconds) shown in in figure 2.17 we observe the last stage of this two-dimensional simulation that we can reasonably trust. This state contains three very interesting (and quite unique from the previous simulation) features which we will go through: The fully mature billows, which take on a rather unique shape; the location, structure, and shape of the dissipation; and finally, the resulting (partially coherent) structure in the bottom half of the boundary.

First, and foremost, we will consider the shape of the resulting billows in this domain size and asymmetric setup. The resulting billows are almost triangular in shape, with the top of the billows (where they come into contact and are thus influenced by the no-slip boundary) is broad and flat. The core itself remains roughly elliptical, with the east and west edges becoming deformed by contact both with the boundary and with adjacent




Figure 2.16: Density contours (top) and dissipation (bottom) of Kelvin Helmholtz billows in an asymmetric 0.75 cm domain at 1.75 seconds. The simulation displays mature billows interacting with the top boundary and exhibiting limited subharmonic pairing.

billows. In this domain, we have achieved a combination of billow-boundary and billow-billow interactions nearly superimposed, resulting in a unique structure worthy of further examination.

The next feature of note is the dissipation. As expected from intuition and previous experiments the vast majority of dissipation occurs near the top and bottom boundaries




Figure 2.17: Density contours (top) and dissipation (bottom) of Kelvin Helmholtz billows in an asymmetric 0.75 cm domain at 2.0 seconds. Note the strong dissipation present at the far boundary.

of the domain. The earlier times clearly showed that the dissipation is formed directly above and below the billows at the boundary and what significant dissipation exists in the domain's interior (namely in the bottom half of the domain) is advected there from the boundaries. The dissipation near the top boundary, which we had speculated *a priori* would be the more significant location, is confined to the boundary layer, likely by the fact

that the induced flow field causes the billows to remain near the top boundary. In contrast, the dissipation near the bottom boundary is located in nearly unstratified fluid, and thus the precise nature of the velocities and mechanisms present in that region is quite unclear.

This observation brings us smoothly to what is likely the most interesting aspect of this simulation, the formation of coherent structures in the dense fluid region. Our previous simulations (namely the 1 and 2 cm domains) indicated that the generated billows had a propensity for becoming somewhat free-floating within the domain, and becoming ejected from the “near-boundary” region where they were generated. In contrast to that behaviour what we see here is that the billows remain near their initial boundary and instead form coherent structures in parallel to the initially formed billow cores.

Rather than explore these (somewhat surprising) features further here, we will instead put off the full analysis for the three-dimensional simulations, which we expect to provide a more robust picture of the mechanics of this system, and will allow us to say for certain whether these results are truly realistic or if spanwise instabilities will come into effect and disrupt this system in ways we cannot yet predict.

2.3.6 Integrated Measures of Mixing and Dissipation

Before moving on to consider the effects of spanwise instabilities and shear stress on walls we can examine the total energy and dissipation of the 2D simulations.

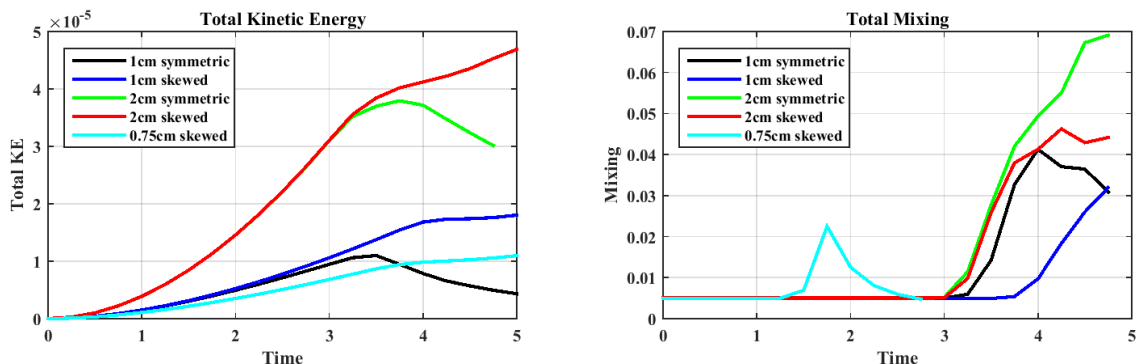


Figure 2.18: Kinetic energy (left) and total mixing, M , (right) of all four one-dimensional simulations. Note that the energy of the symmetric simulations is monotonic, while the skewed simulations reach a peak just as the vertical velocities drive the billows away from the boundary.

Figure 2.18 shows the total kinetic energy and mixing of all 2D simulations. It is first worth noting that the total kinetic energy of all the simulations is roughly equivalent - i.e doubling the domain does not produce energy that is an order of magnitude larger. Also immediately obvious is that the skewed (off-centered pycnocline) simulations have monotonically increasing kinetic energy profiles, while the centered pycnocline simulations reach an energetic peak relatively early. This peak corresponds well with the initial billow formation, indicating that the billow-forming event is more energetic than later interactions.

In contrast, for the off-center pycnocline domains, the billow formation is as energetic (in fact, it is slightly higher than the corresponding centered pycnocline cases), however, the subsequent billow-billow and billow-boundary interactions induce velocities which exceed the initial formation event. This result is corroborated by our dissipation plots (figures 2.8, 2.9, 2.10, and 2.13), which showed that the initial billow formation produced less dissipation than the following billow-billow and (more importantly) billow-boundary interactions.

Finally it is worth noting that the skewed 1cm domain actually surpasses the centered 2cm domain in kinetic energy. Although (as expected) the larger domain does have larger initial formation velocities, the (1cm) asymmetric domain actually yields larger long-term velocities in a more confined setting, a fact that can have interesting implications for mixing and chemical processes. The first order explanation for this effect is that the asymmetric domain allows for more rapid destabilization. Lack of symmetry means that there is a preferred direction for the billows to begin destabilizing. This results in a faster transition to a turbulent regime and, as a result, larger integrated kinetic energy.

Considering next the total dissipation induced by the various flows, we see from figure 2.19 that the total dissipation across all simulations is roughly equal in magnitude. As expected, the highest levels of dissipation occur in the centered pycnocline 1cm domain. It is worth noting that the centered domains, although less energetic in the long run, induce more dissipation than their skewed counterparts. This is due to the fact that the symmetric domains allow for higher velocities to develop when compared to their skewed counterparts, which in turn produce more dissipation when interacting with the boundary layers.

It is immediately obvious from the total dissipation plots (figure 2.19) that the majority of dissipation occurs not within the interior of the domain, but near the boundaries.

Additionally, since in the skewed cases the billow size is limited by the presence of the nearer boundary, we observed less billow-boundary interaction and thus less dissipation. However, these differences do not dramatically alter the magnitude of total dissipation observed - i.e. the presence of boundaries, in combination with shear, will produce some dissipation regardless of the precise configuration of the pycnocline.




Figure 2.19: Total Dissipation (solid lines) and dissipation in the interior - removing the top and bottom 20 grid points - (dashed lines) for the one-dimensional simulations. Notice that the vast majority of dissipation across all simulations occurs within 20 grid points of the boundary.

2.4 Three-Dimensional Simulations

After the initial onset and two-dimensional rollup of KH billows spanwise instabilities rapidly develop and the system becomes active in all three dimensions. The two-dimensional simulations indicate that the 2cm and 1cm domains exhibit very similar behavior. Thus we have chosen to limit our three-dimensional simulations to a 4cm by 2cm by 0.75cm (x, y, and z, respectively) domain. This choice was made for two reasons, one physical,

Table 2.3: Three-Dimensional simulations and their numerical parameters.

| Simulation | Domain Size ($x \times y \times z$) | Grid Size ($x \times y \times z$) |
|---------------------|--|--|
| Symmetric Two-layer | 4 cm \times 2 cm \times 0.75cm | 512 \times 96 \times 128 |
| Skewed Two-layer | 4 cm \times 2 cm \times 0.75cm | 512 \times 96 \times 128 |
| Linearly stratified | 4 cm \times 2 cm \times 0.75cm | 512 \times 96 \times 128 |

one practical. The physical reason is that we wish to study more in depth the effects of boundaries on the evolution of the billows, and the larger domain (i.e. 2cm) demonstrated characteristics such as billow pairing which were not the goal of this investigation. The practical reason is that by reducing the vertical domain size we are able to get higher vertical resolution, the most relevant direction for the boundary-billow interaction, while keeping computation times within reason (in our case under a week using available resources). In order to further develop our understanding of the transition to turbulence in a confined setting we have chosen to focus on three distinct stratifications, a two-layer case with a thin ($< 1mm$) initial pycnocline at either the mid-depth or off-set by 25% from the mid-depth and a linearly stratified case.

The initial evolution of the three cases (not shown) follows the standard ‘billowing’ paradigm, with 2D instabilities forming first (as predicted by Squire’s theorem) and expanding in size to fill the domain. This initial buildup occurs almost identically for linearly stratified and quasi two-layer simulations. The only observable differences between the various stratifications investigated is the timescale of the initial rollup - the two-layer flows developing much faster than the linearly stratified case [41, 77, 76].

Table 2.3 shows the two cases presented in this section, along with relevant numerical parameters such as domain size and grid resolution. The extent in the spanwise dimension is 2cm, which is sufficient to capture all spanwise instabilities that develop within this confined regime (confirmed by trial and error). The domain is periodic in the spanwise direction. For the density isosurfaces below we refer to perturbation from the constant density value initially found in the pycnocline centre.

2.4.1 Symmetric Two-Layer Flow

The first three-dimensional simulation to discuss is the extension of the 1cm symmetric run. Extending the domain into three dimensions, we create a 4 cm \times 2 cm \times 0.75 cm domain (see table 2.3 for details).

As predicted by Squire’s theorem the initial evolution occurs in two dimensions and thus the first instance of the 3D (numerical) experiment (top panel of figure 2.20) corresponds very well with the early to mid stages of the two-dimensional simulations. The evolution follows the classical path of instability formation, growth, and billow formation separated by braids. The fully mature billow state (visible at around 1.5 seconds) then begins to destabilize in the spanwise direction (as predicted by theory [15, 72]) and it is on this state that we wish to perform the bulk of our analysis. The confined domain (0.75cm in the vertical) means that billows quickly grow to fill the entire domain, and do not have the opportunity to self-interact. Unlike the 2 cm case there is no evidence of subharmonic pairing, and thus we are free to study exclusively the effects of confinement on dissipation and the destabilization of the billows in the spanwise direction.

Although the domain is slightly (25%) smaller than the 1cm case, there are few significant differences in the formation or evolution of the initial billows. Additionally the dissipation (discussed below) evolves in the same qualitative manner. The only significant difference is that the smaller domain forms billows slightly slower than it’s larger counterparts; however the difference is not large enough to truly affect computation time, and so we are satisfied with the trade-off.

As predicted from the two-dimensional simulations, the flow remains mostly confined to two dimensions for a significant period of time. It is not until the billows have fully matured and grown to the full size of the domain (approximately 2 seconds into the simulation - bottom panel of figure 2.20) that we see the development of significant spanwise activity. This spanwise instability sets in quite rapidly. By 2 seconds (figure 2.21 we can clearly observe the growth of instabilities in the spanwise direction especially around the billow cores. The dissipation plot shows that while there is some dissipation present in the billow cores and interior of the domain, the majority occurs near the top and bottom boundary, where billow-induced velocities interact with the no-slip boundary conditions.

A key aspect of the simulation that we wish to track is how three-dimensional the flow becomes. As the primary instability, the KH billows, grow and reach peak maturity, these instabilities themselves will destabilize in the spanwise direction. This destabilization is frequently referred to as the spanwise or secondary instability. We wish to determine several factors, such as how active the flow is in the spanwise direction (which tells us to what extent secondary instabilities have developed) as well as how much energy goes into motions in the spanwise direction. This information can help us to understand the importance of three-dimensionalization with respect to mixing and dissipation of the flow. The standard deviation of kinetic energy provides a useful discussion point as it allows us to determine quantitatively both how the quickly the spanwise instability develops, and how much energy it contains.

Figure 2.22 shows the spanwise standard deviation of kinetic energy at 1.75 and 2 seconds scaled by 10% of the maximum kinetic energy at that output. The earlier output (top panel) shows some three-dimensionalization, present mainly in the braids connecting billows, though these locations are highly localized. In contrast, the later output (bottom panel) shows significant three-dimensionalization in the billow cores. This is confirmed by the isosurfaces in figures 2.20 and 2.21.

The now fully mature billows, begin to destabilize rapidly, and by 2.25 seconds (figure 2.23) we can already see that the flow has become significantly three-dimensional. The dissipation at this time shows the significant three-dimensionalization which is occurring both in the billow cores and in the near wall region. Additionally the spanwise instabilities have induced significantly larger dissipation in the billow cores (as denoted by the red regions away from the boundaries). This region of high dissipation near the center is something we did not see in any of the two-dimensional simulations, causing us to conclude that this is a direct result of secondary instabilities. This result confirms our intuition that three-dimensional flows produce more dissipation (and more mixing) than their two-dimensional counterparts.

Although the flow at this time (2.25 seconds) remains quite anisotropic (i.e. the down-channel velocities are significantly higher than the across-channel velocities) we do see the clear transition into a fully turbulent regime, as evidenced by the reduction of the smallest energy containing scales of the system as well as an increase in dissipation.

If we consider the enstrophy (and the 3D enstrophy) at this time, figure 2.24, we observe the formation of tubes of enstrophy corresponding to the internal structure of the billow cores. In contrast, the 3D enstrophy shown in the right panel corresponds to the deviations of total enstrophy in the spanwise direction. As such, the structure is primarily vertical and identifies physical locations where the tubes of enstrophy begin to destabilize. Due to the relatively two-dimensional state of the simulation 3D enstrophy is relatively low.

Another features of interest in figure 2.24 that the isosurfaces of highest enstrophy appear not in the center of the domain, but in the boundary layer - a fact which will be of great interest when we examine top and bottom shear.

As the evolution of the system progresses, we observe a nearly complete spanwise destabilization of the flow. By 3 seconds (figure 2.25) the system superficially resembles isotropic turbulence. The density isosurface (top panel) shows fine-scale features - a characteristic indicative of a fully turbulent flow. In addition to this, the locations of dissipation are no longer strictly confined to the near-wall or in-core regions. Although the down-stream velocities still dominate (mean values of approximately $2 \times 10^{-5} \frac{\text{m}}{\text{s}}$ compared to mean spanwise velocities of $5 \times 10^{-6} \frac{\text{m}}{\text{s}}$) this is primarily due to the large tilt employed.

Once three-dimensionalization is fully developed, and the system has transitioned to a fully turbulent regime, the density structure persists to some degree since mixing is imperfect, but the system is greatly reduced in energy, as discussed below. This period is associated with large spanwise variations in the density and dissipation fields, and a such is worth examining using total and 3D enstrophy (figure 2.26).

The primary feature of total enstrophy - shown in figure 2.26 - is (still) the large tubes associated with vorticity in the billow core. In contrast, the rapid destabilization in the spanwise (the rise of spanwise instabilities) causes 3D enstrophy to identify regions through the cores where instabilities develop. The 3D enstrophy isosurfaces provide useful information only when compared with the total enstrophy plots. By themselves they can identify regions where the dominant enstrophy is in the spanwise direction, however it is only when contrasted with the structure of the enstrophy tubes formed in billows that the full picture is revealed.

Total enstrophy both at 2.25 and 3 seconds (figures 2.24 and 2.26, respectively) showed significant enstrophy (50% isosurfaces) near the top and bottom boundaries directly above and below billow cores. In order to understand this, we must consider the velocity profile near the no-slip boundary as it has been generated by the billow velocities. The flow at the boundary itself is necessarily 0, however, as we proceed from the boundary towards the center of the domain over the center of the KH billows, we arrive in a region where the velocity is large and parallel to the wall. Thus as we leave the boundary the velocity is going to be exactly monotonically increasing, producing vorticity at the boundary, as expected.

A final key feature that is worth discussing for the three-dimensional simulations is the horizontal structure of the shear stresses on the top and bottom boundaries. In the symmetric case, the shear (and hence the shear stress) remains almost identical across the center of the domain, so we present only the shear on the top wall. However it is worth noting that for the asymmetric simulation this will not be the case.

During the initial two-dimensional rollup (figure 2.27) the billows induce significant activity in large portions near the top of the domain and this is expressed in three broad regions of elevated shear. Once the billows have matured to the point of filling the domain and beginning to interact with the boundary (bottom panel of figure 2.27) we see a larger number of thinner bands of high shear.

At later times, once the billows have begin to destabilize (figure 2.28) we see the uniform spanwise structure break down and produce interesting structure in the across-channel direction. While the total velocity gradients are smaller, and thus cause less activity, there is more interesting structure present. Despite this, it is evident that the fundamental na-

ture of the billow-induced shear remains essentially two-dimensional. The shear induced, while not uniform, is structured in bands.

As with the two-dimensional simulations it is worthwhile to consider the total dissipation, mixing, and kinetic energy (figure 2.29). There are several interesting phenomenon evident in this figure. The first is that the kinetic energy, dissipation, and mixing all reach a peak relatively early in the simulation (around 1.25 to 1.75 seconds) and then monotonically decrease. The second feature of note is that the dissipation in the interior of the domain (in green) is significantly smaller than the total dissipation. This is consistent with our physically intuition that the majority of induced dissipation occurs at the boundary-billow interface.

Another, if somewhat more subtle, feature is that the mixing peaks slightly later than the dissipation and kinetic energy. The precise reason for this delay is not entirely clear to the author, however it is evident that the dissipation rises as the kinetic energy increases, and drops rapidly once the fluid velocities decrease. It is worth remarking that at 3s, the time discussed above, and shown in Figure 2.25, the total kinetic energy is around 2% of the maximum, while the mixing continues at around 8% of its peak intensity.

Finally, it is worth noting that all three quantities peak while the simulation is firmly within the two-dimensional regime. This means that maximum velocities, dissipation, and mixing all occur in two-dimensions. This is a somewhat surprising result, as classically it is three-dimensionalization which is thought to be responsible for mixing and dissipation ([15, 60, 61]). As such this is a feature of confined KH billows which appears to deviate dramatically from the classically prescribed regime. It's also important to reiterate that the Reynolds numbers at these times range between 100 and 300 for various simulations (between 0.75 and 2 cm). While this is below values of KH observed in ocean or atmosphere, it is well above the "low-Re" regime in which we expect two-dimensional activity to dominate.

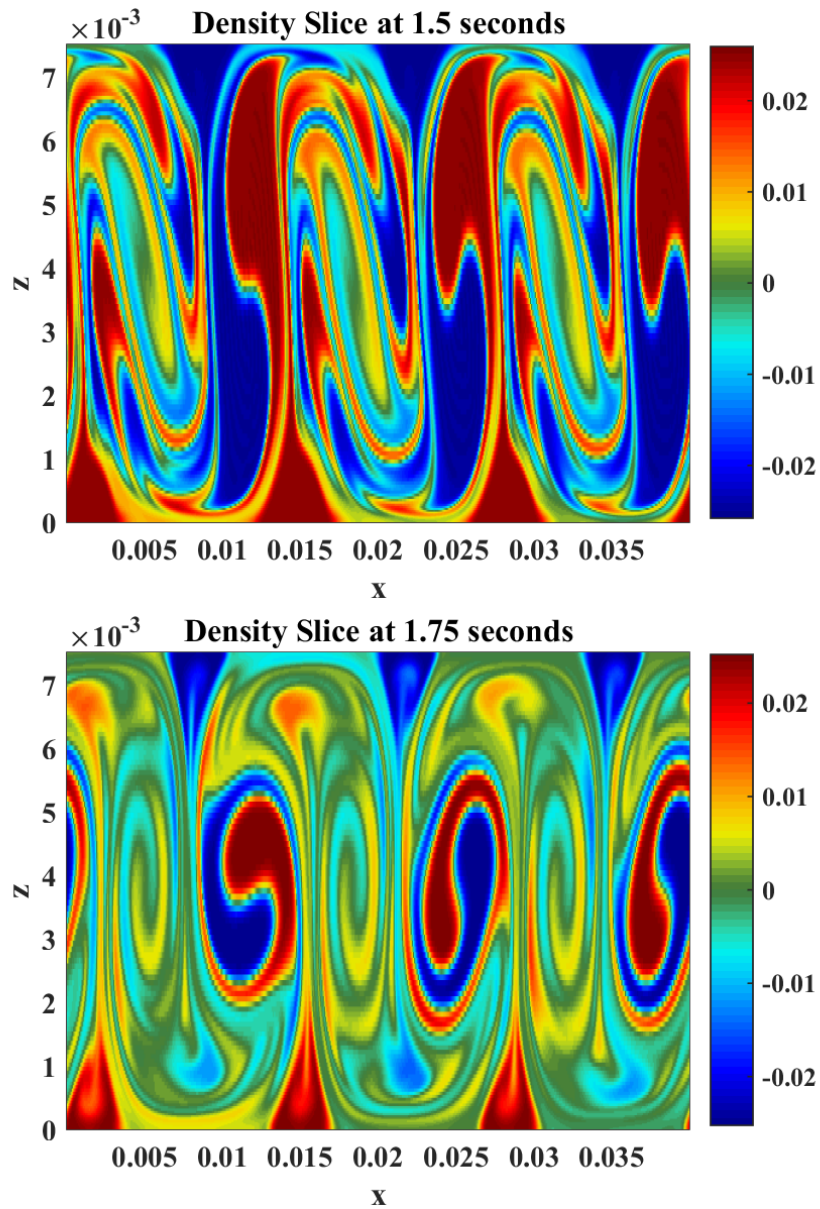


Figure 2.20: Isosurface of density showing $\pm 10\%$ of maximum density perturbation at 1.5 (top) and 1.75 (bottom) seconds. Notice that as predicted the initial rollup (top panel) occurs entirely in two dimensions with minimal activity (equivalent to initial noise) in the spanwise (y) direction until significant overturning has occurred in the billow cores (bottom panel).

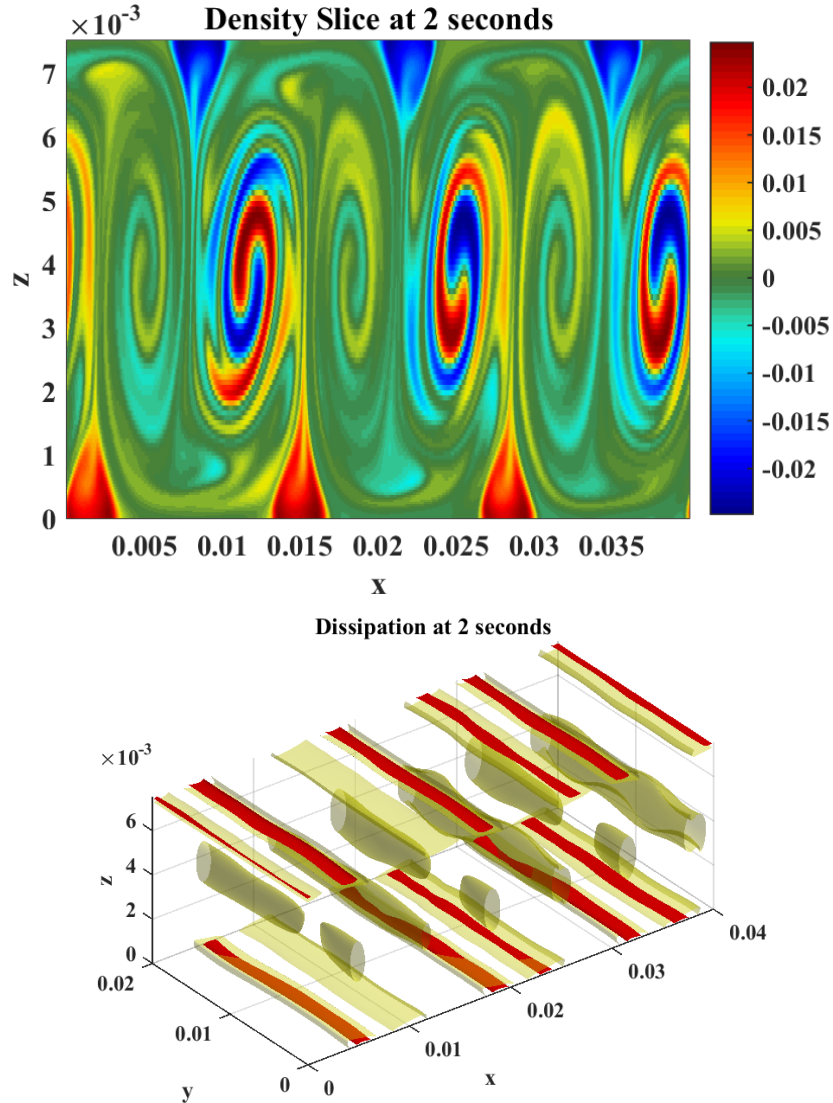


Figure 2.21: Density (top panel showing $\pm 10\%$ of maximum density perturbation) and dissipation (bottom panel showing 10 and 50% of max dissipation) isosurfaces of the symmetric quasi-two layer flow at 2 seconds. At this time the billows have fully matured and begin to destabilize in the spanwise (y) direction. The majority of dissipation (denoted in red) occurs not near the billow cores, but the edge of the domain.




Figure 2.22: Spanwise standard deviation of kinetic energy of the symmetric quasi two-layer flow at 1.75 (top) and 2 (bottom) seconds saturated at 10% of the maximum kinetic energy. As predicted by classical theory, the spanwise (secondary) instabilities develop in the cores and braids of the KH billows where the majority of overturn and mixing occurs.

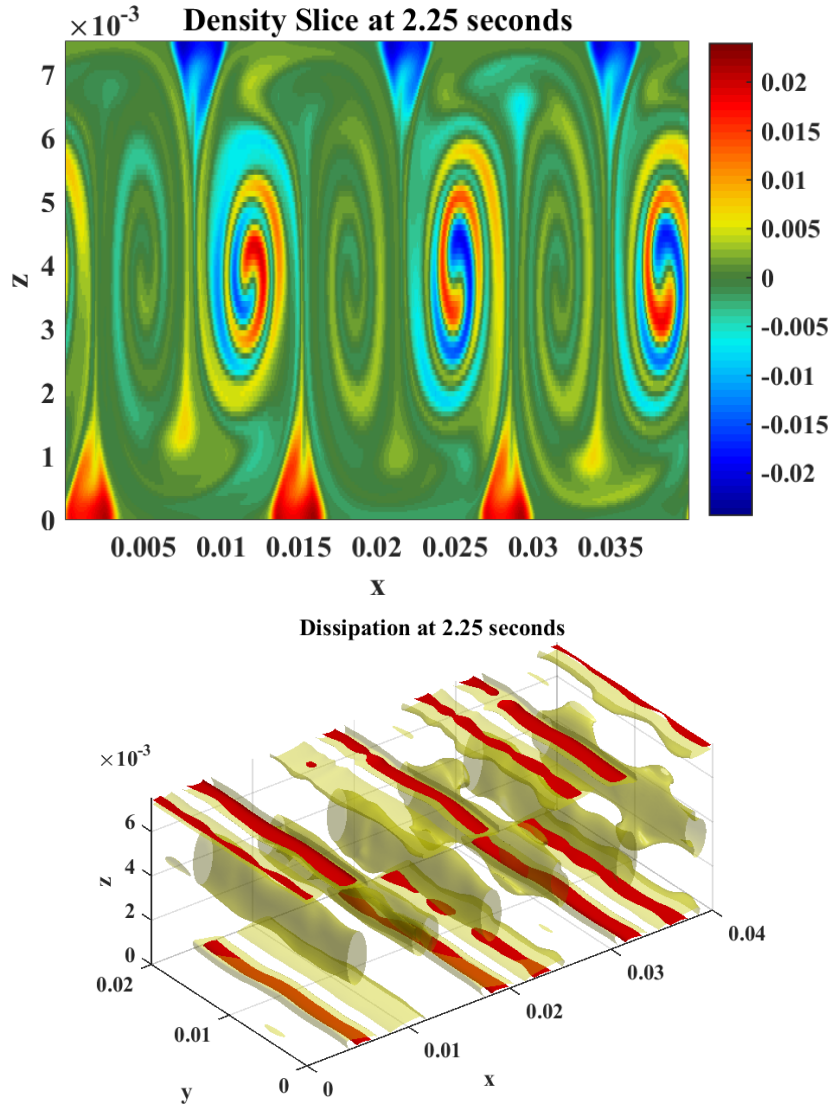


Figure 2.23: Density (top panel showing $\pm 10\%$ of maximum density perturbation) and dissipation (bottom panel showing 10 and 50% of max dissipation) isosurfaces of the symmetric quasi-two layer flow at 2.25 seconds. Note the presence of large regions of high dissipations in the billow cores which did not exist in two-dimensional flows.

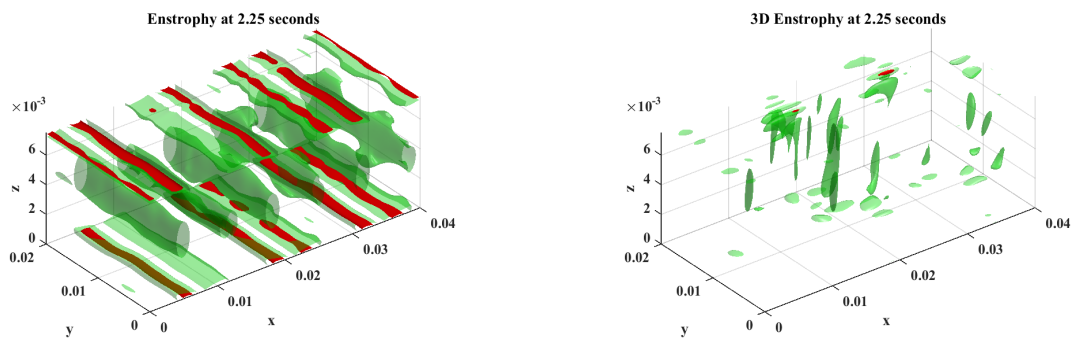


Figure 2.24: Enstrophy (top panel showing 10 and 50% of max enstrophy) and 3D enstrophy (bottom panel showing 10 and 50% of max 3D enstrophy) isosurfaces of the symmetric quasi-two layer flow at 2.25 seconds. Note the tube-like structure of enstrophy corresponding to billow cores, while the 3D enstrophy is relatively low due to the two-dimensionality of the simulation at this time.

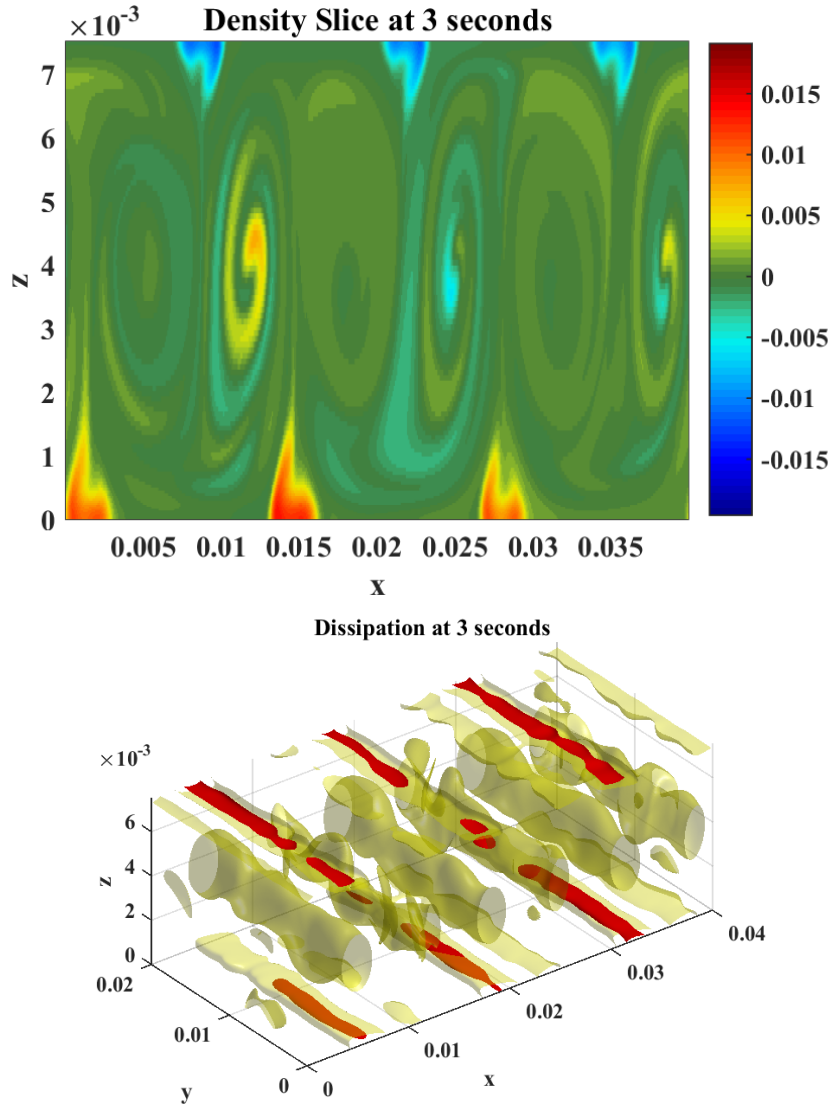


Figure 2.25: Density (top panel showing $\pm 10\%$ of maximum density perturbation) and dissipation (bottom panel showing 10 and 50% of max dissipation) isosurfaces of the symmetric quasi-two layer flow at 3 seconds. The system has destabilized to the point that it now resembles isotropic turbulence.

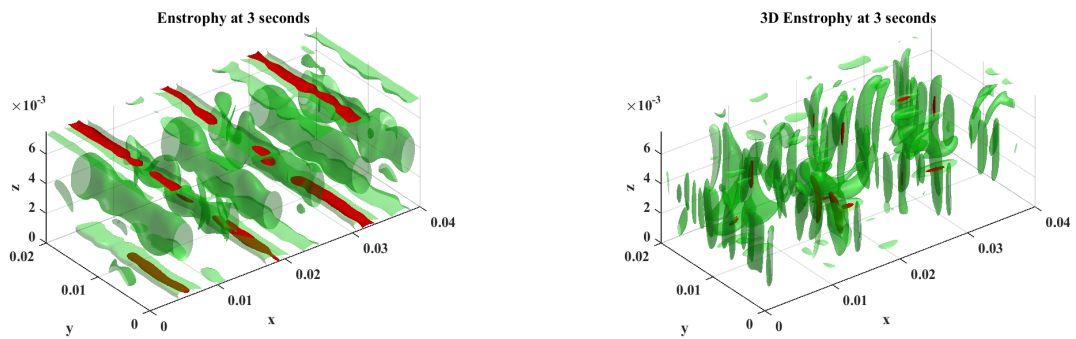


Figure 2.26: Enstrophy (top panel showing 10 and 50% of max enstrophy) and 3D enstrophy (bottom panel showing 10 and 50% of max 3D enstrophy) isosurfaces of the symmetric quasi-two layer flow at 3 seconds. The total enstrophy is dominated by tube-like structures (corresponding to billow cores), while the 3D enstrophy identifies regions of spanwise instability.




Figure 2.27: Shear stress in $x - y$ plane at the top boundary $\mu\sqrt{u_z^2 + v_z^2}$ at 1.25 (top) and 1.5 (bottom) seconds scaled by the maximum local value (0.21 Pa for both outputs). At these times the flow is entirely two-dimensional, and as such the spanwise structure is uniform.




Figure 2.28: $x-y$ plane of shear stress near the top boundary $\mu\sqrt{u_z^2 + v_z^2}$ at 2.25 (top) and 3 (bottom) seconds scaled by the maximum local value (0.072 Pa and 0.038 Pa respectively). The spanwise destabilization appears as wave-like variations in top and bottom shear lines..




Figure 2.29: Summary plot showing total kinetic energy, total mixing, and dissipation (both internal and total) scaled by the maximum values achieved over the length of the simulation. Notice that as predicted the mixing and dissipation both peak during the initial, two-dimensional rollup event. As the system mixes and becomes nearly homogeneous gravitational adjustment stops and all motion ceases.

2.4.2 Asymmetric Two-Layer Flow

The next fully three-dimensional simulation we will explore in-depth is the asymmetric quasi two-layer case. This setup places the pycnocline at three quarters of the domain depth. As with the first three-dimensional simulation examined, we extend the grid in the spanwise direction. The domain and simulation sizes are given in table 2.3.

The very early parts of this simulation correspond almost exactly to the two-dimensional evolution - the billows form and mature in the top half of the domain. The structure is consistent with the classical picture of cores connected by braids, and for the early evolution the presence of no slip boundaries (even so close to the pycnocline) is negligible.

As with the symmetric two-layer case, the initial evolution occurs strictly in two dimensions, with billows forming and growing uniformly in the spanwise direction. The early and mid stages of this simulation (see figure 2.30) once again correspond well to the early two-dimensional runs in section 2.3.5. As with the two-dimensional simulations we see that the thinner top layer causes lower velocities, and thus a slower billow evolution. Thus the timescale for the evolution in this domain is somewhat slower than the symmetric counterpart.

In the symmetric case the billows formed, grew to the size of the domain, and then destabilized and became three-dimensional. In this asymmetric run we have an additional complication, as in two-dimensions we saw the billows “ejected” from the near-top region towards the far boundary. Whether this effect is present in three dimensions, or if the flow destabilizes beforehand and becomes turbulent is something we wish to investigate here.

Considering the density at later times (figure 2.31) we immediately notice that the flow does begin to destabilize upon reaching the nearer (top) boundary. The three-dimensional simulations become too active in the spanwise direction, destabilizing the flow and inducing dramatically different regimes from those observed in two dimensions. This can be seen in the density isosurfaces shown in figure figure 2.31, as well as the standard deviation of kinetic energy in figure 2.32.

The standard deviation of kinetic energy also hints at something unique to the asymmetric two-layer case. Although the physical billows are (mostly) confined to the top half of the domain, there is significant “anti-billow” activity that occurs in the bottom half. These velocities have a large spanwise component - as seen in figure 2.32 and appear to be structurally very similar to the primary billows.

In larger domains (namely 1cm and 2cm two-dimensional simulations) we witnessed significant subharmonic pairing occur once the billow cores had matured. As we saw in the 0.75cm two-dimensional simulations this is not present in this domain, and the




Figure 2.30: Density isosurface at 1.5 and 1.75 seconds showing $\pm 10\%$ of maximum density perturbation. The initial billow evolution is entirely two-dimensional, evolving as predicted by the 2D simulations.

three-dimensional simulations confirm this result. In fact, an argument could be made that the three-dimensional simulations exhibit even less billow-pairing structure than the two-dimensional counterparts in section 2.3.5 which would demonstrate the importance of three-dimensional mechanics such as vortex stretching and spanwise instability as a




Figure 2.31: Density isosurface at 2.125 and 2.25 seconds showing $\pm 10\%$ of maximum density perturbation.

mechanism for the removal of background energy.

The specific mechanics of the top half of the domain are very interesting in their own right. The pycnocline is deformed (as is expected of this type of shear flow) and eventually leads to significant vorticity and overturning in the billow cores. These regions are connected by relatively thin braids, again in agreement with classical theory ([55, 32, 41, 64]).




Figure 2.32: Standard deviation of kinetic energy in the spanwise direction scaled by the local maximum of kinetic energy (saturated at 10% of max). The majority of activity occurs near the center of the domain, where destabilization in the spanwise (y) direction is strongest.

Although the billows appear to remain confined to the top half of the domain there is still considerable activity present in the far (i.e. bottom) boundary. This is particularly evident in the dissipation and enstrophy shown at the later times. The dissipation isosurface in figure 2.33 shows the 10% (in yellow) and 50% isosurfaces of maximum dissipation achieved. This allows us to immediately identify the billow cores, but also to note that although the bulk of the billows is confined to the top half of the domain, there is significant dissipation induced at the far boundary.

