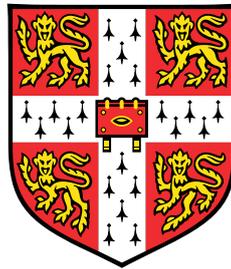


The dynamics and mixing properties of vortex rings obliquely impacting a density interface



Benjamin Michael Jackson

Department of Applied Mathematics and Theoretical Physics
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Pembroke College

December 2021

To my parents, Margaret and Peter, my brother, James, and my sister, Katy.

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and acknowledgements.

Benjamin Michael Jackson

December 2021

Acknowledgements

First and foremost, I wish to thank my supervisor, Professor Stuart Dalziel, for his endless support and guidance. I will always be grateful to Stuart for giving me the opportunity to study at DAMTP as his PhD student. I couldn't have asked for a kinder person to learn from and work with over these past few years. Thank you, Stuart.

I would like to thank Dr. Jamie Partridge, for his assistance and advice regarding experimental procedure and methodology over the course of my PhD studies. Jamie was always there to lend a hand when technical issues arose, and always happy to spare his time. I would also like to thank Dr. Jason Olsthoorn, for helping me get my feet off the ground when I began to embark on my PhD studies.

The experiments presented in this thesis would have not been possible without the help of an extraordinary group of laboratory technicians. Many thanks to David Page-Croft, Colin Hitch, Paul Mitton, Andrew Denson and John Milton, all of whom provided invaluable support for my experimental investigations. I would also like to give thanks to Dr. Mark Hallworth for his assistance throughout my time in the laboratory.

Working at DAMTP has been an absolute pleasure, in large part due to the wonderful group of friends that have been made along the way. I would particularly like to thank Maksim Dodonov and Christopher Howland, who started their PhD journeys alongside me, and my office mates over the years: Amalia Thomas, Jonny Tsang, Katherine Grayson and Georgia Ioannou. I am also deeply grateful to Joana Ferreira, Paul Cosgrove and the wonderful group of friends made at Pembroke College, who were always there to escape from work with.

The last year (and a bit) of completing this thesis, the Earth became gripped by the covid-19 pandemic. I had the unbelievably good fortune of being locked down with the most amazing group of people: Lorena Gazotti, David Kaloper-Mercinjak, Matt Mahmoudi, Saksham Sharma, Alex Waghorn, and last but not least, my wonderful girlfriend Yunhui Zhuang, whom I've been unbelievably lucky to get to know over this past year. I will always fondly remember our Sunday brunches and daily circuits at 6 Grange Rd during this strange time in global history.

Finally, I am deeply thankful to all my family, whose support for me has been unwavering throughout my time at Cambridge, and to Yunhui and her younger brother Huanzhao, for keeping me sane during the final stages of writing this thesis. I am grateful to EPSRC who provided the financial support necessary for me to pursue the research presented in this thesis.

The dynamics and mixing properties of vortex rings obliquely impacting a density interface

Benjamin Michael Jackson

This thesis presents an experimental investigation into the dynamics and mixing properties of vortex rings obliquely impacting the density interface in a two-layer density stratification. For turbulent two-layer zero-mean-shear flows, the classical grid-mixing experiments of J.S. Turner (*J. Fluid. Mech*, 33:639–656, 1968) demonstrated the intermittent interaction of strong, coherent eddy-like structures with the density interface to be a dominant mixing mechanism. Previous studies have made an analogy between this mixing mechanism and the mixing induced by a vortex ring vertically impacting the density interface in an otherwise stationary flow, to study the mixing mechanism in isolation. Our research extends this analogy by investigating vortex rings obliquely impacting the density interface, at propagation angles $\theta_0 \leq 25^\circ$ relative to the vertical.

The dynamics of the ring-interface interaction were explored using a double-pulsed laser system to take two-dimensional planar simultaneous Particle Image Velocimetry (PIV) and Laser-Induced Fluorescence (LIF) measurements. To be able to obtain a high signal to noise ratio in all the observable velocity scales of the flow, we developed a ‘multi-frame’ PIV algorithm that makes use of interrogating pairs of PIV images at several different time intervals apart. This algorithm was used to process our PIV data, and can easily be adapted to process PIV data for other flows with regions of localised turbulence.

Ring–interface interactions are classified as ‘penetrative’ or ‘non-penetrative’, corresponding to whether downward entrainment across the interface occurs or not. For both types of interaction, our PIV/LIF measurements reveal that oblique ring impacts lead to an azimuthally asymmetric production of baroclinic vorticity, triggering instability mechanisms that are not present in the $\theta_0 = 0^\circ$ case. With the aim of investigating the influence of θ_0 on the mixing properties of the ring–interface interaction, experiments were conducted in which a periodic sequence of 600 vortex rings were generated to mix an initially two-layered stratification. Insights made after the completion of these experiments revealed that, in the non-penetrative regime, the system converges to a state where a significant fraction of the total mixing is convective, as opposed to being directly associated with the ring interacting with the interface. It is argued that the high mixing efficiency observed is attributable in large part to the convective mixing, rather than directly to the ring–interface interaction as has been previously reported. These results are discussed and contextualised with previous grid-mixing experiments.

Table of contents

List of figures	xiii
List of tables	xxv
1 Introduction	1
1.1 Context	1
1.2 History of the ring–eddy analogy	2
1.3 Motivation and thesis outline	7
2 Underlying Theory and Literature	9
2.1 Vortex dynamics: fundamental theory	9
2.1.1 Definitions	10
2.1.2 Theorems and governing equations	11
2.2 Vortex rings	15
2.2.1 Axisymmetric models	16
2.2.2 Formation process	20
2.2.3 The Widnall instability and the stability of vortex tubes	25
2.3 Vortex-ring interactions	32
2.3.1 Normal impact on to a no-slip wall	32
2.3.2 Oblique impact onto a no-slip wall	35
2.3.3 Normal impact on to a deformable free surface	36
2.3.4 Oblique impact on to a deformable free surface	37
2.3.5 Normal impact with a density interface	39
2.3.6 Oblique impact with a density interface	41
2.3.7 Vortex rings propagating through a weak, linear stratification	43
2.4 Aspects of stratified mixing	45
2.4.1 The energetics of mixing	45
2.4.2 Turner’s mixing box and the ring–eddy analogy	47

3	Experimental Apparatus, Setup and Methodology	51
3.1	Creating a two-layer stratification	51
3.1.1	Filling the tank	52
3.1.2	Lower-layer filling apparatus	53
3.1.3	Molecular diffusion and interfacial thickness	56
3.1.4	Thermal effects	57
3.2	Vortex rings	59
3.3	Planar data acquisition	63
3.3.1	Refractive index matching	63
3.3.2	Plane illumination	65
3.3.3	Camera setup	70
3.3.4	Particle image velocimetry	72
3.3.5	Laser-induced fluorescence	76
3.4	Vertical profiles	77
3.4.1	Apparatus and setup	77
3.4.2	Calibration	81
3.5	Summary	85
4	A Variation of Multi-Frame PIV for Localised Turbulence	87
4.1	Introduction	88
4.2	Image pair interrogation	90
4.2.1	Measurement error of a single particle trajectory	91
4.2.2	Interrogation windows	93
4.2.3	Measurement error of a group of particle projections within an inter- rogation window	95
4.3	Multi-frame PIV algorithm	97
4.4	Implementation of algorithm	100
4.5	Discussion	108
5	Influence of propagation angle on the ring–interface interaction	111
5.1	Observed measures	114
5.2	Non-penetrating rings at small propagation angles	116
5.2.1	Vertical propagation	117
5.2.2	Oblique propagation: evolution of the primary ring	119
5.2.3	Oblique propagation: azimuthal flow	125
5.3	Penetrating rings at small propagation angles	129
5.3.1	Vertical propagation	130

5.3.2	Oblique propagation: a first glance	133
5.3.3	Oblique propagation: evolution of the primary ring	138
5.3.4	Oblique propagation: secondary structures	142
5.4	Summary	147
6	Vortex-ring-induced mixing of an initially two-layered density stratification	149
6.1	Experimental methodology	151
6.1.1	Setup and procedure	151
6.1.2	Overview of input parameters	154
6.1.3	Data correction and validation	155
6.2	Energetics of sustained periodic forcing	157
6.2.1	Quantifying kinetic energy	157
6.2.2	Quantifying ring-induced increases in potential energy	159
6.2.3	Mixing efficiency measurements	160
6.3	Evolution of stratification for non-penetrative mixing	163
6.3.1	The initial-adjustment phase	163
6.3.2	Ring evolution and mixing in the quasi-steady state	165
6.3.3	Mixing equilibrium in the quasi-steady state	169
6.3.4	Influence of control parameters on the QSS	170
6.4	Summary	173
7	Conclusions	175
7.1	Thesis summary	175
7.2	Turner's mixing box and the ring–eddy analogy	177
7.3	Avenues for future research	179
7.3.1	Multi-frame PIV algorithm	179
7.3.2	Ring–interface interaction dynamics	180
7.3.3	Conductivity probe design	181
7.3.4	Eddy-induced mixing	182
7.4	Final words	183
	References	185
	Appendix A Kelvin Wave properties for different azimuthal wavenumbers	193
	Appendix B Tracer Particle Calculations	195
B.1	Estimation of the Kolmogorov Lengthscale	195
B.2	Order of magnitude calculation of the Stokes Number	196

B.3	Calculation of the settling velocity	196
Appendix C	Analysis of the conductivity probe	197
C.1	Probe circuit	197
C.2	Sampling the density stratification	200
C.2.1	Probe-tip region	200
C.2.2	Inner-probe region	203
C.3	Summary	205
Appendix D	Ring-induced heat fluxes	207

List of figures

1.1	Plot of the regions in the $(\text{Ri}_0, \text{Re}_0)$ parameter space that have been explored in experimental studies in the context of the ring–eddy analogy for vertically propagating vortex rings impacting a density interface. The legend items correspond to studies by Linden (1973), Olsthoorn and Dalziel (2015) and Olsthoorn and Dalziel (2017). The divide between regions A , B and C mark possible transition regions in the nature of the turbulent entrainment induced by the ring–interface interaction.	6
1.2	Diagram of a vortex ring of diameter a propagating obliquely towards a sharply stratified density interface at propagation speed U and propagation angle θ_0 . Here, the illustrated two-dimensional cross-section coincides with the central plane of the flow. Richardson and Reynolds numbers Ri_0 , Re_0 are defined based on bulk properties of the vortex ring and the fluid density of each layer. The parameters presented here will be used throughout this thesis.	8
2.1	Illustration of a portion of a vortex tube, with one vortex line marked on the vortex tube surface. Velocity and vorticity vectors are drawn to aid the visualisation of the flow associated with the vortex tube.	12
2.2	(a) Diagram of the cross-section of an axisymmetric vortex ring, with radius a , circulation Γ_0 and propagation speed U . The vortex core has core size b and vorticity distribution $\omega(\mathbf{x})$. The core region D has been drawn with circular cross-section for simplicity. (b) Sketch of the relationship between the Cartesian and cylindrical coordinate systems.	16
2.3	(a) Plots of the core boundary for Norbury’s rings for a range of values of α , duplicated from Norbury (1973). (b) Diagram of Maxworthy’s bubble entrainment model, where the left-hand side (i) prohibits the shedding of vorticity into the bubble exterior, and the right-hand side (ii) allows for the formation of a wake behind the vortex ring where deposited vorticity can annihilate. This diagram is duplicated from Maxworthy (1972).	19

2.4	Snapshots of the ring-formation process, using dye to visualise a central plane of the vortex ring. The tube diameter is 5 cm and the piston speed is 4.6 cm/s, giving Reynolds number $Re_p = 2300$. Figure reproduced from Didden (1979).	22
2.5	Taking $c = 1$, plots of (a) (2.34) and (b) (2.35) for the ring core volume (dashed) and total volume carried by the ring (solid), non-dimensionalised by the plug volume. The graphs are plotted for Norbury rings, using tabulated values from Norbury (1973). Plot duplicated from Linden and Turner (2001).	24
2.6	(a) Illustration of the Widnall instability with 8 peaks and wavenumber k . The unperturbed ring, whose centreline is indicated by the blue dashed circular line, has core diameter b and ring diameter a . (b) Photograph of a vortex ring exhibiting the Widnall instability, with 9 peaks. (c) Photograph of the Widnall instability viewed from the side. Photographs reproduced from Krutzsch (1939).	25
2.7	For $m = 1$, using red and blue for retrograde and cgrade waves respectively: (a) Taking $\kappa = 1$, plots of the left-hand side (black) and right-hand side of the dispersion relation (2.40) as a function of ξ . (b) For $n = 0, 1, 2$, plots of the frequencies $\omega_n(\kappa, 1)$ given by (2.41) as a function of κ . Arrows are used to indicate the direction of increasing n	28
2.8	Illustration of the Crow Instability. Figure modified from Leweke et al. (2016). Each vortex tube diverges along the stretching plane of the strain field induced by the other.	30
2.9	Sketches of the early stages of a ring interacting with a no-slip wall, for normal and oblique impacts. (a) As the ring approaches the wall, its radius increases, as indicated by the dashed arrows. (b) Proximity of the ring with the wall leads to the generation of secondary vorticity at the wall. (c) For oblique impacts, secondary vorticity is generated to a greater extent by the side of the ring nearer the wall.	34
2.10	Photographs of a dye visualisation of a vortex ring impacting a wall at an oblique angle, taken from a single experiment with $(Re_0, \theta_0) = (600, 38.5^\circ)$. The ring is moving towards the reader, with Region A of the ring impacting the wall first. The dimensionless time $\tau = t/t_0$ is normalised by an advective timescale $t_0 = D_0/U_0 = 1.716$ s, where D_0 is the diameter of the tube outlet and U_0 is the ring propagation speed. Note that $\tau = 0$ corresponds to Figure 4a of Lim (1989). Photographs selected from Figure 4 of Lim (1989). . . .	35

- 2.11 Schematic illustrating the transformation of surface parallel and surface normal vorticity components during the vortex connection process. Duplicated from Zhang et al. (1999). Note that the terms ‘vortex stretching’ and ‘vortex turning’ used here relate respectively to the vorticity parallel and vorticity normal components of the vortex stretching term in the vorticity evolution equation given by Equation (2.17). 38
- 2.12 Horizontal slices plotting the vertical velocity of the ring–interface interaction, demonstrating the modal structure of the stratification instability. Cases (a), (b) and (c) correspond to $(\text{Ri}_D, \text{Re}_D) = (2.4, 1600)$, $(1.7, 2400)$ and $(0.98, 1600)$ respectively, and have respective mode numbers 8, 10 and 8. Figure duplicated from Olsthoorn and Dalziel (2017). 41
- 2.13 Selected photographs of the vortex ring–interface interaction for $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2070, 4.6, 0^\circ)$, taken from a dye visualisation experiment where the lower layer is dyed green to visualise the distortion of the density interface. Here, the ring diameter $a = 4.9$ cm and the ring propagation speed $U = 42.2$ cm s^{-1} . Small perturbations to the flow in the upper layer were found to heavily compromise the symmetry of the instability. The time interval over which the four photos were taken is approximately 2.5 seconds. 42
- 2.14 (a) Illustration of the experimental setup used by Maxworthy (1977) to observe vortex rings propagating through a linear stratification. A plot of the ambient density profile $\rho_\infty(z)$ and a sketch of the ring diameter $D(z)$ as it propagates through the stratification are both provided. (b) Sketch of the vortex ring after it has propagated well into the linearly-stratified region. Three distinct regions emerge: the core of the ring, a mixed intermediate region, and the ambient. Vorticity is produced baroclinically at the interfaces between adjacent regions. Figures reproduced from Maxworthy (1977). . . 44
- 2.15 Plots of the measured mixing rate as a function of Richardson number. (a) Experimental results when a single layer stirred. Filled black circles and empty black circles correspond to experiments where heat and salt are the stratifying agents respectively. (b) Experimental results when salt is the stratifying agent. White empty dots and black filled squares correspond to experiments where one layer is stirred and both layers are stirred respectively. Graphs reproduced from Turner (1968). 48
- 3.1 Sketch of the lower-layer height as a function of time during the lower-layer filling process. 53

3.2	Sketch of the apparatus used to fill the larger tank. (a) View of the apparatus positioned in the tank. (b) View of the underside of the apparatus from directly beneath it. (c) Side view of the apparatus at the bottom of the tank.	54
3.3	(a) Example of a typical vertical profile of the temperature stratification, taken a few minutes after the lower-layer filling process had been completed. (b) The temperature of the laboratory measured using an RC-5 data logger (Elitech) for a week in September 2020. The blue lines correspond to 8am for each day.	59
3.4	(a) Simple illustration of the two ring tubes used, which both have an inner diameter $D_p = 39\text{mm}$. Shading is used to indicate the cylindrical geometry of the tubes. (b) A photograph of the angle setter. The ring tube in use is secured to the angle setter, which we use to set the ring propagation angle θ_0 .	60
3.5	(a) The non-dimensionalised displacement profile of the bike pump traverse used to generate vortex rings. (b) Calibration curves for the straight (S) and angled (A) ring tubes, where PIV measurements were used to relate the ring propagation speed U to the traverse time T_s . Across all calibration measurements, precision errors for U were less than $\pm 0.6\text{ mm/s}$	62
3.6	Diagram illustrating the deflection of a light ray at the boundary between two isotropic media, as described by Snell's law given by Equation (3.4).	64
3.7	Plots of the density of sodium chloride (NaCl) and sodium nitrate (NaNO ₃) solutions at 20°C against the refractive index, n under green light. Data obtained by Lefauve (2018).	65
3.8	Schematic of the planar PIV/LIF setup. (a) Front view; the central plane of the light sheet is aligned with the axis of the ring tube. (b) Plan view; two cameras (labelled 'A' and 'B') used to take measurements are positioned to have overlapping fields of view of the light sheet. For ease of interpretation, we have omitted illustrations of the ring tube, angle setter and tank filling apparatus in the plan view.	66
3.9	Side-view illustration of the three-dimensional light sheet projected onto a camera sensor and quantised onto the sensor array (defined in Section 3.3.3). The light sheet thickness Δy_L has an hour-glass like shape. The z -position of the minimum thickness $\Delta y_{L\text{min}}$ is adjustable by refocusing the laser optics.	68
3.10	Schematic of the timing sequence between the camera shutters opening and closing for frame capture and the laser pulses, where the laser-pulse offset δt is the time interval between consecutive laser pulses and δT is the inverse sampling frequency of the camera.	70

3.11	Close-up photograph of tracer particles that have settled to their neutral-buoyancy level in a linear stratification. Particles are seen to settle in two bands, with peaks at $\rho_{p-} = 1.023 \text{ g/cm}^3$ and $\rho_{p+} = 1.025 \text{ g/cm}^3$ and each with an approximate range of 10^{-3} g/cm^3	74
3.12	LIF snapshot of the ring–interface interaction for $(\text{Ri}_0, \theta_0) = (1.18, 15^\circ)$ during the transition to turbulence. The fluorescing dye clearly marks large disconnected regions of upper and lower layer fluid in the lower layer indicate significant out-of-plane motion. The instability structures in convectively unstable regions of the flow are also clearly visible, illustrating the presence of small length scales and high in-plane velocity gradients.	78
3.13	Schematic of the conductivity probe used in this thesis.	80
3.14	Close-up photograph of the probe tip and thermistor submerged in a fluid sample with the magnetic puck stirring the fluid.	82
3.15	(a) Calibration curve relating the thermistor voltage output to the temperature of the sampled fluid. (b) Example set of the data points taken for a probe/thermistor calibration.	83
4.1	Diagram illustrating how the parameters \bar{T} , ΔT associated with an image pair $\{\mathcal{I}_i, \mathcal{I}_j\}$ are defined, where image \mathcal{I}_k is a snapshot of the flow at time T_k when the laser pulses. In the illustrated case, $(i, j) = (2n + 2, 2n + 5)$. . .	91
4.2	Particle pattern-matching in two interrogation windows, separated by a time ΔT . In this example, the measured pixel displacement is $(\delta X_A, \delta Y_A) = (6, 0)$. The normalised absolute difference signal is shown on the right, where (dX_A, dY_A) is the pixel deviation from the measured pixel displacement. . .	94
4.3	Diagram illustrating the relationship between the sensor array A and the interrogation array A_I . Interrogation points \mathbf{X}_I , which are spaced a pixel distance δ_I apart, are marked with a red cross. A single interrogation window of size $n_I \times n_I$ pixels is shaded blue. The displacement measurements $\delta \mathbf{X}_I$ obtained from interrogating a pair of PIV images are taken to be evaluated on the interrogation array.	95
4.4	Diagram illustrating how, for a PIV image sequence $\Phi(\delta T, \delta t)$, the parameters $\bar{T}_m^{(k)}$ and $\Delta \bar{T}^{(k)}$ are defined for $k < K$. In the illustrated case, $\Delta T^{(k)} = 4\delta T$ and $\Delta \bar{T}^{(k)} = 2\delta T$	98
4.5	A pseudocode for the multi-frame algorithm discussed in this chapter. . . .	99

- 4.6 (a) Example of a displacement array obtained using the largest image-pair time step we employ using our PIV algorithm. The rectangle drawn encloses the region to be interrogated at the next smallest time step. The scalar field is the curl of the pixel-displacement field, denoted ω_y . (b) LIF image of the ring-interface interaction at the instance in time corresponding to the measured displacement array in Figure 4.6 (a). 102
- 4.7 Following the displacement measurements shown in Figure 4.6 (a), from left to right, displacement measurements produced utilising the image-pair time steps $\Delta T^{(2)}$, $\Delta T^{(4)}$ and $\Delta T^{(6)}$. For each set of displacement measurements, the scalar background is the curl of the pixel displacement field. 104
- 4.8 Plots of $W_{n,ij} = k_{n,ij} + \alpha_{n,ij}$, indicating both the image-pair time steps utilised at each interrogation point to construct the PIV velocity field given in (b), and the weighting constants $\alpha = \alpha_{n,ij}$, defined in Equation (4.11). Plots are (a) before, and (b) after, three filters are applied to the initial guess of the W_n array in the following order: a 3×3 median filter, a 3×3 maximum filter, and a 5×5 low-pass convolution filter. 107
- 4.9 (a) Example of a velocity field constructed using the algorithm presented in this Chapter, at the time coinciding with the LIF image given in Figure 4.6 (b). (b) For comparison, a velocity field evaluated at the same instance in time, utilising only the smallest available time step in all regions of the flow. Labels correspond to features identified in the text. Plots (c) and (d) correspond to (a) and (b) respectively, using a reduced scale in ω_y to illustrate the reduction in RME achieved using our algorithm in slower-moving regions of the flow. For all plots, a subsample of the constructed velocity vectors are printed, with the largest velocity vectors corresponding to approximately 11.5 cm/s. 109
- 5.1 Diagram of a vortex ring of diameter a propagating obliquely towards a sharply stratified density interface (at height $z = H_2$) at propagation speed U and initial propagation angle θ_0 . Here, the vortex-core cross-sections C_1 and C_2 coincide with the central plane of the flow. Richardson and Reynolds numbers Ri_0 , Re_0 are defined based on bulk properties of the vortex ring and the fluid density of each layer. Note that the tube height $z = H_B$ is out of view in this diagram. 112
- 5.2 Diagrams illustrating the definitions of the propagation angle $\theta_p(t)$ and core angle $\theta_c(t)$. Note that before the ring interacts with the density interface, the two angles are equal, with $\theta_p(t) = \theta_c(t) = \theta_0$ 116

- 5.3 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (1810, 3.09, 0^\circ)$. The time separation between consecutive snapshots is 0.5 seconds, with the first frame taken at time $t = 1$ s. Note that the advective timescale $a/U = 1.33$ s. 118
- 5.4 LIF images at (a) $t = 2.0$ s and (b) $t = 3.5$ s, overlaid with a subsample of 50% the PIV algorithm-constructed velocity vectors, and ω_y contours, with the blue and red contours respectively representing anti-clockwise and clockwise in the central plane of the flow. The in-plane vorticity contours are drawn for $|\omega_y| = 5, 15, 25$ and 35 s^{-1} . Note that the largest velocity vectors in (a) correspond to a speed of 59.2 mm/s. 119
- 5.5 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (1810, 3.09, 5^\circ)$. The time separation between consecutive snapshots is 0.5 seconds, with the first frame taken at time $t = 1$ s. Note that the advective timescale $a/U = 1.33$ s. 120
- 5.6 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (1810, 3.09, 10^\circ)$. The time separation between consecutive snapshots is 0.5 seconds, with the first frame taken at time $t = 1$ s. Note that the advective timescale $a/U = 1.33$ s. 121
- 5.7 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (1810, 3.09, 20^\circ)$. The time separation between consecutive snapshots is 0.5 seconds, with the first frame taken at time $t = 1$ s. Note that the advective timescale $a/U = 1.33$ s. 122
- 5.8 Plots of the observed measures of the vortex ring introduced in Section 5.1, in addition to the non-dimensionalised penetration depth $Z(t)/a$ into the lower layer. For each figure, data are plotted between $t = 0$ and $t = 4$ seconds. Vertical dashed lines are drawn in plots (a) to (d), marking the instances in time that correspond to the first seven LIF snapshots presented in each of Figures 5.3, 5.5, 5.6 and 5.7. In plots (e) and (f), filled circles mark the position of the ring centroids and ring centre every 0.5 seconds. The legend corresponding to all plots is given in (b). 124
- 5.9 LIF snapshots at $t = 4$ s for the cases (a) $\theta_0 = 10^\circ$ and (b) $\theta_0 = 20^\circ$, overlaid with velocity vectors and contours of $Q_{2D} = 1, 2, 3, 4$ and 5 s^{-2} . Vortex S and cores C_1 and C_2 are labelled in both figures where they are made reference to in the text. 126

- 5.10 LIF snapshots at (a) $t = 5.5$ s for $\theta_0 = 5^\circ$, and (b) $t = 6.5$ for $\theta_0 = 10^\circ$. Snapshots are overlaid with velocity vectors and contours of $\nabla \cdot \mathbf{u}_{2D} = 5, 10, 15, 20$ and 25 s^{-1} . Circles are drawn to clearly identify the vortex core S_2 128
- 5.11 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2140, 2.24, 5^\circ)$. Each snapshot is labelled with its associated instant in time. Red circles are drawn around the recoil ring cores exhibited in the first snapshot for clarity. 129
- 5.12 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 0^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.90$ s. 131
- 5.13 LIF images at (a) $t = 2.00$ s, (b) $t = 2.75$ s and (c) $t = 3.50$ s, overlaid with (sparsely drawn) velocity vectors and ω_y contours, with the blue and red contours respectively representing anti-clockwise and clockwise in the central plane of the flow. The in-plane vorticity contours are drawn for $|\omega_y| = 2, 15, 25$ and 35 s^{-1} . Black crosses are drawn to identify the in-plane stagnation point separating the ring-entrained fluid and ring fluid ejected by the surrounding baroclinic vorticity. 132
- 5.14 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (3270, 0.98, 0^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.74$ s. 133
- 5.15 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 5^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.90$ s. 135
- 5.16 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 10^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.90$ s. 136
- 5.17 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 20^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.90$ s. 137

5.18	LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (3270, 0.98, 5^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.74$ s.	138
5.19	LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (3270, 0.98, 10^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.74$ s.	139
5.20	LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (3270, 0.98, 20^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.74$ s.	139
5.21	Plots of the observed measures of the vortex ring for the cases corresponding to $(\text{Re}_0, \text{Ri}_0) = (2680, 1.48)$, in addition to the non-dimensionalised penetration depth $Z(t)/a$ into the lower layer. For plots (a), (c) and (d), data are plotted up to the stage where either C_1 or C_2 can no longer be reliably identified in the central plane of the flow. In plots (e) and (f), filled circles mark the position of the ring centroids and ring centre every 0.5 seconds. The legend corresponding to all plots is given in (b).	140
5.22	Plots of the in-plane vorticity ω_y overlaid with velocity vectors, LIF snapshots overlaid with vorticity contours, and LIF snapshots overlaid with $\nabla \cdot \mathbf{u}_{2D}$ contours for the times $t = 2.6$ s, 3.0 s, 3.4 s and 3.8 s. The magenta and cyan $\nabla \cdot \mathbf{u}_{2D}$ contours respectively represent flow converging on to and diverging away from the central plane.	144
5.23	Close-up LIF snapshots of the small secondary vortices that develop on the leading edge of the BVC as it begins to destabilise for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 20^\circ)$. The time interval over which the snapshots are taken correspond to the panels presented in Figure 5.22 (c-d).	145
5.24	Selected LIF snapshots from the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2960, 1.18, 15^\circ)$, with consecutive snapshots separated in time by 0.125 seconds. In this case, there is a clear development of secondary vortices along the leading edge of the destabilising BVC, as well as disconnected regions of upper-layer fluid indicating significant out-of-plane flow.	147

6.1	Representative examples of the density profiles taken after every 100 ring generations for the cases in the (a) non-penetrative, and (b) penetrative, regimes respectively. Arrows are included to indicate the evolution of time. The selected experiments are labelled A0 and Q0 in Table 6.2, corresponding to $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2960, 6.30, 0^\circ)$ and $(2960, 1.18, 0^\circ)$ respectively.	150
6.2	Diagram of the experimental setup of a typical periodic-mixing experiment.	152
6.3	Representative plot of the vortex-ring kinetic energy measurements δKE_i made over a sequence of 100 consecutive PIV velocity fields, plotted against the non-dimensionalised distance $(\bar{z}_i - z_0)/a_0$ between the vortex ring and the ring-tube outlet. A least squares approach is used to fit a line to the data, indicated by the blue dashed line. For each ring case, δKE is taken to be the value of the fitted line evaluated at $(\bar{z}_i - z_0)/a_0 = 3$	158
6.4	Plots of $\Sigma\delta\text{PE}_R^{(n)}$ against the number of ring generations, n for the cases (a) $\text{Ri}_0 = 6.30$ and (b) $\text{Ri}_0 = 0.98$. Colours correspond to the legend given in (b).	161
6.5	Plots of averaged mixing efficiencies from each experiment listed in Table 6.2 against Ri_0 . (a) Plots of η_{200}^{600} against Ri_0 , where $\eta_c = 0.42$ is the mixing efficiency determined by Olsthoorn and Dalziel (2015). (b) Plots of η_0^n against Ri_0 , where $n = 80$ and $n = 40$ for the penetrative and non-penetrative regimes respectively. Colours for (a) and (b) correspond to the legend given in (b). (c) Plots of η_0^n against n for experiments A0 and Q0, corresponding to the non-penetrative and penetrative regimes respectively. Standard error bars are included in all plots to illustrate the uncertainty associated with our measurements.	162
6.6	Non-dimensionalised density profiles selected to visualise the evolution of the periodic-mixing system for case A0, corresponding to $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2960, 6.30, 0^\circ)$. Here, the non-dimensionalised tube height $h_B = H_B/L = 0.75$, and unmixed upper and lower-layer densities are $\rho_1 = 1.0017\text{g/cm}^3$ and $\rho_2 = 1.0481\text{g/cm}^3$ respectively.	164
6.7	Non-dimensionalised density profiles after every 60 ring generations for the cases (a) B0 and (b) D0. Red profiles mark the transition between the IAP and the quasi-steady state.	165

6.8	(a) The average change in transitional Richardson number δRi_T per ring generation, against θ_0 . (b) Plot of $Ri_T^{(n)}$ against n for the cases corresponding to $\theta_0 = 0^\circ$, demonstrating the typical evolution of $Ri_T^{(n)}$ in the non-penetrative regime. (c) The product of Ri_0 and the non-dimensionalised interfacial height change $A\delta H_2/V_R$ in the quasi-steady state per ring generation, against θ_0 . Standard error bars for (a) and (c) are included.	168
6.9	Non-dimensionalised density profiles after 120, 300 and 480 ring generations for the cases (a) (A0, A20), (b) (B0, B20), (c) (C0, C20), and (d) (D0, D20) respectively. Black and red profiles correspond to the cases where $\theta_0 = 0^\circ$ and $\theta_0 = 20^\circ$ respectively.	171
6.10	Plots of (a) $Ri_{g,U}^{(n)}$, and (b) $\text{std}(Ri_{g,U}^{(n)})$, against ring generation n for $\theta_0 = 0^\circ$. Cases A-D correspond to the labels given in Table 6.2 and the black profiles in Figures 6.9 (a)-(d).	172
6.11	A simplified sketch of the horizontally-averaged density profile $\bar{\rho}(z)$ during the lifespan of a vortex ring mixing event after the IAP. Each vortex ring entrains fluid from the lower layer, after which convective mixing occurs as the ring-transported and ring-mixed fluid rise to find their new neutral buoyancy levels.	174
7.1	Sketches of the entrainment events in (a) grid-mixing experiments, and (b) ring-mixing experiments after the initial-adjustment phase. Horizontally-averaged density profiles $\bar{\rho}(z)$ are provided for both sketches.	178
C.1	Illustrations of the conductivity probe tip, secured to the inner and outer electrodes that form the conductivity probe. (a) The shaded regions illustrates the different regions that form the circuit-closing fluid, where the conductivity of the fluid plays a role in setting the electrical conductivity of the circuit. (b) Illustration of the lengthscales L_k, D_k that set the electrical resistivity R_k for each of the circuit-closing regions of fluid. For a full schematic of the conductivity probe, see Figure 3.13.	198
C.2	Plot of the electrical conductivity γ against the density ρ_{20} of NaCl solution at 20°C , covering the range of ρ_{20} used in our experiments. Data plotted from values tabulated by Weast et al. (1984).	200

C.3	(a) From our point-sink model, plot of the dimensionless elevation height ζ normalised by the height ζ_0 obtained when the fluid parcel makes contact with the point sink. The plot is given against the non-dimensionalised time $\tau = f(Q, \zeta)$, given by Equation (C.5). (b) Plots of the probe-measured voltage profiles $C_V^{(M)}$, taken before the start of an experiment for three values of U_S^\downarrow . Plots are given against the height $(z - H_2)$ [cm] relative to the prescribed interface.	202
C.4	Schematic of a radially confined axisymmetric jet flow. Figure duplicated from Revuelta et al. (2002).	204
C.5	(a) Plot of the derived salinity excess $S^{(n)}$ normalised by the mean derived salinity excess \bar{S} against the traverse number n . (b) Plot of the potential energy gain $PE^{(n)} - PE^{(0)}$ in mJ against the traverse number n . Both plots are taken from a single experiment, where alternating probe/thermistor profiles are taken at traverse speeds $U_S^\downarrow = 10\text{mm/s}$ (blue) and $U_S^\downarrow = 5\text{mm/s}$ (red). . .	206
D.1	Plot of L_ϕ for NaCl solution at 20°C against fluid density and the molality. Data plotted from Tables 22A and 22B of Clarke and Glew (1985).	208

List of tables

2.1	Table presenting propagation velocity U , kinetic energy E_k and impulse I_z for the Hill's spherical vortex and Lamb's thin-core cross-section models.	18
2.2	Table of studies referenced in Section 2.3 that investigated the dynamics of vortex rings impacting either no-slip surfaces, free surfaces ('Free surf.') or density interfaces ('Dens. int.'). Where applicable, the range of Re_0 , θ_0 and Ri_0 investigated for each reference are given. Starred (*) parameters under the Re_0 and Ri_0 columns correspond to alternative definitions of the Reynolds and Richardson numbers that incorporate the ring circulation Γ	33
3.1	Table presenting the molecular diffusivities of chloride (Cl^-), sodium (Na^+) and nitrate ions (NO_3^-), and the molecular diffusivity of rhodamine 6G dye (R6G). Data obtained from Weast et al. (1984) and Gendron et al. (2008).	57
4.1	Table of the input parameters used to determine which image pairs are interrogated in the implementation of our multi-frame algorithm. The n th image pair $\{\mathcal{I}_{A_n}, \mathcal{I}_{B_n}\}^{(k)}$ interrogated at the k th image-pair time step $\Delta T^{(k)}$ is taken to be evaluated at the time $\bar{T}_n^{(k)}$. Interrogation points with pixel-distance measurements $ (\delta X_I, \delta Z_I)_{n,ij}^{(k)} > s^{(k)}$ are interrogated further, utilising the time step $\Delta T^{(k)}$	101
5.1	Table of parameters for the experiments investigated in this chapter. The values and precision errors for U and a/U are determined from our bike pump calibration, discussed in Section 3.2, noting that measurements of the ring diameter were in the range $a = 49.2 \pm 1.0$ mm.	114

- 6.1 Table presenting the ring propagation speeds U , Reynolds numbers Re_0 and kinetic energy δKE for the vortex rings generated in our periodic-mixing experiments. The values and precision errors for U are determined from our bike pump calibration, discussed in Section 3.2. We label each ring case from 1 to 4 for ease of reference. 154
- 6.2 Table of parameters for the experiments presented in this chapter. A dividing line separates the experiments associated with non-penetrative and penetrative regimes (characterised by Ri_0). The ring cases correspond to vortex-ring parameters given in Table 6.1. The labels associated with each Ri_0 are used for reference throughout this chapter. 155

Chapter 1

Introduction

1.1 Context

Vortices are almost ubiquitous in fluid dynamics, appearing in a wide range of industrial, astrophysical, environmental and biological settings. Since the development of the vorticity equations by Helmholtz (1858) over 160 years ago, the wealth of research devoted to the subject of vortex dynamics has led to a rapid advancement in our understanding of many fluid dynamical phenomena, including atmospheric processes, boundary-layer separation and free-surface interactions, to name but a few. Such advances have enabled us to make accurate weather predictions, travel by air and improve the efficiency of a many industrial processes, with technologies such as centrifugal separators.

Some of the greatest challenges facing fluid dynamicists today concern the nature of turbulent flows. As suggested by Davidson (2013), turbulent flows can be classified for incompressible fluids as '*a spatially complex vorticity field which advects itself in a chaotic manner via the Biot-Savart law*'. Such flows can be further complicated by a number of physical features, including non-uniformities in the density stratification, the geometry of domain boundaries and their response to large-scale rotation. An important environment in which these three complicating features play a significant role is the oceans. This has made their accurate modelling incredibly challenging, even with the aid of today's most advanced super computers.