As previously mentioned, three-dimensional simulations differ drastically from two dimensions. The presence of vortex stretching allows for fundamentally different regimes to exist in 3D than those confined to two dimensions. If we consider the enstrophy isosurfaces at later times (figure 2.34) we notice immediately that the regions of highest enstrophy occur near the boundaries. The presence of no-slip boundaries causes large shear near the boundary, which is reflected in significant enstrophy. This is an easily explained effect and is useful for helping identify regions of high shear near boundaries. One feature of note about this shear is that it corresponds with the along-channel (x) location of billow cores, indicating that the billows near the top of the domain induce large velocities even at the far boundary.

Another interesting feature of the enstrophy plots is that at these later times we begin to see regions of high enstrophy near the center of the domain. This isosurface allows us to immediately identify that the center of the domain contains regions of high enstrophy. However, we are left with the question of what the source of this region was - whether the

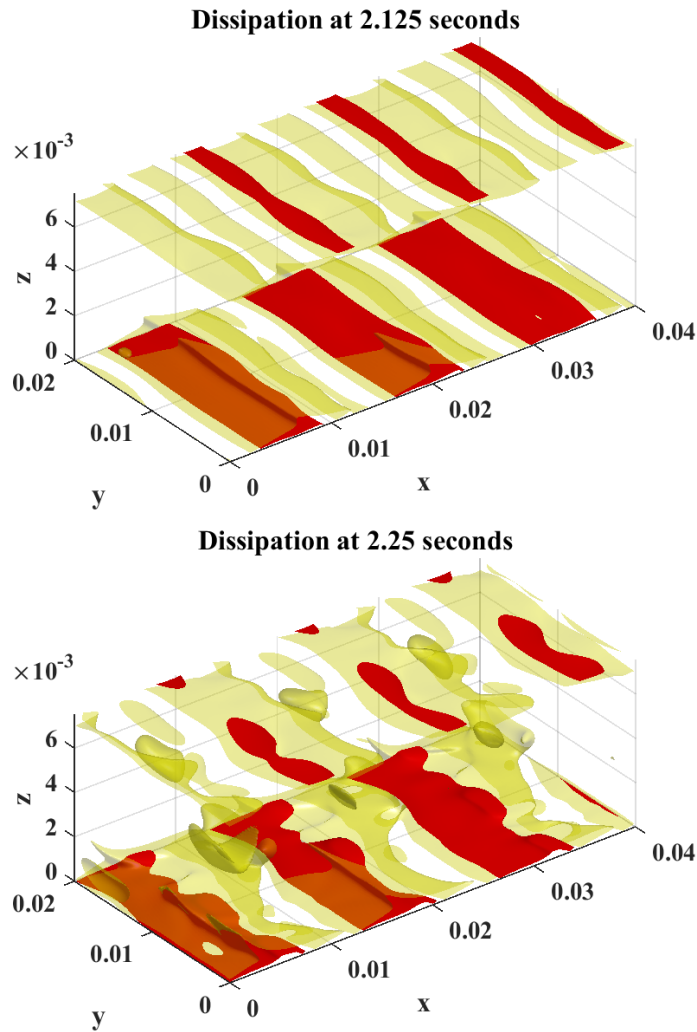


Figure 2.33: Dissipation isosurface at 2.125 and 2.25 seconds showing $\pm 10\%$ of maximum density perturbation.

enstrophy is a result of billow core (or braid) activity, or if it was advected into the center from the boundary layer). To that end we can examine a vertical slice through the center of the domain in figure 2.35.

Due to the relatively two-dimensional nature of the simulation at this time, the 3D enstrophy provides minimal new information. The earlier time (top-right panel showing

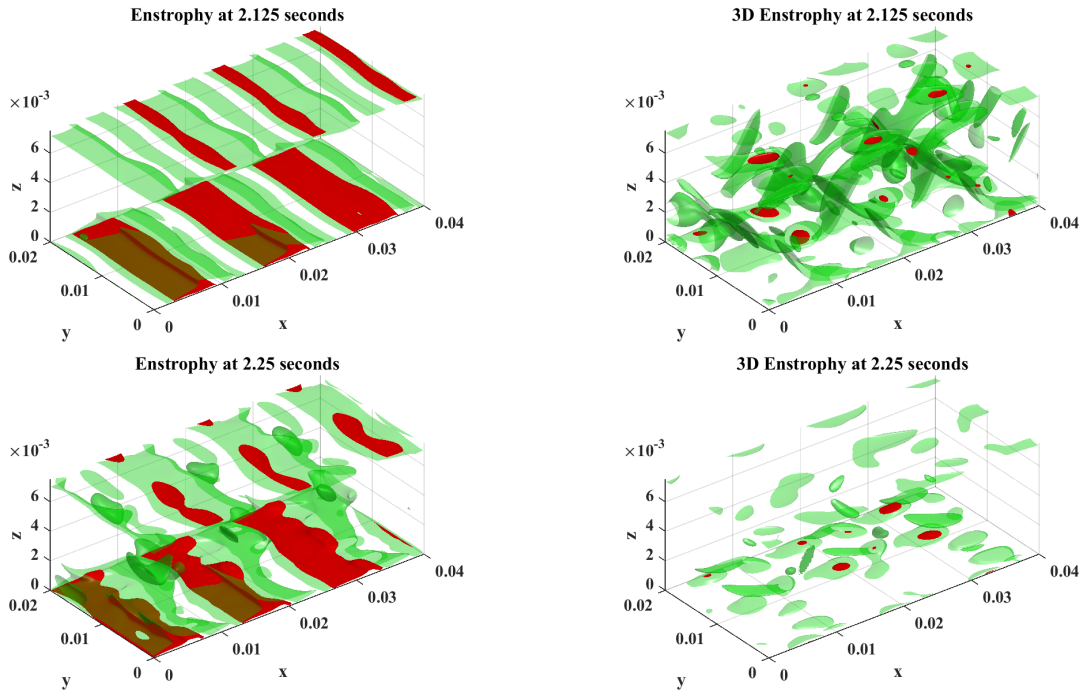


Figure 2.34: Enstrophy (left) and 3D enstrophy (right) isosurface at 2.125 and 2.25 seconds showing 10% (in green) and 50% (in red) of maximum enstrophy or 3D enstrophy, respectively. The highest regions of enstrophy are located at the far (bottom) boundary. We begin to see some regions of high enstrophy near the center of the domain. The structure of 3D enstrophy appears quite regular at the early times (due to lack of spanwise structure) becoming more interesting as the system destabilizes.

2.125 seconds) displays interesting structure, however the values of 3D enstrophy at this time are clearly very low, making this plot intriguing, but not very useful. In contrast the later time (bottom-right panel, showing 2.25 seconds) contains significantly less activity, highlighting only a few regions in the boundary, however looking at the total enstrophy plot we are forced to conclude that this is a much more physically meaningful plot, showing that the regions of significant destabilization occur near the boundary - with minimal interesting activity in the domain center.

Figure 2.35 shows the a vertical slice of density (left panels) and enstrophy (right panels) at 2, 2.125, and 2.25 seconds into the simulation through the (spanwise) center. The left panels show that the already mature billows continue to destabilize in the spanwise

direction until the resulting slice loses the traditional billow structure. Interestingly, the two layers remain relatively distinct, with mixing occurring relatively slowly across the sharp density gradient.

The density slices also reveal the presence of pseudo-billows created in the deep layer of the domain (near the bottom boundary). These “anti-billows” are created by the KH billow velocities around the pycnocline and are responsible for a significant amount of shear (and consequently enstrophy) near the bottom no-slip boundary.

The vertical slices of the enstrophy field yield an interesting, but not unexpected result. Due to the high amount of shear caused by the billows, the enstrophy production (top panels of figure 2.35) is dominated by interaction with the no-slip boundaries near the top and bottom of the domain. As the flow becomes more active, vertical (z direction) velocities become more prominent. This causes the vorticity produced in the boundary-layer to be advected into the center of the domain. The braids between billow cores act like guides for the vorticity, forcing it into the center of the domain.

One interesting takeaway from this result is the fact that the regions of high enstrophy in the domain’s center were extremely easy to identify using semi-transparent isosurfaces (as in figure 2.34). Unfortunately, these figures, while arousing our interest, did not sufficiently explain the formation or evolution of this newfound vorticity.

Figure 2.36 shows the total kinetic energy, dissipation (total and internal), and mixing throughout this simulation. In this asymmetric case there are several key differences from the symmetric three-dimensional simulation. The first feature of note is that the kinetic energy continues to increase throughout the simulation time. Although the increase is not monotonic (we see a slight decrease in overall kinetic energy when three-dimensionalization takes place around 2.5 seconds), there is a distinct upward trend. This is due to the fact that the dominating force in these simulations, the tilted gravity term, continues to act strongly as the induced mixing is not vertically uniform.

Whereas in the symmetric (and later linear) simulations the fluid is rapidly mixed to a near-uniform density, in this simulation the mixing occurs in stages. The initial billows are formed in the top half of the domain, and so it is that region which is mixed first. The resulting state is still partially stratified, and thus further shear is induced by the tilted gravity term. This transition can be clearly seen in the (spanwise) mean density plots in figure 2.37.

The next interesting aspect of the line plot in figure 2.36 is the double peak of dissipation and mixing. The first thing to note is that the first peak occurs around 1.75 seconds, which corresponds to the location of the initial two-dimensional rollup event. This event sees the maximum peak in mixing, and the smaller peak in dissipation. This is particularly

interesting as it implies that the more significant mixing event occurs primarily in two dimensions. The second peaks occur around 2.5 seconds, when the flow is transitioning to a fully three-dimensional (pseudo-turbulent) regime. This event sees a maximum of dissipation, and the smaller mixing event.

The double peaks in figure 2.36 are indicative of a fundamental regime change from the shear-dominated two-dimensional billow formation event around 1.5 seconds to the quasi-turbulent semi-mixed state around 2.5 seconds. The dramatic mixing near the end of this simulation results in a quasi-linear state. This state, which arises naturally from the gradual mixing in the asymmetric simulation is indicative of a transition from a sharp pycnocline to linearly stratified regime, making the linear stratification worth considering at this time.




Figure 2.35: Vertical ($x - z$) slices through the spanwise center of the domain of density (left) and enstrophy (right) at 2.25, 2.375, and 2.5 seconds. As the billows destabilize and secondary braid and core instabilities form enstrophy is generated first near the boundaries then near the center. Despite the largest density fluctuations being confined to the top half of the domain, the bottom boundary experiences much greater dissipation and enstrophy. Additionally the generation of “anti-billows” near the bottom boundary provides a mechanism for interaction at a distance.




Figure 2.36: Summary plot showing total kinetic energy, total mixing, and dissipation (both internal and total) scaled by the maximum values achieved over the length of the simulation. Notice that unlike the symmetric case the asymmetric simulation experiences two distinct peaks in all the measured quantities, with mixing achieving its maximum during the initial (two-dimensional) event, while kinetic energy and dissipation peak during the second (more three-dimensional) event.

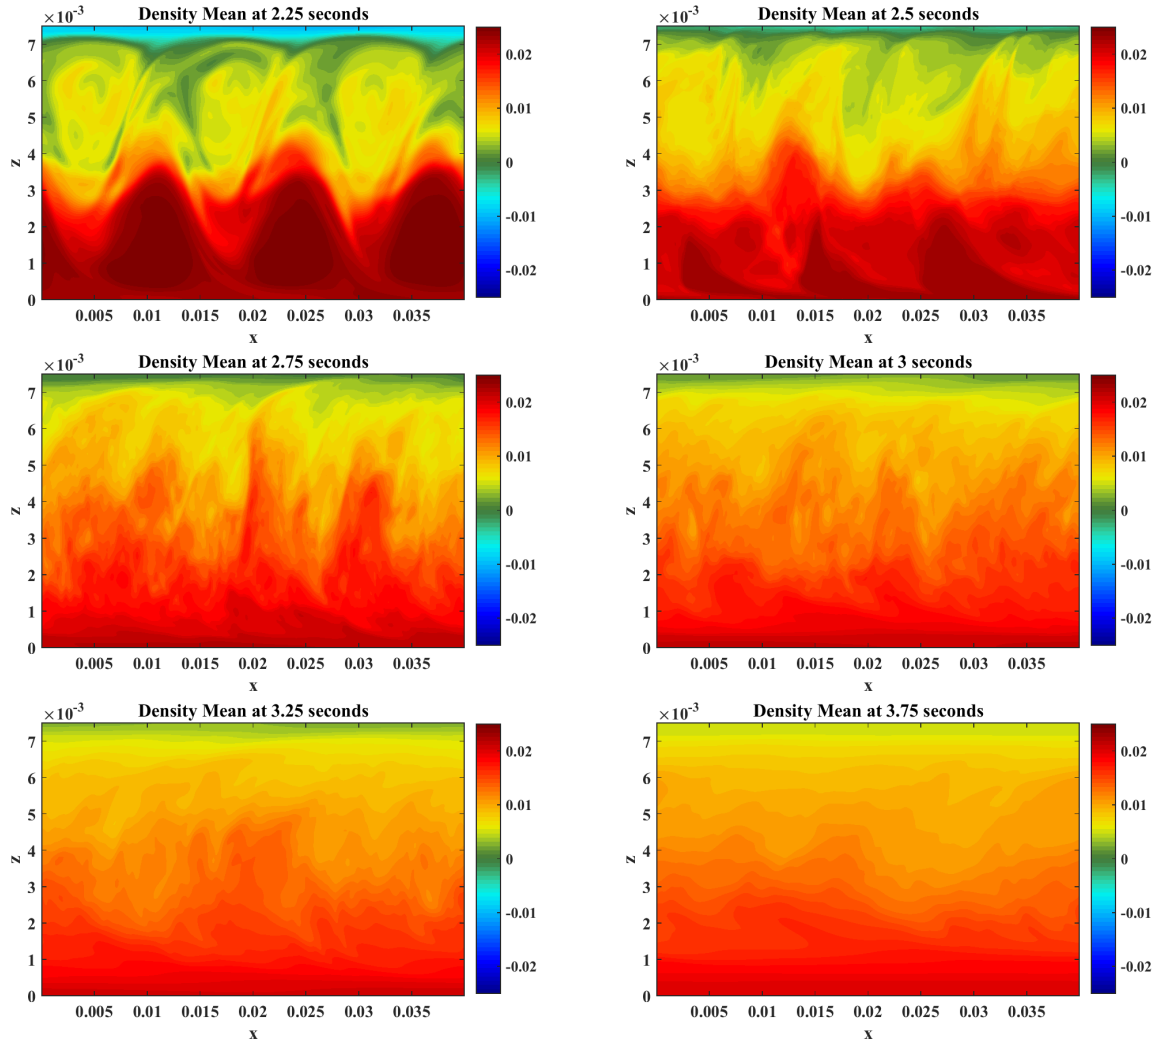


Figure 2.37: Spanwise-averaged density at later stages of the simulation. The average state transitions from a sharp pycnocline to a pseudo-linear regime. This figure provides a qualitative measure of the mixing which takes place in this simulation, whereby the sharp divide between light and heavy fluid is smoothed out.

2.4.3 Linear Stratification

The initial evolution of the instability follows the standard ‘billowing’ paradigm. This initial buildup occurs almost identically for linearly stratified or quasi two-layer simulations. The only observable differences between the the various stratifications investigated being the timescale of the initial rollup - the two-layer flows developing much faster than the linearly stratified case. This is easily explained since an initially sharp pycnocline yields much faster velocities from gravitational adjustment (light fluid travelling upwards en masse and heavy fluid travelling downwards). These faster velocities produce stronger shear and thus the instability forms much more quickly in 2 layers as opposed to linear stratification. However, after the initial billowing event, and a following period of quiescence, a second billowing event is observed.

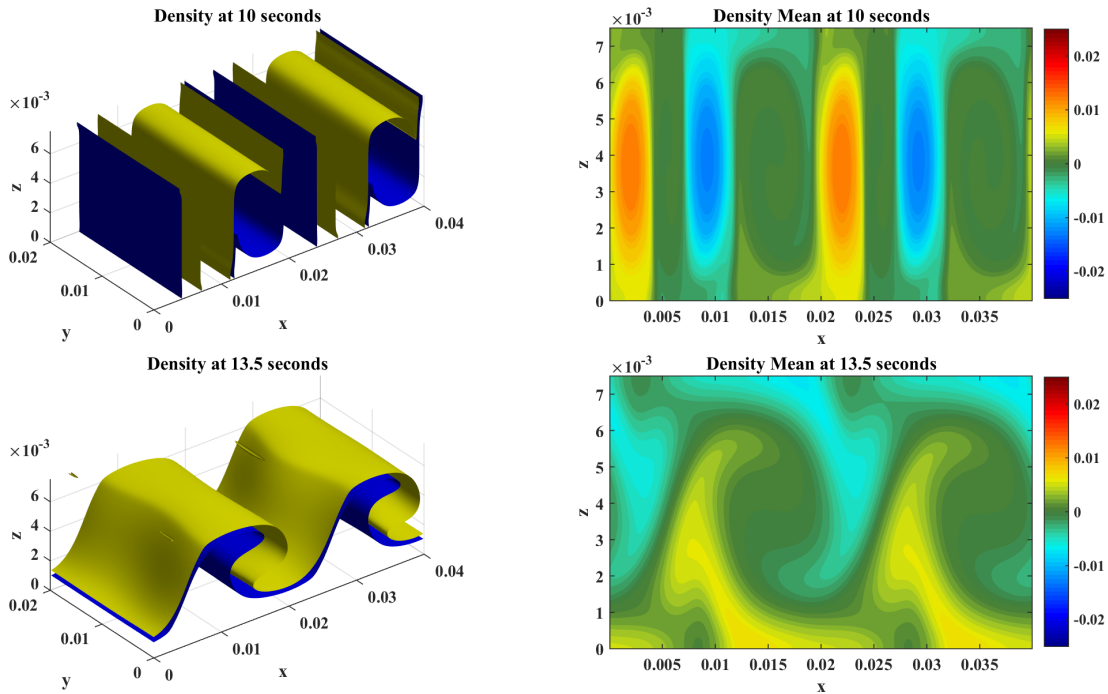


Figure 2.38: Density isosurfaces (left) showing $\pm 10\%$ of max perturbation (ρ') and spanwise mean (right) at 10 (top) and 13.5 (right) seconds. This early part of the simulation serves to demonstrate the drastically different timescale of the linear stratification as well as the consistency with classical theory in describing the initial billow rollup and formation event.

Figure 2.38 shows the isosurfaces and spanwise-averaged density through the domain

for two distinct billowing episodes in the linearly stratified simulation. The initial billow formation occurs much in the same way as the quasi two-layer case, reaching the mature state shown in the top panels of figure 2.38. After a period of quiescence, new disturbances form, grow and eventually form two large billows shown in the bottom panels of figure 2.38.

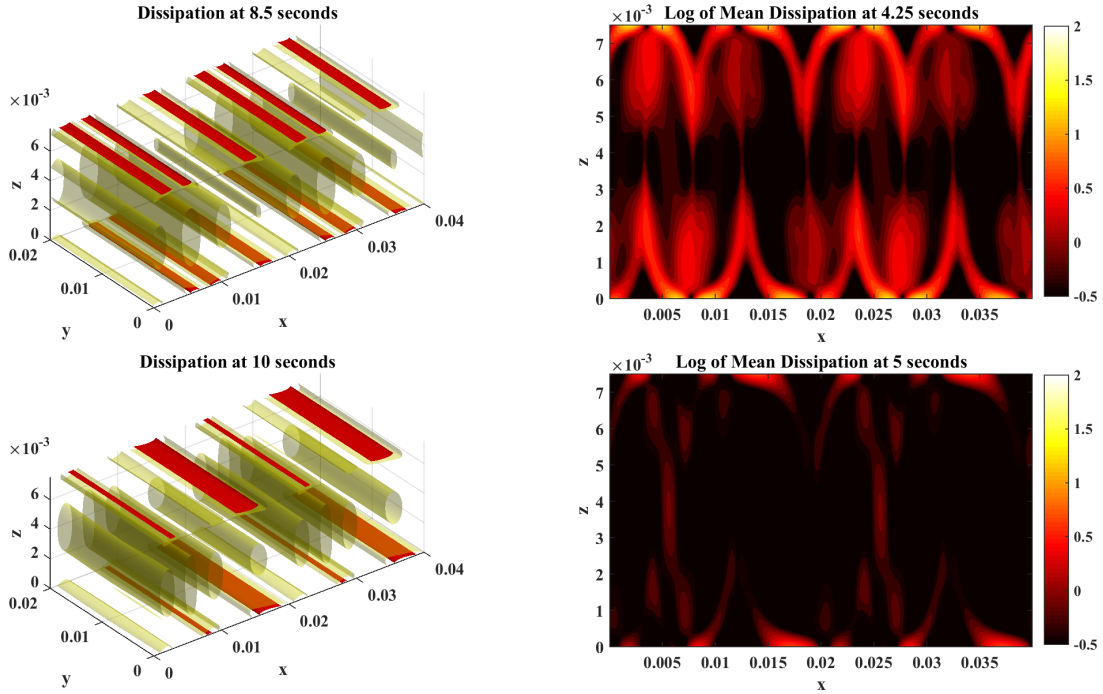


Figure 2.39: Dissipation isosurfaces (left) at 10 and (yellow) and 50 (red) percent of local max. Mean dissipation (right) scaled by one quarter of local maximum dissipation. Both figures are presented at 8.5 (top panels) and 10 (bottom panels) seconds. Early dissipation (top panels) - caused by the initial shear and broadening of the pycnocline - quickly gives way to boundary dominated dissipation once the billows have matured (bottom panels).

If we consider the locations of dissipation throughout the early part of the simulation (see fig. 2.39) we notice immediately that the majority of dissipation occurs near the boundaries (as expected). At early times the next highest regions of dissipation occur near the braids of the initial billows, where the local shear is highest in the interior. Finally there is some dissipation that occurs in the interior of the billows, but it is relatively much weaker than the other two regions described above. This corresponds well to the

classical description of billow formation and evolution, as modified by the specifics of our experimental setup. Classical theory dictates that the braids must contain the highest local shear (and thus the highest dissipation) of the interior of the domain. However, as we may expect, the interaction with no-slip boundaries introduces a physical location where the fluid undergoes extremely large, highly localized velocity changes. Thus we fully expected the near boundary regions to dominate the dissipation plots, especially once the billows have matured and reached the full size of the domain.

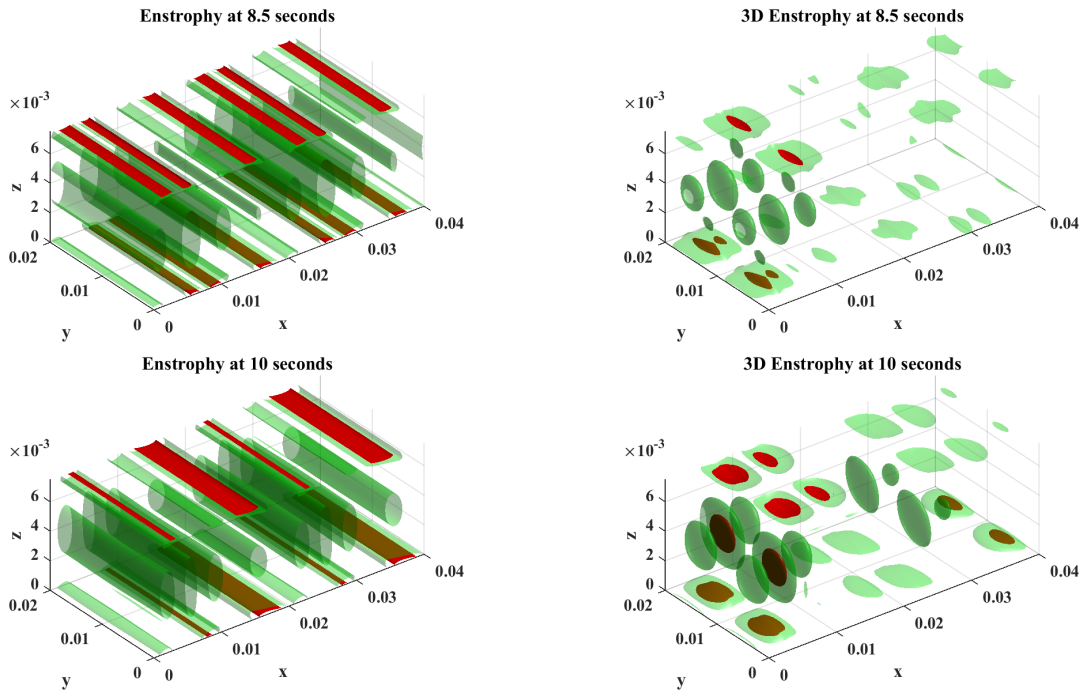


Figure 2.40: Enstrophy (left) and 3D enstrophy (right) isosurfaces at 10 (green) and 50 (red) percent of local max at 8.5 (top panels) and 10 (bottom panels) seconds. Mean dissipation (right) scaled by one quarter of local maximum dissipation. The enstrophy at these times is dominated by the billow cores and spanwise nearly constant, making the 3D enstrophy interesting to look at but not meaningful.

The enstrophy at these times, figure 2.40 - corresponding to the initial rollup - are predictably two-dimensional and as such the enstrophy appears to focus entirely on the newly-formed billow cores. The 3D enstrophy at these times once again appears to contain some interesting structure, however we must remember that the lack of spanwise variation

means that this structure is artificial (likely arising from initial noise seeding) and thus not physically meaningful.

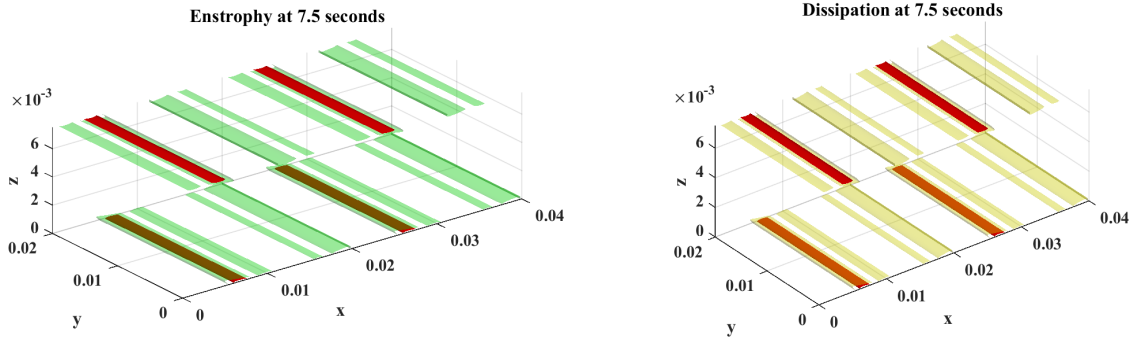


Figure 2.41: Enstrophy (left) and dissipation (right) isosurfaces (10% of local maximum in green and yellow, respectively, and 50% in red) at 7.5 seconds. The vast majority of activity is clearly confined to the top and bottom boundary layers. The induced dissipation and vorticity occur only at the top and bottom boundaries.

We considered isosurfaces of enstrophy. The enstrophy, \mathcal{E} , at early times corresponded almost exactly to the dissipation, and as such we will omit showing all the early time figures here, however, figure 2.41 shows the enstrophy (top) and dissipation (bottom) as a representative example of results at an early time. Note that the two fields are indistinguishable at this time. At later times we will be interested in the spanwise component of enstrophy as a measure of three-dimensionalization.

This, relatively early, part of the simulation shows that the finite amplitude two-dimensional billows begin to show some destabilization in the spanwise direction. While at these times, this spanwise structure has a negligible effect on the features of the flow, it is worth noting its presence.

Progressing to later times, we observe the billows grow to the size of the domain, destabilize in the spanwise direction and collapse into a pseudo-mixed state (see figure 2.42). At this point there is little to no trace left of the original billows. It is worth noting at this time that a significant amount of mixing has taken place. The mean density plot, saturated at the original maximum values, shows that the flow has become significantly mixed. This is also evidenced by the total mixing line in figure 2.44. The final nail in the coffin of the original billows is the total absence of significant enstrophy in the interior of the domain.

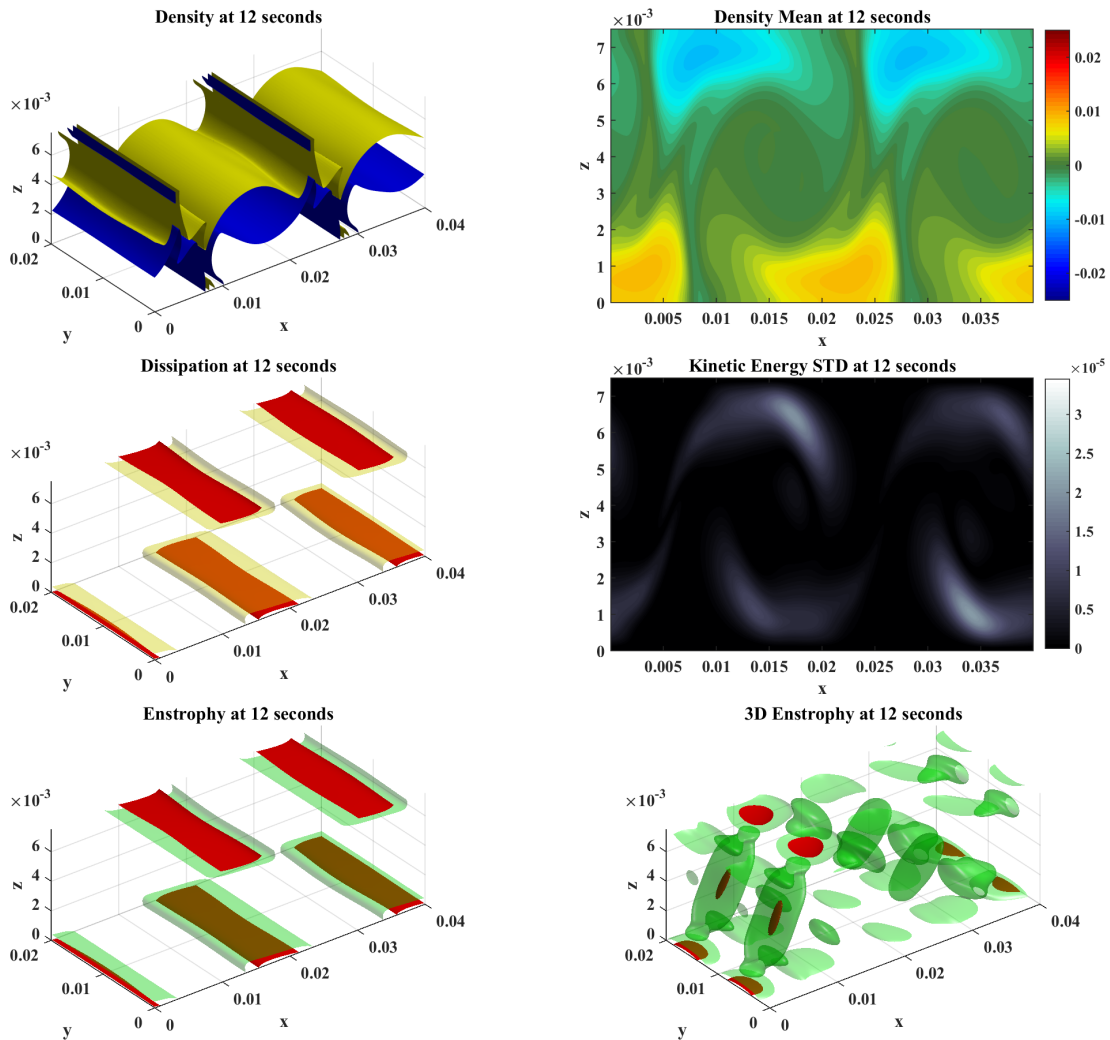


Figure 2.42: Density isosurfaces (top-left), spanwise-averaged density (top-right), dissipation isosurface (middle-left), standard deviation of kinetic energy saturated at 10% of max (middle-right), and estrophy and 3D enstrophy isosurfaces (bottom left and right) at 12 seconds. Note that by this time the linear simulation has reached a quasi-mixed state, with no distinct billows or braid-core structures visible any longer.

Once a significant amount of mixing and three-dimensionalization has occurred in the simulation the flow becomes more regular (through viscous dissipation). However here we

witness something unique to the linearly stratified case (figure 2.43). The flow experiences a second (albeit weaker) instability and rollup. Interestingly, this secondary instability is inherently three-dimensional, with spanwise velocities consistently contributing a significant amount to the kinetic energy of the system.

Though the averaged density does not display the classical billow shapes, it is clear from the isosurfaces that the flow has developed a secondary instability. The standard deviation of kinetic energy (bottom-right panel) clearly demonstrates that this secondary instability is fundamentally three-dimensional in nature. Finally, the dissipation isosurfaces (middle and bottom left panels - again at 10 and 50%) show that the majority of intense dissipation and enstrophy during the secondary event occurs strictly near the boundaries. Somewhat interestingly, and more importantly contradictory, is that the 3D enstrophy does appear to have some structure in the interior. Although the values of enstrophy are clearly very low in the interior, the regular structure suggests some innate three-dimensionalization - a fact which is visually confirmed by the small deformations of the density isosurfaces in the spanwise direction.

To quantify this result, the mean streamwise and across-channel velocities during the initial rollup event are on the order of $7 \times 10^{-3} \frac{m}{s}$, whereas the spanwise velocities are on the order of $6 \times 10^{-5} \frac{m}{s}$ - less than one percent of the horizontal velocity. In contrast, during the second billowing event (around 10-14 seconds) the typical stream wise velocities are around $8 \times 10^{-3} \frac{m}{s}$, while the across channel and spanwise velocities are comparable, around $8 \times 10^{-4} \frac{m}{s}$. These values indicate a fundamentally three-dimensional regime. The much larger average (almost 1cm/s) stream wise velocities are a result of the nature of the experiment. The flow must be primarily dominated by the lighter fluid moving upstream, and the heavier moving downstream, a factor which creates shear and is the driving force behind billow formation. Thus we are not surprised by the 10 to 1 ratio of along channel to across channel velocity.

Fig. 2.44 shows the total viscous dissipation (scaled by the maximum achieved over the length of the simulation) and spanwise standard deviations of kinetic energy and density. This plot demonstrates a key aspect of the simulation. Namely, the bulk of the viscous dissipation occurs during the initial (two-dimensional) billowing event. This is also where the majority of the mixing takes place. It should be clear from the spanwise standard deviation of kinetic energy and density histories (blue and black lines) that this event takes place almost entirely in two dimensions. The secondary and tertiary rollup events can be seen in the standard deviation of the kinetic energy history (blue) at 14 and 20 seconds, respectively. These secondary rollup events are inherently three-dimensional, as can be seen from the density isosurfaces in the lower panels of fig. 2.38. However, while the secondary (and tertiary) rollup events produce regions of intense local dissipation, it

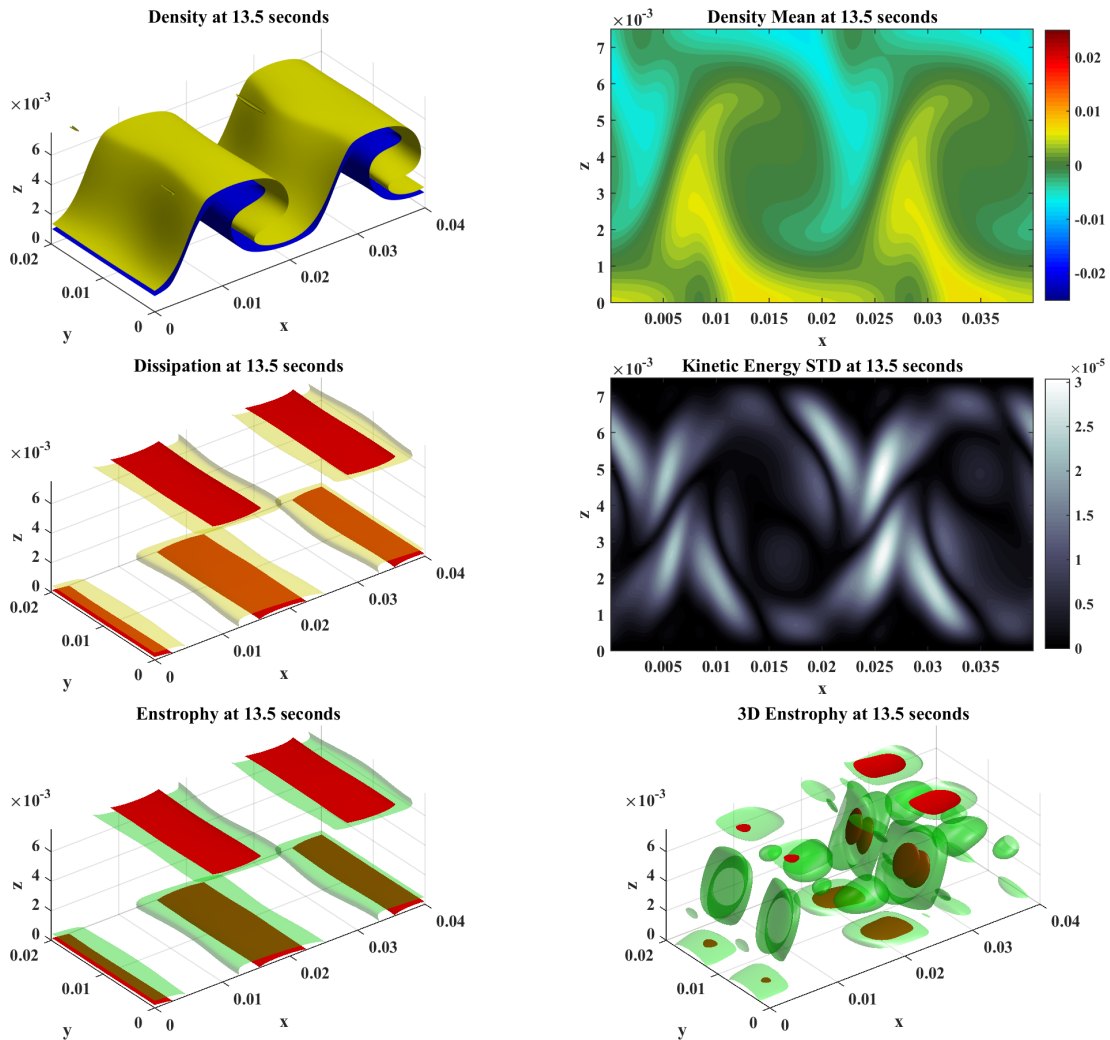


Figure 2.43: Density isosurfaces (top-left), spanwise-averaged density (top-right), dissipation isosurface (middle-left), standard deviation of kinetic energy saturated at 10% of max (middle-right), and enstrophy and 3D enstrophy isosurfaces (bottom left and right) at 13.5 seconds. This corresponds well to the peak of the secondary instability which develops in the linearly stratified case.

is clear that the system undergoes dramatically less mixing at these subsequent events.

An important aspect of this system is that while there are several discrete billow-forming




Figure 2.44: Summary plot showing total kinetic energy, total mixing, and dissipation (both internal and total) scaled by the maximum values achieved over the length of the simulation (top) and a semi-log plot of the same variables (bottom). The majority of the interesting dynamics (dissipation, mixing, and the bulk of the enstrophy generation) occurs during the initial rollup event, around 3.5 seconds. Notice, however, the secondary event at 12-14 seconds which causes a significant rise in relative mixing as well as a small bump in dissipation and kinetic energy.

events, the system is changed drastically and irreversibly after the first event. The first billow formation event occurs largely in two dimensions, and is responsible for the vast majority of the mixing in the system.