One aspect of turbulence in density-stratified flows is the mixing of scalar fields, such as salinity or temperature. Scalar mixing consists of two physical processes that occur simultaneously: the stirring of scalar iso-surfaces, and the enhancement of diffusion as a result of the stirring-induced intensification of the scalar gradients and the associated temporary increase in interfacial area over which diffusion can occur. Mixing of the density field can be quantified as the irreversible increase in gravitational potential energy that results

from a given energy input. This quantification allows for the efficiency of a mixing event to be defined, as the ratio of the irreversible increase in potential energy to the total energy input into the system (Winters et al., 1995). While evaluating the mixing efficiency η of discrete events in closed systems is generally achievable, as common definitions are inherently non-local, it becomes much more challenging to appropriately characterise the mixing efficiency for a number of different system types. Examples of such system types include coastal formations such as estuaries where there is significant spatio-temporal variability in the dynamics and mixing of the flow (Geyer et al., 2008), and Rayleigh-Taylor-type setups where measurements of η can depend on ‘dynamically isolated’ regions of the fluid domain (Davies Wykes et al., 2015).

A particular challenge of ocean modelling, and one that motivates this thesis, is improving our ability to account for the spatio-temporal variability of the mixing efficiency of different dynamical processes that play a significant role in the energetics of the ocean. A substantial majority of the energy input into the oceans is contributed by lunisolar tides and shear-driven forcing by surface winds (Wunsch and Ferrari, 2004). Determining the magnitude at which these energy inputs drive vertical fluxes of heat, salt and momentum in the upper-mixed layer of the ocean is a major component in ocean circulation models (Ferrari and Wunsch, 2009). Appropriate parameterisation of these fluxes is essential for modelling the ocean state under different conditions, which would be beneficial for assessing the potential future impact of different climate-change scenarios (Ivey et al., 2008).

An important component of the community-wide effort in pinning down the intricacies of turbulent mixing in density-stratified flows has been the carrying out of small-scale laboratory experiments in tanks, usually with either linear or two-layer salt stratifications. Such experiments have the advantage of being able to isolate particular features of turbulent mixing for their concentrated analysis, such as the formation of layers developing from an initially linear density stratification (Park et al. (1994), Holford and Linden (1999), Oglethorpe et al. (2013)), or the turbulent entrainment across a sharply stratified density interface (Turner (1968), Woods et al. (2010), Shravat et al. (2012)). This thesis contributes to this endeavour by considering the dynamics and mixing properties of vortex rings obliquely impacting a sharply stratified density interface.

1.2 History of the ring–eddy analogy

Vortex rings are toroidal fluid structures consisting of concentrated vorticity. In their simplest form, they are axisymmetric and propagate along their symmetry axis, under their own self-induced velocity field. Vortex rings naturally form in a wide and surprising range

of different contexts, including the diastolic phase of the cardiac cycle (Töger et al. (2012)) and the acceleration of charged ions travelling through superfluid helium (Rayfield and Reif (1964)). Academic interest in vortex rings (and vortex pairs) for high Reynolds number flows increased considerably in the early 1970s, motivated in large part by furthering the understanding of the roll-up process of wing-tip vortices and the evolution of the wake that flying aircraft leave behind (Crow (1970), Moore and Saffman (1972), Didden (1979)). The relatively simple geometry and compact nature of vortex rings has lent themselves to serving as a useful case study for a wide variety of fluid dynamical phenomena, including vortex–vortex interactions, boundary-layer separation and turbulent entrainment. This has led vortex rings to being a relatively well-understood vortical structure (though there is still much more to be learned).

Initial motivation for the application of vortex rings to the problem of stratified turbulent mixing can be attributed to the grid-mixing experiments of Turner (1968) and the observations of Maxworthy (1972). Turner investigated the mixing induced in a two-layered density stratification by a vertically oscillating grid of horizontal bars. Separate experiments were conducted using either salt or temperature as the density-stratifying agent, and oscillating grids placed in either one or both of the fluid layers, away from the interface. For the range of Richardson numbers considered, Turner observed that the grid-induced mixing sharpened the density interface, and the rate of mixing in one layer was independent of whether stirring was being induced in the other layer. Turner attributed the latter observation to statistical independence of the largest eddies that interacted with the interface, concluding that these eddies must be responsible for the majority of mixing that occurs.

Experiments were conducted by Maxworthy (1972) to clarify the structural characteristics of vortex rings in real flows, and the discrepancies that emerge when compared to classical, idealised vortex-ring models. Maxworthy highlighted the similarity between the observed entrainment processes of the vortex ring and those at a turbulent–laminar interface, leading Maxworthy to suggest that vortex rings could be treated as idealised eddies to investigate the mechanisms that allow turbulent interfaces to remain sharp. This, along with the grid-mixing observations from Turner (1968), motivated Linden (1973) to make use of this ring–eddy analogy by considering isolated vortex rings vertically impacting a density interface, as a model for the turbulent entrainment process observed by Turner (1968). Note that, throughout this thesis, we refer to the entire evolution of the interaction between the vortex ring and the interface as the ‘ring–interface interaction’ and the associated mixing as a ‘mixing event’.

For experimental studies, the ring–eddy analogy is attractive for three reasons. Firstly, the flow conditions that initiate the ring formation process can be made repeatable using simple experimental technologies, allowing for the production of highly reproducible eddy

structures. Secondly, for suitable Reynolds number, vortex rings remain laminar as they propagate from the ring generation site to the mixing site, retaining their structure and losing only a small amount of energy in their wake or to viscous dissipation. Experiments are therefore deterministic up to the stage where significant mixing begins to occur. Thirdly, theoretical models for axisymmetric vortex rings in homogeneous fluids allow for bulk properties of experimentally produced vortex rings to be estimated by measuring easily observable quantities, such as the propagation speed of the ring or the ring diameter. For example, Norbury rings (Norbury, 1973) can be compared with the experimentally produced vortex rings to provide an estimate for the amount of kinetic energy being delivered to the mixing site.

In addition to conducting a series of shadowgraph experiments, Linden (1973) developed a theoretical model to estimate the entrainment rate of an ensemble of independent eddies mixing at a density interface. An entrainment law for the ring-induced mixing was derived by assuming the rate at which energy is made available for mixing scales with the buoyancy response timescale of the perturbed interface during the ring–interface interaction. The entrainment law is given by

$$\frac{u_e}{U} \sim \text{Ri}_0^{-3/2}, \quad (1.1)$$

where u_e is the entrainment velocity, U is the propagation speed of the ring and Ri_0 is the bulk Richardson number, given by

$$\text{Ri}_0 = \frac{\rho_2 - \rho_1}{\rho_1} \frac{ag}{U^2}, \quad (1.2)$$

for ring diameter a and upper and lower-layer densities ρ_1 , ρ_2 respectively. The relation given by Equation (1.1) is consistent with the observed entrainment law by Turner (1968) when salt was used as the stratifying agent.

It has only been in the last few years that the ring–eddy analogy has been revisited. The mixing induced by the periodic generation of vertically propagating vortex rings in an initially two-layered stratification was investigated experimentally for $4 \lesssim \text{Ri}_0 \lesssim 12$ by Olsthoorn and Dalziel (2015). It was found that the initial interface remained intact and was sharpened, tending to a constant thickness. While direct measurement of the entrainment velocity was not possible due to the imposed period of consecutive ring generations, the change in height of the interface was used to determine an entrainment volume, which was shown to have an inverse dependence on Ri_0 . The mixing efficiency η of the system was found to be independent of Ri_0 over the range of Ri_0 considered, with $\eta \approx \eta_0 = 0.42$. Though there was an uncertainty of approximately 20% in estimating the kinetic energy of the ring and hence η_0 , this did not effect the main result, which was the independence of η on Ri_0 . The estimate

for η_0 was supported by Olsthoorn and Dalziel (2018), who constructed a one-dimensional model for the evolution of the energetics of the periodic-ring-mixing system, which found values similar to $\eta = 0.42$ for the mixing efficiency.

To understand how the dynamics of the ring–interface interaction allowed the mixing efficiency determined in Olsthoorn and Dalziel (2015) to be so high, three-dimensional velocity fields of the ring–interface interaction for vertically propagating vortex rings were constructed by Olsthoorn and Dalziel (2017) using a novel PIV approach. Three cases were considered, parameterised by

$$(\text{Re}_D, \text{Ri}_D) = (1600, 0.98), (2400, 1.7), (1600, 2.4). \quad (1.3)$$

Here, $\text{Ri}_D = D\text{Ri}_0/a$ and $\text{Re}_D = D\text{Re}_0/a$, where the bulk Reynolds number Re_0 is defined as

$$\text{Re}_0 = \frac{Ua}{\nu}, \quad (1.4)$$

where ν is the kinematic viscosity. Note that in Olsthoorn and Dalziel (2017), for practical considerations, the Reynolds and Richardson numbers were constructed using the diameter D of the tube from which vortex rings were generated, rather than the ring diameter a . It was noted however that $a/D \approx 1.25$ for all cases, consistent with the experiments of Maxworthy (1977).

The cases given by Equation (1.3) were classified as deeply penetrative, partially penetrative, and non-penetrative respectively, according to the maximum penetration depth into the lower layer achieved by the vortex ring. A Crow-like instability with modal structure was found to develop in the vorticity baroclinically produced by the distortion of the interface. By considering different Re_0 , Olsthoorn and Dalziel (2017) demonstrated that the dominant mode of the instability was controlled by the bulk Reynolds number, with modes 8 and 10 emerging as dominant for $D\text{Re}_0/a = 1600, 2400$ respectively. It was argued that the instability mechanism was responsible for the high mixing efficiency of the ring-induced mixing, as it allowed for the efficient transfer of the kinetic energy of the ring to smaller lengthscales before the turbulent breakdown of the instability structure.

For the cases of $(\text{Re}_0, \text{Ri}_0)$ considered, Olsthoorn and Dalziel (2017) determined a timescale τ_I associated with the growth of the Crow-like instability and showed this to depend inversely on Ri_0 ,

$$\frac{U\tau_I}{a} \sim \frac{1}{\text{Ri}_0}. \quad (1.5)$$

The relation given by (1.5) indicates that, for $\text{Ri}_0 \gtrsim 1$, the rate at which the instability grows will exceed the buoyancy-response rate of the interfacial recoil, allowing for the stratification

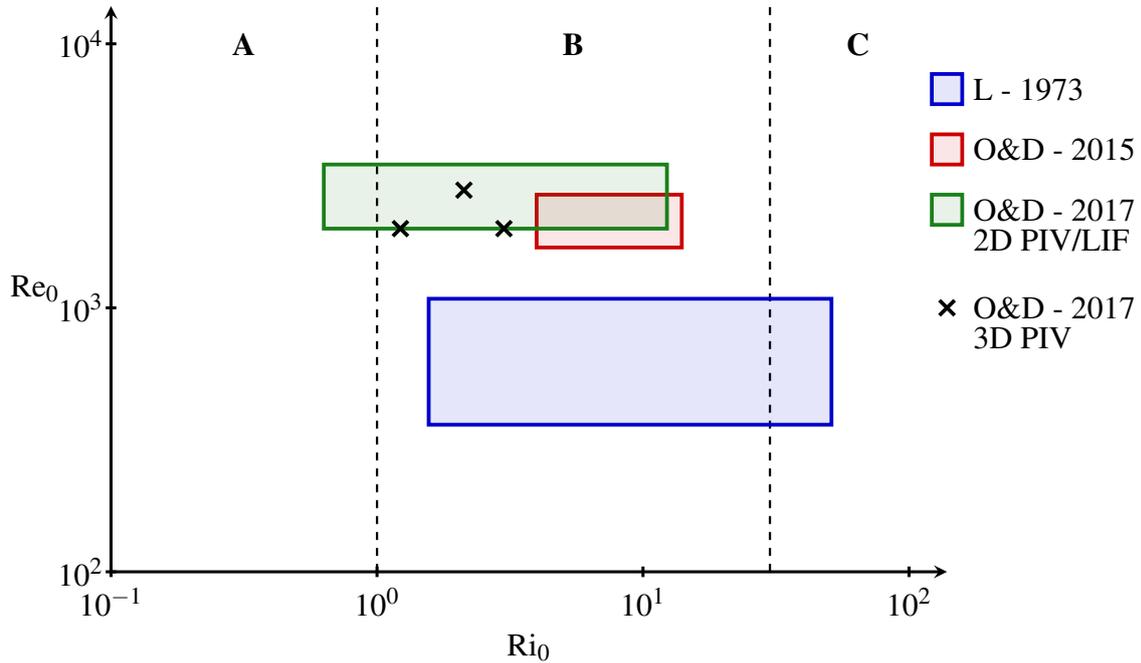


Fig. 1.1 Plot of the regions in the (Ri_0, Re_0) parameter space that have been explored in experimental studies in the context of the ring–eddy analogy for vertically propagating vortex rings impacting a density interface. The legend items correspond to studies by Linden (1973), Olsthoorn and Dalziel (2015) and Olsthoorn and Dalziel (2017). The divide between regions **A**, **B** and **C** mark possible transition regions in the nature of the turbulent entrainment induced by the ring–interface interaction.

instability to develop. Conversely, for $Ri_0 \lesssim 1$, the interface would rebound before the stratification instability is able to develop. It was suggested by Olsthoorn and Dalziel (2017) that the subsequent dynamics for $Ri_0 \lesssim 1$ may be responsible for a reduction in the mixing efficiency, though this was not investigated.

In summary, Figure 1.1 plots the regions of the (Ri_0, Re_0) parameter space where the impact of a vertically propagating vortex ring at a density interface has been investigated experimentally in the context of using vortex rings to model the turbulent entrainment process at a sharply stratified interface. The divide between regions **A** and **B** represents a transition at $Ri_0 \sim 1$ where, as Ri_0 increases, the Crow-like instability identified by Olsthoorn and Dalziel (2017) is allowed to grow before the buoyancy response of the interfacial recoil. The divide at $Ri_0 \sim 30$ between regions **B** and **C** highlight a transition for high Ri_0 indicated by Fernando and Long (1985) for grid-generated mixing, where for sufficiently large Ri_0 , the sole mixing mechanism at the interface is wave-breaking. In this large Ri_0 regime, the largest turbulent eddies are effectively flattened by buoyancy force, arresting the splashing mechanism of the ring–interface interaction described by Linden (1973).

1.3 Motivation and thesis outline

In this thesis, we extend the analogy made between vortex rings and turbulent eddies by investigating the family of oblique ring impacts onto the density interface. Previous studies employing the ring–eddy analogy have restricted their attention to vortex rings propagating normal to the density interface. For the range of (Re_0, Ri_0) considered by Olsthoorn and Dalziel (2015, 2017), their results suggest that it is the timescale for the growth of the Crow-like stratification instability non-dimensionalised by the advective timescale a/U that sets the entrainment coefficient, as opposed to the similarly non-dimensionalised timescale of the buoyancy response suggested by Linden (1973). From physical intuition this is perhaps not so surprising, as it is the stratification instability that appears to be providing the pathway for energy to be efficiently transferred from larger to smaller scale structures, where diffusive processes become more effective at converting mechanical energy into internal energy.

Assuming the stratification instability mechanism is primarily responsible for the mixing energetics, then the usefulness of the ring–eddy analogy for studying eddy-induced mixing is limited if attention is restricted to the case of normal ring impacts. This is because the azimuthal symmetry of the ring–interface interaction leads to a stratification instability that appears to be unique in the broader class of eddy–interface mixing events, thus the associated mixing properties are unlikely to be representative of general eddy–interface mixing events. This motivates our investigation of oblique ring impacts. By breaking the azimuthal symmetry, a wider range of dynamics will be exhibited in the ring–interface interaction, which we will show to have important consequences for the mixing that occurs. By considering this larger class of ring–interface interactions, further insights into the nature of eddy-induced mixing events can be made that the ring–eddy analogy has been unable to provide in previous studies.

Figure 1.2 sketches a vortex ring obliquely impacting a sharply-stratified density interface. We define the propagation angle, θ_0 , as the angle the central axis of the vortex ring makes with the vertical. The dimensionless parameters that govern the evolution of the ring–interface interaction are taken to be Re_0 , Ri_0 , and θ_0 . There are four additional parameters that can play a role in governing the evolution of the ring–interface interaction: the Formation number, F_p , which controls the core-thickness of the vortex ring relative to the ring diameter, the Schmidt number $Sc = \nu/D_\rho$ (where D_ρ is mass diffusivity), the thickness $\delta z/a$ of the density interface, non-dimensionalised by the ring diameter, and the distance between the ring-generation site and the density interface. Our experimental procedure ensures that variation in these parameters across the experiments conducted for this thesis is small enough not to compromise comparisons of different experiments with the same values for (Re_0, Ri_0, θ_0) .

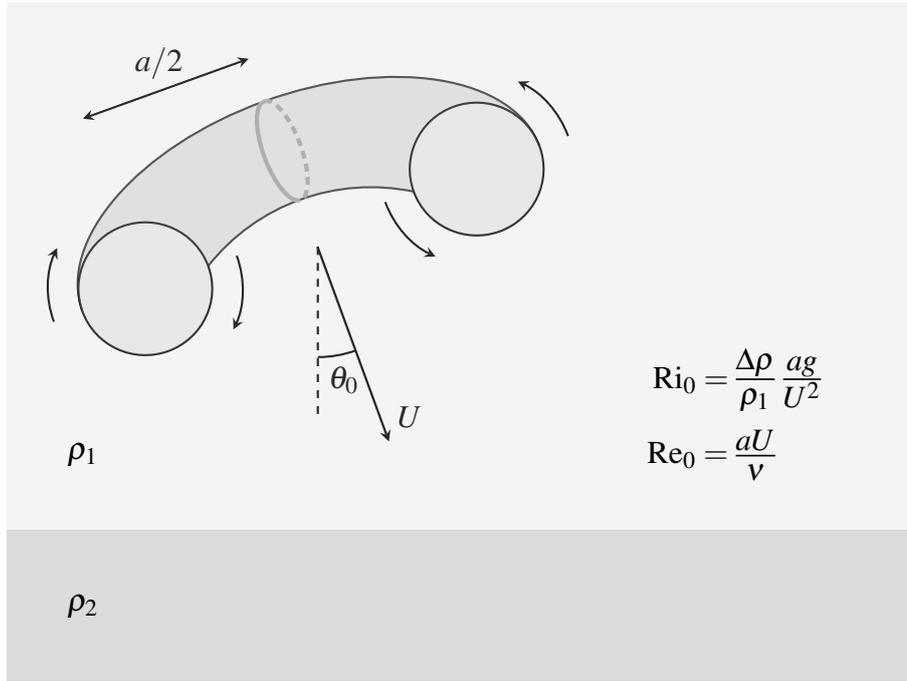


Fig. 1.2 Diagram of a vortex ring of diameter a propagating obliquely towards a sharply stratified density interface at propagation speed U and propagation angle θ_0 . Here, the illustrated two-dimensional cross-section coincides with the central plane of the flow. Richardson and Reynolds numbers Ri_0 , Re_0 are defined based on bulk properties of the vortex ring and the fluid density of each layer. The parameters presented here will be used throughout this thesis.