The locations, timing, and magnitude of dissipation are especially important when considering the mixing and three-dimensionalization in the system. It is clear from the line

plot in fig. 2.44 that the dissipation takes place during the early stages of the simulation - in the initial rollup event. What is less clear, but corresponds to our intuition, is that the majority of mixing also takes place during this initial event. The initial rollup event occurs primarily in two dimensions (very little spanwise velocity) and induces the greatest amount of dissipation and mixing. This is clearly visible via the red and green lines in fig. 2.44. The dissipation and mixing peak sometime during the initial billowing event (entirely in two-dimensions). It is worth noting that the mixing peaks slightly after the dissipation. The most likely explanation for this is that the billows quickly grow to the size of the (vertical) domain and interact with the boundary to enhance dissipation, while the magnitude of the density gradient peaks during the period of strongest billow formation sometime after.

2.4.4 Discussion

The fully three dimensional simulations presented above have several notable features worth discussing. The first is that even on sub-centimeter scales and in the presence of boundaries the initial evolution of the billows (the formation event) occurs exactly as described by classical theory of shear instability and transition. Moreover, the three-dimensional results agree almost exactly with our two-dimensional simulations up to the point when the flow destabilizes and becomes fully 3D. This means that for the study of the initial rollup and billow formation (even in the presence of boundaries), two-dimensional simulations are more than sufficient. Additionally, this is important because it indicates that the initial formation is largely unaffected by the boundary layers that form as the fluid accelerates.

The initial billow formation event occurs primarily in two dimensions (with very low spanwise velocities) and yet is responsible for the majority of the mixing and dissipation throughout the simulation. While two dimensional mixing layers have a rich mathematical literature (see [14, 32, 67, 100] for experimental and numerical discussions) and some application in low Reynolds number fluids (such as two-dimensional turbulence in soap films [35, 78]), it is a central aspect of the theory of stratified shear transition that the onset of three dimensionalization leads to a large increase in the rate of mixing. The linearly stratified simulation presented above is in blatant contradiction of this commonly held view ([15, 72]) that mixing is a three-dimensional phenomenon. As we can see in the line plots (see figures 2.29, 2.36, and 2.44) the mixing and dissipation occur strongly during the initial billowing event across our simulations, but is most dramatic in the linear case where the primary event is responsible for dramatically more mixing than subsequent instabilities. While the secondary and tertiary events in the linear case are fully three dimensional in nature - as opposed to the entirely two-dimensional initial rollup - but

there is very little mixing or dissipation induced in the interior of the domain after the initial billow formation event which occurs (mostly) in two dimensions.

We have seen from these simulations of quasi two-layer and linearly stratified flow that the evolution of Kelvin-Helmholtz billows follows the classically prescribed path [14]: braids and cores develop first in two dimensions, instabilities arise in the spanwise direction and the flow becomes fully three-dimensional (and quasi-turbulent). An important, if somewhat overlooked, feature of confined billows is that the majority of mixing and dissipation in the simulation takes place during the initial rollup event, occurring primarily in two dimensions. The subsequent billowing events, which are intrinsically three-dimensional from the onset, do continue to contribute some mixing of the fluid as well as the generation of regions of high dissipation near the boundaries, however these are dwarfed in intensity and effect by the initial events.

The other result of note is that unlike the classical Kolmogorov-inspired view of fully developed turbulence, the linearly stratified simulation is dominated by successive discrete events that mix fluid, dissipate energy and cascade energy to shorter scales. Interestingly, the secondary and tertiary rollup events do not appear to become energetic enough to excite large, channel filling billows. Moreover, after each billowing event the system seems to settle down for a finite interval of time before the next event takes place. As expected, with each event the system becomes more mixed, the stratification is decreased, and the subsequent rollup events are less violent than the ones preceding. This is fundamentally different from the quasi two-layer simulations we have performed in which the initial billowing event led to strong three-dimensionalization and mixing so that the system settled to something close to a quasi-steady, mixed state.

An interesting visualization phenomenon we discovered was in attempting to acquire a measure of three-dimensionality through 3D enstrophy. We discovered that separately enstrophy or 3D enstrophy yielded some useful information (such as the locations of vortex-dominated billows or regions of destabilization, however when combined it allowed us to gather a much more complete picture of the state of the dynamics of the system. These two metrics when combined demonstrated where in the flow the enstrophy dominated and how susceptible it was to destabilization.

Finally, while the effects of confinement on the linearly stratified shear flow are less prominent than in the two-layer cases, it is clear (see fig. 2.42 and 2.43) that interaction with the boundary layer does occur and induces significant regions of dissipation (and to a lesser degree, mixing). It is interesting to consider how the physically dominant interaction with the boundaries is expressed in more sophisticated analysis techniques used for turbulent flows (namely the Q and R metrics), and this will be taken up in the

following chapter.

Prior to continuing on to examine the behaviour of our confined billows using turbulent metrics, it may be worth pausing to consider the implications of the behaviour we have observed here. The aspect of our investigations which makes them fundamentally unique is the proximity of the no-slip boundaries and the change in flow dynamics that arises from this interaction with the boundary layer. Classical turbulence theory [25, 55] tells us that the boundary layer thickness scales as the square root of the far-field velocity (\sqrt{U}). This means that for flows with a vertical length scale on the order of this thickness either due to extreme confinement (as we have explored here) or large velocities we expect the presence of boundaries to impact the evolution of the flow.

In the case of KH billows we have seen that the initial billow formation occurs in a manner that is mostly independent of the presence of the boundaries, but later aspects, such as the secondary (spanwise) destabilization of billows is fundamentally altered by the interaction with the boundary. Turning to the physical implications of this result we may consider several real-world scenarios which could produce similar (though likely not identical) results.

Chief among our considerations is the natural exchange flow problem (e.g. that leads to the formation of salt wedges in rivers) where high density (salty or cold) fluid intrudes into a lighter fluid region, whether caused by wind-driven effects as in the Hamilton harbour [56], or by effects of the tide such as in the Frasier estuary [91, 103]. A key difference between this type of (fairly well-studied) effect and the one we have presented here is the nature of the instability. In the papers mentioned above all the measurements indicate that the proximity of boundaries is a possible source of altered dynamics, however it is in a way that is different from the confined simulations presented herein. All three papers discuss the effects of what is termed “bottom friction” on the reduced exchange flow and pointed to an large amount of mixing which cannot be predicted by a linear theory such as the Taylor Goldstein equations. Additionally the resulting instabilities differ somewhat from classical KH billows, likely due to the close proximity of a single boundary which will tend to prevent the formation of fully mature billows and instead cause the type of “one-sidedness” which is a hallmark of the salt wedge exchange flow.

It is not immediately clear whether our investigations contain a direct parallel to ocean or lake-scale measurements, however it is certainly worth further thought to determine if there are lessons learned from our very small scales which can be easily translated to large scale behaviour.

Before concluding this chapter we will return to the question of the somewhat unexpected behaviour found in the asymmetric 1 and 0.75 cm simulations. Both of these

domains experienced the creation of a region of near-constant fluid which appeared to exist alongside the primary billows. The exact nature of these regions was unclear, though it appeared to have some relation to the production of dissipation and enstrophy at the far boundary. This region demonstrates an area which merits closer inspection, but also identifies a clear weakness in our analysis of the system: we are unable at a glance to determine the dominant mechanics throughout the domain. Additionally we currently have no good way to identify the creation of coherent structures (which is what this region will be shown to be in the following chapter). As such, we will present a suite of tools, using tertiary variable analysis, of classifying flow topology, identifying regions of interest, and facilitating the analysis of complex flows.

Chapter 3

Q and R: Tertiary Variable Analysis

3.1 Analysis Tools

Direct numerical simulation (DNS), and our chosen tool for it (SPINS), is a very powerful tool which allows for the accurate simulation of two and three-dimensional fluid flows (in our case studies satisfying the Boussinesq approximation). One of the biggest advantages of SPINS is that the simulations generate full density, velocity, and pressure fields for the flow at any specified output time in a wide variety of geometric configurations and boundary condition combinations. This does, however, present us with a very large amount of information which must be presented in the most efficient and informative manner possible.

To this end, this chapter deals with Q and R , a set of tertiary variables, as a tool for the analysis of a particular fluid flow and its origin, as well as links to currently established two-dimensional turbulence literature (namely the Okubo-Weiss criterion). We begin by introducing the tertiary variables as they are derived from the eigenvalues of the velocity gradient tensor (following [25]). We then proceed to apply these tools to some canonical flows, such as simple two-dimensional channel flow and the more complex Burgers and Taylor-Green vortices. We then proceed to a discussion of visualization tools for Q and R and come up with a (somewhat heuristic) approach to beginning our analysis. Finally, we will present several case studies, namely the fully three-dimensional Kelvin-Helmholtz billow simulations from the previous chapter, and a new case study, the Rayleigh-Taylor of a falling plug of dense fluid. For the latter case we demonstrate the value of $Q - R$ analysis, when applied to a system we know very little about *a priori*.

We discussed in the introduction (section 1.5.1) the nature of the tertiary variables Q and R and their use as a diagnostic tool for the assessment of a flow. We will now focus on the use of these turbulent metrics as a means of analyzing the flow. However, first we will begin with a brief discussion of the nature of Q and R as a metric, as well as some of the drawbacks of Q vs R plots and analysis tools.

3.2 Nonuniqueness

When first considering the tertiary variables Q and R it is natural to wonder whether Q and R are capable of uniquely specifying a flow type. To this end we present a simple sample flow as a counter-example. The purpose of this exercise is to demonstrate, using a toy problem, the limitations of Q - R and the need for metrics beyond Q and R in order to fully characterize the flow.

We begin with the definition of the invariants:

$$\begin{aligned} P &= \lambda_1 + \lambda_2 + \lambda_3 = \text{trace}(A) = u_x + v_y + w_z = 0 \\ Q &= \lambda_1\lambda_2 + \lambda_2\lambda_3 + \lambda_1\lambda_3 \\ R &= \lambda_1\lambda_2\lambda_3, \end{aligned}$$

where λ_i is the i th eigenvalue of the velocity gradient tensor. For simplicity (and without loss of generality) let $\lambda_1 = 1$. This simplifies our invariants to

$$\begin{aligned} -1 &= \lambda_2 + \lambda_3 \\ Q &= \lambda_2 + \lambda_2\lambda_3 + \lambda_3 = \lambda_2\lambda_3 - 1 \\ R &= \lambda_2\lambda_3. \end{aligned}$$

This leads to two interesting results. The first is that we obtain the relationship $Q = R - 1$. The second is that we can write

$$R = -\lambda_2(1 + \lambda_2).$$

These two properties combined imply that multiple eigenvalues (and thus multiple flows) can yield the same sets of Q - R data. For example all flows for which $\lambda_2 = 0$ yield $R = 0$ and $Q = -1$. Similarly all flows for which $\lambda_2 = -1$ yield $R = 0$ and $Q = -1$. While this example is somewhat contrived, it does definitively show that while Q and R can be used to characterize a flow, they do not uniquely define it. Namely, two similar flows may have different Q vs R profiles if their underlying mechanics are different, while two very different flows can have nearly identical Q vs R profiles if the dominant system mechanics are the same.

3.3 Canonical Flows

In order to gain a better understanding of how the tertiary variables Q and R can be used to analyze a complex flow we will apply this method of analysis to some standard flows. We begin by investigating simple 2D flows in order to gain intuition on the derivation and predictive ability of the tensor invariants before considering more complex analytic solutions, such as the Burgers and Taylor-Green vortices.

3.3.1 Two-Dimensional Analysis

The tertiary variables P , Q , and R are defined as the invariants of the velocity gradient tensor, $A_{ij} = \frac{\partial u_i}{\partial x_j}$, whose eigenvalues (λ_i) satisfy the characteristic equation

$$\lambda^3 + P\lambda^2 + Q\lambda - R = 0.$$

This characteristic equation only arises in three dimensions. In 2D we are left with the equation

$$\lambda^2 + P\lambda - Q = 0,$$

which has only two invariants. This means that as written, QR analysis can only be applied to 3D flows. In order to get around this we will consider not only the algebraic definition of Q , but also the expansion of the meaning of the invariants in terms of the shear and vorticity.

Since the velocity gradient matrix can be written as the sum of the rate of strain and rotation matrices as

$$A_{ij} = S_{ij} + W_{ij},$$

we can write the invariants [25] Q and R as

$$\begin{aligned} Q &= -\frac{1}{2}S_{ij}S_{ji} + \frac{1}{4}\boldsymbol{\omega}^2 = \frac{1}{2}(a^2 + b^2 + c^2) + \frac{1}{4}\boldsymbol{\omega}^2, \\ R &= \frac{1}{3}(S_{ij}S_{jk}S_{ki} + \frac{3}{4}\omega_i\omega_j S_{ij}) = \frac{1}{3}(a^3 + b^3 + c^3) + \frac{1}{4}\omega_i\omega_j S_{ij}. \end{aligned}$$

Here a , b , and c are the eigenvalues of the strain tensor, S_{ij} . Thus Q allows us to determine the relative strength of the strain and vorticity present.

In two dimensions we only get access to two invariants, P and Q . The first invariant, P , is still equal to the trace of the matrix A , and as such $P = \nabla \cdot \mathbf{u} = u_x + v_y = 0$. Since

we only have two invariants in two dimensions Q takes the place of R in three dimensions as the coefficient of the ones (or the identity matrix) in the characteristic polynomial.

By the Cayley-Hamilton theorem [68] we know that the matrix must satisfy its characteristic equation. As such in 2D we have

$$A^2 - QI = 0.$$

The trace of this equation gives

$$Q = \frac{1}{2}A_{ij}A_{ji},$$

Which is the same result obtained in three dimensions for the variable Q . This means that in two dimensions our Q variable can be written as

$$Q = -\frac{1}{2}S_{ij}S_{ji} + \frac{1}{4}\omega^2 = \frac{1}{2}(a^2 + b^2) + \frac{1}{4}\omega^2,$$

Which once again gives the balance between shear and vorticity.

In the oceanography (and two-dimensional turbulence) literature there is a parameterization used to characterize the topology of the flow field. By looking at the shape of trajectories of two particles initially close together, Okubo speculated by means of linear analysis that the relationship could be defined via the eigenvalues of the velocity gradient tensor [69]. In 1991 Weiss [98], analyzing two-dimensional turbulence, developed a criterion for the flow topology based on the sign of

$$W = s_n^2 + s_s^2 - \omega^2,$$

where s_n , s_s , and ω are the normal and shear components of strain and the vorticity respectively defined by

$$\begin{aligned} s_n &= \frac{\partial u}{\partial x} - \frac{\partial v}{\partial y}, \\ s_s &= \frac{\partial v}{\partial x} + \frac{\partial u}{\partial y}, \\ \omega &= \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}. \end{aligned}$$

The parameter, W , is called the Okubo-Weiss parameter, and separates the flow into two regions: vorticity-dominated ($W < 0$) and strain-dominated ($W > 0$). For practical (experimental) applications the critical value used is sometimes set as a threshold parameter, W_0 , equal to a scaling of the standard deviation of W , σ_W .

In two dimensions using subscript notation for partial derivatives, we have the invariant $Q = A_{ij}A_{ji} = \det(A) = -u_x^2 - u_y v_x = -v_y^2 - u_y v_x = -\frac{1}{2}u_x^2 - \frac{1}{2}v_y^2 - u_y v_x$. In contrast, the Okubo-Weiss parameter can be written as

$$\begin{aligned} W &= (u_x - v_y)^2 + (u_y + v_x)^2 + (u_y - v_x)^2 \\ &= u_x^2 + v_y^2 - 2u_x v_y + 4v_x u_y \\ &= 2\left(\frac{1}{2}u_x^2 + \frac{1}{2}v_y^2 + u_y v_x\right) + 2(u_x v_y + v_x u_y) \\ &= -2Q + 2(-Q) \\ &= -4Q. \end{aligned}$$

As we have clearly demonstrated above, the Okubo-Weiss parameter is simply a scaling of the second invariant, Q , in two dimensions. This invariant is used to determine the relative strength of strain and vorticity, which can in turn be used to characterize the flow topology. It may be worth remarking that although the 2D literature (see [40, 39, 43, 44] and references therein) refers to the regimes of W (and thus Q) as being able to categorize flow topology, there is really very little in the way of real topological differences present in two dimensions. The only regimes worth considering are whether the fluid is undergoing predominantly strain, vorticity, or neither. In three dimensions the added invariant R adds significant complication (see the intro to this chapter for details) in analyzing the types of dominant mechanisms, such as vortex stretching or compression, plane shear or line shear, and other, more difficult to categorize, regimes.

Despite the relative simplicity of two-dimensional flow topology it is worth considering what this parameterization says about some canonical flows. Thus we will investigate the irrotational and solid-body vortices, internal wave in a channel (sometimes called a translating baroclinic vortex) as well as canonical three-dimensional flows such as the Burgers and Taylor-Green vortices.

3.3.2 Irrotational Vortex

The irrotational vortex is one of the first simple examples taught in any introductory fluids course. This simple potential flow is a vortex which has no vorticity (by definition of a potential flow), or put another way all the vorticity is concentrated at a singularity (commonly placed at the origin), a non zero shear stress, but a zero net viscous force [55]. This provides us with an excellent test case, as the diagnostic, Q , should be able to clearly show that there is no vorticity, and only strain present in the resulting velocity field.

The streamfunction is given by

$$\psi = -\frac{\Gamma}{2\pi} \ln(\sqrt{x^2 + y^2}),$$

which yields the velocity field

$$u = -\frac{\Gamma y}{2\pi(x^2 + y^2)} \quad v = -\frac{\Gamma x}{2\pi(x^2 + y^2)}.$$

We can now compute the velocity gradient matrix:

$$A = \begin{pmatrix} \frac{\Gamma xy}{\pi(x^2+y^2)^2} & \frac{-\Gamma(x^2-y^2)}{2\pi(x^2+y^2)^2} \\ \frac{-\Gamma(x^2-y^2)}{2\pi(x^2+y^2)^2} & -\frac{\Gamma xy}{\pi(x^2+y^2)^2} \end{pmatrix}.$$

Which results in the relatively simple calculation

$$Q = -\frac{\Gamma^2}{4\pi^2 (x^2 + y^2)^2}.$$

This result tells us something that we pretty much expected, namely that Q is strictly negative. There is no vorticity, as can be verified by $\omega = u_y - v_x = 0$ (as is expected of the irrotational vortex), and thus the only mechanism present is strain (or shear).

3.3.3 Solid Body Rotation

A solid body vortex can be created by steadily rotating a cylindrical tank containing viscous fluid. Once the transients die out, the resulting velocity is proportional to the radius of the streamlines. If the angular velocity of the rotation is ω_0 the velocity field is given by

$$u = -\omega_0 y \sqrt{x^2 + y^2} \quad v = \omega_0 x \sqrt{x^2 + y^2}.$$

This yields the velocity gradient matrix

$$A = \begin{pmatrix} -\frac{\omega_0 xy}{\sqrt{x^2+y^2}} & -\frac{\omega_0(x^2+2y^2)}{\sqrt{x^2+y^2}} \\ \frac{\omega_0(2x^2+y^2)}{\sqrt{x^2+y^2}} & \frac{\omega_0 xy}{\sqrt{x^2+y^2}} \end{pmatrix}.$$

While somewhat messy, this gives the relatively neat result that $Q = \det(A)$ is given by

$$Q = 2\omega_0^2(x^2 + y^2).$$

In contrast to the irrotational vortex above, since Q is strictly positive we have an example of a flow with vorticity and no shear - as expected from our initial setup.

Now that we have confirmed that our analysis is capable of correctly identifying strictly rotational or strictly shear regimes we can consider a more complex canonical example.

3.3.4 Internal waves in a Channel

A baroclinic torque occurs when the gradients of pressure and density are not collinear. As a result, waves in a stratified fluid produce vorticity [55], while plane internal waves propagating in an arbitrary direction have the exotic property that energy propagates in a direction that is perpendicular to the motion of individual crests, internal waves propagating horizontally in a channel are the simplest example of an internal wave since they are in some sense analogous to surface waves. The vertical structure of the velocity field can be solved for from a simple linear eigenvalue problem. The flow in a channel is confined to the x and z directions. The presence of no flux boundary conditions at $z = 0$ and $z = H$ allows for a separation of variables for the vertical and horizontal structure, as well as a solution for the dispersion relation. When the fluid is linearly stratified the solution can be given in closed form. Due to the baroclinic generation of vorticity, this type of flow is also sometimes referred to as a translating baroclinic vortex.

The streamfunction is given by

$$\psi = \cos(kx - \sigma t) \sin(mz),$$

where N_0 is the constant buoyancy frequency, k and m are the horizontal and vertical components of the wavenumber vector, respectively, and $\sigma = \frac{kN_0}{\sqrt{k^2+m^2}}$ is the dispersion relation. The boundary conditions imply that m is quantized as $m = n\pi/H$.

Differentiating the stream function yields the velocity field

$$\begin{aligned} u &= m \cos(kx - \sigma t) \cos(mz), \\ w &= k \sin(kx - \sigma t) \sin(mz). \end{aligned}$$

Thus we get the velocity gradient matrix

$$A = \begin{pmatrix} -km \sin(kx - \sigma t) \cos(mz) & -m^2 \cos(kx - \sigma t) \sin(mz) \\ k^2 \cos(kx - \sigma t) \sin(mz) & km \sin(kx - \sigma t) \cos(mz) \end{pmatrix}.$$

The eigenvalues are given by

$$\begin{aligned} \lambda_1 &= \frac{1}{2} \left(k^4 (\cos(kx - \sigma t))^2 (\sin(mz))^2 + 2k^2 (\cos(kx - \sigma t))^2 (\sin(mz))^2 m^2 \right. \\ &\quad \left. + (\cos(kx - \sigma t))^2 m^4 (\sin(mz))^2 - 4 (\cos(kx - \sigma t))^2 k^2 m^2 \right. \\ &\quad \left. - 4k^2 m^2 (\sin(mz))^2 + 4k^2 m^2 \right)^{\frac{1}{2}}, \end{aligned} \tag{3.1}$$

and

$$\begin{aligned} \lambda_2 = & -\frac{1}{2} \left(k^4 (\cos(kx - \sigma t))^2 (\sin(mz))^2 + 2k^2 (\cos(kx - \sigma t))^2 (\sin(mz))^2 m^2 \right. \\ & + (\cos(kx - \sigma t))^2 m^4 (\sin(mz))^2 - 4 (\cos(kx - \sigma t))^2 k^2 m^2 \\ & \left. - 4k^2 m^2 (\sin(mz))^2 + 4k^2 m^2 \right)^{\frac{1}{2}}. \end{aligned} \quad (3.2)$$

Now we can calculate Q exactly via the determinant of A , or the formula $Q = \frac{1}{2}(-a^2 + b^2) + \omega^2$, where a, b are the eigenvalues of the strain tensor, S , to find that

$$Q = -k^2 m^2 (\cos(mz)^2 - \cos(kx - \sigma t)^2) \quad (3.3)$$

This representation of Q allows us to observe that the vertical structure, given by $-k^2 m^2 \cos(mz)^2$ provides a strictly negative contribution to Q - denoting rate of strain, while the horizontal component $k^2 m^2 \cos(kx - \sigma t)^2$ is a positive contribution, denoting vorticity. From this analysis it should be clear that this relatively simple system undergoes both regions of fine rate of strain, and finite vorticity distributed throughout the domain.

We can gain some physical intuition for this result from the equations of motion. If we linearize for simplicity of algebra, we get the velocity evolution equations:

$$\begin{aligned} u_t &= -\frac{p_x}{\rho_0}, \\ w_t &= -\frac{p_z}{\rho_0} - \frac{\rho' g}{\rho_0}. \end{aligned}$$

Cross-differentiating, we can derive the vorticity evolution equation:

$$(u_z - w_x)_t = \omega_t = \frac{\rho'_x g}{\rho_0}. \quad (3.4)$$

The above equation clearly shows that the rate of change of vorticity depends strictly on the horizontal changes in ρ' .

Thus, computing the Q variable provides a slight twist on the classical vorticity analysis. The added benefit is that in higher dimensions (and more complex flows) $Q - R$ analysis provides significantly more information, as we will see when examining more complex 3D canonical flows, as well as in the case studies to follow.

3.3.5 Burgers Vortex

The first canonical three-dimensional flow we wish to examine here is the Burgers vortex. This flow is radially symmetric and exhibits the mechanism of vortex stretching (in the z direction). The primary mechanisms at work are a concentration of vorticity in a narrow column near the $r = 0$ axis balanced by viscous diffusion, which tends to disperse vorticity in the radial direction.

This example is an exact solution to the steady Navier-Stokes equation and is thought to be an intrinsic part of the classical turbulent cascade, occupying the region where the Kolmogorov $-\frac{5}{3}$ law is effective [25].

The velocity field for Burgers vortex is given in cylindrical coordinates - (ρ, θ, z) - by

$$u_\rho = -\frac{\alpha}{2}\rho, \quad (3.5)$$

$$u_\theta = \frac{\Gamma}{2\pi\rho} \left(1 - \exp \left[-\frac{\alpha\rho^2}{4\nu} \right] \right), \quad (3.6)$$

$$u_z = \alpha z. \quad (3.7)$$

This velocity field has several key parameters. The first is Γ , the strength of the vortex. The second is α , which gives a measure of the radial stretching term - i.e. how much vortex stretching occurs. Finally we have the kinematic viscosity, ν , which gives a measure of the diffusion present.

Since the velocity field can be written in the closed form above, we can compute the velocity gradient matrix exactly as

$$A = \begin{pmatrix} -\frac{1}{2}\alpha & -\frac{1}{4}\frac{\Gamma}{\pi\rho^2\nu} \left(-\alpha\rho^2 + 2e^{\frac{1}{4}\frac{\alpha\rho^2}{\nu}}\nu - 2\nu \right) e^{-\frac{1}{4}\frac{\alpha\rho^2}{\nu}} & 0 \\ -\frac{1}{2}\frac{\Gamma}{\rho^2\pi} e^{-\frac{1}{4}\frac{\alpha\rho^2}{\nu}} \left(e^{\frac{1}{4}\frac{\alpha\rho^2}{\nu}} - 1 \right) & -\frac{1}{2}\alpha & 0 \\ 0 & 0 & \alpha \end{pmatrix}$$

We can now write $A = S + W$, the symmetric and anti-symmetric stress and rotation matrices and compute the eigenvalues S :

$$\begin{aligned} a &= \alpha \\ b &= \frac{1}{8} \frac{1}{\pi\nu\rho^2} \left(\Gamma\alpha\rho^2 e^{-\frac{1}{4}\frac{\alpha\rho^2}{\nu}} - 4\pi\alpha\nu\rho^2 + 4\Gamma\nu e^{-\frac{1}{4}\frac{\alpha\rho^2}{\nu}} - 4\Gamma\nu \right) \\ c &= -\frac{1}{8} \frac{1}{\pi\nu\rho^2} \left(\Gamma\alpha\rho^2 e^{-\frac{1}{4}\frac{\alpha\rho^2}{\nu}} + 4\pi\alpha\nu\rho^2 + 4\Gamma\nu e^{-\frac{1}{4}\frac{\alpha\rho^2}{\nu}} - 4\Gamma\nu \right) \end{aligned}$$

Note that one of the eigenvalues corresponds to the stretching term, α , which is intuitively obvious since the only motion in the vertical direction is governed by this term.

Additionally, we can compute the enstrophy of the system:

$$\omega^2 = \frac{1}{32} \frac{\alpha^2 \Gamma^2}{\pi^2 \nu^2} \left(e^{-\frac{1}{4} \frac{\alpha \rho^2}{\nu}} \right)^2.$$

From these quantities Q and R can be computed exactly.

$$Q = \frac{-1}{8\nu \pi^2 \rho^4} \left(6 \alpha^2 \pi^2 \rho^4 \nu + \Gamma^2 e^{-\frac{1}{2} \frac{\alpha \rho^2}{\nu}} \alpha \rho^2 - \Gamma^2 e^{-\frac{1}{4} \frac{\alpha \rho^2}{\nu}} \alpha \rho^2 + 2 \Gamma^2 e^{-\frac{1}{2} \frac{\alpha \rho^2}{\nu}} \nu - 4 \Gamma^2 e^{-\frac{1}{4} \frac{\alpha \rho^2}{\nu}} \nu + 2 \Gamma^2 \nu \right) \quad (3.8)$$

$$R = \frac{-\alpha}{8\nu \pi^2 \rho^4} \left(-2 \alpha^2 \pi^2 \rho^4 \nu + \Gamma^2 e^{-\frac{1}{2} \frac{\alpha \rho^2}{\nu}} \alpha \rho^2 - \Gamma^2 e^{-\frac{1}{4} \frac{\alpha \rho^2}{\nu}} \alpha \rho^2 + 2 \Gamma^2 e^{-\frac{1}{2} \frac{\alpha \rho^2}{\nu}} \nu - 4 \Gamma^2 e^{-\frac{1}{4} \frac{\alpha \rho^2}{\nu}} \nu + 2 \Gamma^2 \nu \right) \quad (3.9)$$

It is worth noting here that since the Burgers vortex is azimuthally symmetric there is no θ dependence and Q and R depend only on the radius, ρ . As such this makes the analysis of the Burgers vortex essentially one-dimensional, which is a significant simplification.

While providing an analytical solution, the above result is not immediately helpful in identifying regions where certain topologies of the flow will dominate. Taking some representative values of $\alpha = 2 \frac{1}{s}$, $\Gamma = 15 \frac{m^2}{s}$, and $\nu = 10^{-3}$ we can consider the behaviour of Q and R as we move away from the origin. Our understanding of Burgers vortex says that we expect to see vortex stretching dominate near the center of the domain, with more shear-dominated flow as ρ approaches infinity.

Figure 3.1 shows a plot of Q and R for a range of ρ . Immediately obvious is that we see only two regimes present. First, near the origin, Q and R are both large and positive, implying a region of vortex stretching, in accord with the postulated physical balance that the Burgers vortex is seeking to demonstrate. Interestingly, we also see that the other dominant region has values of Q and R which are both negative (although still significant in magnitude). This regions corresponds to bi-axial shear. At first this may seem like a contradiction, since the flow field which we defined is designed to produce vortex stretching; however, it is clear after some thinking that the regions further from the origin would experience greatly weakened vorticity, with the only real result being that flow coming in close to parallel will tend to get very slightly stretched in the vertical direction.




Figure 3.1: Line plot of Q (red) and R (blue) versus radius (ρ) for a Burgers vortex with $\Gamma = 15, \alpha = 2, \nu = 10^{-3}$. The region close to the origin (where the interesting dynamics in occurs) correspond to the first and third quadrants of the $Q - R$ plane.

In order to examine the behaviour of this vortex close to the origin we can consider the asymptotic behaviour of the ratio of Q versus R ,

$$\begin{aligned}
\frac{Q}{R} = & \frac{-48 \pi^2 \nu^2 + \Gamma^2}{\alpha (16 \pi^2 \nu^2 + \Gamma^2)} \\
& - 32 \frac{\Gamma^2 \pi^2 \nu}{(16 \pi^2 \nu^2 + \Gamma^2)^2} \rho^2 \\
& - \frac{\Gamma^2 \alpha \pi^2 (112 \pi^2 \nu^2 - 9 \Gamma^2)}{(16 \pi^2 \nu^2 + \Gamma^2)^3} \rho^4 \\
& + O(\rho^6).
\end{aligned} \tag{3.10}$$

This even asymptotic series allows us to observe several key features about the interaction between Q and R in the Burgers vortex. First is the fact that since ν is very small for commonly occurring fluids ($\sim 10^{-6} \frac{m^2}{s}$ for water) the order one term, $O(1) = \frac{-48 \pi^2 \nu^2 + \Gamma^2}{\alpha (16 \pi^2 \nu^2 + \Gamma^2)} \sim \frac{1}{\alpha}$, is governed primarily by the vertical stretching factor, α . Another feature of note is that

the ratio between Q and R remains almost constant for small values of ρ . The implication is somewhat more complex than the previous observation. However, we can clearly see that since the ratio is positive we will be restricted to the two regimes we expected to see, vortex stretching and bi-axial strain.

Since we have a pretty clear separation of scale between the kinematic viscosity, $\nu \sim O(10^{-3})$ and the other factors (roughly $O(1)$) it makes sense to perform a Taylor series expansion in ρ and ν .

$$\begin{aligned} \frac{Q}{R} = & \alpha^{-1} - 64 \frac{\pi^2 \nu^2}{\Gamma^2 \alpha} - 32 \frac{\rho^2 \nu \pi^2}{\Gamma^2} - 9 \frac{\alpha \pi^2 \rho^4}{\Gamma^2} + 1024 \frac{\nu^4 \pi^4}{\alpha \Gamma^4} - 2 \frac{\alpha^2 \pi^2 \rho^6}{\Gamma^2 \nu} \\ & + 1024 \frac{\nu^3 \pi^4 \rho^2}{\Gamma^4} - \frac{235 \alpha^3 \pi^2 \rho^8}{576 \nu^2 \Gamma^2} + 544 \frac{\alpha \pi^4 \nu^2 \rho^4}{\Gamma^4} - 16384 \frac{\nu^6 \pi^6}{\alpha \Gamma^6} + O(\rho^8). \end{aligned} \quad (3.11)$$

This expansion shows clearly that the first term, $\frac{1}{\alpha}$ is the most important since the interesting behaviour is vorticity stretching, which is only relevant for small ρ .

The relative simplicity of this flow has allowed us to observe the resulting $Q - R$ behaviour in a controlled environment. The next logical step is to consider a three dimensional flow which is more complex.

3.3.6 Taylor-Green Vortex

The Taylor-Green (TG) vortex is an exact solution to the Navier-Stokes equations. It is an unsteady flow of a decaying vortex and is often used as a test case in numerical models due to its convenient Cartesian formulation in a periodic cube.

The original paper by Taylor and Green [36] yields the velocity field (at $t = 0$):

$$\begin{aligned} u &= A \cos(ax) B \sin(by) C \sin(cz), \\ v &= A \sin(ax) B \cos(by) C \sin(cz), \\ w &= A \sin(ax) B \sin(by) C \cos(cz). \end{aligned}$$

With the added condition (from incompressibility) that $Aa + Bb + Cc = 0$.

This formulation is useful as an initial setup for numerical models in a box (cube). A version of the TG vortex which satisfies periodic boundary conditions and is a widely used initial state for numerically simulated turbulence (see [7, 6, 28] for examples of simulations

of a TG vortex using DNS and LES methods) is given by:

$$\begin{aligned} u &= \sin(x) \cos(y) \cos(z), \\ v &= -\cos(x) \sin(y) \cos(z), \\ w &= 0. \end{aligned}$$

By eliminating any vertical motion in the initial state we severely limit the analysis available using $Q-R$ techniques. However, it is worth noting that this system is effectively two-dimensional (despite the z -component) and as such a brief consideration of this initial state using Q can still yield interesting results.

Once again we can compute the velocity gradient tensor:

$$A = \begin{pmatrix} \cos(x) \cos(y) \cos(z) & -\sin(x) \sin(y) \cos(z) & -\sin(x) \cos(y) \sin(z) \\ \sin(x) \sin(y) \cos(z) & -\cos(x) \cos(y) \cos(z) & \cos(x) \sin(y) \sin(z) \\ 0 & 0 & 0 \end{pmatrix}$$

We note that one of the rows of this matrix is uniformly 0, and as such expect that the determinant will be 0 - confirming our suspicion that any analysis using R is rendered moot - and that one of the eigenvalues is 0 as well. While we can break this matrix into a symmetric and antisymmetric part, it is far simpler to compute the eigenvalues of A (λ_i) and construct Q via the relationship $Q = -\frac{1}{2}(\lambda_1^2 + \lambda_2^2 + \lambda_3^2)$.

The eigenvalues of A are

$$\begin{aligned} \lambda_1 &= 0, \\ \lambda_2 &= \cos(z) \sqrt{\cos(x)^2 - \sin(y)^2}, \\ \lambda_3 &= -\cos(z) \sqrt{\cos(x)^2 - \sin(y)^2}, \end{aligned}$$

and thus the variable Q (equivalent to the Okubo-Weiss parameter) reads

$$Q = -\cos(z)^2(\cos(x)^2 - \sin(y)^2). \quad (3.12)$$

This simple calculation, combined with our understanding of the limitations of this starting state, demonstrates several interesting features about this vortex. First is that the sign of Q changes throughout the domain (how many times it changes can be determined by increasing or decreasing the coefficients a, b, c). This sign-change occurs in all three directions, thus ensuring that the domain is rich with enstrophy and strain-dominated regions. Additionally, it is worth pointing out that although we do not have any initial

vertical velocities, the magnitude and sign of Q will change based on the value of z . As the simulation is initialized with some perturbation this means the fluid will quickly destabilize and transition to turbulence. The transition regions (regions where the flow goes from strain to vorticity dominated) exist throughout the domain and so we expect that vertical velocities will quickly be induced.

While this initial state is dynamically rich (and obviously complex enough to induce turbulence and thus can be used to verify numerical models), we wish to study the TG vortex more generally. A more dynamically interesting sample vortex can be created by setting the velocities:

$$\begin{aligned} u &= \sin(x) \cos(y) \cos(z), \\ v &= -2 \cos(x) \sin(y) \cos(z), \\ w &= \cos(x) \cos(y) \sin(z). \end{aligned}$$

This yields the gradient matrix

$$A = \begin{pmatrix} \cos(x) \cos(y) \cos(z) & -\sin(x) \sin(y) \cos(z) & -\sin(x) \cos(y) \sin(z) \\ 2 \sin(x) \sin(y) \cos(z) & -2 \cos(x) \cos(y) \cos(z) & 2 \cos(x) \sin(y) \sin(z) \\ \sin(x) \cos(y) \sin(z) & -\cos(x) \sin(y) \sin(z) & \cos(x) \cos(y) \cos(z) \end{pmatrix}.$$

The eigenvalues of this matrix (λ_i) can be analytically computed using symbolic computation, but are too lengthy and reproducing them here serves no purpose. We can, however, use them to find an expression for $Q = -\frac{1}{2}(\lambda_1^2 + \lambda_2^2 + \lambda_3^2)$ and $R = \det(A) = \lambda_1 \lambda_2 \lambda_3$:

$$\begin{aligned} Q &= -\cos(x)^2 \cos(y)^2 - 4 \cos(x)^2 \cos(z)^2 - \cos(y)^2 \cos(z)^2 + 2 \cos(x)^2 - \cos(y)^2 + 2 \cos(z)^2, \\ R &= -2 \cos(x) \cos(y) \cos(z) (\cos(x)^2 + \cos(y)^2 + \cos(z)^2 - 2). \end{aligned}$$

Analyzing the Q term is nontrivial without considering specific values of x, y, z ; we can, however, point out that once again Q will change sign numerous times throughout the domain. The expression for R will also vary in sign and, significantly, will do so in places independent of the sign of Q .

To demonstrate this, we can consider a single line through the domain at $x = \pi/4, y = \pi/4$. This line results in the (very simple) equations

$$\begin{aligned} Q &= -\frac{1}{2} \cos(z)^2 + \frac{1}{4} \\ R &= \sin(z)^2 \cos(z)^2 \end{aligned}$$




Figure 3.2: A line plot showing Q (top-left) and R (top-right) versus height and Q vs R (bottom) in the Taylor-Green vortex for a vertical line at $x = \pi/4, y = \pi/4$. This relatively small sample demonstrates the presence of all regions of interest in the $Q - R$ plane.

Figure 3.2 shows the Q vs z and R vs z plots between 0 and π . In this single line through the domain we already clearly see that all four regions of interest, $Q, R > 0$, $Q, R < 0$, $Q < 0 R > 0$, and $Q > 0 R < 0$, occur.

The closed-form solution of $Q - R$ for this initial state is relatively easy to compute (using a symbolic computation tool like Maple or Mathematica), and quickly yields confirmation that this initialization will quickly yield very complex and turbulent flow. This type of analysis can also be reapplied to a late stage of the simulation in order to compare the $Q - R$ profiles and determine whether the evolved system has preferentially strengthened any of the regimes - i.e. if a single mechanism becomes more dominant, as we expect with vortex stretching.