The rest of this thesis is organised as follows. In Chapter 2, we review theory and literature relevant for understanding the dynamics and mixing properties associated with oblique ring–interface interactions. Then, in Chapter 3 we discuss in detail our experimental setups and methodologies for the experiments reported in this thesis. To improve the dynamic-velocity-range of double-frame single-pulse PIV measurements, we have developed an algorithm that locally optimises the time separation between camera frames to reduce the local relative measurement error. This algorithm is introduced in Chapter 4. To the best of our knowledge, to date there has been no systematic attempt to describe dynamically the family of interactions resulting from a vortex ring impacting a sharply-stratified density interface. In Chapter 5, simultaneously obtained PIV and LIF data are presented of the ring–interface interactions for a range of parameters in (Re_0, Ri_0, θ_0) -space. The mixing properties of oblique ring–interface interactions is then investigated in Chapter 6, by conducting periodic-mixing experiments similar to those conducted by Olsthoorn and Dalziel (2015). Finally, in Chapter 7 we give a final discussion on the findings in this thesis and potential avenues for future research.

Chapter 2

Underlying Theory and Literature

Following our introductory chapter, which provided the historical context and motivation for this thesis, Chapter 2 reviews theory and literature important for understanding the dynamics and mixing properties of oblique ring–interface interactions. In Section 2.1, we cover fundamental theory and governing equations related to rotational and Boussinesq flows. Then, in Section 2.2, we give an overview of theoretical models for vortex rings, as well as covering literature related to their formation process and their stability. As there has been no previous systematic attempt to describe the dynamics of oblique ring–interface interactions, in Section 2.3 we review the interactions of vortex rings with no-slip walls, free surfaces and density interfaces, to provide insights on the interaction this thesis considers. Finally, in Section 2.4, we review aspects of the stratified turbulence literature related to mixing energetics, as well as discussing in more detail the previous studies using the ring–eddy analogy that were introduced in Section 1.2.

2.1 Vortex dynamics: fundamental theory

This section provides an overview of the mathematical and physical underpinnings of rotational flows that are important in understanding the dynamics of vortex rings. We begin by introducing definitions used to describe rotational flows. We then state classical theorems relating to vortex flows, that we provide the proof for later in the section. We derive the vorticity equation for incompressible, boussinesq fluids from the Navier–Stokes equation. Finally, we discuss the contribution of terms that arise in the vorticity equation and how they affect the vorticity distribution. Much of the theory reviewed in this section can be found in Chapter 1 of Green (1995).

2.1.1 Definitions

By Helmholtz's decomposition theorem, any fluid flow in 3D space $\mathbf{x} \in \mathbb{R}^3$ with velocity field $\mathbf{u}(\mathbf{x}) \in \mathbb{R}^3$ (which is twice-continuously differentiable and decays faster than $|\mathbf{x}|^{-1}$ as $|\mathbf{x}| \rightarrow \infty$) can be written as

$$\mathbf{u} = \nabla\varphi + \nabla \times \boldsymbol{\psi}, \quad (2.1)$$

where φ is the 'scalar potential' and $\boldsymbol{\psi}$ is the 'vector potential'. This decomposition yields

$$\nabla \cdot \mathbf{u} = \nabla^2\varphi, \quad \boldsymbol{\omega} = \nabla \times \mathbf{u} = -\nabla^2\boldsymbol{\psi}. \quad (2.2)$$

From Equation (2.2), it follows that φ satisfies Laplace's equation for incompressible flows. The second relation of Equation (2.2) relates the 'vorticity', denoted $\boldsymbol{\omega}$, to the vector potential. As the curl of the velocity field, the vorticity evaluated at a point \mathbf{x} is oriented with the local rotational axis of the flow, with its magnitude proportional to the local rotation rate. When there is no vorticity in a flow, the flow is said to be 'irrotational'.

If the vorticity field is known, it is possible to invert the second relation of Equation (2.2) using the 'Biot-Savart law' to obtain the rotational component of the velocity field,

$$\mathbf{u}_{\text{rot}}(\mathbf{x}, t) = \frac{1}{4\pi} \int_V \frac{\boldsymbol{\omega}(\mathbf{x}', t) \times (\mathbf{x} - \mathbf{x}') dV'}{|\mathbf{x} - \mathbf{x}'|^3}. \quad (2.3)$$

Note that the decomposition given by Equation (2.1) does not uniquely define φ and $\boldsymbol{\psi}$. Therefore the Biot-Savart law only defines the velocity to within an irrotational flow.

Closely related to the vorticity is the 'fluid circulation', Γ , an integral parameter given by

$$\Gamma = \oint_{\partial S} \mathbf{u} \cdot d\mathbf{l}, \quad (2.4)$$

where Γ is evaluated on the closed contour ∂S bounding a surface S . By applying Stokes' theorem in Equation (2.4), the circulation can be equivalently expressed as the integral of the vorticity component normal to the surface S around which Γ is evaluated, i.e.

$$\Gamma = \int_S \boldsymbol{\omega} \cdot \mathbf{n} dS, \quad (2.5)$$

where \mathbf{n} is the normal pointing outwardly from the surface. The representation of Γ given by Equation (2.5) motivates the interpretation of the vorticity as a circulation density. The circulation is one of a number of parameters that are used to characterise vortical structures.

As the vorticity field is divergence free, in an analogous way to the streamlines of a fluid flow, we can define a 'vortex line' to be a curve that is everywhere tangent to the local

vorticity field. Related to this notion is a ‘vortex tube’, defined as the set of all vortex lines passing through a simply connected surface in space. Thus at all points \mathbf{x}_S on the surface of a vortex tube, we have $\boldsymbol{\omega} \cdot \mathbf{n}(\mathbf{x}_S) = 0$. The term ‘vortex filament’ is commonly used to describe a vortex tube with cross-sectional area that is small relative to the other lengthscales exhibited in the flow.

2.1.2 Theorems and governing equations

With the definitions above, we now review four classical theorems relating to rotational flows, which we provide the proofs for sequentially in this section. These theorems are all attributable to Helmholtz (1858). The first theorem states:

- T1. Vortex tubes cannot be open ended: they must terminate at a boundary (solid or fluid-fluid interface), or form a closed loop in the fluid.

The second to fourth are Helmholtz’s vortex theorems, which hold for inviscid fluids under conservative body forces, where variations in the fluid density can be neglected. These can be stated as:

- T2. The circulation around a vortex tube, evaluated by a simple, closed contour that traverses the tube on its surface, is independent of the contour that is chosen. Moreover, this circulation does not vary in time.
- T3. In the absence of boundaries, fluid material elements that are initially irrotational will remain irrotational.
- T4. Vortex lines are material lines, convected with the local fluid velocity.

Theorem T1 is set apart from the others as it can be deduced solely from kinematic principles. In other words, this theorem requires no constraints on the flow dynamics to maintain its validity, as will be demonstrated shortly. For Helmholtz’s theorems (T2-4), despite the required constraints on the flow dynamics, it can be helpful to think of them as approximately true when the effects of viscosity and rotational body forces are negligible.

Theorem T1 can be derived by integrating $\nabla \cdot \boldsymbol{\omega}$ over a finite volume V with bounding surface ∂V . The divergence theorem gives

$$\int_V \nabla \cdot \boldsymbol{\omega} dV = 0 \Rightarrow \int_{\partial V} \boldsymbol{\omega} \cdot \mathbf{n} dS = 0. \quad (2.6)$$

Consider now a portion of a vortex tube with volume V , outer surface S_0 and cross-sectional

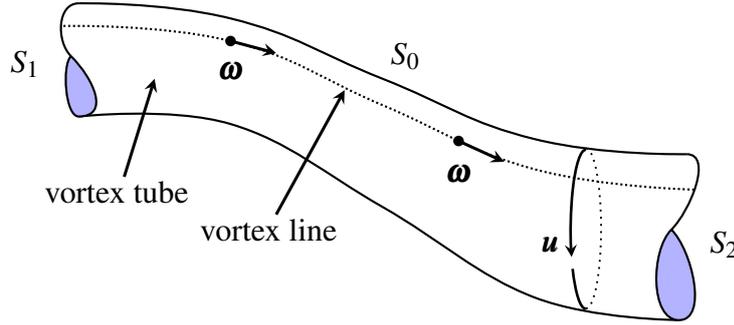


Fig. 2.1 Illustration of a portion of a vortex tube, with one vortex line marked on the vortex tube surface. Velocity and vorticity vectors are drawn to aid the visualisation of the flow associated with the vortex tube.

surfaces S_1, S_2 at its ends, as illustrated in Figure 2.1. Noting that $\boldsymbol{\omega} \cdot \mathbf{n} = 0$ on S_0 , Equation (2.6) gives

$$\int_{S_1} \boldsymbol{\omega} \cdot \mathbf{n} dS + \int_{S_2} \boldsymbol{\omega} \cdot \mathbf{n} dS = 0. \quad (2.7)$$

As \mathbf{n} is the outward pointing vector, from (2.5) we can write $\int_{S_1} \boldsymbol{\omega} \cdot \mathbf{n} dS = \Gamma_1$, $\int_{S_2} \boldsymbol{\omega} \cdot \mathbf{n} dS = -\Gamma_2$ in Equation (2.7), thus giving

$$\Gamma_1 = \Gamma_2. \quad (2.8)$$

This completes the proof for theorem T1 — if a vortex tube were open ended, the circulation conservation law given by Equation (2.8) would necessarily be violated. Thus vortex tubes must terminate at boundaries or form closed loops. The conservation law also demonstrates that as the cross-sectional area of the vortex tube decreases, the average vorticity on that cross-sectional surface must increase, and vice versa.

The proof above demonstrates that the circulation evaluated around a vortex tube is independent of the integration contour, as required for the first part of theorem T2. To prove the latter part of this theorem, we derive the material evolution equation for circulation. We begin with the momentum equation for incompressible, Newtonian fluids subject to a conservative body force,

$$\rho \frac{D\mathbf{u}}{Dt} = -\nabla P + \rho \nabla \zeta + \mu \nabla^2 \mathbf{u}, \quad (2.9)$$

where ρ is the fluid density, $D/Dt = \partial/\partial t + \mathbf{u} \cdot \nabla$ is the material time derivative, P is the fluid pressure, ζ is the body force scalar potential and μ is the dynamic viscosity. Dividing Equation (2.9) by ρ and integrating over a closed contour C , applying the chain rule gives

$$\frac{D\Gamma}{Dt} = - \oint_C \frac{\nabla P}{\rho} \cdot d\mathbf{l} + \oint_C \nabla \zeta \cdot d\mathbf{l} + \oint_C \nu \nabla^2 \mathbf{u} \cdot d\mathbf{l} + \oint_C \mathbf{u} \cdot \frac{D(d\mathbf{l})}{Dt}. \quad (2.10)$$

As C is closed, the second integral on the right-hand side of Equation (2.10) is zero by Stokes' theorem. By considering the evolution of a material loop over a small time interval, it can be shown that $D(d\mathbf{l})/Dt = d\mathbf{u}$. Therefore, as \mathbf{u} is single-valued on C , the last integral in Equation (2.10) is also zero, giving

$$\frac{D\Gamma}{Dt} = - \oint_C \frac{\nabla P}{\rho} \cdot d\mathbf{l} + \oint_C \nu \nabla^2 \mathbf{u} \cdot d\mathbf{l}, \quad (2.11)$$

where $\nu = \mu/\rho$ is the kinematic viscosity. For barotropic flows, we have $P = P(\rho)$, thus the first integral on the right-hand side of Equation (2.11) will vanish as $\nabla(P)/\rho$ is single-valued on C . If the flow is inviscid, we have $\nu = 0$ and the last integral will also vanish. Therefore, provided the flow is inviscid, barotropic and subject only to conservative body forces, the circulation around a material loop is shown to be constant, with Equation (2.11) reducing to

$$\frac{D\Gamma}{Dt} = 0. \quad (2.12)$$

Theorem T2, along with Equation (2.12), is commonly referred to as Kelvin's Circulation Theorem (Thomson (1868)), despite Helmholtz's proof published ten years earlier. From Equation (2.11), we see that the effects of viscous dissipation and baroclinicity (the generation of vorticity through misalignments between pressure and density gradients) are capable of causing the circulation around a vortex tube to vary in time.

To prove theorems T3 and T4 we first derive the vorticity equation. We start with the conservation of momentum equation given by (2.9), where we take the body force to be gravitational (giving $\zeta = -gz$). The density and pressure can be decomposed into a horizontally and temporally averaged hydrostatic component and a perturbation to the hydrostatic regime, given by

$$P = P_h(z) + p'(\mathbf{x}, t); \quad \rho = \rho_0 + \Delta\rho_h(z) + \rho'(\mathbf{x}, t); \quad \frac{dP_h}{dz} = -g(\rho_0 + \Delta\rho_h). \quad (2.13)$$

Substituting Equation (2.13) into Equation (2.9) gives

$$\rho_0 \frac{D\mathbf{u}}{Dt} + (\Delta\rho_h + \rho') \frac{D\mathbf{u}}{Dt} = -\nabla p' - g\rho'\hat{\mathbf{z}} + \mu\nabla^2 \mathbf{u}. \quad (2.14)$$

The form of Equation (2.14) is convenient for applying the Boussinesq approximation, which we adopt throughout this thesis. The Boussinesq approximation assumes that variations in density are small, and the effect of density variations on a fluid's inertia are small relative to the buoyancy force, which remains dynamically significant. Mathematically, the Boussinesq

approximation can be expressed as

$$\Delta\rho_h, \rho' \ll \rho_0, \quad \left| \frac{D\mathbf{u}}{Dt} \right| \ll g. \quad (2.15)$$

Applying the Boussinesq approximation (2.15) allows us to neglect the second term on the left-hand side of Equation (2.14). Then, by dividing through by ρ_0 we obtain

$$\frac{D\mathbf{u}}{Dt} = -\frac{1}{\rho_0}\nabla p' - \frac{g\rho'}{\rho_0}\hat{\mathbf{z}} + \nu\nabla^2\mathbf{u}, \quad (2.16)$$

where ν is assumed to be constant.

Noting that $(\mathbf{u} \cdot \nabla)\mathbf{u} = \nabla(\frac{1}{2}\mathbf{u}^2) - \mathbf{u} \times \boldsymbol{\omega}$ and $\nabla \times (\mathbf{u} \times \boldsymbol{\omega}) = \boldsymbol{\omega} \cdot \nabla\mathbf{u} - \mathbf{u} \cdot \nabla\boldsymbol{\omega}$, we obtain the vorticity equation for an incompressible flow by taking the curl of Equation (2.16), giving

$$\frac{D\boldsymbol{\omega}}{Dt} = (\boldsymbol{\omega} \cdot \nabla)\mathbf{u} + \frac{g}{\rho_0}\hat{\mathbf{z}} \times \nabla\rho' + \nu\nabla^2\boldsymbol{\omega}. \quad (2.17)$$

The last term on the right-hand side of Equation (2.17) represents the viscous diffusion of vorticity, which has greatest impact where vorticity gradients are at their highest. The penultimate term represents ‘baroclinic’ torque, generating vorticity when pressure gradients and density gradients are misaligned. In many stratified flows and environments, baroclinicity is an important mechanism for the generation of vorticity.

The first term on the right-hand side of Equation (2.17) represent the enhancement of vorticity due to vortex stretching. To develop intuition for this mechanism, consider an inviscid, homogeneous flow, so the last two terms of Equation (2.17) can be neglected. The vorticity equation then becomes

$$\frac{D\boldsymbol{\omega}}{Dt} = (\boldsymbol{\omega} \cdot \nabla)\mathbf{u}. \quad (2.18)$$

This evolution equation for $\boldsymbol{\omega}$ is of the same form as that for the length of an infinitesimal segment, l , of a fluid material line. It therefore follows that the following relation must hold for all time,

$$l = C\boldsymbol{\omega}, \quad (2.19)$$

where C is a constant (see Section 1.3.4 of Green (1995)).

From Equation (2.19) we see that stretching a segment of fluid with vorticity increases the vorticity magnitude of the segment. This result is consistent with our understanding of vortex tubes given by theorem T1. In order to maintain uniformity of circulation along a

vortex tube, the vorticity must intensify where the cross-sectional area of the tube decreases, i.e. in regions of the vortex tube that are being stretched.

We are now in a position to prove theorems T3 and T4 (albeit for flows under gravitational body forces rather than a general conservative body force, but the chain of logic is more or less identical). Theorem T3 is deduced from Equation (2.18). For flows with uniform density, we see that if a material element is initially irrotational, then initially we have $D\boldsymbol{\omega}/Dt = 0$, thus vorticity cannot be generated and the material element will remain irrotational for all time. This theorem only holds true in the absence of boundaries. At a boundary, it is possible to obtain a singularity in the $\nu\nabla^2\boldsymbol{\omega}$ term of Equation (2.17), which would act as a source of vorticity.

Theorem T4 is deduced from Equation (2.19). If there are no density variations in the flow, then vortex lines are coincident with material lines for all time, being convected with the local fluid velocity. For such flows, the topology of vortex tubes is preserved: closed vortex loops remain closed and knotted vortex structures remain knotted in the same way (Moffatt (1969)).

For homogeneous flows, the action of viscosity is necessary for the topology of vortex lines to change, with the rate of deviation between material lines and vortex lines proportional to the component of vorticity diffusion perpendicular to the material lines, $\nabla_{\perp}^2\boldsymbol{\omega}$, as demonstrated in Kida and Takaoka (1994). This mechanism is responsible for many vortex interaction phenomena observed in turbulent flows.

2.2 Vortex rings

In the previous section we reviewed fundamental properties of vortex dynamics that are important for understanding the vortical flows that are presented in this thesis. In this section, we begin our discussion of vortex rings. In general, the term ‘vortex ring’ is used to describe closed vortex tubes enclosing a region of concentrated vorticity, with the vortex lines comprising the tube selected using some threshold of the vorticity magnitude. In their simplest form, vortex rings are axisymmetric, toroidal fluid structures with azimuthal vorticity that propagate in the direction of their symmetry axis, under their self-induced velocity field. A diagram of an axisymmetric vortex ring is given in Figure 2.2(a).

This section focuses on vortex rings in homogeneous fluids. We begin by presenting theoretical models of axisymmetric vortex rings. We then continue with a discussion on the formation process of vortex rings in laboratory experiments and related models for their development. We end this section with a discussion on the instability mechanisms exhibited by vortex rings.