One caveat to this analysis is that, as we have seen above, while Q and R may be straightforward to compute (either analytically or numerically) the resulting relationship can be quite complex and as such some *a priori* considerations as to the visualizations of these relationships may help to make further analysis go more smoothly.

3.4 Visualization Tools

The tertiary variables Q and R contain a great deal of information to analyze. This means that choosing the type of visualization we use for Q vs R plots is fundamentally related to what we wish to show. For the purposes of completeness we have considered the a wide variety of available techniques for visualization, both in the physical and the $Q - R$ plane. These techniques are detailed in the following sections along with their major advantages and detriments.

Although Q and R present valuable congregate information, it is difficult to know *a priori* how the Q vs R picture represents the dynamics of the system. As such, we have begun our investigations by considering a sample which we already know a great deal about - namely the symmetric shear flow case from the KH chapter (section 2.4.1). The goal is to see how Q and R can help to capture the relevant dynamics, and to determine if any new information can be acquired.

Since Q and R are intrinsically three-dimensional variables, the various plots will be presented at a point in time at which the instability is mature and the flow is three dimensional (around 3.5 seconds). At this time we are confident that the flow has attained sufficient three-dimensionalization to merit investigation using techniques designed for turbulent flow.

Total Q vs R




Figure 3.3: A Q vs R scatter plot with an overlaid line (in black) at $D_A = 0$ showing a quasi-turbulent simulation.

The first plot we present is the canonical total Q vs R scatter plot. This is a plot of all the calculated Q and R values, at all points in the simulation. This is the canonical figure and represents the categorization of the flow (as described in Ooi et al. [70]). Fundamentally this is the most intuitive way to display these variables as it shows us, via which regions of the plane are active, the type of flow we are seeing. In this case (figure 3.3), we see the flow shaped like a “ice cream cone”, indicative of a state close to isotropic turbulence (as described in [25]).

While this is the standard and in some ways sufficient manner of displaying these variables, it is critically lacking in certain key features which are accessible to us. For instance, due to the presence of no-slip boundaries we expect the nature of our flow to be fundamentally different in the boundary layer, versus near the pycnocline. As such we may wish to consider a subset of Q and R , that in the boundary layer, and that in the interior.

Q vs R Slices

The above discussion naturally leads us to the next logical method of displaying our tertiary variables, namely showing ‘slices’ or subsets of the domain which are of special interest to us. For the KH billow example above it is pretty clear from our previous discussion (section 2.4) that the regions of interest are near the boundary and near the center; however, for

other flows we may wish to examine a three-dimensional subset of the data which contains only features of interest.

The most comprehensive way to display subsets of the data is to choose a direction along which the nature of the flow changes and show slices of Q vs R at that level. Figure 3.4 shows the symmetric KH simulation sliced in the vertical (z) direction. Additionally, we have colour-coded the various regions of flow. So the region of positive R above the line $D_A = 0$, representing a dominance of vortex stretching, is coloured blue, vortex sheeting (compression along one axis, stretching along other two axes) is coloured red, Plane shear is coloured green and line shear is coloured yellow.




Figure 3.4: A layered Q vs R plot showing the Q vs R scatter plot at fixed heights (z 's throughout a domain). Superimposed at each level is the line $D_A = 0$ (in black) ranged between the local minimum and maximum values of R . The points in the scatter plot are coloured by quadrant (corresponding to flow topology as seen in figure 1.7)

An advantage of this approach is that it allows us to consider how the Q vs R slices change in structure as you traverse the domain (in a chosen direction). In our example

from the KH chapter we see that the majority of vortex stretching and planar shear occurs near the boundary of the domain, which is to be expected. In contrast, the interior is much less active, but displays significant features of classical turbulence, namely activity in the first and third quadrants - vortex stretching and plane shear.

While the plot in figure 3.4 contains a great deal of information; it is difficult to discern the relevant information without access to the full dataset (i.e. the ability to rotate or restructure the data). As such, a more useful version of the above plot is to show the individual slices of Q vs R at specific heights. This is precisely what is presented in figure 3.5, showing the near wall (left panel) and near-center (right panel) behaviour of the fluid. If we ensure that the axes are set identically (and symmetrically) for all plots we can clearly see that the boundary layer is more active in all four quadrants, but that the center of the domain adheres better to the classically prescribed ‘ice cream cone’ shape for homogeneous, isotropic turbulence.




Figure 3.5: Q vs R slices at two simulation times. The topological regions (as seen in figure 1.7) are coloured blue for vortex stretching, red for vortex compression, green for planar shear, and yellow for axial strain.

Although selecting the Q vs R profile at a specific height (or region) yields the most useful local results, it is important to determine before hand which regions are the most active (or the most interesting dynamically). This leads to a great deal of pre-processing required before the Q vs R plots can be used to analyze the flow dynamics.

3.4.1 Multi-Region Scatter Plots

Since we are interested in the flow behaviour in certain regions (and how specific sub-domains relate to the total) it is worthwhile attempting to simultaneously plot the Q - R

scatter plane in multiple regions of interest. Figure 3.6 shows the Q-R plane at several horizontal slices through the domain. This allows us to visualize how the center of the domain (in blue at height $z = 3.5\text{mm}$) differs from the boundary layers (in green and red for top and bottom, respectively).




Figure 3.6: Q vs R scatter plots at various levels - near bottom boundary (red), top boundary (green), and near the domain center (blue).

This figure allows us to demonstrate a key feature of the flow - namely that the boundary layer contains much more extreme values of Q and R than the interior of the domain. Additionally, the two boundary layer plots demonstrate a slight asymmetry between the top and bottom of the domain, a fact which maybe worth further exploration. Note that the order of plotting is important to ensure slices with smaller Q and R ranges are not completely hidden.

3.4.2 Summary Histograms

The next method of visualization relies on the fact that the total Q vs R scatter plots will demonstrate the shape of the simulation, but not the breakdown of how many sets of points lie in the “extreme Q ” and “extreme R ” regimes. To this end it is worth considering a two dimensional histogram of the Q - R plane divided in to 10 by 10 or 20 by 20 regions.

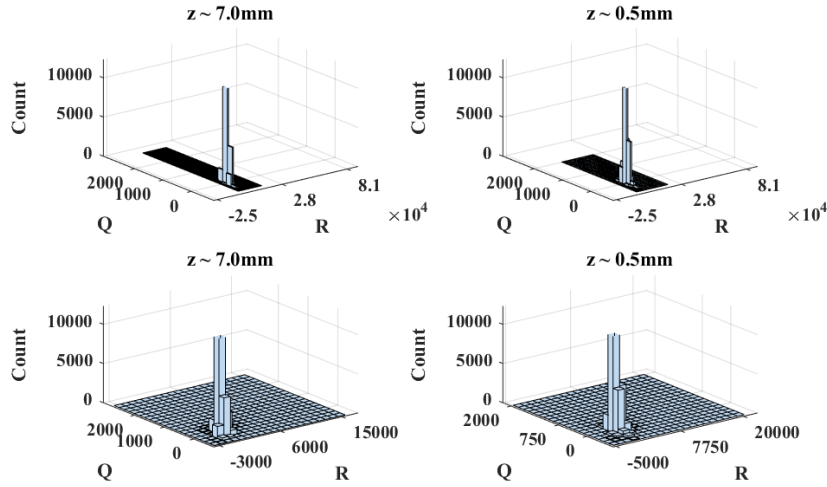


Figure 3.7: Histogram at 2 outputs (left and right) showing an overview (top panels) and zoomed in (bottom panels).

Figure 3.7 shows a histogram of the Q-R plane. Notice that the majority of points lie very near the origin. This is to be expected as only regions of significant enstrophy or shear correspond to extremes of Q and R. By their nature, we are only interested in rarely occurring regions - i.e. those of extreme values. As such the histogram offers somewhat limited value since it invariably identifies the near-origin regions for the eye.

As with the total Q vs R plots we can attempt to take horizontal slices through the domain in order to determine the behaviour of the system in a fixed region. Figure 3.7 shows 2 such slices both for the entire Q-R plane. In this figure we zoom in to show only the active portion near the origin. In both cases the results do not immediately point to any useful characteristic. Although visually appealing, these plots are less informative than the Q-R scatter plots.

One important characteristic that we cannot glean from these figures is that the number of “extreme” points in the simulation, which are dwarfed by the less-active points. By their nature Q and R tend to be nearly gaussian in distribution about 0 [16, 70]. This means that a histogram or other summing plot will necessarily overlook the outlier, or “extreme” points, which are the ones indicative of interesting flow topology.

3.4.3 Augmented Primary Variable Plots

The final way to use the Q-R plane is as a means of augmenting primary variable data. Figure 3.8 shows a (vertical, i.e x-z) slice of density for our sample output (left) along with the same plot overlaid with several diagnostic characteristics. This density plot shows significant overturning and activity at this time. This allows the observer to immediately identify regions of interest such as billow cores, billow braids, or regions with strong density overturns.

On top of this density plot (right), we have overlaid the line $D_A = 0$ in black, which shows where the flow transitions from vortex-dominated to strain-dominated flow (i.e. where the primary topology changes). Additionally, we have added regions of extreme Q and R in white (first quadrant) and cyan (third quadrant). This addition allows us to instantly answer the question of what regime is most significant in locations of interest (again such as cores, braids, or overturn regions).

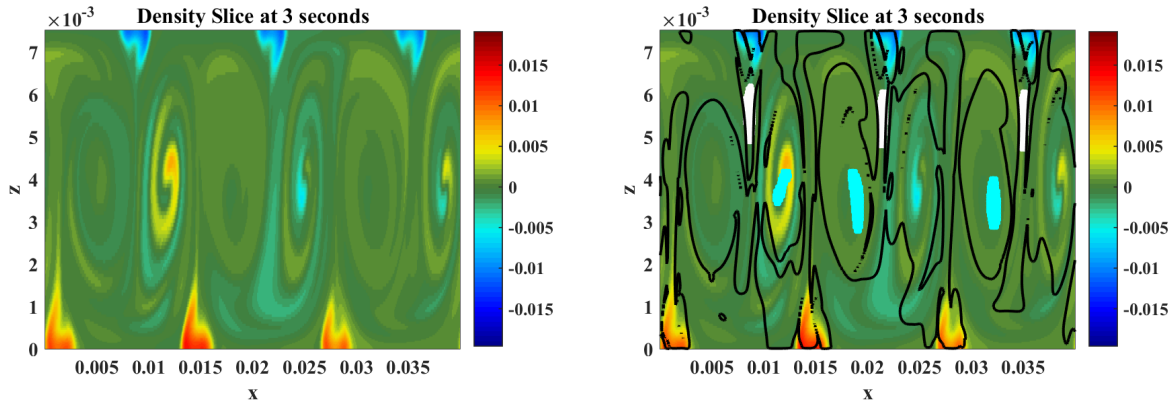


Figure 3.8: Slice of density at 3 seconds (both panels) into the symmetric two-layer simulation. Superimposed in the right panel are the contours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). The cyan points indicate regions of vortex stretching, while the white regions indicate high planar shear.

The advantage of this type of figure is that it demonstrates both what the state of the flow is, using primary variables, as well as what mechanism (shear or vortex stretching) dominates. However, it is important to realize that while it is possible to convey a very large amount of information, the price to pay is in complexity of the resulting figure. Additionally, marked regions can obscure features of the regions of interest which we may wish to examine.

3.5 Case Studies

In the previous sections we have successfully developed an understanding of how Q and R relate to physical processes and developed some useful visualization tools to help with the interpretation of the massive amount of data provided by $Q - R$ analysis. Thus armed we can try to apply this system to numerically generated data in order to see if we can gain new insight into the mechanisms and flow topology present.

In the previous chapter we presented a detailed analysis of the formation and evolution of Kelvin-Helmholtz (KH) billows in a confined environment. This provides a natural launching point for our analysis for two reasons. First, we are already relatively familiar with the dynamics of the evolution, and much of classical theory can be validated using this new approach. Secondly, the initial evolution occurs strictly in two dimensions - so the initial analysis can be performed using strictly Q , allowing us to draw on the Okubo-Weiss literature for the first part of the simulations.

Once we have established this methodology on a classical flow such as KH billows we can consider how it applies to a numerically simulated test case whose dynamics are not immediately obvious to us. The case example which we will use is a confined Rayleigh Taylor instability [52].

3.6 Kelvin-Helmholtz Billows

Our investigation into the formation and evolution of KH billows demonstrated a clear adherence to Squire's theorem: that an initial instability in a fluid always occurs in two dimensions. In all of the simulations we ran, the primary billow formation events always occurred uniformly in the spanwise direction. Accordingly, the initial analysis of the formation events can be performed using strictly Q - corresponding to 2D analysis using the Okubo-Weiss parameter. However, once the flow becomes more complex and destabilizes in the spanwise direction we will be greatly interested in understanding the dominant flow topologies during the turbulent transition period.

Since the initial formation is primarily two-dimensional, and the 2D simulations were an excellent approximation of the initial fully 3D runs, we will be applying the $Q - R$ metrics only to the three-dimensional simulations. The result is that the initial analysis can be performed in 'quasi-2D' - i.e. using primarily Q and drawing upon the Okubo-Weiss criterion for analysis. The subsequent breakdown of the flow in the spanwise direction and the transition to quasi-turbulent flow will require a more *ad hoc* approach. We will

begin by using augmented density plots to identify regions of interest, focussing on two regions specifically, the domain interior and the region(s) near the boundaries. For the symmetric and linear case, the two boundaries behave almost identically, however for the asymmetrically placed pycnocline there exist a near and far boundary, which will be worth considering separately.

3.6.1 Symmetric Two-Layer Flow

The first case study that we will analyze is the symmetric quasi two-layer problem. During our previous analysis of this case we determined that mixing (defined as the regions of high $\nabla\rho$) occurred primarily while the flow was two dimensional (i.e. very little spanwise variation). Additionally we were interested in regions of high dissipation and high shear since it is here that we expect to find the bulk of the dissipation ($\varepsilon = 2\nu e_{ij}e_{ij}$)

We begin by considering the simulation at relatively early times, before destabilization in the spanwise direction. The two-dimensional nature of the early simulation results in two important factors. First, since the spanwise velocities and gradients are very small the resulting Q and R fields are low in magnitude. This means that comparisons with later times will be more difficult, however since we are fundamentally interested in the relative strength of the Q and R fields (since we assume a pseudo-normal distribution, see [69, 98] for original derivation and [44] for standard use in experimental work), we can make a meaningful assessment of the locations of extreme Q .

We can begin our analysis by considering a relatively well-understood part of the simulation, the initial billow formation. As previously discussed this rollup occurs almost entirely in two dimensions, and as such presents a perfect opportunity to display the utility of the density slices superimposed with extreme values of Q and R . Following the Okubo-Weiss related oceanic literature we consider “extreme” values to be those greater than approximately one standard deviation of Q ([44] actually uses 0.2 standard deviations, which we felt highlighted too many regions in most of our very active results). The result, figure 3.9, shows the billows forming as well as the induced locations of high vorticity (in cyan) and high shear (in white). The black contours are lines where $D_A = 0$ i.e. where the eigenvalues change form from purely real to complex.

The distribution of Q relies on the balance between vorticity and strain in the flow field to be investigated. While we have no mathematical assurance that such a balance will be Gaussian in nature, in practice the resulting fields are generated from white (Gaussian) noise, and are empirically observed to follow a roughly Gaussian structure. This is

confirmed in part by the Okubo-Weiss literature, which effectively uses a two-dimensional equivalent of Q and treats it as normally distributed about 0 [40, 39, 43, 44, 69, 97, 98].




Figure 3.9: Slice of density at 1.25 seconds into the two-layer simulation. Superimposed are the contours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). The regions of high vorticity (cyan) correspond well to the billow cores and regions of high strain (white) correspond to the braids.

Figure 3.9 shows a density slice at 1.25 seconds with superimposed contours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). Note immediately that the regions of high vorticity (in cyan) and high planar shear (in white) are immediately tagged as the billow cores and connecting braids, respectively. This plot provides verification of our methodology and allows us to confirm that these turbulent metrics are tagging the relevant regions of interest. The natural next step is to consider the later, more turbulent states of the two-layer simulation. In our previous investigations we determined that the billows grew to occupy the entire domain, at which point interaction with the boundary layer (where viscosity becomes a dominant factor) causes significant

dissipation and shear stress to occur. We now expect these boundary regions to be quickly identified as shear-dominated and tagged in white.




Figure 3.10: Slice of density at 1.5 seconds into the two-layer simulation. Superimposed are the contours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). Once the billows have grown to fill the domain the regions of extreme Q are in the billow cores, where vorticity clearly dominates, and near the boundary, where interaction with the no-slip boundaries induces significant shear - and dissipation.

Figure 3.10 shows a slice of the two-layer simulation at 1.5 seconds, once the billows have grown (remaining two-dimensional) to occupy the entire vertical domain. Already we can see regions of high plane shear (in white) near the boundaries. These regions will produce intense local dissipation in the fluid. Additionally, the billow cores remain very active and are tagged by the cyan high Q and R region to be dominated by vortex stretching. This analysis corresponds well to our understanding of the dynamics of this relatively early time evolution of the system.

Considering later times we will wish to verify two key aspects. First that the metrics Q and R are capable of identifying the dominant mechanisms in the simulation. Secondly to determine if we are able to correlate regions of Q and R activity with dissipation, something we expect to have an important impact on the system due to the presence of no-slip boundaries.




Figure 3.11: Slice of density at 1.625 seconds into the two-layer simulation. Superimposed are the contours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). The simulation is at a state just before three-dimensionalization begins. The domain contains many regions of overturn and significant activity. The $Q - R$ markers have identified several regions of interest, such as the billow cores, where the overturn is most dramatic, braids, and regions where billows come into contact with the no-slip boundary layer near the top and bottom of the domain.

Observing an augmented slice at 1.625 seconds (figure 3.11) shows clearly that the region is rife with density overturns. Moreover, the billow shapes are significantly affected by the presence of boundaries. A useful features of the augmented plot is the cyan regions,

representing vortex stretching, clearly marking the centers of the billows. Additionally it is certainly worth observing that the regions of most intense plane shear (tagged in white) occur both near the boundaries at this time, as well as in the billow braids. This implies that billow self-interactions remain an important dissipation-generating mechanism even at this stage of the simulation, however, it is obvious both here and from our previous discussion, that the majority of dissipation is generated near the boundaries - specifically near the top boundary in this figure - to be advected towards the center of the domain, as we previously observed.

If we consider the vertical structure of Q vs R (top panel of figure 3.12) we notice immediately that the majority of first quadrant activity (in blue) takes place near the center of the domain, away from no-slip boundaries. Meanwhile, the distribution of extreme third-quadrant activity (in green) is located primarily near the boundaries. So while the most extreme regions of shear are tagged somewhat near the braids, the majority of the expected dissipation-inducing domain is located near the no-slip boundaries, as expected by our intuition and analysis in the previous chapter.

Speaking more broadly, the first and second quadrants indicate a predominantly vorticity-dominated flow, while the third and fourth quadrants indicate a shear-dominated flow. Thus figure 3.12 also informs us that at this time in our simulation the center of the domain is primarily vorticity dominated, while the near-boundary regions are shear-dominated, as expected. This is again in agreement with our previous analysis.

As we saw in the previous chapter, the system quickly destabilizes in the spanwise direction, with instabilities forming and rapidly growing in the y direction. The flow quickly becomes fully three-dimensional – with spanwise velocities on the order of 10% of the maximum downstream velocities, similar to vertical velocities – and thus we can apply our $Q - R$ analysis in earnest.

Due to the large spanwise variation in the density and velocity field, vertical slices of the domain no longer present an accurate picture of the simulation state. Thus we will focus on three-dimensional analysis tools, such as the vertical slices of $Q - R$ and Q vs R plots at interior and boundaries.

Figure 3.13 shows the Q vs R scatter plot for three separate layers with fixed z , near the top and bottom boundaries, in green and red respectively, and near the center of the domain, in blue. As we saw from the 3D plot (3.12) the regions of negative Q (corresponding to shear-dominated flow) occur predominantly near the boundaries, while the regions of positive Q (corresponding to vorticity-dominated flow) occur near the center of the domain, where the billow cores dominate. Together with the figure 3.12 this figure provides an overview of the dominant mechanisms at this snapshot in time. There is a significant

Layered Q vs R plots at 1.625 seconds




Figure 3.12: Q vs R scatter plot taken at each horizontal level ($z = z_i$) displayed in an (R, Q, z) space at 1.625 seconds. The black lines correspond to $D_A = 0$ between the minimum and maximum local (at fixed z) values of R . Note the high amount of activity in the first and second quadrants, corresponding to large enstrophy near the center of the domain.

amount of information present in these two figures, some aspects, such as the fact that the vorticity dominates near the center while shear dominates near the boundary is obvious from the previous plots and related discussion. Somewhat more subtle is the observation that the $Q - R$ distribution does not lie along the line $D_A = \frac{27}{4}R^4 + Q^3 = 0$. While we may be tempted to conclude something about this, it is more accurate to note that this means that R does not vary nearly as much as Q . The implication that Q varies greatly while R variations are comparatively small demonstrates that the dominant features of the flow are likely constrained to two-dimensions at this time.

Moving on to later times (figure 3.14) we see a distinct trend emerge. The first quadrant region, indicative of vortex stretching, grows in importance, and by 2.25 seconds it is by




Figure 3.13: Q vs R plot at three vertical levels, near the bottom boundary (red), center (blue, and bottom boundary (green) at 1.625 seconds with the black line corresponding to $D_A = 0$. Note that the majority of vorticity is generated near the interior of the domain, while the shear occurs primarily near the (no-slip) boundaries.

far the dominant features of the plot. In contrast the shear regions (quadrants 3 and 4) remain relatively constant. So as the system destabilizes and transitions to turbulence the shear mechanics continue as before, however they are dwarfed in importance by the newly dominant vortex stretching mechanism. What we see above is a stark demonstration of the importance of vortex stretching in the transition to turbulence. This mechanism, once activated, consumes large amounts of energy, quickly becoming the most important aspect of the energetics of the system.

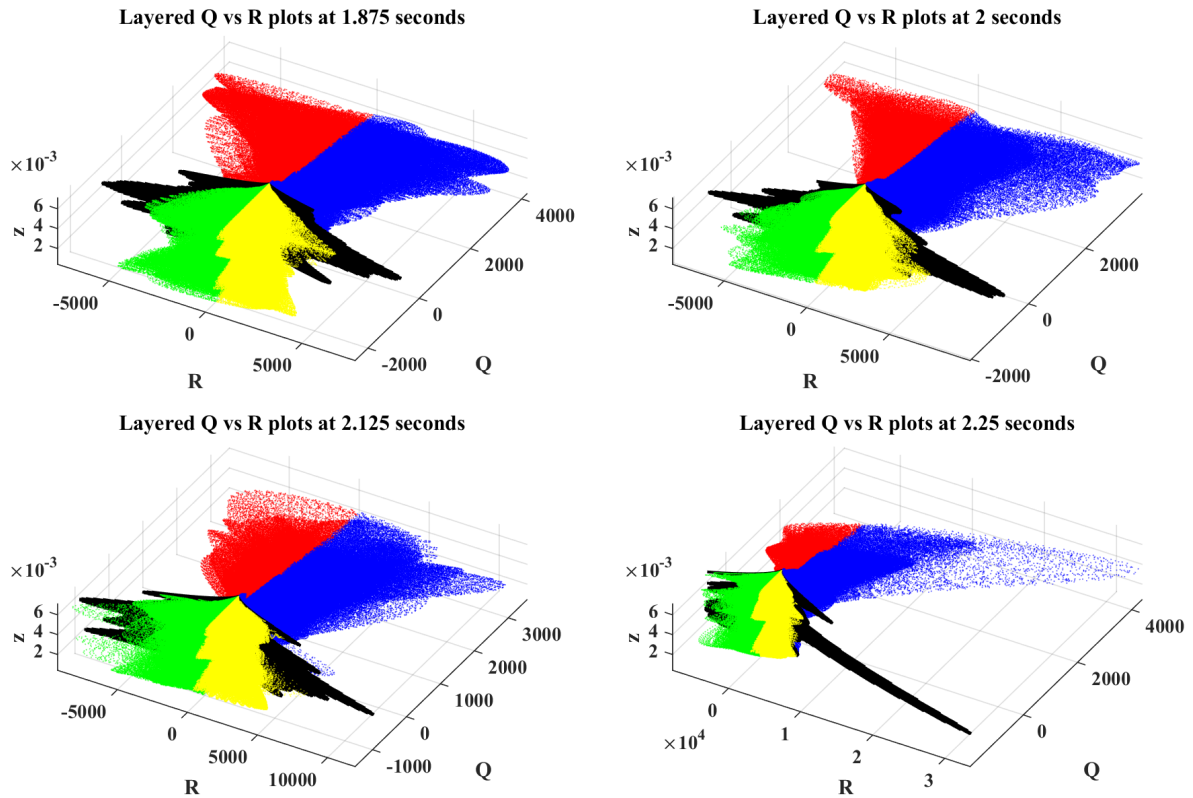


Figure 3.14: Vertically layered Q vs R plots between 1.875 and 2.25 seconds. This period shows the transition between two-dimensional rollup to fully-developed billows which begin to destabilize in the spanwise.

Figure 3.15 shows Q vs R at the same levels as before for the later simulation times. As in the 3D plots we can clearly see that the first quadrant, or vortex stretching region, is by far the most important. What is quite clear in this figure is that the center of the domain is strongly vorticity-dominated, while the boundaries experience large shear early on (as evidenced by the large negative Q regions in the top-left panel) which then diminishes in importance as the flow becomes more turbulent. Another interesting aspect is the increase in vorticity in the bottom boundary (in red) at 2.125 and 2.25 seconds. This result is expected, as the nature of the velocity profile at the boundary (zero at the boundary itself and increasing as you move towards the center of the domain) is conducive to vorticity production. Intuitively we expect this to be a relatively minor source of vorticity compared

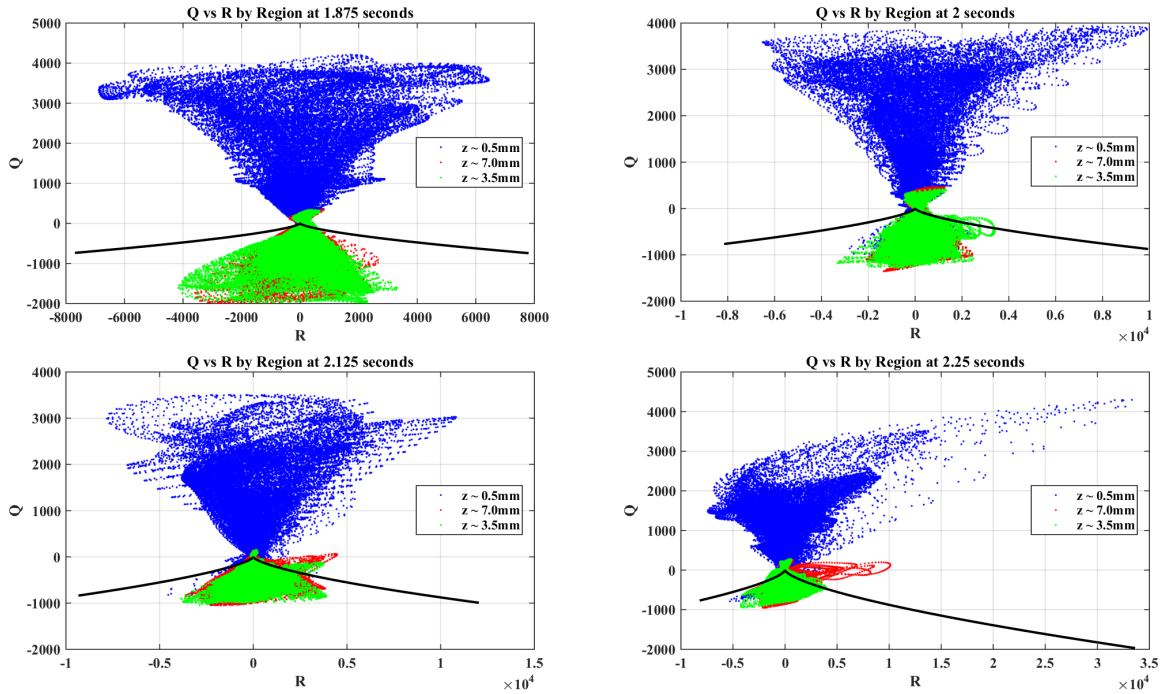


Figure 3.15: Q vs R plots by region (near boundary - red and green - and near center - blue) between 1.875 and 2.25 seconds. This period shows the transition between two-dimensional rollup to fully-developed billows which begin to destabilize in the spanwise.

to the rollup of the billow cores, however it is good to see that this mechanism is not overlooked in our Q vs R diagnostics.

A slice at this time, figure 3.16, shows that indeed the center of the domain is dominated by billow core vorticity, while the shear of the braids is now dwarfed by the boundary-related shear, as was demonstrated by the Q vs R plots above. This billow-boundary interaction is the behaviour that we were first expecting when we began to consider confined KH billows, so it is useful to see it clearly and logically identified by $Q - R$ analysis.

As we observed in the previous chapter, the system continues to destabilize in the classically prescribed manner; spanwise instabilities grow, reaching approximately 10% of downstream velocities (equivalent to maximum vertical speeds) and causing the original billow structure to break down completely. The accompanying dissipation is generated primarily at the boundaries and is likely generated at the shear. For a more detailed analysis see section 2.4.1 of this thesis. We need not reiterate the classical evolution




Figure 3.16: Slice of density at 2.25 seconds into the two-layer simulation. Superimposed are the contours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white). The cyan regions once again identify the original billow cores, while the white regions identify shear-dominated parts of this flow, such as the braids, near $x = 0.015$ and the billow-boundary interaction regions near the top and bottom of the domain.

here again, however it is worth considering the state achieved before the simulation is terminated.

At 4.5 seconds (figure 3.17 we observe that the flow has lost almost all aspects of billow-braid structure. The $\pm 10\%$ density isosurface shows very little structure in either the down-channel or along-channel directions. The resulting state is closer in appearance to stratified turbulence than to any coherent structure observed up to this point. A worthwhile question to ask at this state is whether the $Q - R$ analysis will capture this transition to a pseudo-homogeneous turbulent regime, or whether the dynamics of the billow-braid system continue to dominate despite the state of the system.

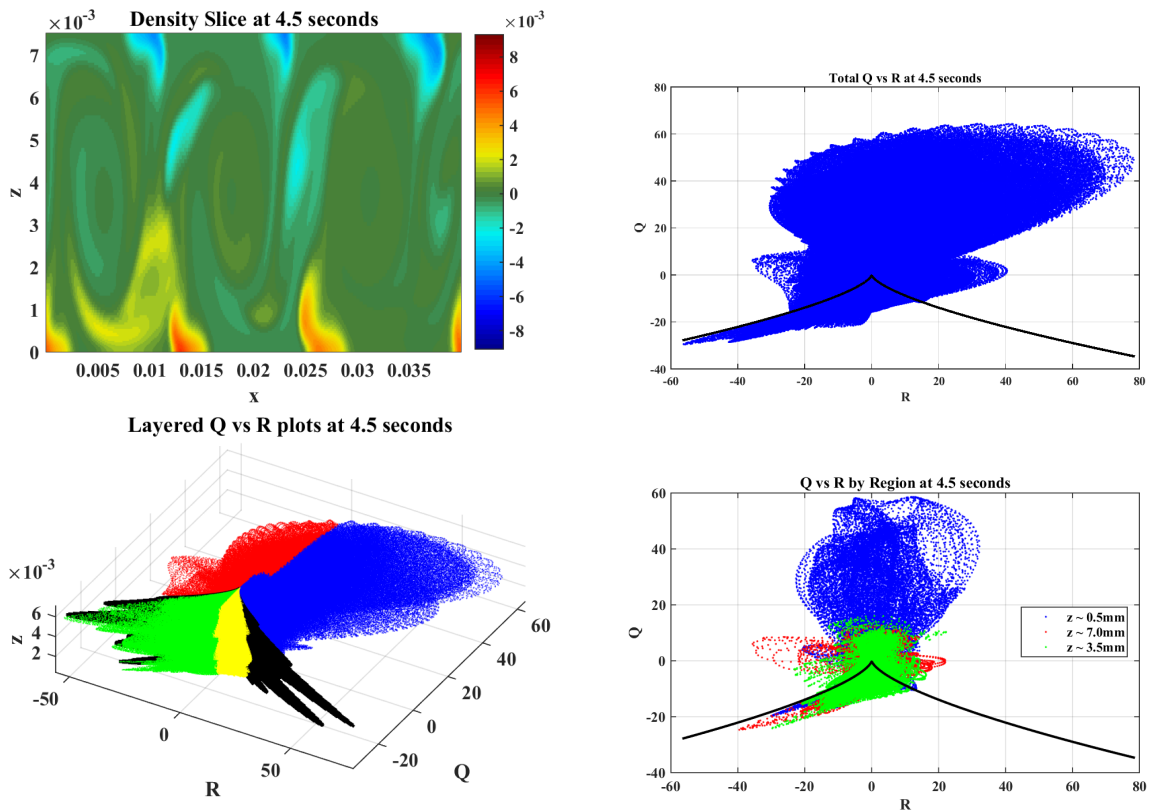


Figure 3.17: Various $Q - R$ related plots at 4.5 seconds showing density isosurfaces, total Q vs R plots, layered Q vs R plots, and Q vs R plotted by region in the top left, top right, bottom left, and bottom right, respectively. The overall impression is that the state of the simulation corresponds well to classical turbulence, and indeed we can note the canonical “ice cream cone” shape begin to emerge.

Figure 3.17 shows the density isosurface (for reference), as well as total, layered, and regional Q vs R plots (top-right, bottom-left, and bottom-right, respectively) at 4.5 seconds. The first observation we can make is that while the range of Q and R is greatly reduced, the total Q vs R plot displays both the classical ice-cream cone structure indicative of isotropic turbulence and several secondary characteristics which are much more difficult to quantify. It is likely that the region around $Q = 60$ corresponds to the vorticity component of the ice-cream cone. What is much less clear is the underlying dynamics of the region around $Q = 0$. While we are well within the confines of the vorticity dominated region,

the low values of Q mean that we have no clear indication of the primary vorticity-shear balance [25]. The bottom-right panel, showing Q vs R by region, clearly shows that this region around $Q = 0$ is confined to the top and bottom boundaries. This is something of a relief, as it now allows us to separate the aggregate Q vs R plot into three separate profiles, each of which resembles a slightly differently shaped ice-cream cone structure.

The logical extension of this observation is that the domain can be separated into three vertical sections, the top and bottom of which exhibit much larger values of R (and thus more axial stretching and, combined with large local values of Q more vortex stretching) while the region passing through the vertical center (captured in the regions plot in the bottom right of figure 3.17) exhibits large vorticity, but significantly less axial vs bi-axial action. These observations serve as one of our most critical clues that although the density (and enstrophy and dissipation) field hint at near-homogeneous behaviour there are fundamental differences in the dominant dynamics throughout the domain.

Finally, it is worth noting that the values of $Q - R$ in the third quadrant remain close to the line $D_A = 0$, remaining strictly below this line - indicating that the dominant mechanism at those points is plane shear. This adherence to the separation line of the tail is indicative of the ice-cream cone structure, and tells us that the flow topology of the region, while within the regime of shear-dominated flow, is not firmly so. This provides a hint at the relative instability of this state, though Q and R theory does not yield the structure of the (presumably small) perturbation necessary to induce a change in flow topology.

Having completed our analysis of the symmetric two-layer flow the natural extension we wish to consider is, as before, the asymmetric domain and linear stratification. The aim is to determine whether $Q - R$ analysis can yield any insight into the differences that arise due to different stratification.

3.6.2 Asymmetric Two-Layer Flow

The next simulation we will attempt to analyze with our tertiary variable toolkit is the asymmetric two-layer flow. The domain remains the same as in the previous section, but the pycnocline is moved to $0.25H$, where H is the domain height.

Figure 3.18 shows an augmented density slice at an early time (1.625 seconds) with regions of high vorticity (cyan) and high shear (white) overlaid. Once again the $Q - R$ analytics have clearly identified the regions of intense vorticity in the billow cores, as well as the high-shear regions in the braids and near the closer boundary. It will be tremendously




Figure 3.18: A density slice at 1.625 seconds with regions of vortex stretching (cyan) and planar shear (white) as well as the line $D_A = 0$ (black) overlaid. We immediately observe that this early output follows the classically prescribed structure and mechanisms with vorticity dominating the billow cores and shear dominating in the braids. What is novel is the presence of shear where the billows come near the no-slip boundary.

useful to be able to track the regions of high shear and vorticity as the billows continue to grow and move away from the nearer boundary.

Figure 3.19 shows augmented density plots from 1.75 to 2.125 seconds. These outputs are particularly interesting because they clearly show something that our initial analysis could only hint at: the relatively constant density region of fluid in the bottom half forms “anti-billows” - regions conforming to the billow core separated by braid structure of classical Kelvin Helmholtz instability [14, 15, 72]. What is particularly interesting is that these anti-billows experience the most intense vorticity at 1.875 seconds, and maintain the largest values of both shear and vorticity for the following output times shown.

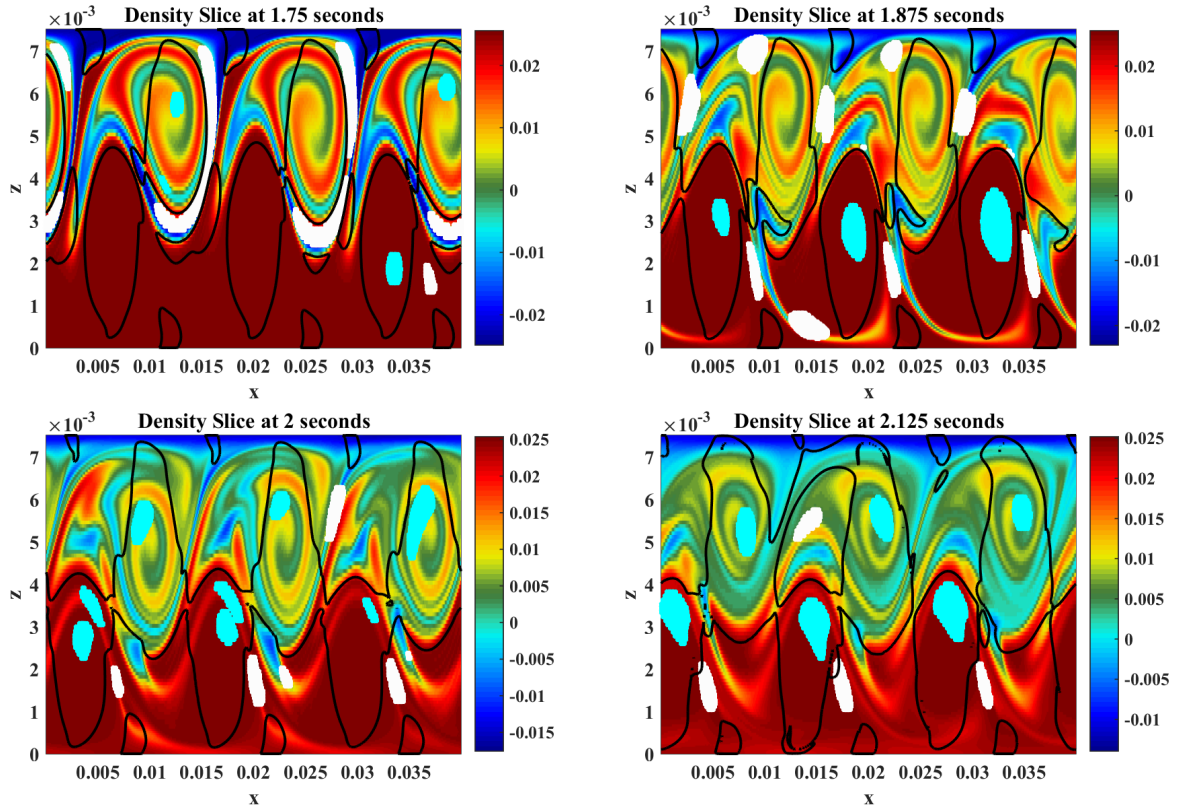


Figure 3.19: Augmented density slices with extreme regions of $Q - R$ in the first (cyan) or third (white) quadrants representing vortex stretching or shear dominated flow, respectively between 1.75 and 2.125 seconds showing the evolution of confined billows in an asymmetric pycnocline and the formation of the “anti-billow” structures in the relatively density constant bottom half of the domain.

The “anti-billow” region is not significantly stratified, however there is no mistaking this new instability for anything but KH billows. It is also worth remarking here that the lines of $D_A = 0$ successfully identify the “anti-billow” regions throughout the simulation, once again confirming that the dominant mechanisms change in character within and without the billow core.

One final point worth noting about this analysis is something we did not expect, namely that the regions of most intense shear are found not near the boundaries but in the billow and anti-billow braids. This diminished importance of the boundaries is both surprising

and insightful as it demonstrates that the billow growth follows a path of least resistance (i.e. inducing sympathetic flow in the ‘empty’, or unstratified region below rather than forced interaction with the far boundary layer by the billows themselves).

Layered Q vs R plots at 1.875 seconds




Figure 3.20: Vertically layered Q vs R plot at 1.875 seconds showing the Q vs R structure at constant heights (z) during the critical point of anti-billow formation.

The next logical step is to consider more fully the vertical structure at 1.875 seconds (the formation time of the anti-billow region). Figure 3.20 shows vertical slices of Q vs R at 1.875 seconds. The black curves, which show the line $D_A = 0$ between the min and max local values of R provide a simple way to visualize where the regions of most intense activity occur. We can therefore conclude that the most interesting regions dynamically are the upper billow region, the center of the domain, and the lower, anti-billow regions.

It is now worth considering the spanwise instabilities generated at this time. In order to consider the rise and formation of these instabilities it is best to consider not composite plots such as Q vs R , but more traditional spatial metrics such as dissipation or enstrophy.

Before we move on to consider these it is worth reiterating that QR -augmented plots were essential in helping us identify the regions of interest, but may not be the best tools for followup analysis.

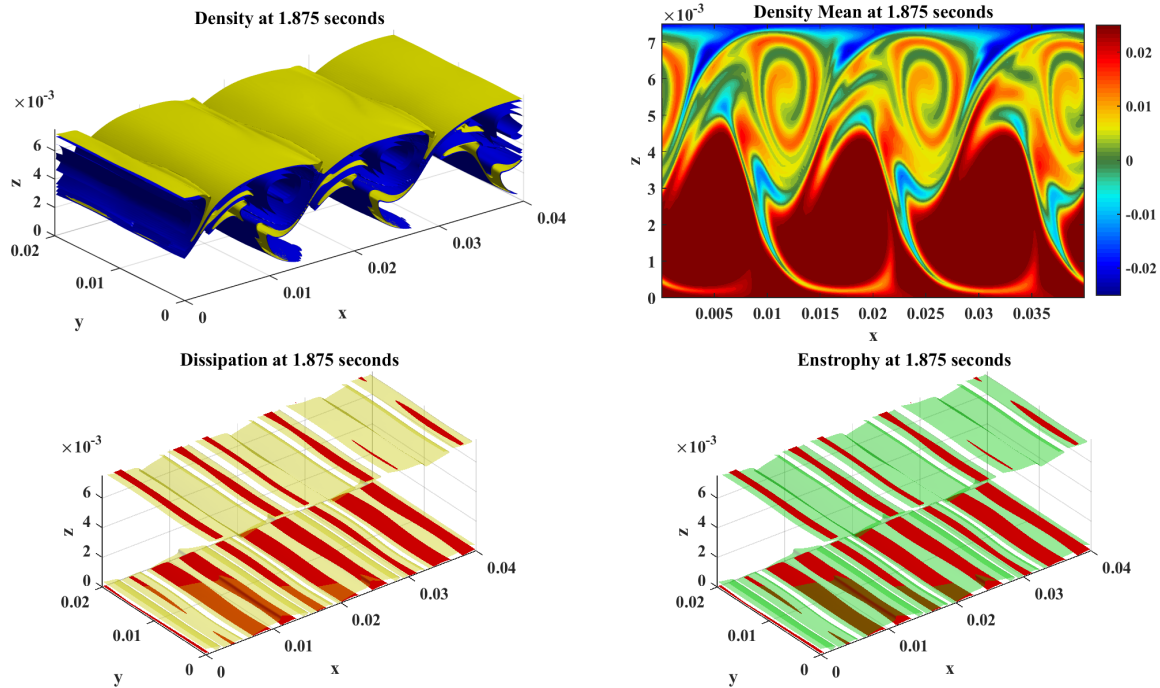


Figure 3.21: Density isosurface (top-left), mean density (top-right), dissipation isosurface (bottom-left), and enstrophy isosurface (bottom-right) at 1.875 seconds. Note the obvious presence of dissipation in the anti-billow regions, but a total lack of enstrophy.

Figure 3.21 shows our standard analysis tools at 1.875 seconds including density (top-left), dissipation (bottom-left), and enstrophy (bottom-right) isosurfaces. Additionally I have included the spanwise-averaged density field (top-right). Clearly visible in the dissipation plot is that there is a strong dissipation response in the anti-billow region in the bottom half of the domain. Interestingly, the enstrophy plot does not tag this region whatsoever. Additionally, as expected, the density isosurface misses the anti-billow region entirely, showing only the small overturning due to entrainment of the former pycnocline by the anti-billows. Finally we note that at this time there are few spanwise fluctuations either in any of the observed fields.

Constructing the same figure at 2.125 seconds we note that the spanwise instability

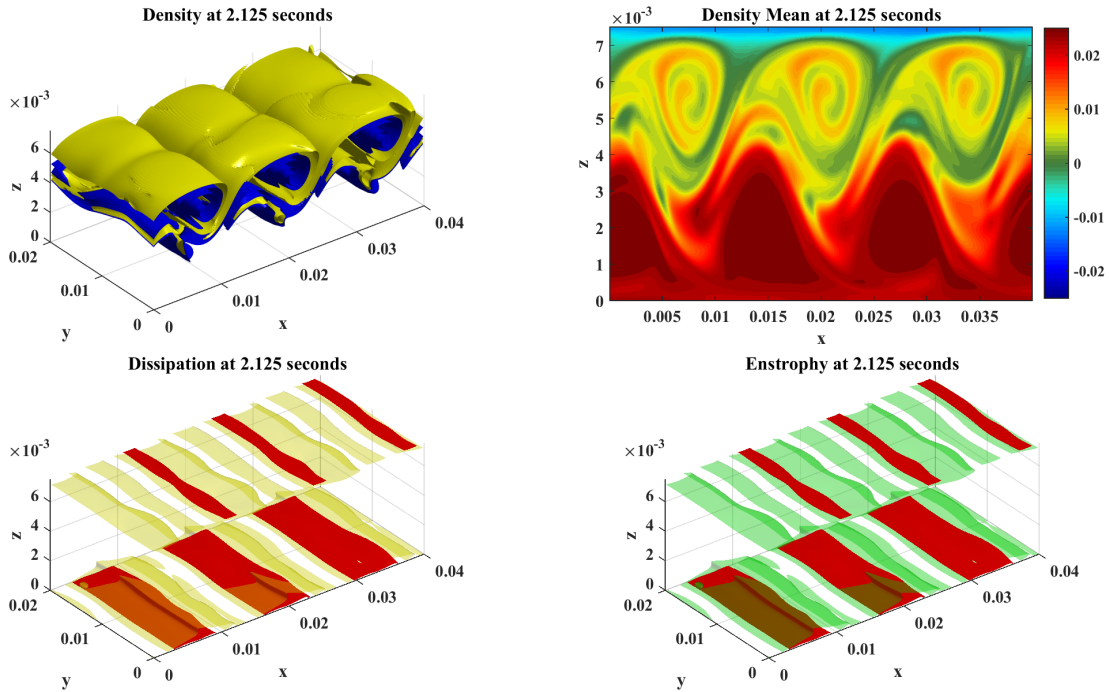


Figure 3.22: Density isosurface (top-left), mean density (top-right), dissipation isosurface (bottom-left), and enstrophy isosurface (bottom-right) at 2.125 seconds. The spanwise destabilization is well under way in both the billow and anti-billow regions.

has begun to materialize in earnest. In the density isosurface (top-left) we see a large (domain-scale) instability form in the billows in the top half of the domain. Additionally the dissipation isosurfaces shows significant across-channel deviations at all heights. These changes mark the beginning of the end for the coherent billow structures in our system. Within 0.25 seconds the system will destabilize and turbulent mechanics, and the attendant energy cascade, will begin to dominate. Nevertheless, the $Q - R$ metrics and the plot types we have developed to examine them have provided a useful tool to probe the mechanics of the onset of instability.

Returning to the QR analysis of the above times we can consider the total, layered 3D, and regional Q vs R plots at 1.875 seconds, figure 3.23. These analysis tools show that the $Q - R$ profile, and thus the flow topology, is most active near the boundaries for regions of high shear. This lines up well with our understanding of what should be happening at this output time, as the boundary layer and interactions with the no-slip boundary should

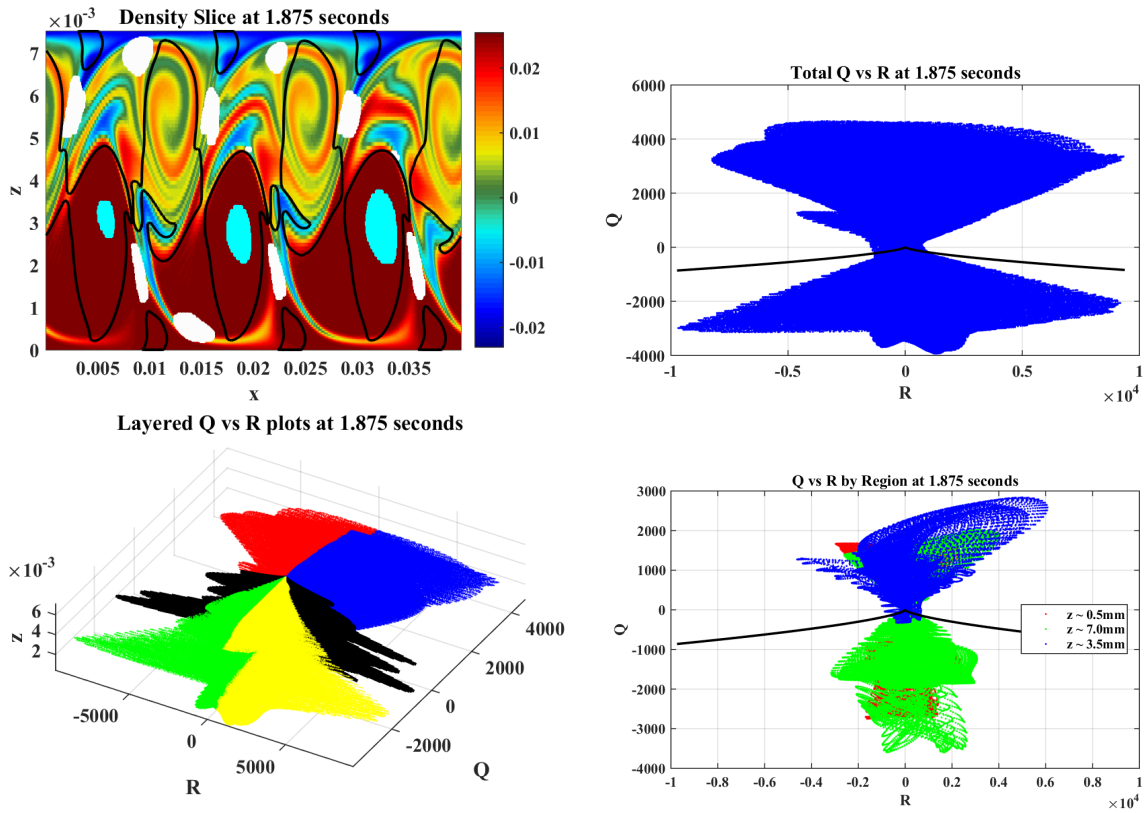


Figure 3.23: Various $Q - R$ related plots at 1.875 seconds showing density isosurfaces, total Q vs R plots, layered Q vs R plots, and Q vs R plotted by region in the top left, top right, bottom left, and bottom right, respectively. The flow topology is clearly dominated by shear in the boundary and vorticity near the domain center.

be the primary cause of shear generation. In contrast, the center of the domain (the blue region in the bottom-right panel) is clearly vorticity dominated, and it is precisely here that we see the formation of the anti-billow regions discussed previously. The magnitude of the enstrophy near the vertical center of the domain (the tops of the anti-billows) is also clearly marked by the cyan regions in the top left panel, indicating that this is where the enstrophy is strongest and the Q vs R plot experiences extreme values (higher than our selected threshold).

Figure 3.24 shows the same series of plots at 2.25 seconds. The anti-billows at this time continue to evolve alongside the primary billows, with the highest vortex stretching

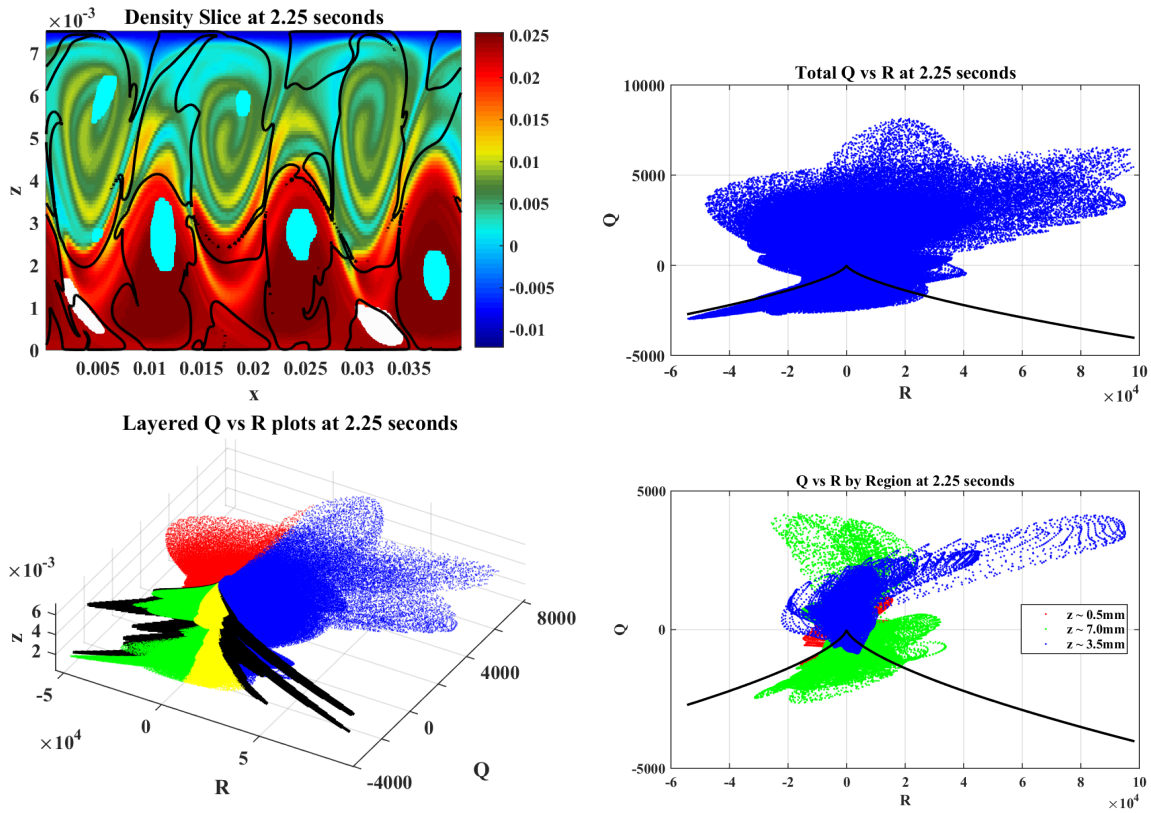


Figure 3.24: Various $Q - R$ related plots at 2.25 seconds showing density isosurfaces, total Q vs R plots, layered Q vs R plots, and Q vs R plotted by region in the top left, top right, bottom left, and bottom right, respectively. Two important features of note is the fact that the most extreme vortex stretching happens near the domain center, and the region near the bottom (far) boundary is by far the very active, providing a unique superimposed profile.

actually taking place in the cores of the anti-billows (as marked by the cyan regions in the top-left panel). Additionally, the total Q vs R plot shows a large amount of both enstrophy and strain, a fact which is not surprising considering the double set of billow cores, and corresponding double interactions with the boundary layer at the top and bottom. From the augmented density plot (top-left panel) we note the influence of mixing near the top boundary, and the very coherent structure of the billows and anti-billows.

Turning to the layered Q vs R plot, we note that the vertical structure has several

interesting characteristics. Firstly it is clear (from the black line in the bottom-left panel) that the vast majority of planar shear (third quadrant) occurs at or near the boundaries. In contrast, the majority of enstrophy is clearly located in two places in the interior (corresponding to the centers of the primary and anti-billow cores). Considering the breakdown of Q vs R by region (bottom-right panel of figure 3.24, we see that the center of the domain is very heavily dominated by enstrophy, with almost no relevant shear, while the regions near the top *and* bottom boundaries both experience significant shear and vorticity. Finally, it is worth noting that the small perturbation from the main ice-cream cone pattern around $Q = 0, R > 0$ is the result of action in the boundary layer, which contains both enstrophy and shear elements in various parts, but exhibiting strong axial tendencies. This is characteristic of the extremely complex interaction of the destabilizing billows and the no-slip boundaries which make this simulation unique.

The next reasonable step in our investigations is to consider the linearly stratified case from the previous chapter, which had a much longer time scale and spawns much weaker destabilization.

3.6.3 Linear Stratification

The previous analysis (section 2.4.3) did a thorough job of documenting the rise, peak, decline, and resurgence of the KH instability in the linear stratification. The simplified summary is that since not enough mixing occurs during the primary billowing event, the system settles somewhat and experiences a secondary Kelvin Helmholtz-like instability. What we hope to gain using our $Q - R$ analysis is a better understanding of what the primary (or dominant) mechanisms are during each of these stages (primary event, settling, secondary event). We wish to recall ahead of time that the initial billowing event was much more energetic than the subsequent settling and re-billowing events, and as such we suspect that the dominant mechanics of the system may have changed.

The first thing we shall do is consider the Q vs R diagram at the peak of the initial mixing state (9 seconds simulation time) in figure 3.25. Immediately clear is the fact that the range of Q in these plots is significantly larger than the range of R . Since they are usually comparable quantities this leads us to conclude that the flow we see here, the classical billowing, does not experience strong variations in the axial or bi-axial motion, but instead alternates between regions of strong vorticity (in the billow cores) and strong shear (in the braids). The next observation is the vertical structure observed in the bottom-left panel indicating that the majority of shear occurs nearer to the boundary, while the vorticity occurs near the domain center, again as expected.

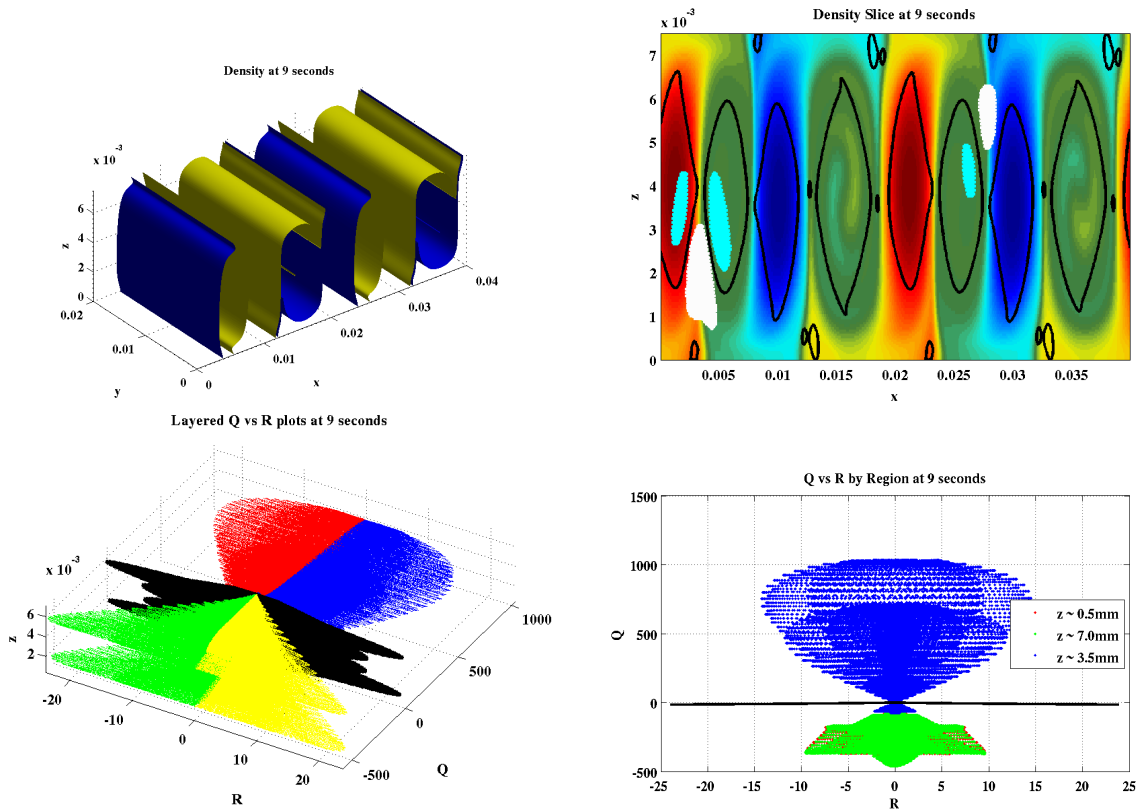


Figure 3.25: Density isosurface (top-left), augmented density slice (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 9 seconds. The linear case evolves in much the same manner as the quasi two-layer case, albeit much more slowly.

Thus far, we have established that the early stage of the linear simulation conforms to our standard picture of KH instability. The next aspect of this simulation worth examining is how the decay and calming of the primary instability appears in $Q - R$ space. From the previous analysis (section 2.4.3), we know that the system will experience some mixing (primarily in two dimensions), which eventually settles to a quasi-stable stratified state.

We next consider the Q vs R plots near the end of the initial instability in figure 3.26, when the flow has settled down and before spanwise instabilities have developed. The flow at this time is vorticity-dominated, however there is as yet almost no spanwise variation in the structure of the density profile and the original billow shape is severely deformed. The structure of the total and layered Q vs R plots indicate that vorticity is the primary

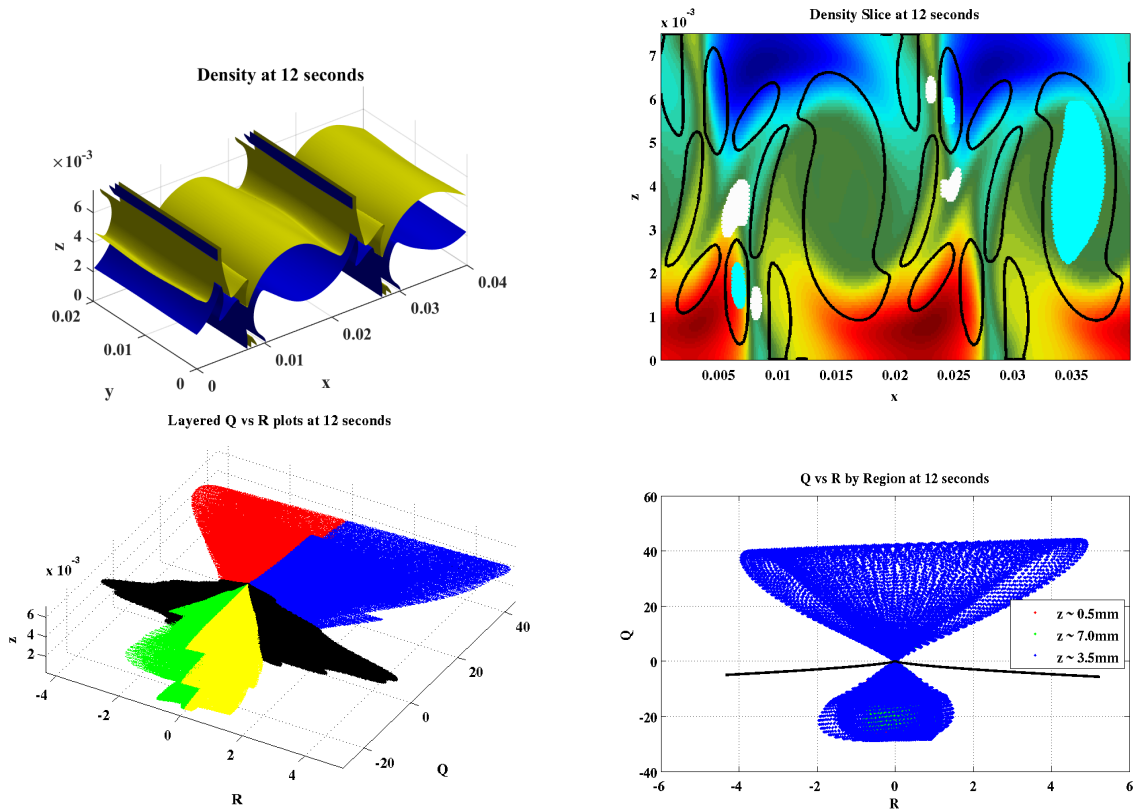


Figure 3.26: Density isosurface (top-left), augmented density slice (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 12 seconds. The simulation at this time is still fundamentally two-dimensional with vortex stretching the dominant factor present.

mechanic and the range of R indicates that there is very little three-dimensionalization at this time.

These later times, shown in figures 3.27 and 3.28, demonstrate the evolution of the linearly stratified system. The secondary rollup, which begins around approximately 14 seconds is accompanied by increased vortex sheeting (second quadrant) and axial shear (fourth quadrant) activity in the $Q - R$ plane. This transition is clearly marked in the total Q vs R plots in figure 3.27 and can in fact be useful as an indicator of when the classical billow formation begins and ends.

One final thing to note about these plots is that the ranges of Q and R give a rough

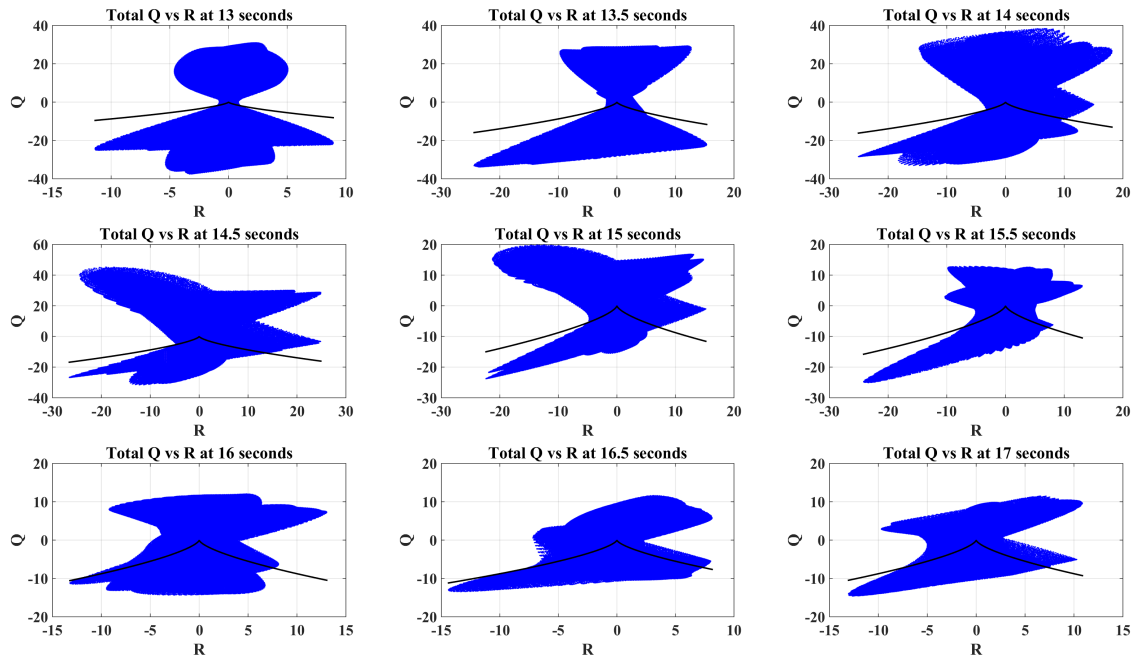


Figure 3.27: Total Q vs R plots between 13 and 17 seconds showing the initial calming period as well as the rise of the second instability around 16 seconds.

indication of the amount of energy in the system. The initial rollup is quite energetic, with values of Q reaching $O(10^3)$, while the calming period and secondary rollup only reach as high as $O(10)$. This drastic decrease in the values of Q corresponds with our earlier analysis of the rapidly decreasing kinetic energy budget, as the system becomes more mixed and more energy is lost to the mixing process.

Examining the system more closely during the secondary rollup event, figure 3.29, we notice that the $Q - R$ profile somewhat resembles the classical ice cream cone pictures, however there are slight deviations indicating the presence of significant vortex sheeting (second quadrant activity). This corresponds well with the mechanisms of billow formation, which involve largely spanwise-invariant vortices.

The above investigations in $Q - R$ space have served to efficiently identify several key aspects of this simulation: The primary mechanisms during the initial rollup events, a rough estimate of the energy present during each instability (quite high initially, and substantially lower for later times/later instabilities), and a (quite surprising) finding that later simulation times greatly resemble the expected profile for homogeneous isotropic

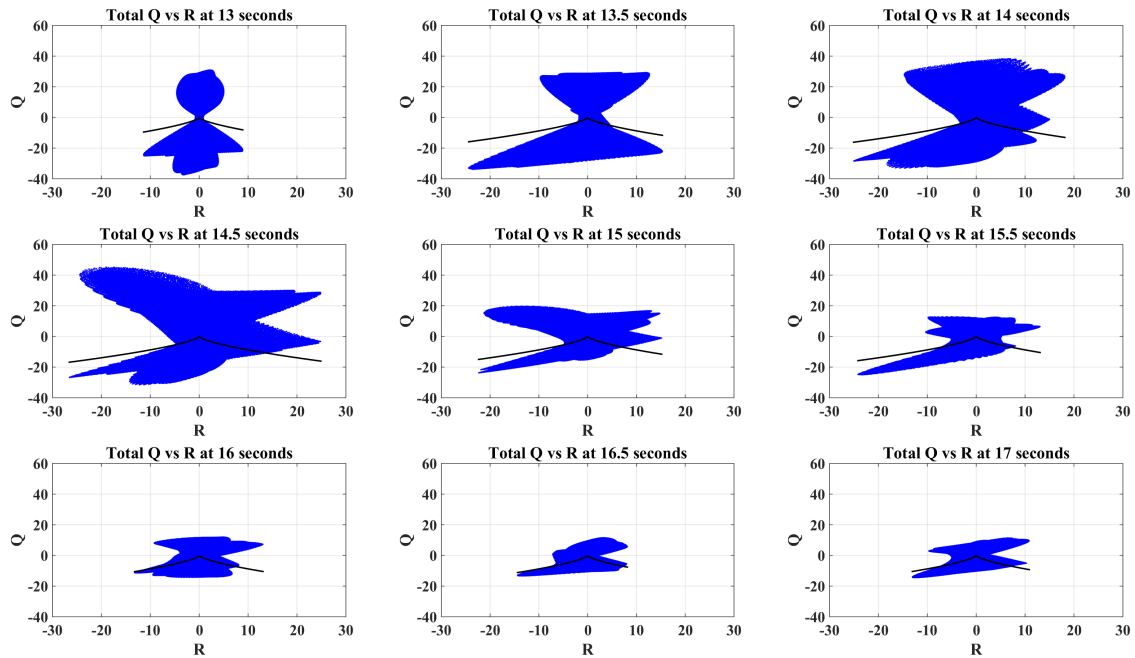


Figure 3.28: Total Q vs R plots between 13 and 17 seconds showing the initial calming period as well as the rise of the second instability around 16 seconds. Identical to figure 3.27 but with axes fixed for better comparison or the energy in the system.

turbulence. These observations serve to validate the use of $Q - R$ space as an initial diagnostic metric, serving as an efficient means to overview the evolution of the flow and to help identify spatial regions where certain elements of interest, such as vorticity or shear, dominate. Additionally combined with some basic information about the flow itself, it allows us to tag certain $Q - R$ profiles and associate them with known instabilities, such as the KH billowing event, for observation in future research.

One obvious criticism of the above considerations is that we have examined a flow which we already knew quite a bit. It's worth reminding ourselves that $Q - R$ analysis did help to tag an interesting new features (namely the 'anti-billow' regions, however, in order to truly validate this method for the analysis of instabilities we shall consider the application of the $Q - R$ analysis to a problem which we have not yet discussed, Rayleigh-Taylor instability (again in the presence of boundaries).

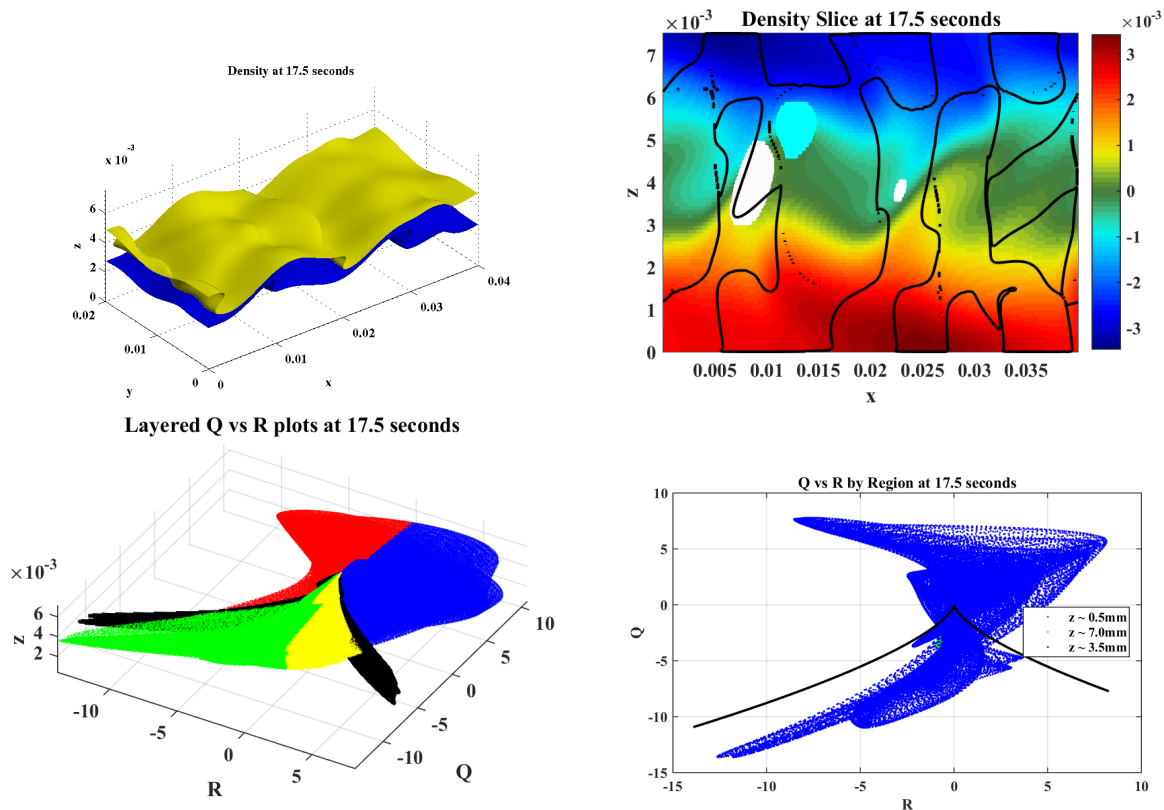


Figure 3.29: Density isosurface (top-left), augmented density slice (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 17.5 seconds during the secondary rollup event.

3.6.4 Rayleigh-Taylor Instability

The Rayleigh-Taylor (henceforth RT) instability, in which a layer of denser fluid initially over- lies a lighter layer, has been well-studied (starting with Lord Rayleigh himself [87] and Taylor [90]). Applications in fluid dynamics, oceanography, and even astrophysics have made the inviscid case very well-studied both experimentally [10, 11, 13, 20, 21, 19, 23, 48, 66, 104] and numerically [24, 57, 58, 5]. In a 2015 paper [52] we sought to numerically investigate the development of an RT instability in viscous fluids that interact with a boundary, for physically realistic parameters (by increasing the Schmidt number). A full discussion of the history of RT instability, especially as it relates to our high-resolution numerical work, can be found in [52]. We will now proceed to offer a brief explanation of

the numerical setup, and then apply our $Q - R$ analysis.

The dimensions of the tank are 8cm by 2cm by 2cm for the stream-wise, spanwise, and cross-tank dimensions, respectively. The experiment was set up so that gravity pointed in the positive x direction. The high density plug initially spanned the width of the tank, was placed near the top of the tank (centered 2 cm below the top of the computational domain), and extended over 1 cm in the streamwise direction. Periodic boundary conditions were adopted in the along channel (x) and spanwise (y) directions, and no slip boundary conditions were adopted in the across channel (z) direction. The periodic boundary conditions in the along channel direction mimic the behaviour of a long, open tube, with continuous dynamic pressure and velocity distributions across the boundary. The Schmidt number (giving the ratio between viscous and molecular diffusion rates, $\frac{\nu}{D}$) is 10, and the Reynolds number, $\frac{UL}{\nu}$ is approximately 600.

The result is a plug of dense fluid “falling” through the lighter surrounding fluid, as shown in figure 3.30. The no-slip boundaries at $z = 0$ and $z = 1$ produce a region where viscosity is the dominant factor, causing significant amount of heavy fluid to be left behind the main plug. Additionally, the shape of the falling plug is quite unique, with the dominant spanwise deformation highly dependent on Schmidt number - as discussed in depth in [52].

For the purposes of our QR investigation, there is quite a bit of activity here that we can examine. First, it will be interesting to see what the dominant mechanisms during the evolution of this plug drop are, a task for which $Q - R$ analysis is well-suited. Secondly, it is worth asking if how the mechanisms change between the near wall and near-center regions. To this end we will once again be considering layered Q vs R plots at constant z . Hopefully this will let us examine the effect that the boundary layer has on the evolution and three-dimensionalization of this flow.

We can begin our examination by considering the total Q vs R plots as the system evolves, shown in figures 3.31 and 3.32. There are several features of these plots which we can immediately use to help categorize the flow and direct further explorations.

First we can note the range of Q and R at the various outputs. Since the fluid is constantly accelerating in this simulation, we expect the range of the tertiary variables to increase monotonically with time. This is what we see, for the most part, however there is a notable exception between 2.4 and 2.6 seconds, where the range of both Q and R briefly decrease. This change is quite significant - 2.5×10^4 at 2.4 seconds down to 1.5×10^4 a decrease of roughly 30% in the outlier values - and identifies a potential area of interest, as it may indicate a shift in the mechanisms that dominate the flow’s evolution.

The next observation of note is the fact that the overall profile of the flow shifts from clearly shear-dominated at the very early times, to more complex at 1.2 seconds, and




Figure 3.30: The density evolution of the RT instability with axes set to scale. The evolution is shown from 0 to 1.4 seconds in 0.2 second increments. The total density difference shown is 2%.

finally transitioning over to the much more standard ice-cream cone shape we have come to expect from mature, turbulent flows. The early stages of the simulations (1.3 and 1.4 seconds specifically) deviate from the turbulent profile we have come to expect, another indication that there exists some fundamental shift in the flow dynamics around 2.4 seconds which is well worth further consideration.