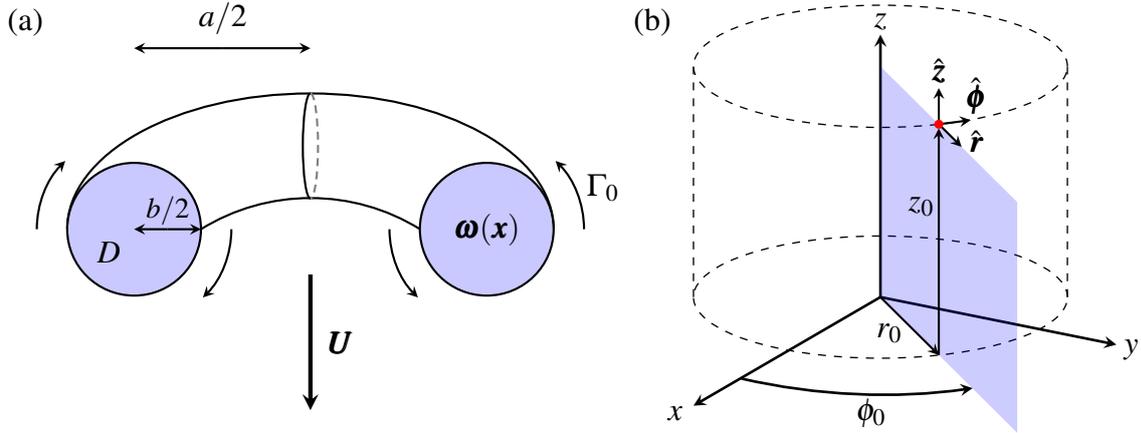


Fig. 2.2 (a) Diagram of the cross-section of an axisymmetric vortex ring, with radius a , circulation Γ_0 and propagation speed U . The vortex core has core size b and vorticity distribution $\omega(\mathbf{x})$. The core region D has been drawn with circular cross-section for simplicity. (b) Sketch of the relationship between the Cartesian and cylindrical coordinate systems.

2.2.1 Axisymmetric models

Much of the theory discussed in this subsection can be found in Akhmetov (2009) and Green (1995). For axisymmetric flows, it is convenient to use a cylindrical coordinate system (r, ϕ, z) , with the z -axis coinciding with the axis of symmetry. A sketch of the coordinate system is given in Figure 2.2(b), including the associated unit vectors $\hat{\mathbf{r}}$, $\hat{\boldsymbol{\phi}}$ and $\hat{\mathbf{z}}$. Note that the flows discussed in this subsection are without swirl ($\partial/\partial\phi = 0$, $u_\phi = 0$), giving $\boldsymbol{\omega} = \omega(r, z)\hat{\boldsymbol{\phi}}$.

In an inviscid, homogeneous fluid, the vorticity equation is of the form given by Equation (2.18). This allows us to reduce Equation (2.18) and the incompressibility condition to give conservation equations

$$\frac{D(\omega/r)}{Dt} = 0, \quad (2.20)$$

$$\nabla \cdot \mathbf{u} = \frac{\partial u_z}{\partial z} + \frac{1}{r} \frac{\partial (ru_r)}{\partial r} = 0. \quad (2.21)$$

The incompressibility condition (2.21) allows us to introduce an axisymmetric stream-function ψ defined by $\mathbf{u} = \nabla \times (\psi/r)\hat{\boldsymbol{\phi}}$, giving

$$u_z = \frac{1}{r} \frac{\partial \psi}{\partial r}, \quad u_r = -\frac{1}{r} \frac{\partial \psi}{\partial z}, \quad \omega = -\nabla^2 \left(\frac{\psi}{r} \right) = -\frac{1}{r} \left(\frac{\partial^2 \psi}{\partial z^2} + \frac{\partial^2 \psi}{\partial r^2} - \frac{1}{r} \frac{\partial \psi}{\partial r} \right). \quad (2.22)$$

If the vorticity field is known, Equations (2.22) allow us to apply the Biot-Savart law (given by Equation (2.3)) to evaluate the streamfunction in terms of the vorticity,

$$\psi(\mathbf{x}) = -\frac{r}{4\pi} \int \frac{\boldsymbol{\omega}(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} dV'. \quad (2.23)$$

In addition to the circulation Γ , for axisymmetric flows without swirl there are two further invariants of the flow that can be simply defined in terms of $\boldsymbol{\omega}$ and ψ : the hydrodynamic impulse \mathbf{I} and the kinetic energy E_k . In general, these invariants are given by

$$\mathbf{I} = \frac{1}{2} \int_V \rho \mathbf{x} \times \boldsymbol{\omega} dV, \quad E_k = \int_V \rho \mathbf{u} \cdot (\mathbf{x} \times \boldsymbol{\omega}) dV. \quad (2.24)$$

For a uniform density fluid, substituting $\boldsymbol{\omega}$ and ψ reduces Equations (2.24) to

$$I_z = \pi \rho \int_0^\infty \int_{-\infty}^\infty \omega r^2 dz dr, \quad E_k = \pi \rho \int_0^\infty \int_{-\infty}^\infty \psi \omega dz dr. \quad (2.25)$$

Note that, while the derivation of I_z in Equation (2.25) is straightforward, the derivation for E_k is less trivial. The derivation for E_k can be found in Section 3.11 of Saffman (1995).

The circulation can be evaluated by taking the integration contour to be the boundary of a semi-infinite (r, z) plane of constant $\phi = \phi_0$ (shaded blue in Figure 2.2(b)). This leads to a simple expression for the circulation,

$$\Gamma = \int_0^\infty \int_{-\infty}^\infty \omega dz dr. \quad (2.26)$$

If the vorticity distribution is entirely contained within a finite volume, the integral limits of the invariants can be replaced with finite limits, providing the integration volume encloses all of the vorticity.

The invariants given by Equations (2.25) and (2.26) restrict the way in which the flow can evolve, thus providing some insight into the flow behaviour in these idealised fluid environments. Note that, in the presence of viscosity, only the impulse I_z remains invariant - kinetic energy is lost through viscous dissipation and there can be a reduction in circulation due to the diffusion of vorticity across the central axis of the ring, leading to its annihilation.

The discussion up to this point has concerned general inviscid, axisymmetric flows without swirl. We now consider the particular case of a vortex ring. From Equation (2.20) it follows that the quantity ω/r is constant for material elements in the flow. In the reference frame moving with the vortex ring, if the flow is steady then the streamlines of the flow do not evolve in time. In this case, the quantity ω/r depends only on the streamfunction ψ ,

Model	U	E_k	I_z
Hill	$\frac{1}{5} \frac{\Gamma_0}{R}$	$\frac{2}{35} \pi \rho \Gamma_0^2 R$	$\frac{2}{5} \pi \rho \Gamma_0 R^2$
Lamb	$\frac{\Gamma_0}{4\pi a} \left(\ln \left(\frac{8a}{b} \right) - \frac{1}{4} \right)$	$\frac{\Gamma_0^2 a \rho}{2} \left(\ln \left(\frac{8a}{b} \right) - \frac{7}{4} \right)$	$\pi \rho \Gamma_0 a^2$

Table 2.1 Table presenting propagation velocity U , kinetic energy E_k and impulse I_z for the Hill's spherical vortex and Lamb's thin-core cross-section models.

giving rise to the steadiness condition,

$$\frac{\omega}{r} = f(\psi). \quad (2.27)$$

This condition can be substituted into the expression for ω given in Equation (2.22) to obtain a PDE in the streamfunction ψ ,

$$\frac{\partial^2 \psi}{\partial z^2} + \frac{\partial^2 \psi}{\partial r^2} - \frac{1}{r} \frac{\partial \psi}{\partial r} + r^2 f(\psi) = 0. \quad (2.28)$$

To make progress in solving Equation (2.28), it is necessary to prescribe the function $f(\psi)$. In doing so, one is making assumptions about the vorticity distribution in the flow. In the classical models of Hill (1894) and Lamb (1932), it is assumed that the vorticity is concentrated in a core region D (shaded blue in Figure 2.2(a)) with characteristic length scale b . It is also assumed that in the interior of D , the function f is constant. This can be written as

$$\frac{\omega}{r} = f(\psi(\mathbf{x})) = \begin{cases} A & \mathbf{x} \in D \\ 0 & \mathbf{x} \notin D \end{cases}. \quad (2.29)$$

An analytic solution to Equations (2.28) and (2.29) was discovered by Hill (1894), who took D to be a sphere of radius R . In the limit of a thin vortex ring with circular cross-section, where $b/a \ll 1$ and $\ln(b/a) \sim 1$, an approximate solution was found by Lamb (1932), correct to $O(b/a)$. Derivations for these solutions can be found in Akhmetov (2009). Expressions for U , E_k and I_z are presented in Table 2.1 in terms of a prescribed circulation Γ_0 and the length scales associated with Hill's and Lamb's models. These expressions demonstrate that the vortex ring dynamics have a significant dependence on the geometry of the vortex core, even in the absence of viscosity.

Using the same form of vorticity distribution given by Equation (2.29), Norbury (1973) performed numerical calculations to determine a family of steady inviscid vortex rings of varying aspect-ratio, of which Hill's and Lamb's vortex rings are end members. These

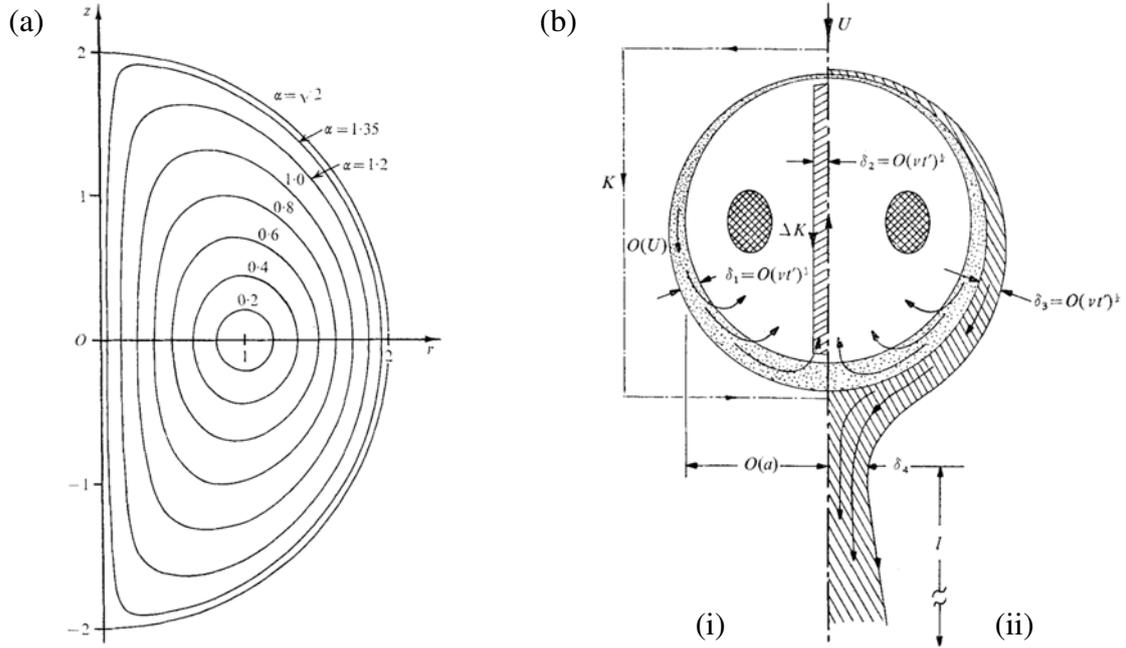


Fig. 2.3 (a) Plots of the core boundary for Norbury's rings for a range of values of α , duplicated from Norbury (1973). (b) Diagram of Maxworthy's bubble entrainment model, where the left-hand side (i) prohibits the shedding of vorticity into the bubble exterior, and the right-hand side (ii) allows for the formation of a wake behind the vortex ring where deposited vorticity can annihilate. This diagram is duplicated from Maxworthy (1972).

rings are parameterised by a non-dimensional mean core radius $0 < \alpha \leq \sqrt{2}$, defined by $A_D = \pi L^2 \alpha^2$, where A_D is the cross-sectional area of D and L is the vortex ring radius. Norbury rings provide a useful comparison to experimentally produced vortex rings, allowing quantities such as the kinetic energy and impulse to be estimated when their precise value cannot be determined. Plots of the core cross-sectional shape are given in Figure 2.3(a).

When viscosity is introduced, the assumption that the flow remains steady in these models can no longer be made. For instance, if Hill's spherical vortex is to satisfy the Navier-Stokes equation at $r = R$, vorticity must be allowed to diffuse at the vorticity discontinuity. For thin-core models, Saffman (1970) extended Lamb's model to include viscosity and predicted a slowing of the ring, given by

$$U = \frac{\Gamma_0}{4\pi a} \left[\ln \left(\frac{8a}{b_e} \right) - 0.558 + O \left(\frac{\nu t}{a^2} \right)^{1/2} \right], \quad (2.30)$$

where $b_e \sim (\nu t)^{1/2}$ is an effective core radius. Saffman demonstrated that if the thin-core approximation is retained for a viscous vortex ring, a contradiction can be reached at large times by considering the invariants of the flow. The circulation would remain constant as the

core-concentrated vorticity cannot self-annihilate at the central axis of the flow. The impulse $I_z = \pi\rho\Gamma_0 a^2$ given in Table 2.1 would also remain constant, yet experimental observations demonstrate that the vortex ring radius a increases in time. The conclusion reached is that vorticity cannot be confined to a thin core in vortex rings and the circulation must be allowed to decrease.

The last model presented in this section, derived by Maxworthy (1972), predicts the formation of a wake behind the vortex ring, a feature that is observed in experimentally produced vortex rings. Maxworthy assumes that the fluid volume travelling with the vortex ring remains similar at all times and that the vorticity is continuously distributed throughout this volume (referred to by Maxworthy as the ‘bubble’). If one assumes that no vorticity is shed into a wake (as illustrated in Figure 2.3 (b)(i)), the impulse contained within the bubble volume remains constant at all times. In this case, by considering the bubble boundary layer diffusing vorticity into the outer irrotational fluid and the inevitable re-entrainment of this newly contaminated fluid back into the vortex ring, one can derive the following power laws:

$$V \sim t, \quad a \sim t^{1/3}, \quad U \sim t^{-1}, \quad \Gamma \sim t^{-2/3}. \quad (2.31)$$

It is the last scaling in (2.31) where a contradiction is reached. As vorticity is confined to the bubble, the circulation can only decrease due to the annihilation of vorticity along the portion of the central axis contained within the bubble volume. This annihilation would lead to an alternative scaling for the circulation, $\Gamma \sim t^{-1/3}$, which is inconsistent with Maxworthy’s model. To resolve this, the vorticity must be allowed to cancel across a longer portion of the central axis. By allowing some vorticity to be shed into a wake behind the bubble where it can self-annihilate (as illustrated in Figure 2.3 (b)(ii)), Maxworthy obtains $\Gamma \sim t^{-2/3}$, thus recovering self-consistency of his entrainment model.

2.2.2 Formation process

One of the primary limitations of the models outlined in Section 2.2.1 is the uncertainty associated with the assumed vorticity distribution of the vortex core. In the absence of instabilities and other disturbances, the vorticity distribution of experimentally produced vortex rings is determined by two mechanisms: the ring-formation process and the diffusion of vorticity after its formation. Whilst the latter is at least described simply by the vorticity equation (2.17), the former is far more complex and varied.

There are a number of techniques used to produce vortex rings (examples are outlined in Shariff and Leonard (1992)). For each case, the ring vorticity distribution is controlled by the flow conditions local to the ring generation site. In particular, the impulse generation

method for producing the ring, the geometry of any confining boundaries at the generation site and the local fluid stratification will all play a significant role in determining the final ring vorticity distribution. Each case demands its own analysis to better understand the structure of the vortex rings that are produced.

In this subsection, attention is restricted to the ring generation technique we employ in the experiments conducted for this thesis, which is to generate rings at the outlet of a cylindrical tube with constant circular cross-section. In particular, we review previous research relating the control parameters to the properties of the vortex ring that is produced. The reader is referred to Section 4.1 of Green (1995) and Section 2 of Shariff and Leonard (1992) for further discussion of the ring-formation process.

The production of a vortex ring is initiated when a piston-like mechanism delivers an impulse to the fluid within the tube, causing some of this fluid to be ejected from the tube outlet. When this happens, a cylindrical vortex sheet containing azimuthal vorticity forms in the viscous boundary layer on the inside of the tube walls, due to the no-slip boundary condition on the walls. As the vortex sheet exits the tube, it separates at the edge of the tube outlet and rolls up into the developing vortex-ring core, continually entraining irrotational fluid from its surroundings during the process. This leads to the vortex rings increasing in diameter as it develops. The formation process ends shortly after the piston stops, when vorticity ceases to be fed to the vortex ring. Dye visualisation snapshots of the formation process are given in Figure 2.4, reproduced from Didden (1979).

It is the nature in which the vortex sheet rolls up into the vortex ring-core that determines the vortex-core structure. Defining a Reynolds number $\text{Re}_p = U_p D_p / \nu$ for the formation process, where U_p is the piston speed and D_p is the tube diameter, Saffman (1978) argues that, for $\text{Re}_p \gg 1$, the effect of viscosity is small on the formation process, with its main effect being to remove a singularity at the centre of the developing vortex core. Then, on dimensional grounds, it would follow that

$$\frac{2a}{D_p} = f\left(\frac{L_p}{D_p}\right), \quad \frac{\Gamma}{U_p L_p} = g\left(\frac{L_p}{D_p}\right), \quad (2.32)$$

for a vortex ring produced from a circular tube, where L_p is the total piston stroke length travelled over the formation process and a is the ring radius (as defined by Figure 2.2 (a)).

The precise form of the functions f and g will depend on both the evolution of the piston velocity profile $U_p(t)$, and the geometry of the ring generation site. For a ring generated from a circular hole, the contact angle between the tube interior and the rest of the fluid domain (denoted α in Figure 1 of Gharib et al. (1998)) influences the vorticity distribution of any vortex rings that are generated from that hole. This motivates the distinction between

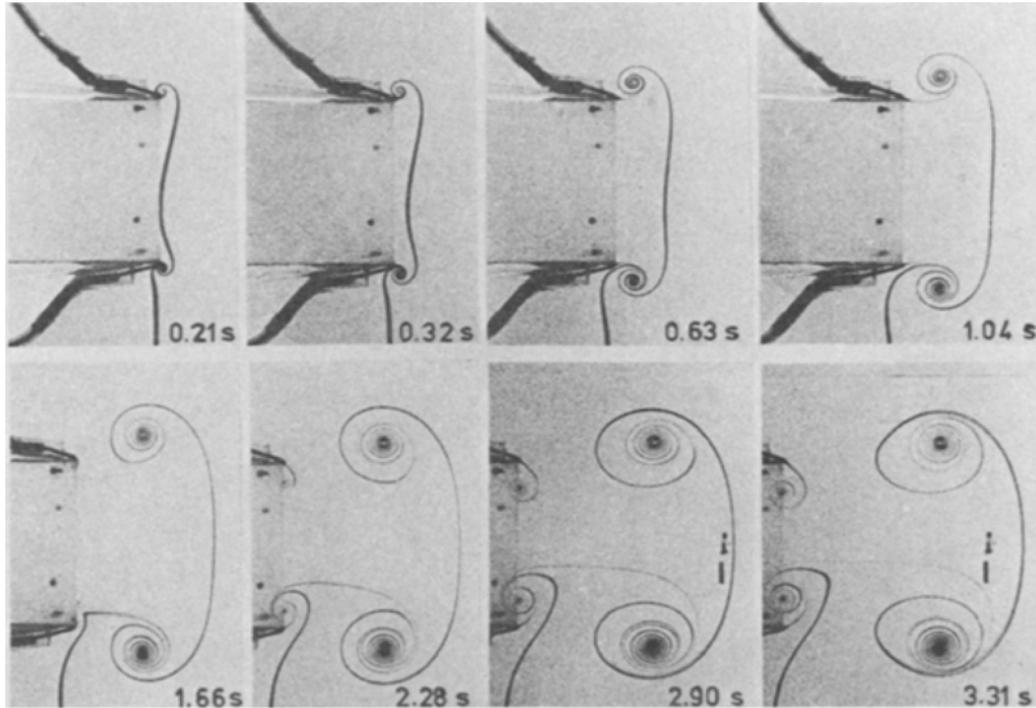


Fig. 2.4 Snapshots of the ring-formation process, using dye to visualise a central plane of the vortex ring. The tube diameter is 5 cm and the piston speed is 4.6 cm/s, giving Reynolds number $Re_p = 2300$. Figure reproduced from Didden (1979).