The last thing of note is again somewhat more subtle. We notice that as the flow first begins to destabilize and develop turbulent features (1.6 to 2.0 seconds) the third quadrant

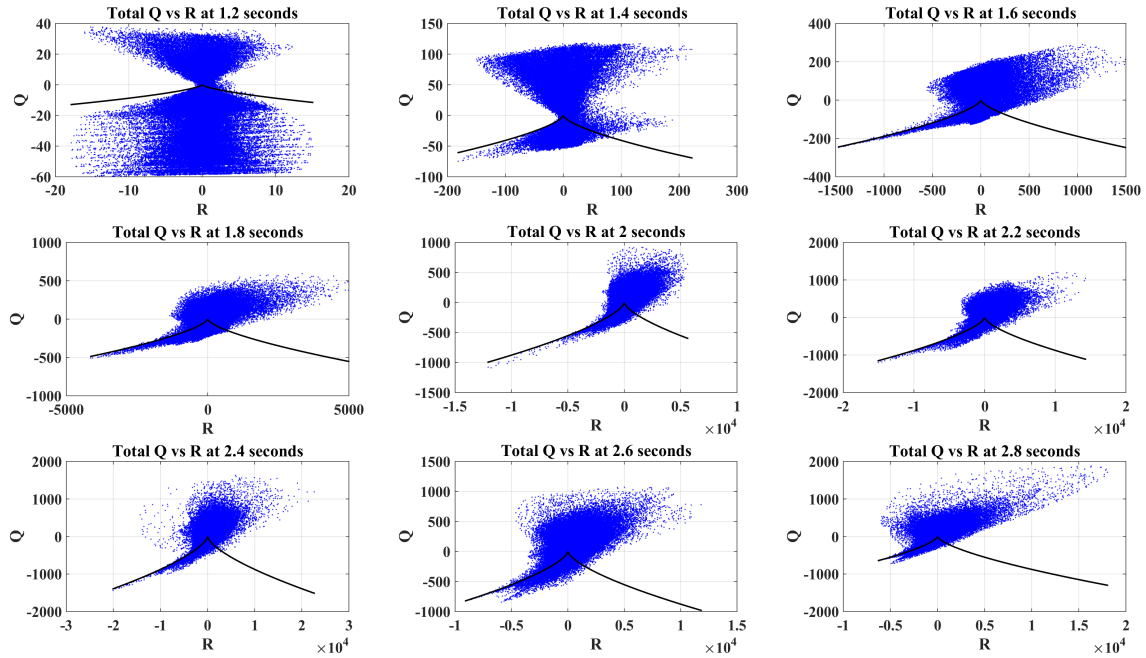


Figure 3.31: Total Q vs R plots between 1.2 and 2.8 seconds showing the entire evolution of the Rayleigh-Taylor instability.

formation begins almost exactly along the separating line, $D_A = 0$. As the flow develops further, 2.2 to 2.8 seconds, we notice the third quadrant behaviour drift away from this line. While the exact implications are not intuitively obvious, the points being scattered along (and on either side) of the $D_A = 0$ line may indicate that the flow frequently shifts between vorticity and strain-dominated regions - although it is important to remember that the locations of these regions need not be adjacent in physical space. This observation is worth keeping a close eye on, as we may wish to try and discern what is causing this behaviour and where, exactly, it occurs.. Importantly, a contour of the line $D_A = 0$ at these times may help us to identify if the flow is rapidly three-dimensionalizing in a small region, or if we have simply overlaid several different regions in this aggregate Q vs R plot.

An examination of the early times of this instability generation, shown in figure 3.33, shows several predictable features about the instability generation. The tagged regions in the augmented plot (top-left panel) show that vorticity is generated in the wake of the instability (around $x = 0.067, z = 0.004$) and shear is generated at the head of the

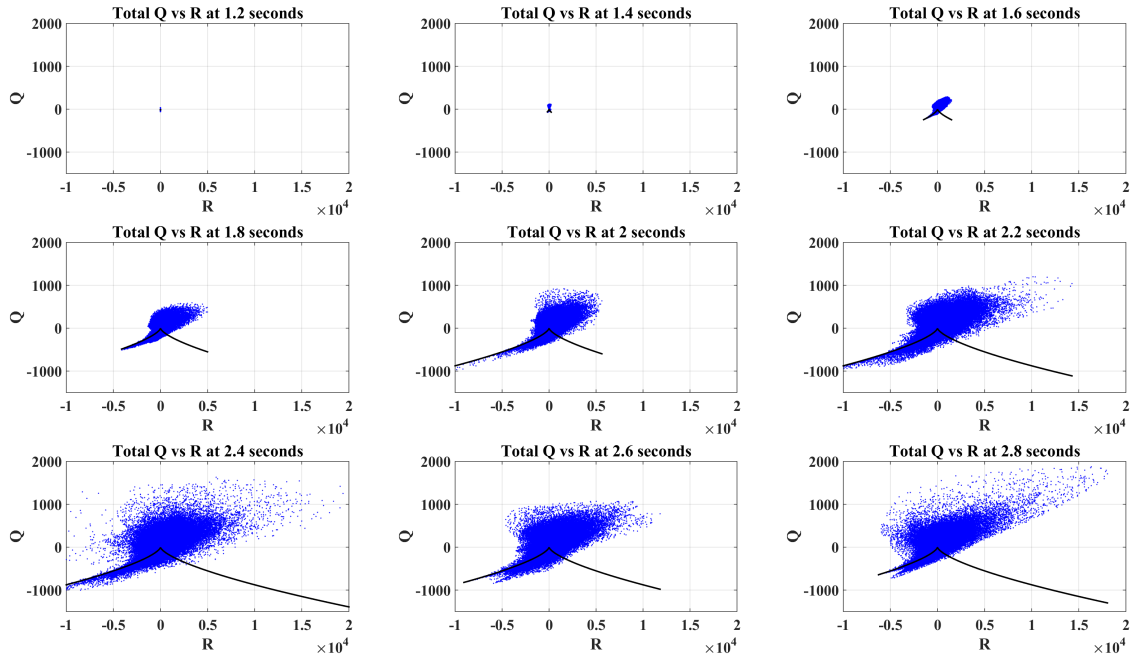


Figure 3.32: Total Q vs R plots between 1.2 and 2.8 seconds showing the entire evolution of the Rayleigh-Taylor instability. This figure is identical in content to figure 3.31 except with axes fixed for easier comparison of the energy in the system.

instability (around $x = 0.067, z = 0.019$) - remember that gravity points in the positive x direction, meaning that the plug of heavy (red) fluid travels to the right. This is consistent with classical RT instability theory [54] and is further indication of the validity of this $Q-R$ methodology. It's worth noting here that the outliers in this simulation i.e. the regions of intense vortex stretching or plane shear) are significantly more localized in the augmented density plot. This could be an effect of the fine structure inside the destabilizing plug or something inherent to RT instabilities in general. This phenomenon is interesting enough to be considered more in depth in future work.

The layered Q vs R plots (bottom-left panel of figure 3.33) demonstrate that at this time the majority of Q and R activity takes place near the boundary. Confirming this we see that the Q vs R by region plots (bottom-right panel) show that the near boundary plots (green and red) are the only ones that display any significant structure. The nature of this structure is of interest to us. First we note that the very-near boundary profile (red) displays significant vorticity, but no real shear. In contrast, the slightly away from

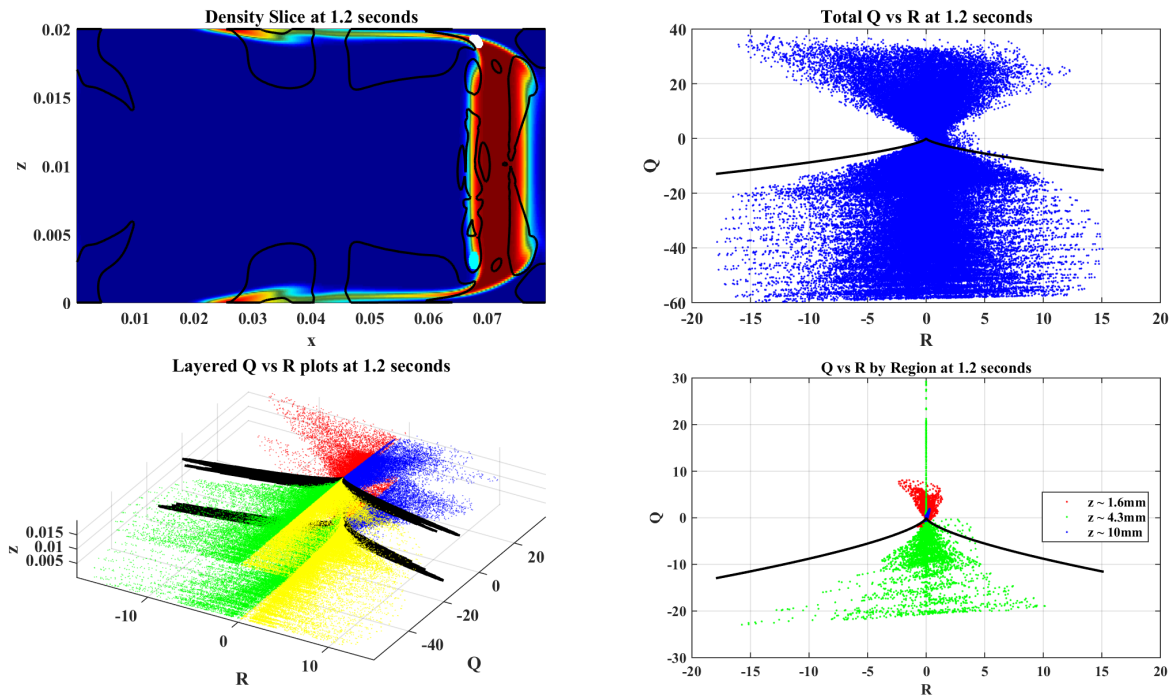


Figure 3.33: Augmented density plot (top-left), total Q vs R plot (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 1.2 seconds during the initial formation of the Rayleigh-Taylor instability.

boundary profile (in green) displays significant shear, but very little vorticity. The dominance of shear away from the boundary is quite easy to explain, the dense plug of fluid drops through the lighter surrounding fluid, causing significant shear as the lighter fluid is forced toward the boundaries. In contrast, we expect the interaction between the velocity field, which will be monotonically increasing as we leave the boundary layer, to produce significant vorticity at the boundary, and this is in fact what we see in the red scatter plot in the bottom-right panel.

Figure 3.34 shows an augmented density slice (through the spanwise center of the domain), a total Q vs R plot, layered Q vs R and Q vs R at various levels, namely inside the boundary layer (red), between the boundary and center (green), and in the center of the domain (blue). What is obvious from the combination of these plots is that at this time the majority of interesting three-dimensionalization and the dynamics corresponding to classical turbulence occur almost exclusively in the boundary layer. The blue scatter

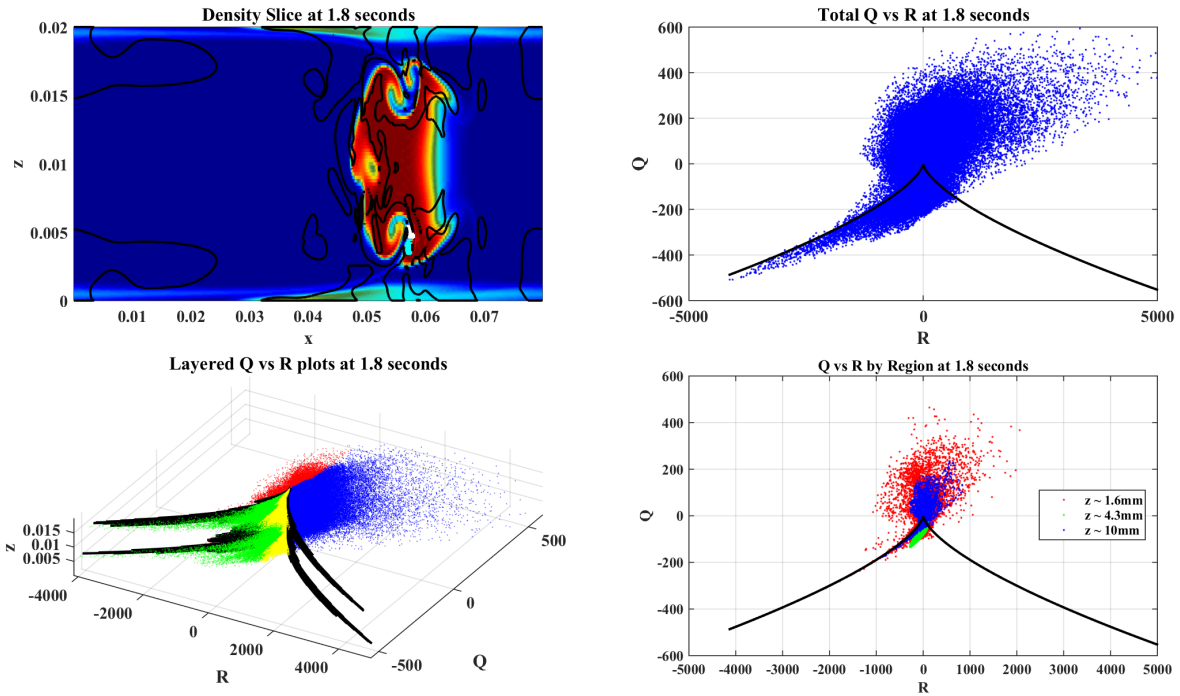


Figure 3.34: Augmented density plot (top-left), total Q vs R plot (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 1.8 seconds. The most active region is still by far in the boundary layer and the overall shape at each level seems to correspond topologically to isotropic turbulence.

plot in the bottom-right panel provides a rough picture of the QR profile near the domain center. It appears to have roughly the same ice-cream cone shape, with vortex stretching and planar shear the dominant flow topologies present at all levels.

As was the case in our analysis of KH billows, the layered Q vs R plots provide the clearest representation of the differences in the profiles at various z values. The black lines, representing the $D_A = 0$ curve, whose range is determined by the local maximum and minimum values of R , clearly show that the system is most active nearest to the boundaries, growing relatively calmer toward the domain center. This is a strange contradiction with the fact that the center of the domain is accelerating the fastest. However, further thought immediately reveals the fact that the no-slip boundaries will cause the highest velocity gradients, due to the presence of viscosity and the rapidly changing nature of the velocity profile at the boundary, this easily explains the much more active boundary region in the

$Q - R$ metric. It also provides a hint as to the later (2.4 second) decrease in the range of Q and R : if the initial drop has completed its interaction with the boundary by this time, we expect significantly smaller velocity gradients, despite increasing absolute velocities.

One final feature stands out quite strongly about the boundary layer. The dominance of vorticity near the boundary is difficult to explain at first, but leads us to an interesting observation. The dominant vorticity production occurs in the near-wall region for vortex lines running in the spanwise direction - we witnessed this at the earlier timestep. As these lines are perturbed, extending them further into the domain, they find themselves in a larger background velocity (since the velocities increase monotonically as we leave the no-slip boundaries. This system stretches the vortices in a way analogous to the hairpin vortex mechanism (see figure 4.13 in Davidson [25]), causing vortex deformation throughout the boundary layer and ensuring large activity in the positive part of the Q vs R plot. We therefore fully expect this mechanism to continue unabated at the boundary for the remainder of the simulation, causing vortex stretching to remain an important component near the boundary.

Although we could continue our analysis of this instability using Q and R indefinitely, we shall have to content ourselves with one final foray, namely into the reduction of Q and R between 2.4 and 2.6 seconds. Observing the state of the system at 2.4 seconds (figure 3.35) and 2.6 seconds (figure 3.36) we notice that by this time the internal instability has clearly destabilized the flow and induced a transition to a turbulent regime at both output times. We observe that the interior of the centre of the instability is rife with pockets of vortex stretching, a key mechanism in the destabilization of this once coherent dense plug. Additionally we see significant regions of high shear interspersed throughout the region (and alternating heavily with the regions of vorticity). As we predicted earlier in this section, the region rapidly alternates between vorticity and shear-dominated regimes, a classic feature of isotropic turbulence. As a result of this in both of the figures the line $D_A = 0$ swerves throughout the domain quite erratically, never taking on any really coherent shapes - indicating that plotting isosurfaces of $D_A = 0$ would not provide any useful diagnostic measure.

What is quite interesting is that we still see significant interaction between the traveling dense fluid (it would be semantically incorrect to continue to refer to it as a plug since it has clearly lost all resemblance to its former shape) and the boundary layer. The lines of $D_A = 0$ easily cross the boundary layer around the central instability. The effects of this are more muddled in the later (2.6 seconds) figure, however, this is likely due to the periodic boundaries and some measure of self-interaction due to the extremely high velocities by this time.

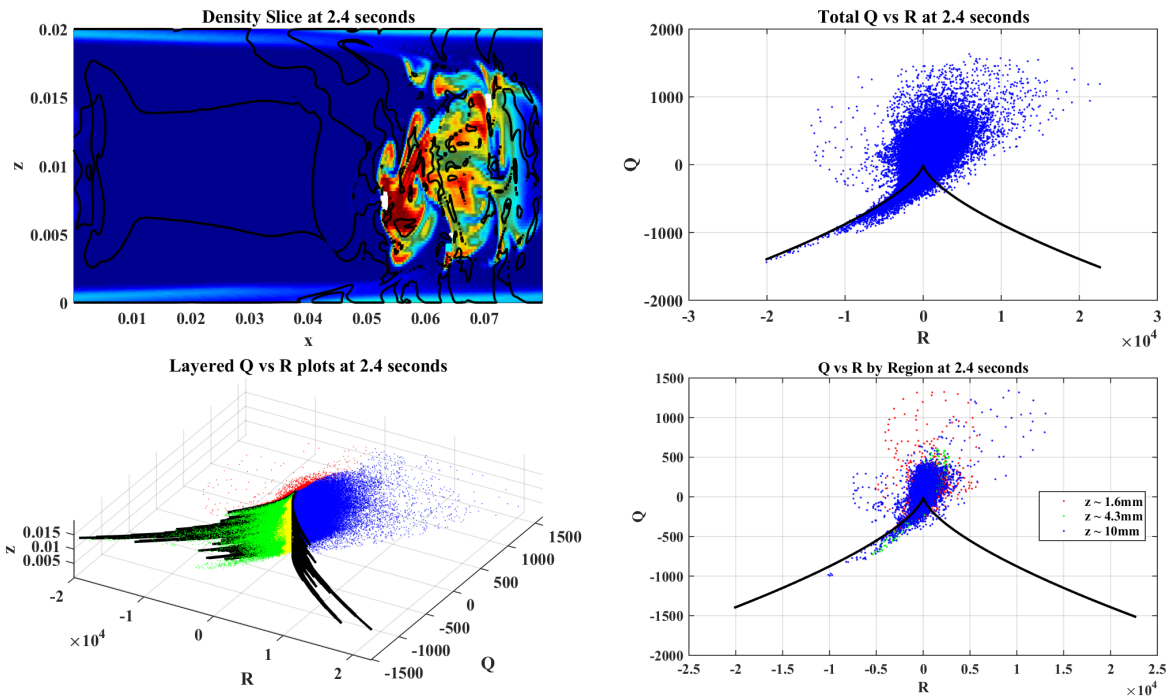


Figure 3.35: Augmented density plot (top-left), total Q vs R plot (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 2.4 seconds. The center has fully destabilized, causing increased velocity gradients and making the center of the domain dominate the Q vs R region plots.

Finally, we return to the question of why the range of Q and R has dropped between these two outputs. Examining the system closely we notice that the destabilization of the dense plug reaches a peak around 2.4 seconds. It would appear that our initial glance (using Q and R) revealed the point in the simulation at which mixing reached a peak in the core. For completeness, we present a summary line plot in figure 3.37.

The kinetic energy, internal, and total dissipation all monotonically increase as the simulation progresses. The mixing, represented by the gradient of density, $\nabla\rho$, appears to reach a maximum at around 2.4 seconds. This corresponds with our identification of a regime change via the range of Q and R .

In all, we have seen how Q and R can help to identify and diagnose several key aspects of a new instability. The initial assessment using $Q - R$ tools led us to identify a possible episode during which a transition in the dynamics takes place. Further analysis led to the

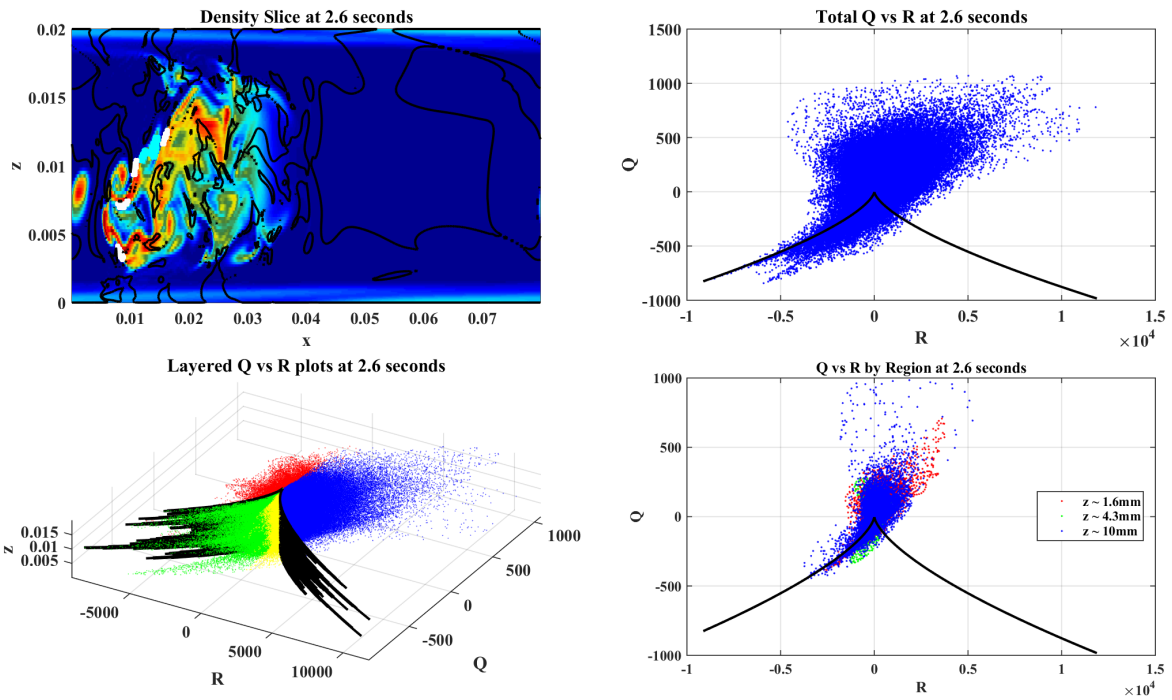


Figure 3.36: Augmented density plot (top-left), total Q vs R plot (top-right), layered Q vs R plot (bottom-left) and Q vs R by region (bottom right) at 2.6 seconds. The plug has completely destabilized and is riddled with alternating regions of high shear and vorticity (in cyan and white on the augmented density plot).

discovery of a relatively unique interaction between the boundary and the velocity profile created by the falling density plug similar in nature to hairpin vortices.




Figure 3.37: Summary plot of Rayleigh-Taylor instability. As expected the kinetic energy and dissipation (both total and internal) rise monotonically as the plug accelerates. In contrast the mixing reaches a peak around 2.4 seconds.

3.7 Discussion

This chapter has presented several interesting facts worth re-iterating here. We begin here, as we began the chapter, by pointing out that there exists a method for analyzing data, namely the tertiary variables Q and R , which requires access to the full velocity field, but yields very useful ways of classifying flow topology. Through analytical explorations and application to standard (canonical) flows such as the Burger’s and Taylor-Green vortices we confirmed that the data sets generated by these tertiary analysis variables can be used to efficiently identify regions of interest, as well as the dominant mechanics in those regions. Although the resulting Q and R relationships are not necessarily sufficient (or unique) in describing a flow, they do allow us to compare the relative dominance of certain mechanisms throughout a particular flow’s evolution. This leads to a brand new way of considering a three-dimensional flow’s evolution, namely the resulting Q vs R profile - either the total profile or the profile only in a region of interest to the investigation.

One interesting facet of our analytical explorations is the unique structure of Q and R in the Burger’s vortex. Since in classical turbulent literature the primary mechanism for energy transfer from large to small scales is vortex stretching we would expect to identify local Burger’s vortex-like structures throughout the domain and it is possible that Q vs R may provide a mechanism for such an analysis.

Having confirmed the validity of our chosen metrics, we then explored several visualization tools. Though we by no means covered the entire range of possible visualizations, we did discover several unique ways to visualize flow topology. First, since the most interesting behaviour (at least for our case studies) involved vortex stretching or planar shear, we found that augmented density plots could serve as a useful way to identify regions of the flow that displayed these (or theoretically other) interesting mechanisms. The next observation of note is that Q and R provide a dramatic amount of information about the flow, and as such one of the most useful ways to use them as diagnostic tools is to limit the regions to which they are applied. By confining ourselves to near-center or near-boundary regions we can immediately identify the dominant mechanics at a given time, in a given place. The obvious downside to this methodology is that you have to know *a priori* where to look, however in our case studies the presence of boundaries provided obvious locations of interest.

Having developed (and confirmed) our chosen graphical tools we then presented four case studies - three Kelvin-Helmholtz simulations presented in the previous chapter, and a brand new Rayleigh-Taylor instability chosen for contrast. Classical isotropic turbulence, to which these metrics were first applied, exhibits a well-established “ice-cream” cone pattern, starting in the third quadrant and moving into the first quadrant through the origin. It is a verification of the universality of Kolmogorov’s theory that this profile, with vortex stretching and plane shear as the dominant mechanics, is manifested in both of our case studies (KH and RT) when the respective instabilities begin to truly three-dimensionalize and the flow transitions to turbulence. At these times, while possibly diminished in energy, the flow becomes closer to classical isotropic turbulence. The Q vs R plots, which are similar for all later times, begin to resemble the classically isotropic turbulent regime.

The next aspect of $Q - R$ analysis worth reiterating is that there is a hidden diagnostic in the range of Q and R at various regions or times. The value of Q and R are determined by the eigenvalues of the velocity gradient matrix. Thus the range of Q and R , and specifically the value of the outliers, gives some indication of the amount of energy in the system. Using this fact we were able to identify when the flow transitioned to a lower-energy regime: in the KH billows once the period of two-dimensional mixing ended; or in the RT instability when the dense plug destabilized, again coinciding with a drop in mixing of the system.

Finally we come to the meat and potatoes of $Q - R$ analysis: identification of flow topology. We have seen in several case studies that the dominant region, and specifically points more than a standard deviation from the mean in $Q - R$ space, identify the dominant flow topology, and thus the mechanics of the system. In several notable cases, such as the identification of vortex stretching as a dominant mechanism near the boundary, we were

able to demonstrate a relationship with classically described behaviour such as hairpin vortices. Additionally, the standard profile of isotropic turbulence, from which these metrics and tertiary variables come to us, makes for a reference point about which we can base our studies of transitional flows, and which creates a fixed point in our analysis of the transition to turbulence.

Q and R offers a way of identifying the flow topology, and by extension the dominant mechanics, of a system; however, it does require access to the full velocity gradient tensor - something which is difficult for experimentalists. The computation of eigenvalues (for generation of Q and R) can be done locally (on a point by point basis), however, this means the computation of 3 eigenvalues from a 3 by 3 matrix makes for a large amount of data to track. Despite the limitations it is my firm belief that as the study of $Q - R$ space develops further we will come to associate specific $Q - R$ profiles with known instabilities - much the same way that we now associate the ice-cream cone with isotropic turbulence - leading to a simple way to characterize flows.

Chapter 4

Extensions

For the final chapter of this work, we will present some extensions of the ideas and results discussed in previous chapters. While somewhat disjointed, this chapter will seek to round out the previous analysis and hopefully serve to help direct future study.

The Kelvin-Helmholtz chapter (2) presents a thorough examination of several confined tilted tube experiments with varying pycnocline locations, however the role of the no-slip boundary, while intuitively obvious, is never fully addressed. Specifically, the formation of the unique ‘anti-billow’ coherent structures from the 0.75 cm, asymmetric pycnocline simulations is worth further examination in order to determine whether it was the confinement or the no-slip boundaries which are responsible for this behaviour.

The next chapter (Q and R , 3) presented a method for the analysis and characterization of flow topology and the identification of dominant mechanics in several numerically simulated flows. The applicability of QR analysis to a different type of numerical simulation (namely Large Eddy Simulations) is investigated in order to demonstrate the methodology’s applicability for a range of numerical methods.

Finally, since the usefulness of Q and R analysis as a visualization and flow characterization tool has been established, we now wish to consider some other visualization and characterization methods from the published literature in order to see how QR analysis measures up.

4.1 KH Extensions

We have been relatively thorough so far in our investigation of confined KH billows. We investigated various domain sizes, pycnocline placements, and even considered the limiting case of the linear stratification.

Arguably, the most interesting result observed occurred in the 0.75 cm domain with the pycnocline placed at 25% of the fluid depth. We observed the formation of coherent billow structures in near constant density fluid, which we labeled as “anti-billows” due to their symmetry with the primary billow structures and their placement in fluid which contained almost no density gradient. A natural question that arises from this result is whether the formation of these anti-billows is the result of the domain size, the presence of boundaries, or the no-slip conditions we have enforced.

The easiest way to determine whether it is in fact the confinement or the no-slip boundaries that is responsible for the formation of anti-billows is to rerun the same simulation (with the same parameters) with free-slip boundaries. This section presents the results of a two-dimensional simulation with the same resolution and parameters as the 2D runs from section 2.3.5 and concludes that it is in fact the no-slip boundaries that are responsible for the formation of anti-billows.

The onset of the KH instability occurs far outside the influence of the boundaries, and as such the no-slip and free slip simulations are identical up to a point. In the no-slip simulation, the initial destabilization occurs around 1.3 seconds, with the formation of classically shaped billows evident by 1.5 seconds. The billows rapidly begin to take up a significant portion (roughly 15%) of the vertical domain. In contrast (see figure 4.1) the free slip simulation experiences a faster initial destabilization (1.2 seconds for roughly the same size and structure billows as before). Additionally, the later outputs clearly show that once the billows have grown to their maximum natural size (roughly 50% of the vertical domain) they immediately begin to experience subharmonic pairing. This result is not surprising, as we saw very similar behaviour in the 2cm tall domain where the billows, less confined but even in the presence of no-slip boundaries, tended to pair and thereby disrupt their classical shapes, just as we see here.

If we next consider the viscous dissipation at these same times, figure 4.2, we see that the only significant dissipation present throughout is during the initial broadening of the pycnocline and initial destabilization event. The dissipation at even earlier times (not shown) is present only around the region where the pycnocline is broadening, in the top half of the domain. This dramatic decrease in overall dissipation is surprising, but not unexpected as it was the boundary layer in the no-slip simulations that contributed the

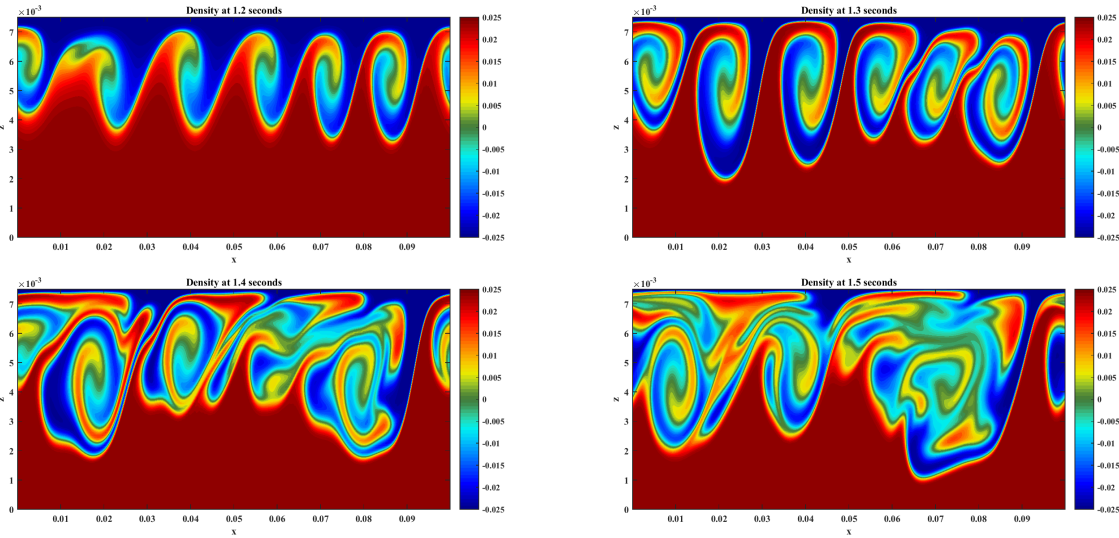


Figure 4.1: Density plots of the asymmetric 0.75 cm simulation with free-slip boundaries at 1.2, 1.3, 1.4, and 1.5 seconds. The initial destabilization occurs slightly faster (1.2 for free slip compared to 1.3 seconds for no-slip boundaries).




Figure 4.2: Dissipation plots (with contours of density in white) of the asymmetric 0.75cm simulation with free-slip boundaries at 1 and 1.1 seconds. The only significant dissipation present is during the initial pycnocline broadening. Once the billow formation begins there is almost no dissipation present.

vast majority of dissipation. Likely there is *some* dissipation that takes place even in the free-slip case, and for this, just as in the interior of the no-slip domain, we must consider a log plot of dissipation in order to see what little structure there is.




Figure 4.3: Log of dissipation plots of the asymmetric 0.75cm simulation with free-slip boundaries at 1.2 and 1.4 seconds. The structure of the dissipation follows fairly closely the shape of the billows, however the magnitude is so weak that the result is nearly negligible.

The structure of what limited dissipation is created in this simulation is shown in figure 4.3. The shape of the dissipation-containing regions closely follows the outer edges of the cores and the braids, however, the magnitude when compared to the dissipation caused by the spreading of the pycnocline (earlier in this simulation) or the dissipation caused by interaction with no-slip boundaries in the other 2D simulations is so small as to be effectively negligible.

Having discounted almost entirely the produced dissipation, the final thing left to investigate in this simulation is the structure of the billows at later times. Figure 4.4 shows density contours of the simulation between 1.6 and 2.0 seconds. It is immediately clear that the primary factor in this simulation is the subharmonic pairing, or billow-billow interactions. What is clearly absent is any form of coherent structure corresponding to the anti-billows we witnessed in the no-slip case. The paired billows are not confined to the top half of the domain and by 1.8 seconds it appears as if significant mixing has taken place, nearly destroying the original billow structures.

We have seen above that given the same initial setup as the no-slip case, a simulation with free-slip boundaries does not produce the coherent structures from its no-slip counterpart. The primary mechanism throughout (once the billows are formed) is billow-billow (as opposed to billow-boundary) interaction, producing a significantly more chaotic system lacking the nice symmetry we witnessed before.

Having performed this experiment we are satisfied that it is in fact the combination of confinement (2 cm domains having produced much less interesting billow-boundary interactions) and no-slip boundaries (free slip boundaries having produced no significant coherent structures) that is responsible for the unique ‘anti-billows’ which were of such interest in previous Chapters.




Figure 4.4: Density plots of the asymmetric 0.75cm simulation with free-slip boundaries at 1.6, 1.7, 1.8 and 2.0 seconds. Note the total lack of coherent structures in this flow with subharmonic pairing dominating nearly ever aspect of the billow evolution.

4.1.1 Reynolds Stress

Another reasonable way in which to consider how the boundary affects the evolution and destabilization of the KH billow regime is to consider the behaviour of the Reynolds stress (the ensemble-averaged stresses) at various proximities to the boundary layer.

The classical Reynolds decomposition consists of splitting the velocity into a mean and perturbation part based on an ensemble of experiments:

$$\mathbf{u} = U + u,$$

where \mathbf{u} is the full velocity, U is the mean velocity, and u is the deviation from the for a specific simulation.

The standard method for computing this decomposition requires an ensemble of simulations; many simulations repeated in parallel with nearly identical starting conditions and averaged to yield the mean flow. Unfortunately we do not possess the resources required to perform a meaningful number of identical simulations. However, it is worth noting that up to the spanwise destabilization of the billows the simulations are primarily two-dimensional. As such, we can treat each slice in the spanwise (y) direction as a separate

two-dimensional simulation and average them to yield a meaningful mean flow. Thus, we can compute a kind of pseudo-Reynolds stress defined as

$$\rho_0 \langle u_i u_j \rangle . \tag{4.1}$$

In this context, this is a symmetric mathematical object (to truly be referred to as a tensor we would have to show invariance under changes of coordinates, but this would require a redefinition of the averaging in the spanwise direction) which can be used to gauge a measure of the anisotropy of a system. The way that we will be using it below, the Reynolds stress gives a measure of deviation from the “mean” state, i.e. the 2D state, thereby allowing us to gauge the rate of spanwise destabilization. We have computed these Reynolds stresses for the asymmetric three-dimensional 0.75 cm tall simulation (as seen in sections 2.4.2 and 3.6.2).

We will consider this simulation between 2.1875 and 2.625 seconds when the destabilization of the primary KH billow is taking place. We can consider the mean velocities, U, V, W , the Reynolds stresses, $\langle uv \rangle, \langle uw \rangle, \langle vw \rangle$, and the horizontally averaged mean velocities and stresses.

Figure 4.5 shows the mean velocities (left panel) as well as the Reynolds stress terms (right panel). As we have seen before the dominant velocities are in the horizontal and vertical directions, thus the greatest variations will be in those velocities, and it is only the $\langle uw \rangle$ stress which is active. This figure shows the structure in the vertical velocities as befitting the internal velocities of the billows, as well as the locations of highest stress, namely near the bottom boundary. While this is useful, a more succinct visualization is to sum the averaged stress (and mean velocities) horizontally to see the total stress as a function of distance from the boundary.

Figure 4.6 shows the horizontally integrated mean velocities and Reynolds stresses at 2.1875 seconds. Note that by far the mean velocities are much more active than the Reynolds stresses. Additionally it is clear from this plot that the horizontal velocities peak around the billow and anti-billow cores (black line) while the spanwise velocities (dashed line) peak near the far boundary. Turning to the Reynolds stresses (blue red and green lines) we see that the only significant activity is in the $\langle uw \rangle$ component. This is easily understood, since the flow is two-dimensional and it is the coupling of vertical and horizontal velocities that produce the only real anisotropy. It is also worth pointing out here that this component is the only one that truly goes negative in any meaningful way.