‘tube’ openings (corresponding to small α) and ‘orifice’ openings (corresponding to small $\pi/2 - \alpha$), introduced and investigated by Pullin (1979).

The relations given by Equations (2.32) suggest that the Formation Number, $F_p = L_p/D_p$, can be used as a control parameter for the structure of generated vortex rings. Related to F_p is the formation time, $t_p = \overline{U}_p t/D_p$, a dimensionless time that varies between 0 and F_p over the ring-formation process. Here, \overline{U}_p is the piston speed time-averaged over the period between the initial and final motion of the piston.

The experimental work conducted by Gharib et al. (1998) investigated the influence of F_p on vortex-ring formation. Gharib found that a limiting formation number $F_p = F_{p,lim}$ exists where the generated vortex ring attains maximal circulation from the formation process. For formation numbers $F_p > F_{p,lim}$, any surplus circulation produced by the formation process would contribute to a trailing jet behind the ring. For tube and orifice openings, Gharib investigated with different piston velocity profiles and exit diameters, finding the limiting formation number for all cases to lie in the range of 3.6 – 4.5. The numerical simulations of Rosenfeld et al. (1998) were presented as a continuation of Gharib’s experimental work. We will not go into the details here, but the findings of these simulations were in broad agreement with the experiments.

Motivated by Gharib's finding of the $F_{p,lim}$, analytic models were developed independently by Mohseni and Gharib (1998) and Linden and Turner (2001) to explain the physical mechanisms that set the value of $F_{p,lim}$. Both models employ a 'before and after' calculation, matching the parameters associated with a plug of fluid (of length L_p , diameter D_p and velocity U_p) ejected from a tube, with the parameters associated with the Norbury family of steady rings (of radius X and velocity U , see Norbury (1973)).

The matching calculations for the non-dimensional ring circulation Γ_R , impulse I_R and kinetic energy $E_{k,R}$ are given by

$$UX\Gamma_R = \frac{1}{2}U_pL_p, \quad (2.33a)$$

$$UX^3I_R = \frac{1}{4}\pi U_pD_p^2L_p, \quad (2.33b)$$

$$U^2X^3E_{k,R} = \frac{1}{8}\pi cU_p^2D_p^2L_p, \quad (2.33c)$$

where Linden and Turner introduced $c < 1$ as the fraction of kinetic energy that is injected from the plug fluid into the ring. Denoting Equations (2.33) as (a), (b) and (c) respectively, by taking (b)/(a) and (c)/(a)² and eliminating X/D_p to combine them, we obtain a condition for the formation number,

$$F_p = \frac{L_p}{D_p} = \sqrt{\frac{\pi}{2}c} \frac{I_R^{1/2}\Gamma_R^{3/2}}{E_{k,R}}. \quad (2.34)$$

Additionally, we obtain a condition relating the volume of the ring to that of the plug fluid V_p by taking (c)/(b)², giving

$$\frac{X^3}{V_p} = \frac{2E_{k,R}}{cI_R^2}, \quad (2.35)$$

where $V_p = \frac{1}{4}\pi D_p^2L_p$. Plots of Equations (2.34) and (2.35) are shown in Figure 2.5, where evaluation of these equations is made for the family of Norbury rings. For this case, all ring parameters are a function of the non-dimensional mean-core radius, α (introduced in Section 2.2.1). The volume of fluid V_c carried within a Norbury vortex core and the volume of irrotational fluid V_e carried along with the core are distinguished.

The model presented by Mohseni and Gharib (1998) employs the Kelvin–Benjamin variational principle for steady-axis touching rings, where an energy constraint limits the maximum energy carried by a vortex ring with touching cores. This method yields a second relation for F_p in addition to Equation (2.34), thus requiring a particular value for α for both relations to be satisfied. Mohseni and Gharib's analysis predicts $\alpha \approx 0.3$, corresponding to $F_{p,lim} \approx 3$. Slightly higher values than the predicted ring velocity yields estimates of $3 \lesssim F_{p,lim} \lesssim 4.5$, consistent with the observations of Gharib et al. (1998).

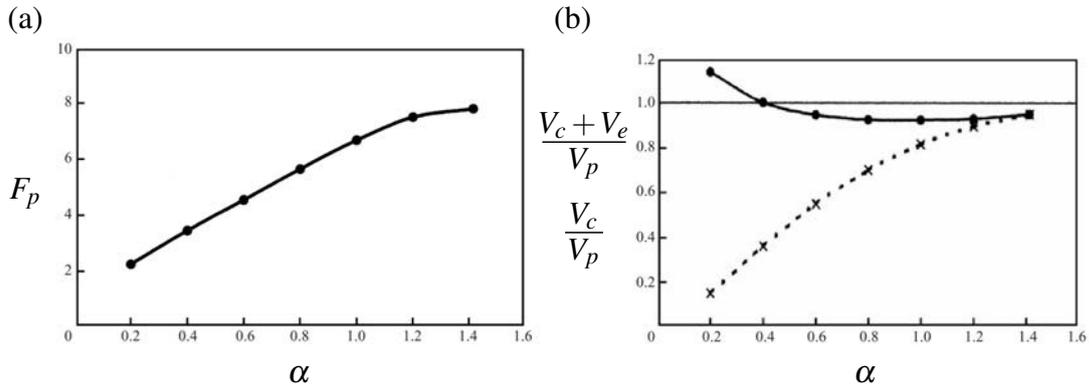


Fig. 2.5 Taking $c = 1$, plots of (a) (2.34) and (b) (2.35) for the ring core volume (dashed) and total volume carried by the ring (solid), non-dimensionalised by the plug volume. The graphs are plotted for Norbury rings, using tabulated values from Norbury (1973). Plot duplicated from Linden and Turner (2001).

However, Mohseni notes that the corresponding value of the Norbury mean-core radius α is considerably smaller than $\sqrt{2}$, indicating a core-touching vortex is not a good approximation to the rings produced at the limiting formation number. A similar observation is made by Linden and Turner (2001) who note that Figure 14(c) of Gharib et al. (1998) shows the vortex ring appears to have two well-separated cores. These observations cast doubt in the validity of applying the Kelvin–Benjamin variational principle in predicting the limiting formation number for this context.

The model presented by Linden and Turner (2001) considers the volumetric constraint given by Equation (2.35), where calculations are plotted in Figure 2.5(b). The calculations demonstrate that for $\alpha \gtrsim 0.42$, corresponding to $F_p \gtrsim 3.5$, the ring is unable to accommodate all the plug fluid. This limiting value is consistent with the observations of Gharib et al. (1998), predicting that for larger formation numbers, some of the circulation-carrying plug fluid will inevitably be left behind in the ring’s wake. For $\alpha \lesssim 0.42$, although the volume of the ring exceeds the plug volume, it is seen in experimentally-produced vortex rings that some of the plug fluid is typically left behind in the wake of the ring (Maxworthy, 1972).

Linden and Turners’ analysis demonstrates that, for vortex rings generated from a circular tube, it is a volumetric constraint that sets the limiting formation number where a single vortex ring is produced, rather than the energy constraint associated with the Kelvin–Benjamin variational principle. Linden and Turner also point out that the matching analysis above could be reused with a more sophisticated model for the ejected fluid, or with experimental data rather than the Norbury family of rings, to make different predictions on the limiting formation number.

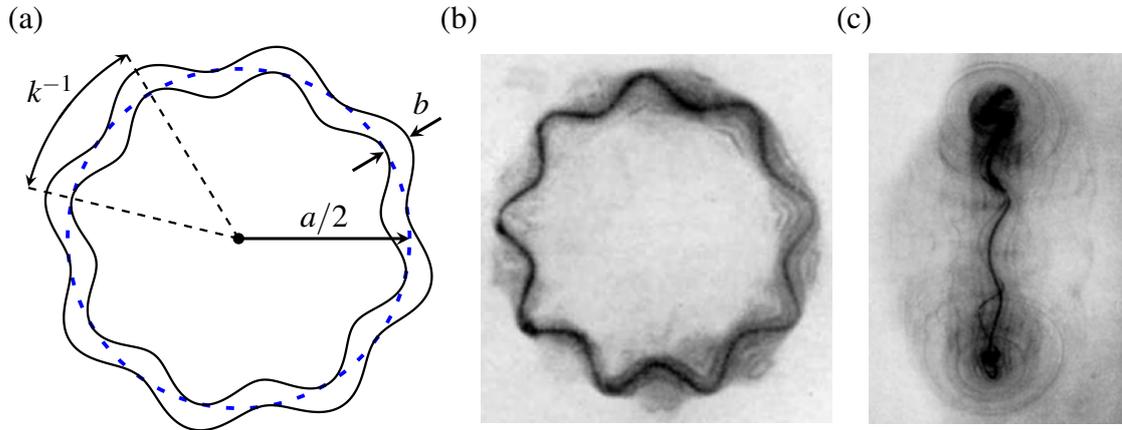


Fig. 2.6 (a) Illustration of the Widnall instability with 8 peaks and wavenumber k . The unperturbed ring, whose centreline is indicated by the blue dashed circular line, has core diameter b and ring diameter a . (b) Photograph of a vortex ring exhibiting the Widnall instability, with 9 peaks. (c) Photograph of the Widnall instability viewed from the side. Photographs reproduced from Krutzsch (1939).

We note that, for other tube/orifice ring generation setups, it is conceivable that the Kelvin–Benjamin energy constraint may become the limiting factor over the volumetric constraint in setting the maximum formation number in producing a single ring. Mohseni and Gharib (1998) speculate, for example, that a piston-tube setup with a time varying exit diameter might be capable of producing thick-core rings similar to Hill’s vortex, in which case the Kelvin–Benjamin variational principle may apply. However, we are unaware of any experimental attempts made to produce thick-core rings with $\alpha \rightarrow \sqrt{2}$ so we make no definitive statement about this possibility.

2.2.3 The Widnall instability and the stability of vortex tubes

The Widnall instability (Widnall and Tsai, 1977) is a three-dimensional, short wavelength, azimuthal instability that can be exhibited by vortex rings. These waves are observed to be the dominant instability in the laminar stage of a vortex ring’s development and are capable of being excited without the forcing of an externally induced flow. The Widnall instability generates waves that deform the centreline of the vortex core and are stationary in the reference frame moving with the bulk motion of the ring. An illustration and photograph of the Widnall instability are given in Figure 2.6.

In this subsection we walk through the theory associated with the Widnall instability so its development can be understood. We begin with a brief historical overview of developments made in understanding the stability of vortex tubes. We then derive the dispersion relation

for small sinusoidal perturbations to a cylindrical vortex column of uniform vorticity. Of the waves permitted by this dispersion relation, attention is then restricted to the mode types that possess the largest growth rates. These modes are then linked to the Crow Instability and Widnall Instability. We end this subsection with a brief discussion on the number of peaks that are expected to appear on the ring core when subject to the Widnall instability.

The earliest, commonly referenced observation of the instability now named after Widnall goes back to Kruttsch (1939). Interest in this phenomenon was rekindled just over 30 years later, with Crow's study on the related problem of the stability of a pair of vortex filaments (Crow, 1970). This later study was motivated by the need to understand the decay of trailing vortices behind aircraft, as this flow could pose a threat to other aircraft. Analytical studies on the stability of cylindrical vortex tubes in the presence of axial flow (Moore and Saffman, 1972) and an externally imposed strain field (Moore and Saffman (1975), Tsai and Widnall (1976)) followed shortly after. Attention was also given to the stability of vortex rings at this time, with analytical (Widnall and Sullivan (1973), Widnall et al. (1974), Widnall and Tsai (1977), Saffman (1978)) and experimental studies (Liess and Didden (1976), Maxworthy (1977)). More recent studies have given considerable insights into the stability of vortex rings in stratified environments and in the presence of boundaries. We defer our discussion on such cases to Section 2.3.

The first analysis on the stability of vortex tubes in inviscid, homogeneous flows can be attributed to Lord Kelvin (Thomson, 1880). For base flows of the form $\mathbf{u} = U_0(r)\hat{\phi}$, Kelvin considered small sinusoidal perturbations flows of the form

$$\begin{cases} \dot{r} = u_r \cos(kz) \sin(m\phi + st), & r\dot{\phi} = U_0 + u_\phi \cos(kz) \cos(m\phi + st), \\ \dot{z} = w \sin(kz) \sin(m\phi + st), & p = P_0 + p_1 \cos(kz) \cos(m\phi + st), \end{cases} \quad (2.36)$$

where $P_0 = \int U_0^2/r dr$ is prescribed by the 0th order radial component of the Euler equation and (k, m, s) are the perturbation wavenumbers, with m restricted to taking integer values. Kelvin gave particular attention to perturbations of a cylindrical column of uniform vorticity 2Ω surrounded by irrotational fluid. The base flow is given by

$$U_0(r) = \begin{cases} \Omega r & r \leq R \\ \Omega a^2/r & r \geq R \end{cases} . \quad (2.37)$$

By using Equations (2.36) and (2.37) to linearise the Euler and continuity equations, it is possible to obtain the second order ODEs governing the evolution of w :

$$\begin{cases} \frac{d^2 w}{dr^2} + \frac{1}{r} \frac{dw}{dr} - \frac{m^2}{r^2} w + \beta^2 w = 0 & r \leq R \\ \frac{d^2 w}{dr^2} + \frac{1}{r} \frac{dw}{dr} - \frac{m^2}{r^2} w - k^2 w = 0 & r \geq R \end{cases}, \quad \text{where } \beta^2 = k^2 \frac{4\Omega^2 - (s + m\Omega)^2}{(s + m\Omega)^2}. \quad (2.38)$$

The general solution of Equation (2.38) can be written as follows:

$$w = \begin{cases} C_1 J_{|m|}(\beta r) + C_2 Y_{|m|}(\beta r) & r \leq R \\ C_3 I_{|m|}(kr) + C_4 K_{|m|}(kr) & r \geq R \end{cases}, \quad (2.39)$$

where C_i are constants, $J_{|m|}$, $Y_{|m|}$ are the order m Bessel functions of the first and second kinds respectively, and $I_{|m|}$, $K_{|m|}$ are the order m modified Bessel functions of the first and second kinds respectively. Regularity at the origin and at infinity demands $C_2 = C_3 = 0$. Then, by imposing continuity of perturbation components of the flow at $r = a$, a dispersion relation is determined,

$$\frac{\kappa^2 J'_{|m|}(\xi)}{\xi J_{|m|}(\xi)} = -\kappa \frac{K'_{|m|}(\kappa)}{K_{|m|}(\kappa)} - \frac{sm\kappa \sqrt{\xi^2 + \kappa^2}}{\xi^2}, \quad (2.40)$$

where $\kappa = ka$ and $\xi = \beta a$. The quantity s is equal to $+1$ or -1 . In the laboratory frame, waves corresponding to $s = 1$ move against the rotation of the column and waves corresponding to $s = -1$ move with the rotation of the column. These waves are referred to respectively as ‘retrograde waves’ and ‘prograde waves’.

For the case $m = 1$, $\kappa = 1$, plots of the left-hand side (black) and right-hand side (red for $s = 1$, blue for $s = -1$) of Equation (2.40) are given in Figure 2.7 (a) as a function of ξ . For each κ , m , the dispersion relation is satisfied for a discrete spectrum of frequencies, $s = \omega_n(\kappa, m)$, $n = 0, 1, \dots$, determined by the intersection points where the two sides of (2.40) are equal.

Returning to Equation (2.38), it follows that the frequencies $\omega_n(\kappa, m)$ are given by

$$\omega_n(\kappa, m) = \Omega \left(-m + \frac{2s\kappa}{\sqrt{\xi_{m,n}^2 + \kappa^2}} \right), \quad (2.41)$$

where $\xi_{m,n} = \xi_{m,n}(\kappa)$ is determined implicitly by Equation (2.40). The function $\xi_{m,n}$ is continuous, with $\xi_n(\kappa) \rightarrow j_{m,n}$ as $\kappa \rightarrow 0$, where $j_{m,n}$ is the n th root of $J_m(\xi)$. The waves permitted by Equation (2.41) are referred to as ‘Kelvin waves’. For the cases $n = 0, 1, 2$, plots

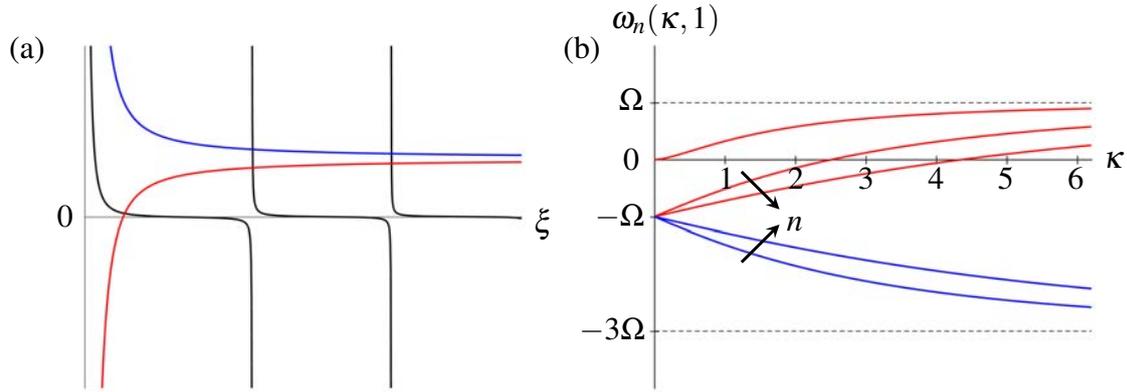


Fig. 2.7 For $m = 1$, using red and blue for retrograde and cograde waves respectively: (a) Taking $\kappa = 1$, plots of the left-hand side (black) and right-hand side of the dispersion relation (2.40) as a function of ξ . (b) For $n = 0, 1, 2$, plots of the frequencies $\omega_n(\kappa, 1)$ given by (2.41) as a function of κ . Arrows are used to indicate the direction of increasing n .

of $\omega_n(\kappa, 1)$ are given in Figure 2.7 (b) for $s = 1$ (red) and $s = -1$ (blue). Arrows indicate the direction of increasing n . Note only the retrograde waves admit a solution for $n = 0$.