The relatively early times chosen in figures 4.5 and 4.6 implies that the Reynolds stresses during this part of the simulation when the initial destabilization is only beginning are largely immaterial. The flow at this time is relatively quiescent in the spanwise direction,

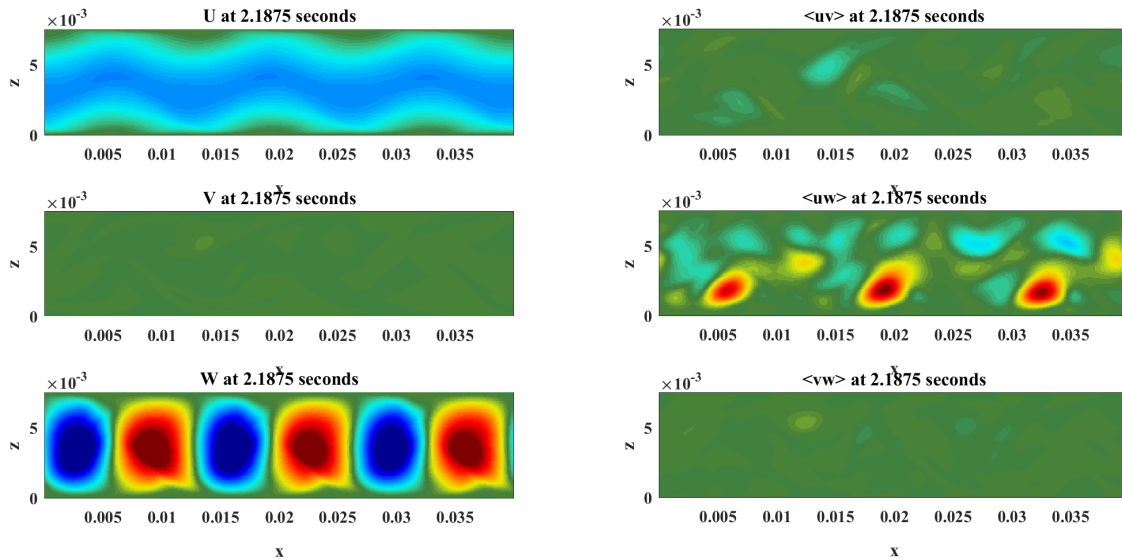


Figure 4.5: Mean velocities (left) and Reynolds stresses (right) at 2.1875s for the asymmetric simulation scaled by the maximum Reynolds stress at this time. The largest mean velocities (and the largest variations causing Reynolds stress) occur in the horizontal and vertical velocities.

with the instabilities only beginning to grow. In order to see how the stresses change we must consider a slightly later time, since the timescale of these secondary instabilities is very short, a snapshot at 2.375 seconds is sufficient.

Figure 4.7 shows the integrated mean velocities and Reynolds stress components at 2.375 seconds. The horizontal velocity is clearly dominated by activity in the bottom two-thirds of the domain where heavy fluid is being pulled downward by gravity. The next striking result is the dramatically different structures of the velocity profiles near the top and bottom boundaries. Specifically we see that the mean vertical velocity, W is nearly zero near the boundaries, rising in the center. This is especially interesting when contrasted with the fact that the horizontal and spanwise mean velocities are very high near the boundaries. Turning to the Reynolds stresses we notice that again (and throughout the destabilization process) the only significant source of anisotropy is the $\langle uw \rangle$ component which goes negative in the top half of the domain and positive in the bottom half. Physically, $\rho_0 \langle uw \rangle$ is the vertical flux of the the part of the downstream momentum that is not invariant in the spanwise direction. What this means is that this




Figure 4.6: Horizontally integrated mean velocities and Reynolds stresses for the asymmetric simulation at 2.1875 seconds. It is clear from the integrated values that the mean velocities are far more important than the Reynolds stresses.

quantity correctly identifies billow cores (and possibly braids) as the regions where the secondary instabilities arise.

A key discernible feature of the Reynolds stress plots is the behaviour near the boundaries. The spanwise activity is clearly dominant near the boundaries (as witnessed by the fine scale structure of the U and V velocities), but the Reynolds stresses at the boundaries are very low. Additionally we see almost no vertical activity in the boundary layer - again this is expected as we cannot flow through the boundary, but it does imply that fluid exchange into and out of the boundary layer will be much less than the activity within it




Figure 4.7: Horizontally integrated mean velocities and Reynolds stresses for the asymmetric simulation at 2.375 seconds. The mean and spanwise velocities are much higher at this time, and it is clear from the spanwise mean velocity (dotted line) and the $\langle uw \rangle$ component of the stress tensor that the spanwise instabilities are much more pronounced.

induced by the background shear and the billow core velocities.

This analysis of Reynolds stresses is a good demonstration of precisely how the system destabilizes via secondary instabilities. While this analysis did shed some interesting light on the near-boundary behaviour of the mean velocities the majority of the information derived we had already acquired via our Q and R analysis. The construction of a true Reynolds stress from an ensemble of simulations is an interesting avenue for future work.

4.2 QR and Large Eddy Simulations

The numerical simulations whose results have been reported in this thesis have been performed using SPINS, a spectral Direct Numerical Simulations (DNS) code. An immediate, and natural, question that should come to mind is why choose a DNS code. Such simulations are notoriously computationally intensive and time-consuming (though it is worth noting that due to our confined setting our simulations fall within the low–medium end of the grid size specturum). If it is the transition to turbulence we are interested in, why not choose an approach more common in the turbulence literature, namely Large Eddy Simulations (LES).

The honest answer to this question is of course that SPINS was available, performed well on the hardware available to us, and proved capable of resolving both the interior of the fluid and the boundary layers near the no-slip boundaries which proved so instrumental in confined KH billows. Additionally, the Reynolds numbers for the flows we were investigating were well within acceptable limits for DNS in general (between 50 and 1000), and SPINS in particular [88]. However, it may be worth asking regardless if the tools for analysis that we have developed here, namely the Q and R metrics from chapter 3, can be applied to other standard techniques in turbulence simulation such as LES.

In general, LES falls under the category of space-averaged equations, since the goal is to separate the flow into two scales and attempt to determine the contribution of the unresolved (subgrid scale) motions in the fluid. Wyngaard in his book, “Turbulence in the Atmosphere” [101] presents a good summary of how space-averaged equations can be used to account for fluxes and discusses the modelling of these fluxes. The goal of LES is to resolve the mean flow and the large, energy-containing eddies exactly [25] while small-scale structures and their influence is accounted for by the specific sub-grid-scale scheme used. In summary, we result with two types of averaging which depends on the spatial scale used (Wyngaard refers to this scale as Δ) relative to the size of the energy-containing eddies present in the flow, ℓ . These categories can be classified as coarse resolution, $\Delta \gg \ell$, or fine resolution, $\Delta \ll \ell$.

A key weakness of LES is the presence of boundaries, where the dynamically relevant mechanics (namely the generated eddies) are very small [25]. Techniques for dealing with boundary-layer dynamics in LES simulations range from mesh refinement near the boundaries (effectively becoming locally DNS) to separating the model of the boundary layer with an eddy-viscosity model and incorporating this in to the greater LES scheme. Both of these methodologies have significant drawbacks, however for the purposes of this section (and the Q and R analysis therein we will be satisfied to use a standard LES-averaging scheme

even near the boundaries in order to see what features (and what dynamics located by Q and R) are preserved or lost.

4.2.1 Simulated Filtering

As was outlined in the introduction, LES is a method of simulating fluids which seeks to separate a given field, f into two parts, the resolved part - denoted by f^r - and the unresolved, or subgrid scale part - f^s ,

$$f = f^r + f^s$$

The specific type of filter used in LES simulations can be applied either in physical space, for example a box car filter, or a Gaussian filter, or in spectral space, for example a wave-cutoff filter (which discards all wavenumbers below some critical value) [25].

Since we have produced several simulations with all components of velocity field $\mathbf{u} = (u, v, w)$ stored, we can apply a filter to our results *a posteriori* in order to simulate the output of an LES simulation of an identical setup.

The filter we have chosen to use here is a Gaussian function given by:

$$G(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}}, \quad (4.2)$$

where σ is the cutoff radius for our filter scale. From this equation $\frac{1}{\sqrt{2\pi}\sigma}$ is the normalization factor. The radius we chose for our subgrid scale is $\frac{L_z}{20}$. We have chosen this σ somewhat empirically as a value which does not remove too much of the fine structure in the simulation while providing sufficiently LES-like results. Since our simulated grid scale, Δ is much larger than the Kolmogorov length scale we can think of this as a type of coarse filtering for dealing with isotropic turbulence.

We will apply this Gaussian filter to several outputs of the three-dimensional simulations presented in sections 2.4.1 and 2.4.2. The process of averaging the domain at every point is computationally very intensive, requiring a triple integral at every point. In order to make this task more feasible we make certain simplifications. First, we will only create filtered velocity outputs at a specific vertical slice. That way we can do our standard Q and R analysis for an augmented plot. A vertical slice should capture enough information about the state of the system, as well as identify the secondary instabilities.

Since the Gaussian filter is only really applicable within a small range of the point in question, we can use a simplified windowing function in order to sum only within a few

dozen points, thus reducing our computation time significantly. The nature of a Gaussian filter is such that the contributions from points in the near vicinity of the central point is dramatically larger (orders of magnitude) than contributions from further away. As a result we need only consider the weighted values of points in the vicinity of our central point (creating a window function). The most dramatic example of this is in the x direction, where we have 512 points. Past 30 points from the point of interest (0.235 cm when compared to a radius of 0.0375 cm) the contributions of the fluid velocities are very close to negligible. Thus by computing our integral over 60 points, rather than the full 512 at each point we save a factor of 10 on our computation time with no loss in filtering accuracy.

The above process results in a set of velocity variables, u^r, v^r , and w^r which can be used to compute the pseudo-LES generated velocity gradient matrix, A^r . From them we can compute the resolved scale eigenvalues, λ_i^r and calculate the resolved scale, or filtered, Q^r and R^r variables.

One final thing worth noting is the accuracy of the filtered derivatives. In order to save on computation time and to create meaningful results we have attempted to use the minimum number of points. As such, in the vertical direction – which uses a Chebyshev grid – we have preserved all points such that the Chebyshev differentiation matrix (and associated spectral accuracy) can still be used [94]. In order to avoid having to compute the entire y domain to get a single slice, we have chosen to use a 4th order central differences scheme, requiring 5 points (technically 4 since the central point is not used). The scheme to find the x (and similarly y) derivative at the point u_i is as follows [18] p49:

$$\frac{\partial u}{\partial x} = \frac{-u_{i+2} + 8u_{i+1} - 8u_{i-1} + u_{i-2}}{12\Delta x} + O(\Delta x^5), \quad (4.3)$$

where the subscript denotes the grid point number and Δx is the grid spacing.

Although we are losing our (quite enviable) spectral accuracy for the derivatives, the resulting fourth order accuracy is sufficient for demonstrating the value of applying Q and R to LES-generated flows.

Before considering the effects of our LES filter on on the flow it is worth examining whether our resulting flow fields have retained the key features of the simulation. Although we do not expect it (since our choice of a Gaussian filter and relatively small filter radius should preserve the velocity field) it is worth convincing ourselves that our averaging process has not dramatically altered the simulation results. To that end, figure 4.8 shows the unfiltered vs filtered fields (left and right panels respectively) for the derivatives of the horizontal velocity, u , (top, middle, and bottom panels for x , y , and z derivatives, respectively). We present here only a single sample time (2.25 seconds). The remaining

outputs were qualitatively verified (not shown) to be structurally similar for the filtered and unfiltered velocity gradient components.

Finally it is worth noting that the algorithm was validated, by testing it with an extremely small σ value (effectively a delta function), which produced velocity fields identical to the original DNS fields, as expected.




Figure 4.8: contour plots of horizontal velocity derivatives ($\frac{\partial u}{\partial x}$, $\frac{\partial u}{\partial y}$, and $\frac{\partial u}{\partial z}$ in top, middle, and bottom panels respectively. Unfiltered (DNS) variables are on the left and filtered (LES) variables are on the right.

4.2.2 0.75 cm Symmetric Domain

The first simulation we will examine with the above-described LES averaging methodology is the symmetric 0.75 cm domain. As we saw in chapters 2 and 3, evolution in this domain follows a standard early KH billow evolution, followed by the rise of spanwise instabilities (around 2 seconds) and eventual destabilization to a fully turbulent regime. Since we are interested primarily in the effects of the filtering on our turbulent metrics Q and R , we will consider only relatively active times (as identified from our earlier analysis).

The first output we will consider occurs at 2.25 seconds. This output time has well formed billows, with spanwise instabilities beginning to significantly affect the structure of the flow. Figure 4.9 shows the unfiltered (top) and filtered (bottom) total Q vs R plots at this time. The effects of the filtering are immediately obvious in this figure. We expected *a priori* that the filtering would smooth out the velocity gradients, thus reducing the outliers in the velocity gradient matrix. This is indeed what we see in figure 4.9, with reduced activity in the extreme ranges of all four quadrants. Since the flow at this time is relatively regular, i.e. not truly what we would consider turbulent, we expect the filtered velocity gradients to remain mostly true to their DNS counterparts.

The most significant difference between the filtered and unfiltered results lies in the 3rd and 4th quadrants. Thus the greatest difference is in reduced shear, both planar (3d quadrant) and axial (4th quadrant). This reduction is relatively easy to explain. The primary sources of planar shear in our KH simulation lie in the braids (thin regions connecting billow cores) and between the billows and the boundary layer. Both of the described regions will be relatively thin, and have high velocity gradients – making them the most vulnerable to a spatial averaging process. This also corresponds to a well-understood weakness of LES – namely that if the scale of the disturbance is approximately equal to or less than the grid scale then the heuristic scheme to approximate subgrid-scale fluxes cannot adequately capture such localized behaviour.

Observing the same system at a later time, figure 4.10, we note the same general trend as before. The filtered Q vs R plot shows a very similar structure, demonstrating on a primarily vorticity-dominated domain, but with an overall reduction of the outliers (which translates to a reduced domain and range). Also the destruction of the strong planar shear aspect in the filtered plot is immediately obvious, and is again in line with our understanding of the characteristics of this simulation (namely regions of very sharp gradients near the boundaries) which are difficult to capture using an LES approach.

Before moving on to consider the asymmetric domain, we will first return to the question of whether the filtered Q and R variables are useful as a diagnostic measure. We have seen




Figure 4.9: Total Q vs R plot from a vertical slice of the 0.75 symmetric domain for unfiltered (top) and filtered (bottom) fields at 2.25 seconds. While the general shape is preserved, the filtered velocity field results in greatly reduced activity in the planar and axial shear regions of the Q vs R plot as well as an overall reduction of extreme values.

(above) that the resulting Q vs R plots convey the majority of relevant information. The next-most useful aspect of our previous analysis was the augmented density plots, and it is these we shall consider further before proceeding.

Figure 4.11 shows a slice of density (unfiltered) with extreme values of Q and R in cyan (1st quadrant, corresponding to vortex stretching) and white (3rd quadrant corresponding




Figure 4.10: Total Q vs R plot from a vertical slice of the 0.75 symmetric domain for unfiltered (top) and filtered (bottom) fields at 2.75 seconds. The general shape is consistent, however we see dramatic reductions in the planar shear region as well as the first quadrant (near $x - axis$) activity.

to plane shear). For the sake of consistency we have used the same threshold in both panels. As we saw in the total Q vs R plots the regions of high vorticity (first and second quadrants in $Q - R$ space) are generally well-preserved. Although we have missed the location of a single billow (right side of the simulation) the remaining billows have been successfully tagged. In contrast, almost none of the regions of high shear have been identified by the

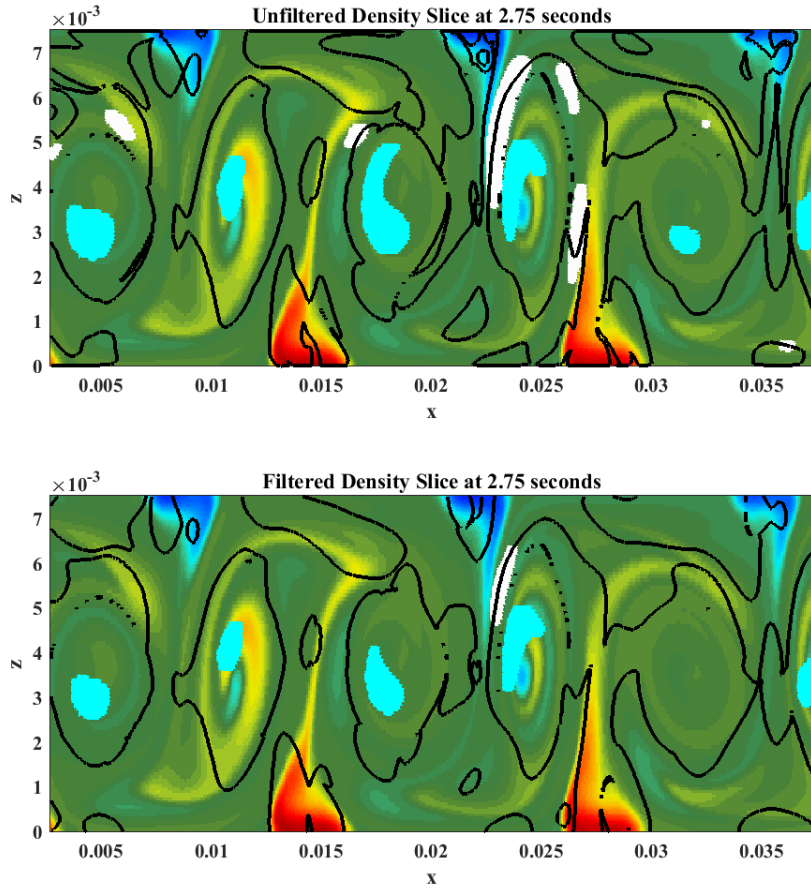


Figure 4.11: Slice of density (unfiltered) at 2.75 seconds into the two-layer simulation. Superimposed are the contours of $D_A = 0$ as well as regions of high positive Q and R (in cyan) and high negative Q and R (white) for the unfiltered (top) and filtered (bottom) variables. At this time in the simulation the flow is primarily two-dimensional. There are as yet no small-scale features and thus we are unsurprised to find the two plots almost identical.

filtered variables.

4.2.3 0.75 cm Skewed Domain

The next simulation we will consider with our averaging, or LES, process is the asymmetric 0.75 cm domain case. This simulation contained the coherent anti-billow structures which were so readily identified using Q and R analysis. This simulation remains relatively two-dimensional – i.e. has spanwise velocities less than 1% of max horizontal velocities – for a significant amount of time, as was discussed in depth in chapters 2 and 3. The first output we shall consider takes place just as the three-dimensionalization is beginning. It is here that we first noticed the distinct ‘anti-billow’ structures.

Figure 4.14 shows total Q vs R at a vertical slice (as in the previous section) at 2.25 seconds. Once again we note that the filtered velocities have led to a $Q - R$ profile which is at first glance very similar to its unfiltered counterpart. There are two significant regions where the filtered and unfiltered domains produce noticeably different results. The first is (as before) the third quadrant region, where planar shear has been greatly reduced in the filtered simulation results. The second, and arguably more interesting, is the loss of the protrusion around $Q = 0, R > 0$. This region, from our previous analysis in chapter 3, corresponded with enstrophy-dominated activity near the bottom boundary of the domain. While disappointing, this is not a surprising result. The near-boundary vortex stretching topography manifests as the hairpin vortex mechanism. A fundamental feature of which is the very small scales of the vortex tubes and the rapidly changing nature of the fluid near the no-slip boundary. This region would be very difficult (arguably impossible) to resolve with an LES model and subgrid scheme and so we are not surprised that the filtered simulation does not show it.

Next we can consider the same output time with augmented density plots (as before). Figure 4.13 shows slices of the filtered and unfiltered variables at 2 times, 1.875 and 2.25 seconds (left and right panels respectively). The earlier time is worth showing here because it was at this time that we initially identified the location and behaviour of the anti-billow structures in the simulation. As we can see from the left panel, the filtered variables are capable of identifying the anti-billow structures with relative ease. The right panel (2.25 seconds) is where we first begin to note differences in the augmented plots.

As the flow becomes more chaotic, and as such smaller-scale, the filtered variables begin to miss the finer features of the simulation: the original billows (top half of the domains) are no longer being correctly identified as regions dominated by vortex stretching; similarly regions of planar shear appear to be shrinking in the filtered simulation, an obvious side-effect of the averaging process. It is still worth noting that the structure of $D_A = 0$ remains consistent, and that the locations tagged by cyan and white in the filtered simulations are not different from those in the original variables. This is a clear indication of the reduction




Figure 4.12: Total Q vs R plot from a vertical slice of the 0.75 asymmetric domain for unfiltered (top) and filtered (bottom) fields at 2.25 seconds. As in the symmetric simulation we note the preservation of the overall structure (with enstrophy clearly dominating the flow topography) but a significant loss of outliers.

of extreme values (which likely come from small-scale characteristics of the flow) in the filtered variables.

The final output we will consider is a relatively mature turbulent flow at 2.5 seconds. Unlike the other outputs in this and the previous section, the flow at this time is quite active. The turbulent cascade is in full force, with large eddies feeding into smaller struc-

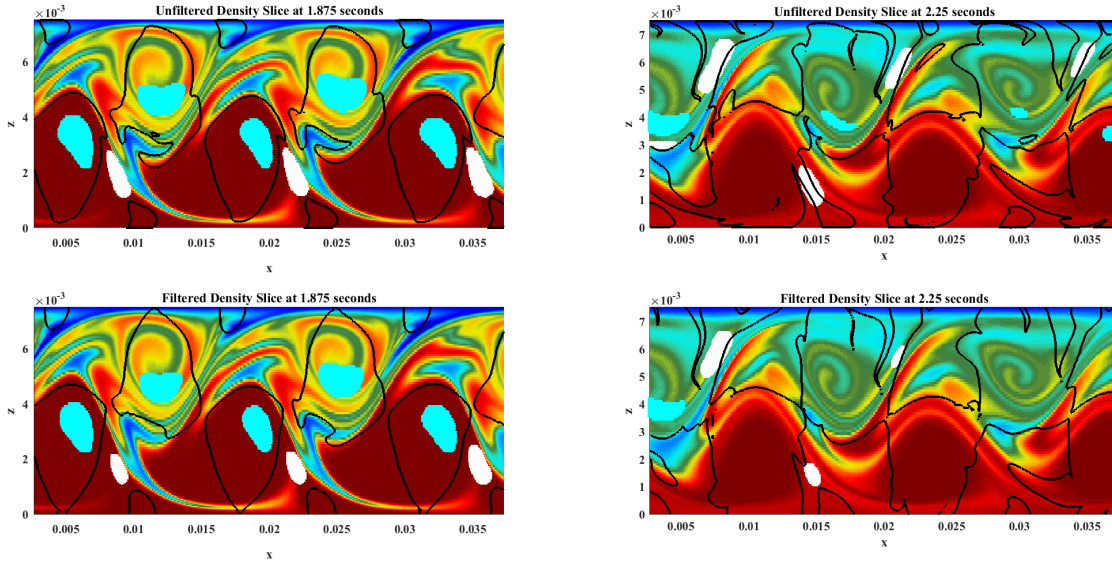


Figure 4.13: Slice of density (unfiltered) at 1.875 (left) and 2.25 (right) seconds into the asymmetric two-layer simulation. Superimposed are the contours of $D_A = 0$ (in black) as well as regions of high positive Q and R (cyan) and high negative Q and R (white) for the unfiltered (top) and filtered (bottom) variables.

tures. As a result, the features of this stage of the simulation are significantly finer, and the filtering process which generates our LES-like variables invariably destroys a large amount of fine structure. Although it is worth noting that the LES variables do manage to re-create the important larger-scale features (as is expected of LES results), the fine structure, and as a result the velocity gradient matrix, is different enough to yield a dramatically different $Q - R$ profile.

We have seen in several of the above simulation results that the data generated via LES methodology (or in this case our pseudo-LES filtering) can be used in conjunction with our $Q - R$ methodology to yield meaningful profiles. The results are most accurate – by which we mean the most true to the DNS-generated Q and R data – when the structures of the simulation are above the grid scale of the LES filter. We have also seen that this tends to happen when the simulation results are more coherent in nature, with the turbulent cascade and interactions with no-slip boundaries causing significant disruptions. This does not preclude the use of Q and R metrics in LES simulations, as we have shown that while




Figure 4.14: Total Q vs R plot from a vertical slice of the 0.75 asymmetric domain for unfiltered (top) and filtered (bottom) fields at 2.75 seconds. At this relatively late, and quite turbulent, time the differences between the unfiltered and filtered variables become more significant.

we may lose the subgrid scale information, we can effectively capture the enstrophy-strain and axial-biaxial balance of the resolved variables quite clearly.

Q and R are a useful and relatively versatile tool which can be applied to any situation where one has access to the full flow field, whether that be generated by numerical simulations such as DNS or LES, or by experiment (such as PTV). The resulting analysis will,

of course, be only as accurate as the variables that are used – there is no magic here, only a good analytical tool for helping to visualize and understand the dominant dynamics of a given fluid flow.

4.3 Visualization Tools

We have given significant thought and attention in this thesis to the subject of concise and relevant flow visualization and diagnostics. Thus far we have concluded that Q and R provide a valuable way to help visualize the relevant flow topology as well as a means to diagnose regions of interest in the flow field. Such an investigation would not be complete without at least some mention or exploration of other techniques discussed in the literature.

There exist dozens (if not hundreds) of different metrics and available schemes to help visualize numerical or physical experiments. Jiang et al. [47] review a variety of successful methods for the identification and visualization of vortices in a complex flow field; we have chosen 2 of these methods (Helicity, and λ_2) to present alongside our analysis in chapter 3. Since it is unrealistic to even list all available methods here we will instead focus in some depth on three different types of visualizations: Visit, a widely-used toolset for the amalgamation and visualization of flow fields; Helicity, a constructed measure for the analysis of flow topology; and λ_2 , a metric derived from the velocity gradient tensor used for the identification of coherent vortex structures.

The aim is to consider some of the best examples from various types of visualizations archetypes: pure visualization, topological identification, and coherent structure analysis. For the sake of standardization we will use the same Kelvin-Helmholtz billow simulations that were presented in chapter 3 of this thesis. At every step we will seek to examine the strong and weak points of each method and attempt to identify areas where it performs better (or worse) than the Q and R metrics we chose to present earlier in this work.

4.3.1 VISIT

A popular and powerful tool that already exists for visualizing computational variables is the plotting toolkit, Visit. This open source tool was first developed as part of the Advanced Simulation and Computing Initiative (ASCI) of the U.S. Department of Energy (DOE) for visualizing the results of complex numerical simulations. This section seeks to create some of the built-in diagnostic plots available in Visit and compare and contrast them to our $Q-R$ space results. Visit is capable of creating standard plots, such as plotting

isosurfaces, slices, and other canonical visualizations, however, we will be more interested in diagnostic tools such as velocity gradient tensor visualization, and it is in fact here that we shall begin our comparison.

For the sake of consistency (as well as not having to describe a brand new simulation) we will examine outputs of the asymmetric 0.75 cm simulation between 1.875 and 2.75 seconds. Due to our access to the full velocity field, and (spectrally accurate) derivatives thereof, we are more than capable of generating the velocity gradient tensor, \mathbf{A} . Once entered into Visit, this tensor can be visualized in terms of the eigenvalues at any given point. This visualization is represented as a disk: Visit finds the eigenvalues of the velocity gradient tensor and visualizes them as an ellipse scaled to be oriented along the principal axes. The connection between eigenvalues and flow topography is a complex one (ergo the creation of tertiary variables such as Q and R) but we shall see that this type of analysis, which would be considered theoretically quite naive, does serve some purpose and is capable of at minimum identifying regions of interest.




Figure 4.15: Isosurfaces of density (2 levels in red and green) with velocity gradient tensor (in blue) overlaid at 1.875 seconds. While the isosurfaces are not capable of identifying the anti-billow regions, the tensor plot shows significant activity in that region.

Figure 4.15 is a Visit-generated plot showing two isosurfaces of density (in green and red) with a tensor plot superimposed in blue. The tensor plot appears as a series of disks with the disk size and orientation determining the magnitude of the eigenvalues at any given point. Although we have access to the eigenvalues at any given point, we have chosen here to display only a total of 300 points throughout the domain as a representative sample. This plot demonstrates two key pieces of information with relative ease: the structure of

the flow field (via the density) and the identification of regions of interest (via the tensor plot).

The location of the primary billows is immediately evident, as is the region (in the bottom half of the domain) where the anti-billows are created. The isosurface is not capable of identifying the fact that there are coherent structures (namely anti-billows) in the bottom half of this simulation, however coupled with the tensor plot we can clearly see that there is significant activity in this region. The tensor plot shows that there is significant activity in several locations throughout the domain. Some features, such as the activity in the primary billow cores and in the braids, is easy enough to identify and classify; others, such as those in the anti-billow region, are only hinted at. This immediately displays a fundamental weakness of this plot, namely that we can identify regions of interest, but without further input (such as the density or enstrophy isosurfaces) the tensor plot does not provide any stand-alone information.




Figure 4.16: Isosurfaces of density (2 levels in red and green) with velocity gradient tensor (in blue) overlaid at 2.3125 seconds. Note the regions of near-constant density with large tensor activity (the anti-billow regions) as well as regions of destabilization where the orientation of the tensor disk has changed to signify a new flow topology.

Considering this same type of plot once the anti-billows have matured and spanwise destabilization begins at a later time in figure 4.16 we note that the contours of density once again serve us well in identifying the destabilization of the primary billows. Again there is a fundamental weakness in the density isosurface's inability to identify the presence or structure of the anti-billows. Interestingly, the tensor plot identifies several regions of

interest. The anti-billows (appearing in this plot once again as conspicuously absent white regions with no tensor activity) are once again identified. The abundance of spanwise instability at this output is also manifested in the tensor plot by some of the disks changing in primary orientation (specifically near the center of the domain where shear velocities are highest). This heuristic analysis is again difficult to quantify – likely indicating a new type of flow topology – but does at the very least clearly identify regions of intense activity and interest in the simulation.




Figure 4.17: Isosurfaces of density (2 levels in red and green) with velocity gradient tensor (in blue) overlaid at 2.5 seconds.

Finally, we consider a relatively mature output time (2.5 seconds shown in figure 4.17) where the coherent structures present at earlier times have destabilized in favour of small-scale turbulent motion. It is clear from the density isosurfaces that there are no more coherent structures present in this simulation. Additionally the fact that the isosurfaces essentially fill the entire domain, makes it clear that mixing has become a dominant factor in the simulation. The tensor plot at this time is heavily obscured by the very active turbulent field. What information we can glean is not directly useful, which presents another weakness of this type of visualization, namely that the more active the flow field, the more difficult the analysis. Examining the tensor field separately (not shown) was attempted, but did not yield any structure or intuitively obvious results. In some sense, the tensor plots require ‘framing’ by a scalar variable in order to be useful.

We have seen from the above simulation outputs that the Velocity tensor plots (of the velocity gradient tensor) when combined with density isosurfaces are capable of providing

an excellent overview plot. This type of visualization is insufficient by itself to classify or analyze a flow, but as a “first-pass” in order to determine what, if anything, of interest occurs in a simulation it is unrivaled. Combined with the relatively simple user interface in Visit as well as the ability to create movies for consecutive outputs this visualization type definitely deserves some place in future analysis.

4.3.2 Helicity

A fundamental aspect of the understanding the evolution of complex flows is the identification of coherent structures – and specifically coherent vortex structures – in a complex flow field. A measure to help visualize flows notably used by Levy et al. [27] for the identification of vortex core lines has been used, primarily in numerical simulations (both DNS and LES) [12, 45, 65, 75, 79], but also in experimental measurements, such as the work of Jang et al. [45] with laser Doppler velocimetry (LDV). In order to identify coherent structures it is assumed that near vortex core regions the angle between the velocity, \vec{u} , and vorticity, $\vec{\omega}$, is small.

This measure, called helicity for reason that will be made clearly shortly, is defined as:

$$H_n = \frac{\vec{u} \cdot \vec{\omega}}{|\vec{u}||\vec{\omega}|}. \quad (4.4)$$

This is the cosine of the angle between the vorticity and velocity vectors. This measure varies between -1 and 1. Thus by locating maximal values of helicity, and tracing the streamlines from the maximal points one can hope to identify coherent vortex structures.

When helicity is ± 1 (the precise value determines the direction of swirl, which is not important in our simulations) the velocity vector points in the same direction as the vorticity and the resulting coherent structure (if there is one) resembles a helix shape, ergo the variable name. It is these helical structures which we will be attempting to identify. To contrast, when the angle between the velocity and vorticity is $\frac{\pi}{2}$ the two are perpendicular and the interaction between velocity and vortex is much more complex. We do not investigate this regime further in this work, however it is worth noting as a possible future direction to explore.

The Helicity measure is purely a geometric one. By scaling out the magnitude of velocity and vorticity we lose two pieces of information that are critical for the diagnosis of flows. While the helicity can be a useful measure, we are also interested not only in the angle between the two vectors, but also in their magnitude. namely, the alignment of vorticity and velocity when both are near zero is not meaningful. Likewise, when one of

the variables is large, but the other is near-zero the helicity measure says very different things, for example large vorticity, helicity of order 1, and low velocity signifies a vortex core which is slowly being advected (or is possibly destabilizing) in the spanwise direction. In contrast, low enstrophy and high kinetic energy is more difficult to visualize, but likely corresponds to vortex sheets [25].

In attempting to apply the Helicity measure to the Kelvin-Helmholtz simulations we noted that the problems discussed above, namely that helicity tends to overvalue regions of low enstrophy and kinetic energy. Plots of isosurfaces (or slices) of H_n yielded interesting results, however did not provide a distinction between important regions (such as previously identified billow cores, and regions where low velocity and vorticity happened to align). In order to help identify the flow topology we have chosen to present the helicity in pairs, coloured by kinetic energy and enstrophy. The colormap used throughout is darkjet, with blue, green, and red representing low, mid, and high values, respectively.

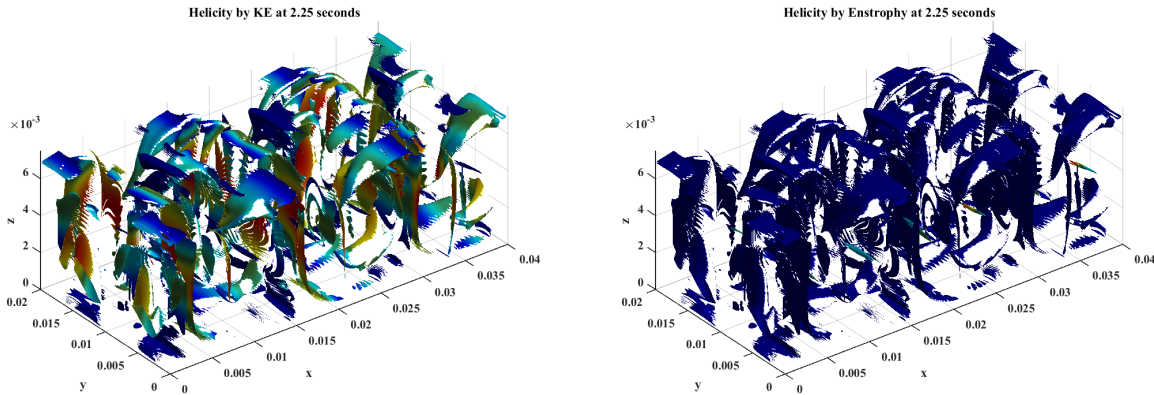


Figure 4.18: Isosurfaces of helicity in the symmetric 0.75 cm domain, $H_n = 0.8$ at 2.25 seconds coloured by kinetic energy (left) and enstrophy (right) with the darkjet colorbar - blue representing low values and red representing high. Notice that the majority of regions identified by helicity have almost no enstrophy.

The output we consider is the symmetric pycnocline 0.75 cm simulation at 2.25 seconds, shown in figure 4.18. We first note in both plots that this isosurface is quite busy, presenting a lot of information. Second, though it is difficult to confirm from the angle at which the isosurface is shown, the helicity has apparently identified the regions around the billows, as well as the centers of the cores themselves. The plot colored by kinetic energy demonstrates that the highest velocities occur near the enter of the domain (at the edges of the billows).

The billows, when they come into contact with the boundary layer, experience greatly reduced velocities – as expected from interaction with the no-slip boundaries. As we saw in chapter 2, it is here that we see the greatest generation of viscous dissipation. In contrast, the isosurface coloured by enstrophy shows only one region of true interest, the center of the billow cores. The simulation at this time is beginning to experience significant destabilization in the spanwise direction, and it is likely this that the helicity in the billow cores is capturing this.

In this simulation, the augmented helicity appears to accurately capture several key characteristics such as the interaction with the boundary layer, destabilization of the billow cores, and the three-dimensionality of the system.

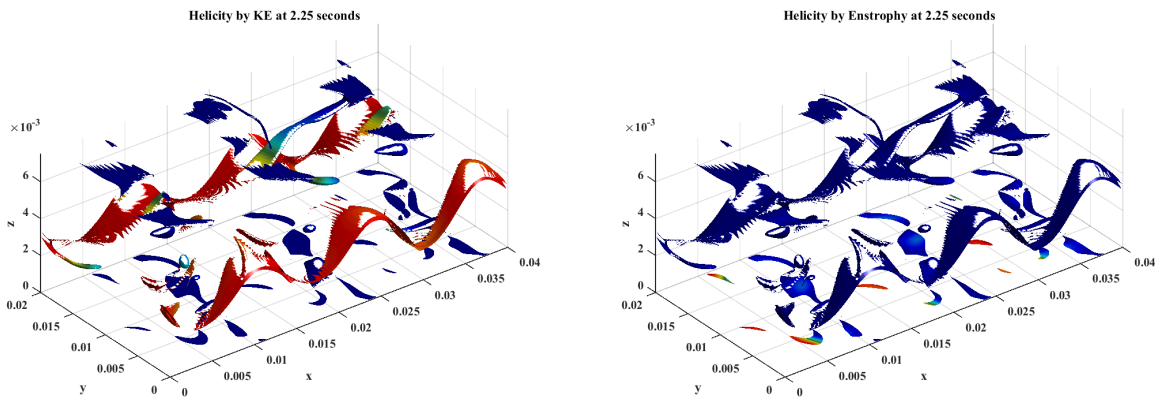


Figure 4.19: Isosurfaces of helicity in the asymmetric 0.75 cm domain, $H_n = 0.8$ at 2.25 seconds coloured by kinetic energy (left) and enstrophy (right) with the darkjet colorbar – blue representing low values and red representing high. The helicity has captured primarily the division between primary and anti billow regions.

The next output we will consider is the same time (2.25 seconds) in the asymmetric pycnocline 0.75 cm domain, shown in figure 4.19. The helicity for this simulation (with identical parameters as before) has captured very different characteristics. First, the regions of highest speeds and aligned vorticity appear to occur in the region between the primary and anti-billows. The regions with lowest KE once again occur near the boundaries, however we do not see the structured helicity we did before, making a complete picture more difficult to ascertain. The enstrophy-coloured plot provides one extremely interesting piece of information, namely that the region of highest enstrophy and helicity occurs near the bottom boundary. We hypothesize that this is associated with the fact

the the vorticity generated near the bottom boundary is pulled away from it and wrapped around the anti-billow regions.

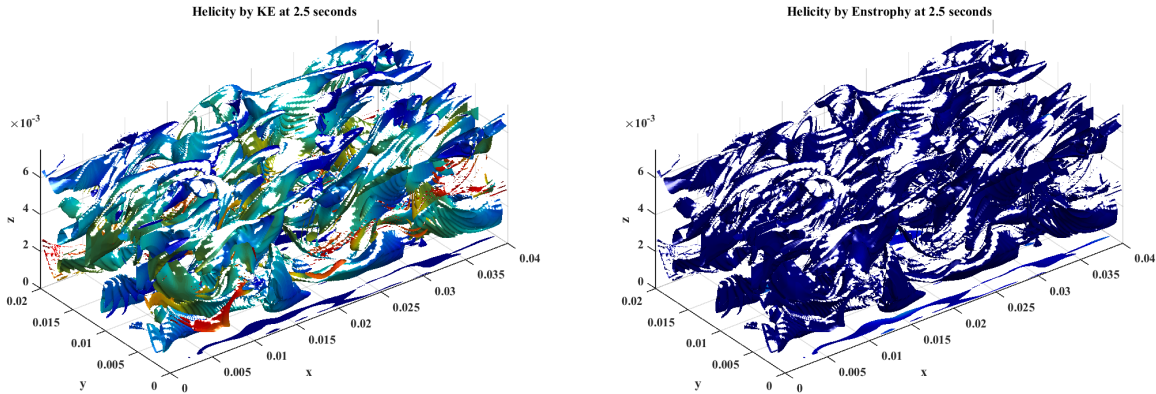


Figure 4.20: Isosurfaces of helicity in the asymmetric 0.75 cm domain, $H_n = 0.8$ at 2.5 seconds coloured by kinetic energy (left) and enstrophy (right) with the darkjet colorbar – blue representing low values and red representing high. The simulation at this time is very active and no longer contains any coherent structures, something the helicity reflects here.