The discussion up to this stage has allowed us to reach Equation (2.41), the general solution to the dispersion relation (2.40), obtained for the vortex considered by Thomson (1880). There are an infinite number of solutions, each of which have their own associated dynamics. For our purposes, it is in our interest to restrict attention to the waves that are most unstable linearly, as these will tend to dominate the development of the flow when it is subjected to small amplitude perturbations.

As m is restricted to taking integer values, it is natural to partition the waves into separate families, each corresponding to a particular value of m , allowing the kinematics of each family of waves to be described separately. Calculations demonstrating the behaviour of Kelvin waves with different azimuthal wavenumbers are given in Appendix A. The waves corresponding to $m = 0$ are ‘varicose modes’. For such waves, there is no azimuthal variation, and the vortex column broadens and narrows with $\cos(kz) \sin(st)$. The waves corresponding to $|m| = 1$ are ‘bending modes’ - these are the only types of Kelvin wave that lead to a deformation of the centreline of the vortex column. As such, one might expect these modes to be the most unstable to the influence of an externally imposed plane strain field. The waves corresponding to $|m| \geq 2$ are ‘fluted modes’. For a given slice $z = z_0$ of the vortex column, fluted modes exhibit m peaks and troughs on the column boundary.

A physical argument is given to demonstrate that our attention can be restricted to the stationary bending modes, corresponding to $|m| = 1$ and $\omega_n \rightarrow 0$. This is the family of waves from which the Crow and Widnall instabilities are generated. Suppose a weak, irrotational

plane strain field of rate $\varepsilon \ll \Omega$, normal to the z -axis, acts on the vortex column described by Equations (2.36) and (2.37). Consider its effect on a cross-section of the vortex column at a fixed z -plane. As $\varepsilon \ll \Omega$, deformations of the cross-sectional surface due to the strain field can be considered as negligible, so the cross-sectional surface remains approximately circular. Then, the main effect of the strain on the cross-sectional surface is to advect it within the z -plane, and its effect on the vortex column is to deform its centreline. As it is only the bending modes that correspond to a deformation of the centreline, it is reasonable to expect that the bending modes will be the most unstable modes to the imposed strain field.

Now, suppose the column experiences a small, purely helical wave disturbance, described by Equation (2.36) for $m = 1$, $s = \omega$. If $\omega = O(1) \gg \varepsilon$, the self-induced rotation of the wave will convect fluid away from the principal straining axis at a sufficient rate to stabilise the wave. Marginal stability of the wave occurs for $\omega = O(\varepsilon)$, and as $\omega \rightarrow 0$, the strain is able to destabilise the wave, stretching the vortex column along the principal straining axis. We can therefore expect that the stationary bending modes will be most unstable to the influence of the strain field.

Returning to Equation (2.41) and Figure 2.7 (b), we find that for all n , there exists a critical wavenumber $\kappa_{c,n} = k_{c,n}a$ such that $\omega_n(\kappa_{c,n}, 1) = 0$. Moreover, from the symmetries of Equation (2.38), we have

$$\omega_n(\kappa, m) = \omega_n(-\kappa, m) = -\omega_n(-\kappa, -m). \quad (2.42)$$

Therefore, for all n , the critical wavenumbers corresponding to the n th bending mode of $m = 1$ and $m = -1$ are equal. Moore and Saffman (1975) used these symmetry properties to demonstrate that, for a vortex column undergoing weak $O(\varepsilon)$ plane strain, a resonance will occur between the n th bending modes corresponding to $(m_1, m_2) = (-1, +1)$ for an $O(\varepsilon)$ band of axial wavenumbers, centred around the critical wavenumber $k_{c,n}$. This resonance corresponds to a superposition of two helical waves with equal amplitude and pitch and opposite chirality, forming a plane sinusoidal wave which is aligned with the stretching plane of the strain field. The analysis of Moore and Saffman (1975) can be generalised to demonstrate resonances between modes with $|m_1 - m_2| = 2$, though, consistently with our previous argument, the $(m_1, m_2) = (-1, +1)$ resonances appear to have the largest growth rates (Kerswell (2002)).

From Figure 2.7 (b), the first three dimensionless critical wavenumbers for $|m| = 1$ are

$$\kappa_{c,0} = 0 \quad , \quad \kappa_{c,1} = 2.504 \quad , \quad \kappa_{c,2} = 4.35. \quad (2.43)$$

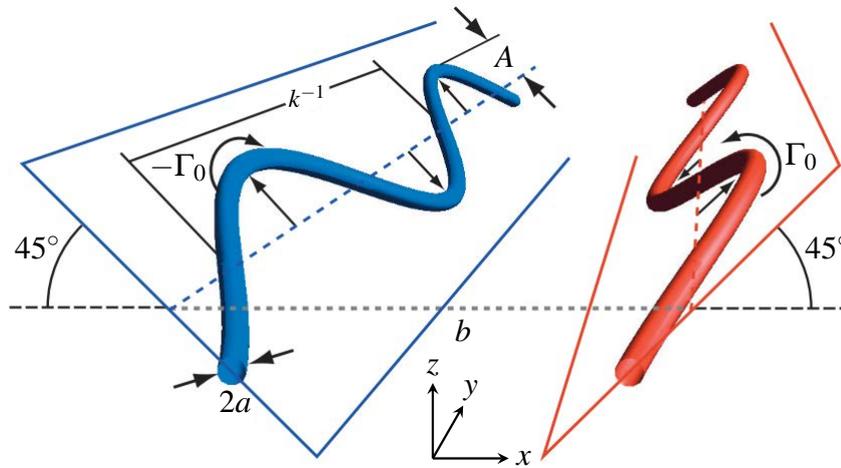


Fig. 2.8 Illustration of the Crow Instability. Figure modified from Leweke et al. (2016). Each vortex tube diverges along the stretching plane of the strain field induced by the other.

Stability analyses taking the long wavelength limit $\kappa \ll 1$ can only predict the instability associated with the zeroth stationary bending mode. The Crow instability for a pair of counter-rotating vortex tubes of equal strength Γ_0 is an example of this (Crow (1970)). An illustration of the Crow instability is given in Figure 2.8. For tubes of radius a separated a distance $b \gg a$ apart, the tubes are unstable to symmetric long wave disturbances due to the weak straining field induced by each tube acting on the other, causing them to diverge along the principal straining planes of the mutually induced strain fields.

In the case of a vortex ring, with core diameter b and ring diameter a (see Figure 2.6 (a)), an additional constraint on the axial wavenumber is that ka must be an integer, so the waves can fit onto the ring. Therefore, unless the ring is thin (with small aspect ratio $\delta = b/a \ll 1$), it is not possible for the zeroth bending mode to become unstable. A stability analysis applying the long wavelength limit to thin vortex rings was carried out by Widnall and Sullivan (1973). However, the instability they predicted occurs at a wavelength that is too short for the long wavelength limit to apply.

An extensive stability analysis for short waves on thin rings of constant vorticity was carried out by Widnall and Tsai (1977), up to $O(\delta^2)$. In their analysis, they demonstrate that the effects of ring curvature are much smaller than the effects of the ring-induced strain field on setting the growth rate and wavenumber of the instability that develops, validating the use of the Kelvin wave analysis. Their analysis is in good agreement with the previous experimental observations presented in Widnall and Sullivan (1973). The instabilities associated with the first and second bending modes are found to have similar

growth rates, suggesting that either mode could be excited on a vortex ring. An observation of a bimodal instability by Maxworthy (1977), where two distinct wavelengths appear simultaneously on the ring, supports the notion that either mode can dominate the instability. However, as the effect of viscous damping will be more significant for shorter wavelength modes, it is the first bending mode that is generally associated with the Widnall instability.

Up to this point, the discussion on stability has been restricted to inviscid vortex tubes of uniform vorticity, surrounded by irrotational fluid. Experimentally produced vortex rings however have peaked, continuous vorticity distributions. Widnall et al. (1974) investigated numerically the dispersion relation for waves on a straight vortex filament with a continuous vorticity profile, given by

$$2\Omega(r) = \begin{cases} (r^2 - a^2)^2 & r \leq a \\ 0 & r \geq a \end{cases} . \quad (2.44)$$

They found that there was a small shift in the critical wavenumbers (for example, $\kappa_{c,1} \approx 2.7$, compared with $\kappa_{c,1} \approx 2.5$ from Equation (2.43)), improving the level of agreement with their previous experimental observations (Widnall and Sullivan (1973)). The example given by Equation (2.44) illustrates that, for peaked vorticity distributions, there will be some deviation of the critical wavenumbers from the predictions made in the Kelvin wave analysis. This is further complicated by the ambiguity in defining a core radius when the vorticity distribution is continuous. Nonetheless, the Kelvin analysis still provides a helpful estimate of where the critical wavenumbers might lie.

A prediction for the number of waves that appears on a vortex ring produced from a tube was made by Saffman (1978). Using the predicted propagation velocity of thin vortex rings of uniform vorticity from Fraenkel (1972), an estimate for the number of waves N on the ring is given by

$$N = kR = \frac{kae^{1/4}}{8} e^{\tilde{U}}, \quad \text{where} \quad \tilde{U} = \ln \frac{8b}{a} - \frac{1}{4}. \quad (2.45)$$

In the experiments of Liess and Didden (1976) and Maxworthy (1977), however, it was observed that, for fixed \tilde{U} , the number of waves would vary by up to a factor of 3 with the piston Reynolds number, defined in Section 2.2.2 as $Re_p = U_p D_p / \nu$. Saffman showed that the predicted number of waves is dependent on the assumed ring vorticity distribution. By assuming the ring vorticity distribution was the same as that derived by Moore and Saffman (1973) for the trailing vortices produced by an elliptically loaded wing, Saffman (1978) obtained a prediction for the number of waves dependent on the piston Reynolds number Re_p and formation number F_p . His prediction was found to agree extremely well with the experimental observations of Liess and Didden (1976) and Maxworthy (1977).

2.3 Vortex-ring interactions

The subject of this thesis is the dynamics and mixing properties of vortex rings obliquely impacting the density interface in a two-layer stratification. In this section, we examine other examples of vortex ring interactions to gain physical insights that may aid our interpretation of the vortex ring interaction this thesis considers. Normal and oblique ring impacts with three different types of interface are considered: a flat, no-slip wall; a deformable free surface where surface tension and surface shear stress can be neglected; and a density interface between two miscible fluids of comparable density. Table 2.2 provides an overview of the ranges (where available) of Re_0 , θ_0 and Ri_0 , investigated in the studies referenced for these three types of interface interactions. Finally, in Section 2.3.7 we briefly examine vortex rings propagating through a weak, linear stratification.

2.3.1 Normal impact on to a no-slip wall

The dynamics of a vortex ring impacting a flat, no-slip wall at propagation angle $\theta_0 = 0^\circ$ (where θ_0 is the angle between the central axis of the ring and the vertical) has received considerable attention in the literature, with many theoretical, numerical and experimental studies (e.g. Walker et al. (1987), Lim (1991), Swearingen et al. (1995) and Cheng et al. (2010)). This ring-wall interaction provides a useful and relatively simple case study to illustrate more general features of evolving vortex structures in wall-bounded turbulent flows. In this discussion the vortex ring is propagating vertically downwards onto a flat, horizontal no-slip wall.

When the ring first feels the presence of the wall, we can treat the flow as inviscid before viscous effects become significant. At this stage, the ring evolves as though it was interacting with its mirror image, through the wall. The vortex ring is stretched by the interaction with the mirror ring, leading to an increase of the ring radius that is indicated by the dashed arrows in Figure 2.9 (a). To preserve the ring volume, this expansion is accompanied by a decrease in the core size, and hence an intensification of the vorticity to preserve the ring circulation.

Inviscid theory predicts that the vortex ring radius increases without bound as the ring moves progressively closer to the wall. This prediction quickly breaks down as viscous forces become significant. As the ring approaches the wall, the no-slip condition induces the production of vorticity within a viscous boundary layer on the wall, directly beneath the ring. This secondary vorticity, of opposite circulation to the ring, is advected outwardly from the central axis of the ring by the ring-induced motion and separates from the viscous boundary layer, after which it begins to surround the ring, as illustrated in Figure 2.9 (b).

Reference	Interface	Re_0	θ_0	Ri_0
Walker et al. (1987)	No-slip	105 – 3000	0°	N/A
Lim (1991)	No-slip	700 – 1070	0°	N/A
Swearingen et al. (1995)	No-slip	645	0°	N/A
Cheng et al. (2010)	No-slip	100 – 1000	0° – 40°	N/A
Ersoy and Walker (1987)	No-slip	Not specified	10° – 30°	N/A
Lim (1989)	No-slip	600	38.5°	N/A
Verzicco and Orlandi (1994)	No-slip	600	0°, 38.5°	N/A
Liu (2002)	No-slip	600, 1000	38.5°	N/A
Couch and Kreuger (2011)	No-slip	2000 – 4000	30° – 87°	N/A
New et al. (2016)	No-slip	2000, 4000	30° – 75°	N/A
Song et al. (1992)	Free surf.	9100 – 32200	0°	N/A
Tyvand and Miloh (1994)	Free surf.	High limit	0°	N/A
Quayan and Chu (1997)	Free surf.	High limit	0°	N/A
Archer et al. (2010)	Free surf.	3700 – 11100*	0°	N/A
Bernal and Kwon (1989)	Free surf.	2660*	0°	N/A
Lugt and Ohring (1994)	Free surf.	100*	45°	N/A
Ohring and Lugt (1996)	Free surf.	100 – 200*	20°, 45°	N/A
Gharib and Weigand (1996)	Free surf.	1150*	83°	N/A
Zhang et al. (1999)	Free surf.	470 – 1570*	10° – 30°	N/A
Dahm et al. (1989)	Dens. int.	2000 – 16000*	0°	4.3 – 275*
Stock et al. (2008)	Dens. int.	High limit	0°	0 – 0.1*
Advaith et al. (2017)	Dens. int.	1350 – 4600	0°	0.1 – 4
Yeo et al. (2020)	Dens. int.	1000 – 4000	0°	Not specified
Linden (1973)	Dens. int.	360 – 1080	0°	1.6 – 51
Olsthoorn and Dalziel (2017)	Dens. int.	1250 – 3500	0°	0.75 – 12.3
Stock (2006)	Dens. int.	High limit	45°	0.1 – 1.0*
Kuehn et al. (2010)	Dens. int.	470 – 1570*	35° – 85°	0.0006 – 3.15

Table 2.2 Table of studies referenced in Section 2.3 that investigated the dynamics of vortex rings impacting either no-slip surfaces, free surfaces ('Free surf.') or density interfaces ('Dens. int.'). Where applicable, the range of Re_0 , θ_0 and Ri_0 investigated for each reference are given. Starred (*) parameters under the Re_0 and Ri_0 columns correspond to alternative definitions of the Reynolds and Richardson numbers that incorporate the ring circulation Γ .

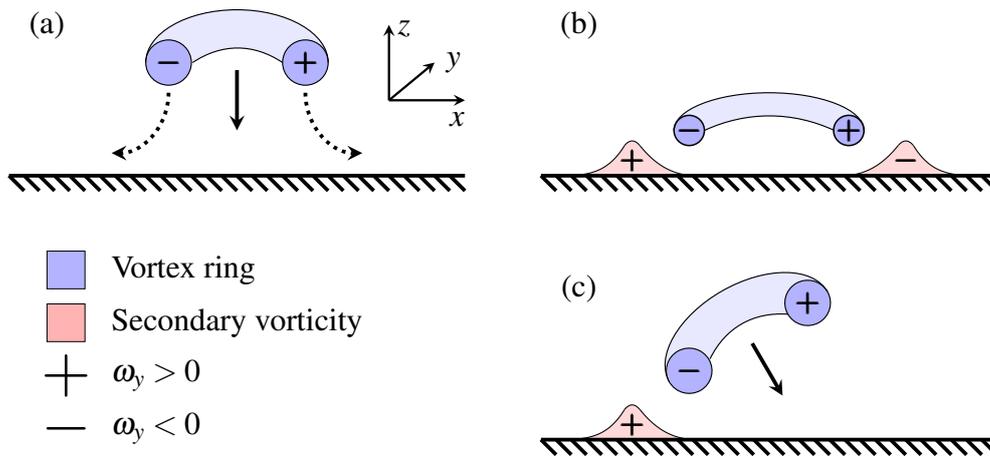


Fig. 2.9 Sketches of the early stages of a ring interacting with a no-slip wall, for normal and oblique impacts. (a) As the ring approaches the wall, its radius increases, as indicated by the dashed arrows. (b) Proximity of the ring with the wall leads to the generation of secondary vorticity at the wall. (c) For oblique impacts, secondary vorticity is generated to a greater extent by the side of the ring nearer the wall.

The surrounding secondary vorticity prevents the unrestricted expansion of the vortex ring predicted by inviscid theory, and causes the ring to rebound slightly from the wall.

The subsequent development of the primary ring and secondary vorticity has a significant dependence on the Reynolds number Re_0 (introduced in Equation (1.4)). Walker et al. (1987) conducted dye visualisation experiments to characterise this flow for $105 < Re_0 < 3000$. For the lowest Re_0 investigated, no significant boundary layer separation appears to occur and the primary ring is dissipated by viscous action of the wall. For $Re_0 \gtrsim 250$, an adverse radial pressure gradient along the wall cause boundary layer separation of the secondary vorticity, which can lead to the formation of a secondary vortex ring. In the absence of instabilities, the secondary ring will propagate around the primary ring before travelling through the primary ring's centre and re-merging with the wall-bounded secondary vorticity. For $Re_0 \gtrsim 470$, azimuthal instabilities develop in the flow. The structure of these instabilities, and their ability to trigger the turbulent breakdown of the flow, both depend on Re_0 . In all cases, the viscous action of the wall and the annihilation of vorticity between the primary and secondary ring structures will promote the dissipation of the flow.