The final output we will consider here is a relatively turbulent result, at 2.5 seconds in the asymmetric pycnocline 0.75 cm domain case, shown in figure 4.20. Recall, from chapters 2 and 3, that the simulation at this time is very active with little in the way of coherent structures and a great emphasis on small-scale motion. This is precisely what the helicity shows at this time. There is no real structure and the kinetic energy distribution is fairly uniform with the exception of lower velocities near the boundaries. The enstrophy-coloured plot (right panel) does not contribute any new information in this case, only confirming that there is no discernible structure in the vorticity or helicity.

We have seen above that helicity, if properly augmented, can be applied as a useful measure. In terms of identifying regions of interest it served as well so our Q and R analysis, however with Q and R we had the added benefit of being immediately able to identify and classify the flow topology.

4.3.3 The λ_2 parameter

The final visualization method we will consider here is the so-called λ_2 methodology. This method is a not-to-distant cousin of $Q - R$ analysis (as we will demonstrate shortly) and is widely used in the identification of coherent structures, primarily in numerical simulations [33, 75, 89], but also in experimental measurements which give access to the full velocity field, such as laser Doppler velocimetry (LDV) [30].

Q and R , the second and third invariants of the velocity gradient tensor, $A = \nabla \mathbf{u}$, are not the only tertiary variables available for analysis. One useful, and frequently used analysis tool is so called λ_2 analysis. First proposed in a 1994 paper in JFM, Jeong and Hussain [46] put forth the idea that there is no good metric (to date) to identify and categorize coherent structures in fluid flows. The paper focusses on the idea of a vortex, discussing several methods which have been proposed in attempting to identify vortex structures such as local pressure minima, pathlines, vorticity magnitude, and even Q and R .

The proposed metric in the Jeong and Hussain paper is the middle eigenvalue of the symmetric tensor (whose eigenvalues will all be strictly real) $\mathbf{S}^2 + \mathbf{\Omega}^2$, where \mathbf{S} and $\mathbf{\Omega}$ are the symmetric and anti-symmetric components of the velocity gradient tensor, $A = \nabla \mathbf{u}$. The result is a measure which indicates the presence of a vortex core when λ_2 is negative. In practice we find that small negative values of λ_2 (in our simulations on the order of 10^{-5}) are not sufficient to identify a vortex; however, values of λ_2 above a certain threshold did prove useful.

Previous applications of λ_2 have been performed for a variety of turbulent regimes. Most involve numerical methods, either DNS or LES [1, 2, 17, 29, 31, 33, 80, 89, 105]. A particularly successful use for λ_2 is the identification of hairpin vortices [17] and vortex identification in wall turbulence [33, 80], however, these metrics have also been applied to large-scale flows such as vorticity in the atmosphere [31] as well as channel flow [1, 105] and other regimes (see Post's article [75] for an excellent overview of the various types and applications of visualizations available including helicity and λ_2).

Figure 4.21 shows an isosurface of λ_2 at 2.375 in the symmetric pycnocline simulation domain coloured by height (z value). Coloring by height provided to yield the clearest images. Large negative values of λ_2 should correspond well to coherent vortex structures and it is here that we hope to use λ_2 to identify the billows which dominate this simulation at this time. This plot shows off very clearly the usefulness of the λ_2 parameter with the regions near the center clearly identifying the billow cores. Instabilities in the billow structure manifest themselves as deformations of λ_2 isosurfaces, corresponding well to our




Figure 4.21: Isosurfaces of λ_2 in the symmetric 0.75 cm domain at 10% of minimum value at 2.375 seconds coloured by vertical height. This plot successfully identifies the billow cores as well as the regions of intense dissipation near the boundaries – located directly above and below the billows.

understanding of this output. The other regions of strong λ_2 visible here are locations where billows are interacting with the boundary layer, directly above and below the billow core formations. These structures, while not coherent in the classical sense, are nonetheless a key location in the simulation due to the local production of vorticity and viscous dissipation at these points.

Figure 4.21 shows the same isosurface levels of λ_2 as before but at a later time (2.875 seconds). The billow structure is still clearly discernible with λ_2 successfully, and consistently, tagging the core regions as well as the high-dissipation regions near the boundary. In both this and the previous output we have seen λ_2 's ability to identify the coherent structure of the billow core. Even with growing destabilization in the spanwise and the rise of secondary instabilities, λ_2 has been sufficient in locating regions of interest. The only concern is the relatively *ad hoc* nature of the threshold we have chosen, namely 10%




Figure 4.22: Isosurfaces of λ_2 in the symmetric 0.75 cm domain at 10% of minimum value at 2.875 seconds coloured by vertical height. The coherent billow cores are still clearly visible, as is the interaction with the no slip boundaries at the top and bottom of the domain.

of minimum value at any given time. It is also important to note that the outputs and simulation results we have discussed thus far have been by far the simplest of all the simulations we have presented. It will be worth observing how this metric stands up to a more complex system such as the asymmetric pycnocline case.

Figure 4.23 shows the isosurfaces of λ_2 at 2.5% of minimum value achieved between 1.875 and 2.5 seconds for the asymmetric pycnocline case. The early time results displayed accurately capture the locations of both the billows and the anti-billows. Also clearly visible (at 2.125 and 2.25 seconds) is the spanwise destabilization in the deformation of the regular tubes from the outputs at 1.875 and 2 seconds. As the system destabilizes further, λ_2 reflects the rise of turbulence via an increase in the small-scale structures present in the λ_2 isosurface. The late times do not show any coherent structures, and in fact it is unclear if the ‘vortices’ tagged here have any structure whatsoever. One final aspect worth noting




Figure 4.23: Isosurfaces of λ_2 in the symmetric 0.75 cm domain at 2.5% of highest minimum value achieved between 1.75 and 2.75 seconds in 0.125 second increments coloured by vertical height. The early times clearly identify both primary and anti-billow locations. Mid times (2.25–2.5 s) show the spanwise destabilization and start of the turbulent regime. Late times (2.5–2.75 s) show the onset and rise of turbulence and the destruction of the earlier coherent structures.

is a strong ongoing presence of λ_2 in the boundary layer. If we recall from chapter 2, the majority of viscous dissipation is generated by the interaction between billows and the no slip boundaries. The structure of this interaction remains relatively coherent until spanwise instabilities begin to dominate the simulation (after 2.5 seconds) and this is precisely what the isosurfaces have identified near the top and bottom of the domain.

The λ_2 parameter is capable of identifying coherent structures; however, in order to

do so we were forced to choose a relatively *ad hoc* level at which to plot isosurfaces. In the symmetric simulation, the optimal level at which to view λ_2 was 10% of the minimum value at any given time. This provided a consistent picture of the structure and location of coherent vortices, namely in the billow centers and at the boundary layer. In contrast a similar measure for the asymmetric simulation (not shown) produced results which did not capture any of the coherent structures which we knew to be present in this simulation. In order to successfully identify the regions of interest we were forced to manually adjust the threshold of λ_2 until the structures became visible. This value turned out to be -2500 (a somewhat arbitrary value with no context here). Additionally, the values of minimum λ_2 varied dramatically in this simulation, and as such the threshold value we chose had to remain consistent. The resulting value was equal to 2.5% of the smallest minimum achieved over the course of the simulation - a result which was not at all obvious *a priori*.

Another important weakness of λ_2 analysis is that, unlike Q and R , it provides no additional information about the flow topology in or around the coherent structure. We know that the structures identified are likely vorticity-dominated, but lack context for the type of mechanisms present (whether it be vortex stretching or sheeting) or the surrounding flow field. Nevertheless, the types of plots shown could be part of a visualization strategy.

A simple recipe is presented here for the analysis of a flow. One fundamental aspect which we cannot remove is the visualization of the flow field as an overview for the state of the simulation, thus we must begin with a density isosurface. A slight modification may be made from the isosurfaces presented here, in that we may wish to display density coloured by a relevant variable, kinetic energy if we are interested in the most active regions, or dissipation in the presence of boundaries. The next step is the identification of regions of high viscous dissipation, something we can achieve either with isosurfaces of dissipation as was used in chapter 2 or with vertical slices of mean and standard deviation of dissipation in the spanwise direction. The final step is to attempt to identify the presence of coherent structures in the simulation, for this we recommend a combination augmented density slices from chapter 3 (which provided a clear identification of the vertical structure and dominant mechanics at a slice) and isosurfaces of λ_2 (which easily identify coherent structures, namely billows, as well as displaying spanwise instabilities in those structure). Both of these require a thresholding factor; in our investigations of KH billows we found that the thresholding for Q vs R plots proved more consistent than λ_2 isosurfaces, however both effectively identified any coherent structures present in the flow.

4.4 Discussion

This final chapter sought to round out our earlier discussions. The three primary sections contained herein focussed on understanding the studies presented in the first two chapters in a wider context.

The Kelvin Helmholtz section explored the importance of no-slip boundaries in the evolution of our simulations, concluding that the boundary layer is a fundamental aspect of the billow confinement, and without it the primary mechanism in the flow is subharmonic pairing (what we called billow-billow interactions). This simulation definitively demonstrated that it was the combination of confinement, domain size, and the no-slip boundaries which was responsible for the unique coherent structures present in our simulations in chapter 2, with no-slip boundaries being the *sine qua non* aspect.

The Q and R section investigated the relevance of the analysis methodology which we had developed theoretically in the first half of chapter 3, and applied in the second half of this chapter, when applied to a different type of numerical scheme, namely large eddy simulations (LES). This section determined that while LES flows lost the finer details of the flow, the $Q - R$ analysis remained viable throughout the simulation with little need for modification. The most dramatic result we observed was that Q and R are very efficient in capturing the small-scale dynamics. Averaged (LES-simulated) turbulent outputs demonstrate that the Q vs R profile changes dramatically in range as well as details of the specific shape. This alteration makes Q and R analysis precisely as trustworthy as the true resolution of the variables used – a result which is not surprising as there is no magical means of generating information which is not resolved.

The final, and certainly most different, part of this chapter considered Q and R as a visualization and coherent structure identification tool and contrasted it with other well known methods such as the Visit suite of plotting tools, helicity, and λ_2 analysis. We determined that while all three have their strengths and weaknesses, our choice of $Q - R$ analysis for the identification of coherent structures and the classification of flow topology was entirely justified in the context of our simulations and analysis goals. The Visit tools are perhaps the most visually and aesthetically pleasing, however, we noted that these visualizations required us to possess a relatively thorough understanding of the simulation state and the dominant mechanics in order to make full use of the produced plots. In contrast, topological measures such as helicity produced clear indications of the locations of coherent structures, and even one type of dominant feature (namely a vortex core). The drawback of helicity was its propensity for misidentifying regions of importance, as the normalization factor weighs regions of high kinetic energy and enstrophy identically to

relatively calm ones. In order to combat this inherent flaw we attempted to superimpose the relevant measures (namely KE and enstrophy) but found the resulting plots to still be lacking as vortex sheeting caused the regions of interest to be dwarfed by noise. Finally, examining a popular (and well-established) metric such as λ_2 allowed us to compare the effectiveness of our methodology – namely Q and R – to a relatively standard method. We concluded that while λ_2 analysis did succeed in identifying coherent vortex structures in the simulation it suffered from two major drawbacks: first the visualization we chose (isosurfaces of λ_2) required an *ad hoc* parameterization and threshold adjustment which varied with simulation type; and second, when used in a highly turbulent output λ_2 did not contribute any new information about the dominant mechanics or the state of the flow.

Hopefully, the preceding discussions have served their purpose. There are several future directions of study which we can follow up on, ranging from a more coherent understanding of the nature of anti-billows, to refining our visualization tools (namely Q and R) and how they can be combined with more standard metrics – one such possibility is the combination of Q vs R and λ_2 plots. Some of these ideas will be presented in the next section along with a discussion as to how these future directions fit into the scope of the current work and the greater body of fluid mechanics knowledge in general.

4.5 Future Directions

Having outlined here some obvious extensions to the work presented in chapters 2 and 3 we may now consider some future directions for research and how they would fit in with the scope of this study. The structure of this section once again follows the three major themes of the thesis, KH billows, QR turbulent metrics, and visualization tools.

4.5.1 KH

The most natural (and obvious) extension to the DNS results presented in chapter 2 is an expansion of the parameter space presented. As was mentioned in the introduction to the two-dimensional simulations in section 2.3, we initially explored several different ratios of momentum to thermal diffusivity (Prandtl or Schmidt number). While these results did not make it into this work, it may be worth considering an expansion of this regime in order to see how it affects the mixing and dissipation in the presence of boundaries. Another natural parameter which can be modified is the initial stratification. We attempted to explore a range of interesting stratifications (such as the symmetric, asymmetric or skewed,

and linear initial stratifications) however this was by no means a complete list. Varying the location of the pycnocline, the total density jump (5% is at the large end of the spectrum for a total density jump), and the overall structure (double or triple pycnocline or a linear stratification over a portion of the vertical extent of the domain, etc) are all valid investigations we may wish to consider in the future.

Another worthwhile consideration is the regime over which the unique coherent structure we observed in the asymmetric pycnocline, 0.75 cm domain case is generated. We determined earlier in this chapter that it was the combined effect of confinement and no-slip boundaries which resulted in the formation and continued existence of this anti-billow structure. In future work we may seek to confirm the domain parameters which lead to this phenomenon such as the required stratification, pycnocline placement, domain size, etc.

As we noted in the introduction (section 1.4), unstable shear flows can fall into two major regimes, Kelvin-Helmholtz instabilities (by far the more studied and classically well-understood instability) and Holmboe instabilities, which occur for sharp density changes if the gradient Richardson number ($J = Ri(z = 0) = \frac{N^2}{(\frac{dU}{dz})^2}(z = 0)$ (some papers use J as the bulk Richardson number, which is not the same as the gradient Richardson number here)) exceeds some critical value – i.e. when the pycnocline is sharp enough. This is a regime which we are capable of stimulating, and in fact some of the suggested parameter space explorations above would almost certainly fall into this regime. A key point is that KH billows behave dramatically differently in a confined domain and in the presence of boundaries, thus it is natural to ask how a related instability, albeit one with a different structure, behaves under similar constraints.

4.5.2 Turbulent Metrics

The turbulent metrics which we explored in chapter 3 proved to be an efficient means of characterizing the destabilization of our sample flows, as well as the identification of regions of interest – such as the coherent structures in the KH simulations. Further work involving these turbulent metrics is needed in order to help complete our understanding of various Q vs R profiles of instabilities and the transition to turbulence. The application of Q and R to a wider range of flows, specifically turbulent transitional flows such as a transitional boundary layer, can help in two ways: it confirms the usefulness of Q and R as a tool for characterizing flow topology and it develops a suite of ‘canonical’ profiles for various instabilities. The goal of building such a toolset is that when examining a new system we can apply Q vs R analysis first and based on the profile determine the dominant mechanics

and the dominant instability (or at least the instability which is most similar dynamically to the one we are observing).

A more analytic problem posed by our exploration of Q and R as a toolset arose from our investigation of the Burgers' vortex in section 3.3.5. We determined that for a Burgers vortex the Q and R distribution versus radial distance from the center of the vortex is very specific. If vortex stretching is the primary mechanism in isotropic turbulence then we expect these iconic Q and R profiles to exist throughout a fully turbulent flow. In isotropic turbulence we expect such vortices (if they exist) to be ubiquitous throughout the domain in various orientations. This means that the identification of the profiles may be locally difficult; however, we have two mathematical advantages: first the tensor invariants should be immune to the orientation of the vortex, and second we have access to the local principal axes due to having to compute eigenvalues and eigenvectors during the computation of Q and R . Such an investigation would lend credence to the generally accepted theory that it is vortex stretching which is the primary mechanism in isotropic turbulence as well as providing a useful means of identifying a local Burgers vortex in more complex flows.

4.5.3 Visualizations

We have compared in this chapter several key methods of visualizations: tensor eigenvalue plots, helicity, and λ_2 . These metrics all have certain advantages and canonical uses. One particularly widespread usecase for the λ_2 parameter which we did not consider here is the identification of hairpin vortices in turbulent boundary layers (see Adrain's 2007 paper [2] for a recent examination of hairpin structure at a boundary). Due to its origins in isotropic turbulence (away from boundaries) we have focussed primarily on the use of Q and R as it relates to the development of turbulent regimes that bear at least some resemblance to isotropic turbulence. One promising area for future investigation is the use of Q and R as a means of identifying hairpin vortices – based on the fundamental flow topology which results in hairpin vortex production near a boundary.

4.6 Concluding Comments

It is my sincere hope that the contents of this thesis have achieved our initial goals of discovering something worthy of further study by placing a well-known phenomenon in a new setting. We set out with the intention of exploring a relatively new facet of a well-known phenomenon – Kelvin Helmholtz billows in a confined setting. Along the way we

discovered a key weakness in our visualization strategy for complex flows, a gap which we attempted to fill by borrowing metrics from the turbulence literature to help characterize the flow field. This tool – Q and R – proved useful in not only identifying the dominant mechanics at any given point, but in locating and understanding coherent structures in the flow. Through this analysis we discovered a brand new coherent structure in confined shear flow, namely the “anti-billows” from sections [2.4.2](#) and [3.6.2](#).

A consistent theme throughout the last four years of research has been the identification of interesting phenomena or novel tools and the subsequent combination of these into something new. This effort manifested here as either new regimes, such as the no-slip boundaries for KH billows, or a new context for a known tool set, such as examining transitional flows with using Q and R . The combination of the known (or well understood) and the unknown (or unfamiliar) has, in my experience, yielded the most rewarding scientific inquiry. It is my belief that by placing the familiar in an unfamiliar context we illuminate aspects of both that we would not have otherwise been available for study, and that is precisely what has kept me interested in this field for the past four years.

References

- [1] ADRIAN, R., CHRISTENSEN, K., AND LIU, Z.-C. Analysis and interpretation of instantaneous turbulent velocity fields. *Experiments in Fluids* 29, 3 (2000), 275–290.
- [2] ADRIAN, R. J. Hairpin vortex organization in wall turbulence. *Physics of Fluids (1994-present)* 19, 4 (2007), 041301.
- [3] BAYLY, B. J., ORSZAG, S. A., AND HERBERT, T. Instability mechanisms in shear-flow transition. *Annual Review of Fluid Mechanics* 20, 1 (1988), 359–391.
- [4] BLUMEN, W., BANTA, R., BURNS, S. P., FRITTS, D. C., NEWSOM, R., POULOS, G. S., AND SUN, J. Turbulence statistics of a Kelvin–Helmholtz billow event observed in the night-time boundary layer during the Cooperative Atmosphere–Surface Exchange Study field program. *Dynamics of Atmospheres and Oceans* 34, 2 (2001), 189–204.
- [5] BOFFETTA, G., DE LILLO, F., AND MUSACCHIO, S. Anomalous diffusion in confined turbulent convection. *Physical Review E* 85, 6 (2012), 066322.
- [6] BRACHET, M. Direct simulation of three-dimensional turbulence in the Taylor–Green vortex. *Fluid Dynamics Research* 8, 1 (1991), 1–8.
- [7] BRACHET, M. E., MEIRON, D. I., ORSZAG, S. A., NICKEL, B., MORF, R. H., AND FRISCH, U. Small-scale structure of the Taylor–Green vortex. *Journal of Fluid Mechanics* 130 (1983), 411–452.
- [8] BROWNING, K. Structure of the atmosphere in the vicinity of large-amplitude Kelvin–Helmholtz billows. *Quarterly Journal of the Royal Meteorological Society* 97, 413 (1971), 283–299.

- [9] BUTLER, K. M., AND FARRELL, B. F. Three-dimensional optimal perturbations in viscous shear flow. *Physics of Fluids A: Fluid Dynamics (1989-1993)* 4, 8 (1992), 1637–1650.
- [10] CABOT, W. Comparison of two- and three-dimensional simulations of miscible Rayleigh–Taylor instability. *Physics of Fluids (1994-present)* 18, 4 (2006), 045101.
- [11] CABOT, W. H., AND COOK, A. W. Reynolds number effects on Rayleigh–Taylor instability with possible implications for type Ia supernovae. *Nature Physics* 2, 8 (2006), 562–568.
- [12] CALA, C., FERNANDES, E., HEITOR, M., AND SHTORK, S. Coherent structures in unsteady swirling jet flow. *Experiments in Fluids* 40, 2 (2006), 267–276.
- [13] CARNEVALE, G., ORLANDI, P., ZHOU, Y., AND KLOOSTERZIEL, R. Rotational suppression of Rayleigh–Taylor instability. *Journal of Fluid Mechanics* 457 (2002), 181–190.
- [14] CAULFIELD, C., AND PELTIER, W. Three dimensionalization of the stratified mixing layer. *Physics of Fluids (1994-present)* 6, 12 (1994), 3803–3805.
- [15] CAULFIELD, C., AND PELTIER, W. The anatomy of the mixing transition in homogeneous and stratified shear layers. *Journal of Fluid Mechanics* 413 (2000), 1–47.
- [16] CHONG, M., PERRY, A. E., AND CANTWELL, B. A general classification of three-dimensional flow fields. *Physics of Fluids A: Fluid Dynamics (1989-1993)* 2, 5 (1990), 765–777.
- [17] CHRISTENSEN, K., AND ADRIAN, R. J. Statistical evidence of hairpin vortex packets in wall turbulence. *Journal of Fluid Mechanics* 431 (2001), 433–443.
- [18] CHUNG, T. *Computational fluid dynamics*. Cambridge University Press, 2010.
- [19] COOK, A. W., CABOT, W., AND MILLER, P. L. The mixing transition in Rayleigh–Taylor instability. *Journal of Fluid Mechanics* 511 (2004), 333–362.
- [20] COOK, A. W., AND DIMOTAKIS, P. E. Transition stages of Rayleigh–Taylor instability between miscible fluids. *Journal of Fluid Mechanics* 443 (2001), 69–99.
- [21] COOK, A. W., AND ZHOU, Y. Energy transfer in Rayleigh–Taylor instability. *Physical Review E* 66, 2 (2002), 026312.

- [22] DA SILVA, C. B., AND PEREIRA, J. C. Invariants of the velocity-gradient, rate-of-strain, and rate-of-rotation tensors across the turbulent/nonturbulent interface in jets. *Physics of Fluids* 20, 5 (2008), 55101–55101.
- [23] DALZIEL, S., LINDEN, P., AND YOUNGS, D. Self-similarity and internal structure of turbulence induced by Rayleigh–Taylor instability. *Journal of Fluid Mechanics* 399 (1999), 1–48.
- [24] DALZIEL, S. B., PATTERSON, M. D., CAULFIELD, C., AND COOMARASWAMY, I. A. Mixing efficiency in high-aspect-ratio Rayleigh–Taylor experiments. *Physics of Fluids (1994-present)* 20, 6 (2008), 065106.
- [25] DAVIDSON, P., ET AL. *Turbulence: an introduction for scientists and engineers*. Oxford University Press, USA, 2015.
- [26] DE SILVA, I., FERNANDO, H., EATON, F., AND HEBERT, D. Evolution of Kelvin-Helmholtz billows in nature and laboratory. *Earth and Planetary Science Letters* 143, 1 (1996), 217–231.
- [27] DEGANI, D., SEGNER, A., AND LEVY, Y. Graphical visualization of vortical flows by means of helicity. *AIAA Journal* 28, 8 (1990), 1347–1352.
- [28] DRIKAKIS, D., FUREBY, C., GRINSTEIN, F. F., AND YOUNGS, D. Simulation of transition and turbulence decay in the Taylor–Green vortex. *Journal of Turbulence* N20, 8 (2007).
- [29] DUBIEF, Y., AND DELCAYRE, F. On coherent-vortex identification in turbulence. *Journal of Turbulence* 1, 1 (2000), 011–011.
- [30] DUWIG, C., FUCHS, L., LACARELLE, A., BEUTKE, M., AND PASCHEREIT, C. O. Study of the vortex breakdown in a conical swirler using LDV, LES and POD. *ASME Paper No. GT2007-27006* (2007).
- [31] FRITTS, D. C., AND ALEXANDER, M. J. Gravity wave dynamics and effects in the middle atmosphere. *Reviews of Geophysics* 41, 1 (2003).
- [32] FRITTS, D. C., PALMER, T. L., ANDREASSEN, Ø., AND LIE, I. Evolution and breakdown of Kelvin-Helmholtz billows in stratified compressible flows. Part I: Comparison of two- and three-dimensional flows. *Journal of the Atmospheric Sciences* 53, 22 (1996), 3173–3191.

- [33] FUREBY, C., AND GRINSTEIN, F. F. Large eddy simulation of high-Reynolds-number free and wall-bounded flows. *Journal of Computational Physics* 181, 1 (2002), 68–97.
- [34] GARRETT, C., AND MUNK, W. Internal waves in the ocean. *Annual Review of Fluid Mechanics* 11, 1 (1979), 339–369.
- [35] GHARIB, M., AND DERANGO, P. A liquid film (soap film) tunnel to study two-dimensional laminar and turbulent shear flows. *Physica D: Nonlinear Phenomena* 37, 1 (1989), 406–416.
- [36] GREEN, A., AND TAYLOR, G. Mechanism of the production of small eddies from larger ones. In *Proc. Royal Soc. A* (1937), vol. 158, pp. 499–521.
- [37] GUALA, M., LIBERZON, A., TSINOBER, A., AND KINZELBACH, W. An experimental investigation on Lagrangian correlations of small-scale turbulence at low Reynolds number. *Journal of Fluid Mechanics* 574 (2007), 405–427.
- [38] HAIGH, S., AND LAWRENCE, G. Symmetric and nonsymmetric Holmboe instabilities in an inviscid flow. *Physics of Fluids (1994-present)* 11, 6 (1999), 1459–1468.
- [39] HALLER, G. Lagrangian coherent structures from approximate velocity data. *Physics of Fluids (1994-present)* 14, 6 (2002), 1851–1861.
- [40] HALLER, G., AND YUAN, G. Lagrangian coherent structures and mixing in two-dimensional turbulence. *Physica D: Nonlinear Phenomena* 147, 3 (2000), 352–370.
- [41] HAZEL, P. Numerical studies of the stability of inviscid stratified shear flows. *Journal of Fluid Mechanics* 51, 01 (1972), 39–61.
- [42] HERBERT, T. Secondary instability of plane channel flow to subharmonic three-dimensional disturbances. *Physics of Fluids (1958-1988)* 26, 4 (1983), 871–874.
- [43] HUA, B., AND KLEIN, P. An exact criterion for the stirring properties of nearly two-dimensional turbulence. *Physica D: Nonlinear Phenomena* 113, 1 (1998), 98–110.
- [44] ISERN-FONTANET, J., FONT, J., GARCÍA-LADONA, E., EMELIANOV, M., MILLOT, C., AND TAUPIER-LETAGE, I. Spatial structure of anticyclonic eddies in the Algerian basin (Mediterranean sea) analyzed using the Okubo–Weiss parameter. *Deep Sea Research Part II: Topical Studies in Oceanography* 51, 25 (2004), 3009–3028.

- [45] JANG, C.-M., FURUKAWA, M., AND INOUE, M. Analysis of vortical flow field in a propeller fan by LDV measurements and LES—part I: three-dimensional vortical flow structures. *Journal of Fluids Engineering* 123, 4 (2001), 748–754.
- [46] JEONG, J., AND HUSSAIN, F. On the identification of a vortex. *Journal of Fluid Mechanics* 285 (1995), 69–94.
- [47] JIANG, M., MACHIRAJU, R., AND THOMPSON, D. Detection and visualization of vortices. *The Visualization Handbook* (2005), 295.
- [48] KIMURA, S., AND SMYTH, W. Direct numerical simulation of salt sheets and turbulence in a double-diffusive shear layer. *Geophysical Research Letters* 34, 21 (2007).
- [49] KLAASSEN, G., AND PELTIER, W. Evolution of finite amplitude Kelvin-Helmholtz billows in two spatial dimensions. *Journal of the Atmospheric Sciences* 42, 12 (1985), 1321–1339.
- [50] KLAASSEN, G., AND PELTIER, W. The onset of turbulence in finite-amplitude Kelvin-Helmholtz billows. *Journal of Fluid Mechanics* 155 (1985), 1–35.
- [51] KLAASSEN, G., AND PELTIER, W. The influence of stratification on secondary instability in free shear layers. *Journal of Fluid Mechanics* 227 (1991), 71–106.
- [52] KOBERINSKI, A., BAGLAENKO, A., AND STASTNA, M. Schmidt number effects on Rayleigh-Taylor instability in a thin channel. *Physics of Fluids (1994-present)* 27, 8 (2015), 084102.
- [53] KOLMOGOROV, A. N. The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers. *Dokl. Akad. Nauk SSSR* 30, 4 (1941), 299–303.
- [54] KULL, H.-J. Theory of the Rayleigh-Taylor instability. *Physics Reports* 206, 5 (1991), 197–325.
- [55] KUNDU, P., AND COHEN, I. *Fluid mechanics. 2004*, vol. 307. Elsevier Academic Press, San Diego, 2008.
- [56] LAWRENCE, G., PIETERS, R., ZAREMBA, L., TEDFORD, T., GU, L., GRECO, S., AND HAMBLIN, P. Summer exchange between hamilton harbour and lake ontario. *Deep Sea Research Part II: Topical Studies in Oceanography* 51, 4 (2004), 475–487.

- [57] LAWRIE, A. G., AND DALZIEL, S. B. Turbulent diffusion in tall tubes. I. Models for Rayleigh-Taylor instability. *Physics of Fluids (1994-present)* 23, 8 (2011), 085109.
- [58] LAWRIE, A. G., AND DALZIEL, S. B. Turbulent diffusion in tall tubes. II. Confinement by stratification. *Physics of Fluids (1994-present)* 23, 8 (2011), 085110.
- [59] LÜTHI, B., TSINOBER, A., AND KINZELBACH, W. Lagrangian measurement of vorticity dynamics in turbulent flow. *Journal of Fluid Mechanics* 528 (2005), 87–118.
- [60] MASHAYEK, A., AND PELTIER, W. The ‘zoo’ of secondary instabilities precursory to stratified shear flow transition. Part 1 Shear aligned convection, pairing, and braid instabilities. *Journal of Fluid Mechanics* 708 (2012), 5–44.
- [61] MASHAYEK, A., AND PELTIER, W. The ‘zoo’ of secondary instabilities precursory to stratified shear flow transition. Part 2 The influence of stratification. *Journal of Fluid Mechanics* 708 (2012), 45–70.
- [62] MENEVEAU, C. Lagrangian dynamics and models of the velocity gradient tensor in turbulent flows. *Annual Review of Fluid Mechanics* 43 (2011), 219–245.
- [63] METCALFE, R. W., ORSZAG, S. A., BRACHET, M. E., MENON, S., AND RILEY, J. J. Secondary instability of a temporally growing mixing layer. *Journal of Fluid Mechanics* 184 (1987), 207–243.
- [64] MILES, J. W. On the generation of surface waves by shear flows Part 3. Kelvin-Helmholtz instability. *Journal of Fluid Mechanics* 6, 04 (1959), 583–598.
- [65] MOFFATT, H., AND TSINOBER, A. Helicity in laminar and turbulent flow. *Annual Review of Fluid Mechanics* 24, 1 (1992), 281–312.
- [66] MUESCHKE, N. J., AND SCHILLING, O. Investigation of Rayleigh–Taylor turbulence and mixing using direct numerical simulation with experimentally measured initial conditions. I. comparison to experimental data. *Physics of Fluids (1994-present)* 21, 1 (2009), 014106.
- [67] NAKAMURA, N. Two-dimensional mixing, edge formation, and permeability diagnosed in an area coordinate. *Journal of the Atmospheric Sciences* 53, 11 (1996), 1524–1537.
- [68] NOBLE, B., AND DANIEL, J. W. *Applied linear algebra*, vol. 3. Prentice-Hall New Jersey, 1988.

- [69] OKUBO, A. Horizontal dispersion of floatable particles in the vicinity of velocity singularities such as convergences. *Deep Sea Research and Oceanographic Abstracts* 17, 3 (1970), 445–454.
- [70] OOI, A., MARTIN, J., SORIA, J., AND CHONG, M. A study of the evolution and characteristics of the invariants of the velocity-gradient tensor in isotropic turbulence. *Journal of Fluid Mechanics* 381 (1999), 141–174.
- [71] ORSZAG, S. A., AND PATERA, A. T. Secondary instability of wall-bounded shear flows. *Journal of Fluid Mechanics* 128 (1983), 347–385.
- [72] PELTIER, W., AND CAULFIELD, C. Mixing efficiency in stratified shear flows. *Annual Review of Fluid Mechanics* 35, 1 (2003), 135–167.
- [73] PIERREHUMBERT, R. Universal short-wave instability of two-dimensional eddies in an inviscid fluid. *Physical Review Letters* 57, 17 (1986), 2157.
- [74] PIROZZOLI, S., AND GRASSO, F. Direct numerical simulations of isotropic compressible turbulence: influence of compressibility on dynamics and structures. *Physics of Fluids (1994-present)* 16, 12 (2004), 4386–4407.
- [75] POST, F. H., VROLIJK, B., HAUSER, H., LARAMEE, R. S., AND DOLEISCH, H. The state of the art in flow visualisation: Feature extraction and tracking. In *Computer Graphics Forum* (2003), vol. 22:4, Wiley Online Library, pp. 775–792.
- [76] RAHMANI, M., LAWRENCE, G., AND SEYMOUR, B. The effect of Reynolds number on mixing in Kelvin–Helmholtz billows. *Journal of Fluid Mechanics* 759 (2014), 612–641.
- [77] RAHMANI, M., SEYMOUR, B., AND LAWRENCE, G. The evolution of large and small-scale structures in Kelvin–Helmholtz instabilities. *Environmental Fluid Mechanics* 14, 6 (2014), 1275–1301.
- [78] RIVERA, M., VOROBIEFF, P., AND ECKE, R. E. Turbulence in flowing soap films: velocity, vorticity, and thickness fields. *Physical Review Letters* 81, 7 (1998), 1417.
- [79] ROTH, M. *Automatic Extraction Of Vortex Core Lines And Other Line Type Features For Scientific Visualization*. Hartung-Gorre, 2000.
- [80] SCHOPPA, W., AND HUSSAIN, F. Coherent structure generation in near-wall turbulence. *Journal of Fluid Mechanics* 453 (2002), 57–108.

- [81] SMYTH, W., CARPENTER, J., AND LAWRENCE, G. Mixing in symmetric Holmboe waves. *Journal of Physical Oceanography* 37, 6 (2007), 1566–1583.
- [82] SMYTH, W., NASH, J., AND MOUM, J. Differential diffusion in breaking Kelvin-Helmholtz billows. *Journal of Physical Oceanography* 35, 6 (2005), 1004–1022.
- [83] SMYTH, W., AND PELTIER, W. The transition between Kelvin-Helmholtz and Holmboe instability: an investigation of the overreflection hypothesis. *Journal of the Atmospheric Sciences* 46, 24 (KH, Holmboe 1989), 3698–3720.
- [84] SMYTH, W., AND PELTIER, W. Instability and transition in finite-amplitude Kelvin-Helmholtz and Holmboe waves. *Journal of Fluid Mechanics* 228 (1991), 387–415.
- [85] SPIEGEL, E., AND VERONIS, G. On the Boussinesq approximation for a compressible fluid. *The Astrophysical Journal* 131 (1960), 442.
- [86] STAQUET, C., AND SOMMERIA, J. Internal gravity waves: from instabilities to turbulence. *Annual Review of Fluid Mechanics* 34, 1 (2002), 559–593.
- [87] STRUTT, J., AND RAYLEIGH, L. Investigation of the character of the equilibrium of an incompressible heavy fluid of variable density. *Proc. London Math. Soc* 14, 1 (1883), 8.
- [88] SUBICH, C. J., LAMB, K. G., AND STASTNA, M. Simulation of the Navier-Stokes equations in three dimensions with a spectral collocation method. *International Journal for Numerical Methods in Fluids* 73, 2 (2013), 103–129.
- [89] SUJUDI, D., AND HAILES, R. Identification of swirling flow in 3D vector fields. Tech. rep., Technical report, Department of Aeronautics and Astronautics, MIT, 1995. AIAA Paper 95-1715, 1995.
- [90] TAYLOR, G. The instability of liquid surfaces when accelerated in a direction perpendicular to their planes. *Proc. R. Soc. London Ser A*, 201 (1950), 192–196.
- [91] TEDFORD, E., CARPENTER, J., PAWLOWICZ, R., PIETERS, R., AND LAWRENCE, G. A. Observation and analysis of shear instability in the fraser river estuary. *Journal of Geophysical Research: Oceans (1978–2012)* 114, C11 (2009).
- [92] THORPE, S. Laboratory observations of secondary structures in Kelvin-Helmholtz billows and consequences for ocean mixing. *Geophysical & Astrophysical Fluid Dynamics* 34, 1-4 (1985), 175–199.

- [93] TRAN, C. V. The number of degrees of freedom of three-dimensional Navier–Stokes turbulence. *arXiv preprint arXiv:0912.2000* (2009).
- [94] TREFETHEN, L. N. *Spectral methods in MATLAB*, vol. 10. Siam, 2000.
- [95] VAN HAREN, H., AND GOSTIAUX, L. A deep-ocean Kelvin-Helmholtz billow train. *Geophysical Research Letters* 37, 3 (2010).
- [96] VAN HAREN, H., GOSTIAUX, L., MOROZOV, E., AND TARAKANOV, R. Extremely long Kelvin-Helmholtz billow trains in the Romanche Fracture Zone. *Geophysical Research Letters* 41, 23 (2014), 8445–8451.
- [97] WANG, Q., AND SQUIRES, K. D. Large eddy simulation of particle-laden turbulent channel flow. *Physics of Fluids (1994-present)* 8, 5 (1996), 1207–1223.
- [98] WEISS, J. The dynamics of enstrophy transfer in two-dimensional hydrodynamics. *Physica D: Nonlinear Phenomena* 48, 2 (1991), 273–294.
- [99] WORTH, N., NICKELS, T., AND SWAMINATHAN, N. A tomographic PIV resolution study based on homogeneous isotropic turbulence DNS data. *Experiments in Fluids* 49, 3 (2010), 637–656.
- [100] WYGNANSKI, I., AND FIEDLER, H. The two-dimensional mixing region. *Journal of Fluid Mechanics* 41, 02 (1970), 327–361.
- [101] WYNGAARD, J. C. *Turbulence in the Atmosphere*. Cambridge University Press, 2010.
- [102] YEUNG, P., AND POPE, S. Lagrangian statistics from direct numerical simulations of isotropic turbulence. *Journal of Fluid Mechanics* 207 (1989), 531–586.
- [103] YONEMITSU, N., SWATERS, G. E., RAJARATNAM, N., AND LAWRENCE, G. A. Shear instabilities in arrested salt-wedge flows. *Dynamics of atmospheres and oceans* 24, 1 (1996), 173–182.
- [104] YOUNG, Y., TUFO, H., DUBEY, A., AND ROSNER, R. On the miscible Rayleigh-Taylor instability: two and three dimensions. *Journal of Fluid Mechanics* 447, 1 (2001), 377–408.
- [105] ZHOU, J., ADRIAN, R. J., BALACHANDAR, S., AND KENDALL, T. Mechanisms for generating coherent packets of hairpin vortices in channel flow. *Journal of Fluid Mechanics* 387 (1999), 353–396.

- [106] ZHU, D. Z., AND LAWRENCE, G. A. Holmboe's instability in exchange flows.
Journal of Fluid Mechanics 429 (2001), 391–409.