The stability characteristics of the secondary ring differ fundamentally from that of an isolated vortex ring. Swearingen et al. (1995) demonstrated using the long wavelength approximation that thin vortex rings are unstable to long wavelength perturbations, which is not the case for isolated vortex rings. The strength of the secondary ring instability is found to be dependent on the relative positions of the two rings and the ratio of the primary and

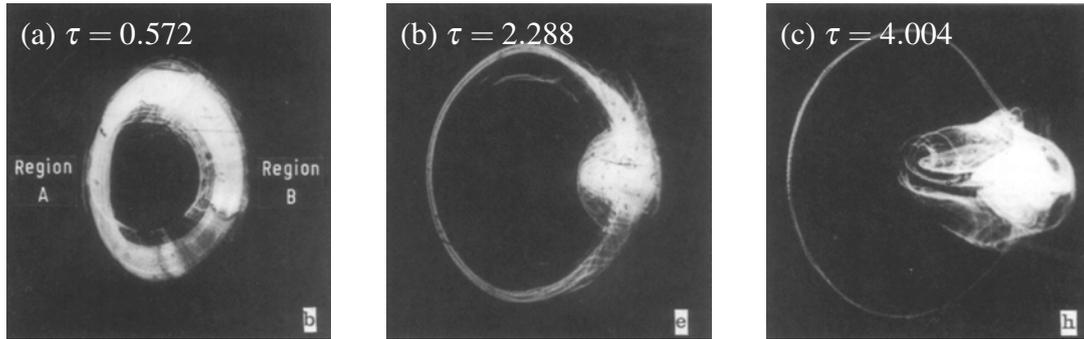


Fig. 2.10 Photographs of a dye visualisation of a vortex ring impacting a wall at an oblique angle, taken from a single experiment with $(Re_0, \theta_0) = (600, 38.5^\circ)$. The ring is moving towards the reader, with Region A of the ring impacting the wall first. The dimensionless time $\tau = t/t_0$ is normalised by an advective timescale $t_0 = D_0/U_0 = 1.716$ s, where D_0 is the diameter of the tube outlet and U_0 is the ring propagation speed. Note that $\tau = 0$ corresponds to Figure 4a of Lim (1989). Photographs selected from Figure 4 of Lim (1989).

secondary ring circulations. This is consistent with the instability of a vortex tube of weaker circulation in a counter-rotating vortex pair (Bristol et al. (2004)).

2.3.2 Oblique impact onto a no-slip wall

In addition to the Reynolds number, the dynamics of a ring obliquely impacting a flat, no-slip wall has a significant dependence on θ_0 , the initial propagation angle that the translational velocity of the ring makes with the vertical before it interacts with the wall. This interaction has received relatively little attention compared to its normal counterpart. Specific cases of the interaction in (Re_0, θ_0) space are presented by Ersoy and Walker (1987), Lim (1989) and Verzicco and Orlandi (1994). Recent studies have explored more extensively the range of dynamics for different cases of (Re_0, θ_0) , using numerical simulations (Liu (2002), Cheng et al. (2010)) and laboratory experiments (Couch and Kreuger (2011), New et al. (2016)).

Here we describe the case of a moderate propagation angle, similar to the case presented by Lim (1989) for $Re_0 \approx 600$, $\theta_0 = 38.5^\circ$. It is worth noting that significant dynamical changes in the behaviour of the interaction were demonstrated by Cheng et al. (2010) for $100 \leq Re_0 \leq 1000$, $0^\circ < \theta_0 < 40^\circ$, and New et al. (2016) for $Re_0 = 2000, 4000$ and $30^\circ < \theta_0 < 75^\circ$. As such, it should be noted that the case presented in this section is only representative of the interaction dynamics for a narrow range of (Re_0, θ_0) .

As the vortex ring approaches the wall, secondary vorticity is generated by the near side of the ring interacting with the wall due to the no-slip boundary condition, as illustrated in Figure 2.9 (c). This secondary vorticity interacts with the ring, leading to some vorticity

annihilation at the near side of the ring and a local increase in the ring core pressure. This results in an azimuthal pressure gradient in the ring core, leading to a transport of the core fluid from the near side to the far side of the vortex ring. The dye visualisations presented by Figures 2.10 (a) - (c) clearly illustrate this transport.

Figure 2.10 (b) indicates the presence of bi-helical vortex lines that wind around the circumferential axis of the ring. These are accounted for by the circumferential variations of the ring-core size and vorticity that develop due to the variation in stretching along the ring. Liu (2002) observes that the helical winding becomes more pronounced for larger values of Re_0 . As the flow continues to evolve, a secondary vortex loop develops and forms around the portion of the ring that was initially furthest from the wall, as shown in Figure 2.10 (c). The pressure-driven flow accompanied with viscous diffusion continues to weaken the primary ring vorticity until it eventually breaks where it first made contact with the wall.

2.3.3 Normal impact on to a deformable free surface

The evolution of a vortex ring impacting a deformable free surface at $\theta_0 = 0^\circ$ has been previously investigated (Song et al. (1992), Tyvand and Miloh (1994), Quayan and Chu (1997), Archer et al. (2010)). Most of the early studies of this interaction were motivated by a growing interest in the satellite detection of ship wakes. More broadly, this interaction provides a useful case study for better understanding of how vortical structures evolve when near a free surface.

Experimental observations of the free surface deformation resulting from the ring impact were made by Song et al. (1992) using shadowgraph imaging and a capacitance probe to quantify the surface deformation. The ring was visualised using hydrogen bubbles. While the ring remains laminar as it nears the free surface, axisymmetric waves are produced on the free surface and propagate radially outwards. Shortly after the ring impacts the free surface, the ring destabilises, breaking down and leading to the formation of several smaller U-shaped vortices which connect to the free surface and also propagate outwards. The instability of the primary ring is associated with the generation of short surface waves with complex three-dimensional structure, and is likened by Song et al. to the Crow instability.

Archer et al. (2010) used numerical simulations to investigate the instability of the primary ring structure. In these simulations, the ring exhibits a small amplitude Widnall instability prior to the ring-surface interaction, as can be expected for isolated rings with suitable Reynolds number. Associated with the initial instability, the radial structure of the ring core comprises inner and outer regions, which undergo small oscillatory displacements in opposite directions to one another, moving along the local principal straining axis. The interaction of the ring with the free surface leads to both the outer core vorticity being deposited onto the

free surface and the rotation of the perturbed inner core structure in the azimuthal plane of the ring. This realignment of the instability structure allows the Crow instability to further destabilise the ring by amplifying the radial perturbations that had initially developed due to the Widnall instability.

2.3.4 Oblique impact on to a deformable free surface

The oblique impact of a vortex ring on to a deformable free surface has been studied primarily to examine the mechanism of a vortex connecting to a free surface (Bernal and Kwon (1989), Lugt and Ohring (1994), Ohring and Lugt (1996), Gharib and Weigand (1996), Zhang et al. (1999)). The connection process is made possible by the free-slip condition, where the kinematic and zero shear-stress boundary conditions permit surface-normal vorticity at the surface. The U-shaped vortices observed by Song et al. (1992) is an example of this process. The connection process is not possible for the case of a no-slip wall, as only surface parallel-vorticity is permitted on the boundary, thus vortex lines are forced to remain closed in the fluid.

Early observations of the ring connection process were made by Bernal and Kwon (1989) using dye visualisation techniques, for rings initially travelling parallel to the free surface. In these experiments, the ring is seen to break at the portion of the ring nearest the free surface, and reconnect to the free surface at its open ends, forming a single U-shaped vortex. This process is elucidated by Gharib and Weigand (1996), who use shadowgraph and PIV to investigate the case where the vortex ring is initially travelling at a shallow attack angle towards the free surface (corresponding to $\theta_0 = 83^\circ$). By considering the pre-impact symmetry plane of the interaction, Gharib and Weigand (1996) demonstrate that production of surface normal vorticity cannot occur on the symmetry plane but must take place on either side. The amplification of secondary vorticity in these side regions coupled with the free-surface condition of zero shear stress (requiring the formation of a boundary layer) leads to the inevitable formation of a U-shaped vortex connecting to the free surface.

The ring-connection process was explained through the surface-boundary-layer structure by Zhang et al. (1999), who used numerical simulations to study the ring impact for propagation angles $\theta_0 \geq 60^\circ$ and small Froude numbers Fr_Γ , given

$$Fr_\Gamma^2 = \frac{\Gamma_0}{g^{1/2}(a/2)^{3/2}}, \quad (2.46)$$

where Γ_0 is the initial circulation of the ring and a is the ring diameter. In this study, a Cartesian coordinate system is used where the portion of the vortex ring nearest the free

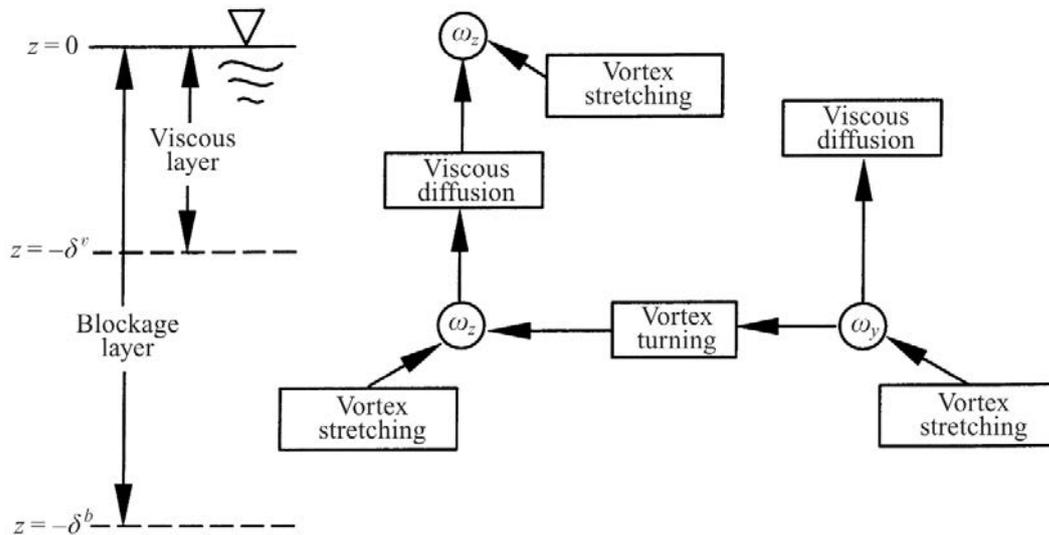


Fig. 2.11 Schematic illustrating the transformation of surface parallel and surface normal vorticity components during the vortex connection process. Duplicated from Zhang et al. (1999). Note that the terms ‘vortex stretching’ and ‘vortex turning’ used here relate respectively to the vorticity parallel and vorticity normal components of the vortex stretching term in the vorticity evolution equation given by Equation (2.17).

surface before impact has its vorticity aligned with the y -axis, and the z -axis is normal to the undisturbed free surface, pointing outwardly.

By considering the surface normal and parallel components of the vorticity equation that are responsible for vortex stretching, vortex turning and vortex diffusion, Zhang et al. (1999) demonstrate that the surface boundary layer consists of two regions. The outer ‘blockage layer, with thickness δ^b comparable to the characteristic length scale of the flow, is where surface normal motion is restricted due to the kinematic boundary conditions, leading to the reorientation of surface parallel vorticity towards the surface normal. The inner viscous layer, with thickness $\delta^v \sim \text{Re}^{-1/2}$, is due to the vanishing of shear stresses required by the dynamic boundary conditions. Here, remaining surface parallel vorticity is rapidly diffused and surface normal vorticity is strengthened, facilitating the vortex connection to the free surface. A schematic illustrating the transformation mechanism of surface parallel and surface normal vorticity is given in Figure 2.11.

The free-surface boundary layer has similar structure to the boundary layer first identified by Hunt and Graham (1978) for free-stream turbulence near a no-slip wall, where surface normal turbulent velocities decay in the outer ‘source’ region of the boundary layer, and turbulent fluctuations decay to zero in the inner viscous region. The similarity between the two boundary layer types is mainly attributable to the restriction of surface normal motion imposed by the kinematic condition near the boundary. Therefore, provided surface

deformations are small, the theory developed by Hunt and Graham (1978) is applicable to turbulent flows near free surfaces and to a lesser extent, fluid-fluid density interfaces.

2.3.5 Normal impact with a density interface

The normal impact of a vortex ring onto a density interface between two fluids has been investigated extensively, motivated by better understanding vortex interactions with stratified environments both in geophysical and industrial flows. Examples include the interaction of wakes produced by marine vessels with the seasonal thermocline, the interaction of convective thermals with atmospheric inversion layers, and the behaviour of turbulent flows at flame fronts (Dahm et al. (1989), Marcus and Bell (1992), Stock et al. (2008), Advaita et al. (2017), Yeo et al. (2020)).

Study of the vortex ring–interface interaction has also been motivated by better understanding the entrainment process and mixing at turbulent density interfaces, using the vortex ring as a model turbulent eddy to isolate the eddy entrainment mechanism (Maxworthy (1972), Linden (1973), Olsthoorn and Dalziel (2015, 2017, 2018)). The ring–eddy analogy also provides the motivation for our study of oblique vortex ring impacts at a sharp density interface. As the ring–eddy analogy was previously discussed in Section 1.2, here we restrict our attention to the dynamics of the induced flow.

The first investigation of the ring–interface interaction was made by Linden (1973), who conducted laboratory experiments using shadowgraph to visualise the flow. Vortex rings were generated in the upper layer and made to propagate vertically downward towards the interface. The focus of this study was to develop a theoretical model for the turbulent entrainment process induced by the vortex ring impinging on the sharp interface, which is discussed in greater detail in Section 2.4.2). Linden demonstrates that the maximum extent to which the ring penetrates the interface is proportional to the square of a Froude number Fr , given by

$$Fr = \frac{U}{(ag\Delta\rho/\rho_1)^{1/2}} \sim Ri_0^{-1/2}, \quad (2.47)$$

where a and U are the diameter and propagation speed of the vortex ring respectively, $\Delta\rho$ is the density step at the interface, ρ_1 is the upper layer density and g is the acceleration due to gravity. This is equivalent to being proportional to the ratio between the bulk kinetic energy of the vortex ring and the potential energy strength associated with the density jump at the interface.

The interaction was reexamined later by Dahm et al. (1989), who used a combination of dye and LIF experiments and numerical simulations to investigate the flow both for ‘thin’

interfaces (where the interfacial thickness δz is small relative to the ring diameter a), and ‘thick’ density interfaces (where $\delta z \sim a$). By considering the vorticity evolution equation on a thin interface, Dahm demonstrates that in the Boussinesq limit, the production of baroclinic vorticity at the interface is governed by the product of the Atwood number, A , and the ratio R of hydrostatic to hydrodynamic pressure gradients at the interface. These are given by

$$A = \frac{\Delta\rho}{\rho_2 + \rho_1}, \quad R = \frac{a^3 g}{\Gamma^2}, \quad (2.48)$$

where ρ_2 is the lower layer density of the ring and Γ is the initial circulation of the ring. The product AR relates to the Froude number introduced in Equation (2.47) by $AR \sim Fr^{-2} \sim Ri_0$, since the ring circulation scales as $\Gamma \sim aU$.

For $AR \gtrsim 1$, Dahm et al. (1989) observes that the ring is unable to penetrate the interface. Baroclinic vorticity generated at the interface pinches off to form a secondary ring around the outside of the primary ring and just above it. The primary ring inducts the secondary ring through its centre, leading to the ring diameter increasing for the primary ring and decreasing for the secondary ring. This observation was also made by Marcus and Bell (1992), who point out the similarity between the no penetration case and the case of a ring normally impacting a no-slip wall. Dahm et al. (1989) observes the development of a wavy instability on the secondary ring due to the compressional and extensional strains exerted by the primary ring, a feature also present in the no-slip wall case.

For small values of AR , the ring is able to penetrate the interface deeply. The vortex ring transports a substantial volume of upper layer fluid into the lower layer region, both via the ring atmosphere and its wake. Horizontal buoyancy gradients at the interface between the transported upper layer fluid and the surrounding lower layer fluid lead to compression of the vortex ring and a rapid production of baroclinic vorticity. This secondary vorticity erodes the primary ring structure, eventually leading to the turbulent breakdown of the ring and a considerable amount of mixing.

Recent advances in technological capabilities have led to the ring–interface interaction being reexamined. In particular, Olsthoorn and Dalziel (2017) used a two-dimensional stereo-PIV technique to reconstruct three-dimensional velocity fields of the ring–interface interaction. This reconstruction was achieved by repeating each experiment multiple times, taking measurements at a different horizontal plane for each experiment. Three cases were studied: $(Ri_D, Re_D) = (2.4, 1600)$, $(1.7, 2400)$ and $(0.98, 1600)$, corresponding to non-penetrative, partially penetrative and deeply penetrative ring impacts. Note that Olsthoorn and Dalziel (2017) use the vortex-ring-tube diameter D rather than the ring diameter a in defining the Reynolds and Richardson numbers, so $Ri_D = DRi_0/a$ and $Re_D = DRe_0/a$.

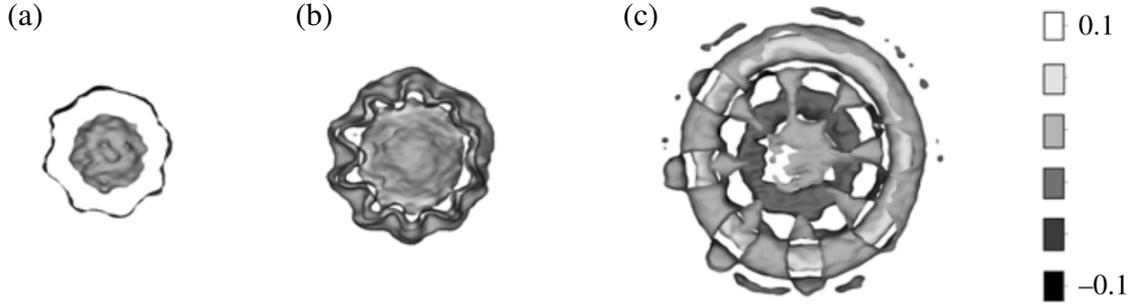


Fig. 2.12 Horizontal slices plotting the vertical velocity of the ring–interface interaction, demonstrating the modal structure of the stratification instability. Cases (a), (b) and (c) correspond to $(\text{Ri}_D, \text{Re}_D) = (2.4, 1600)$, $(1.7, 2400)$ and $(0.98, 1600)$ respectively, and have respective mode numbers 8, 10 and 8. Figure duplicated from Olsthoorn and Dalziel (2017).

These experiments demonstrated that the mode number of the ‘Crow-like’ instability that develops during the interaction is controlled by the Reynolds number, an observation consistent with the prediction of Saffman (1978) for the unstratified case. Figure 2.12 demonstrates this Reynolds number dependence. For all three cases, the instability structure was primarily expressed through the vorticity that had been generated baroclinically. By considering the angular distribution of vertical velocity at different heights, the timescale τ of the instability growth is shown to have an inverse dependence on the Richardson number Ri_0 . The scaling is given by

$$\frac{U\tau}{a} \sim \frac{1}{\text{Ri}_0}, \quad \text{where} \quad \text{Ri}_0 = \frac{g\Delta\rho}{\rho_1} \frac{a}{U^2}. \quad (2.49)$$

For $\text{Ri}_0 \gtrsim 1$, the scaling predicts that the ring instability is able to grow significantly before the interface rebounds, whereas for $\text{Ri}_0 \lesssim 1$, the buoyancy response will inhibit the development of the ring instability. Selected photographs from one of our preliminary dye experiments for the parameters $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2070, 4.6, 0^\circ)$ are presented in Figure 2.13 to showcase the distortion of the density interface in the regime where the Crow-like instability is able to develop.

2.3.6 Oblique impact with a density interface

The case of a vortex ring obliquely impacting a density interface is the subject of this thesis. To the best of our knowledge, there have only been two previously published studies that have considered this interaction, by Stock (2006) and Kuehn et al. (2010).

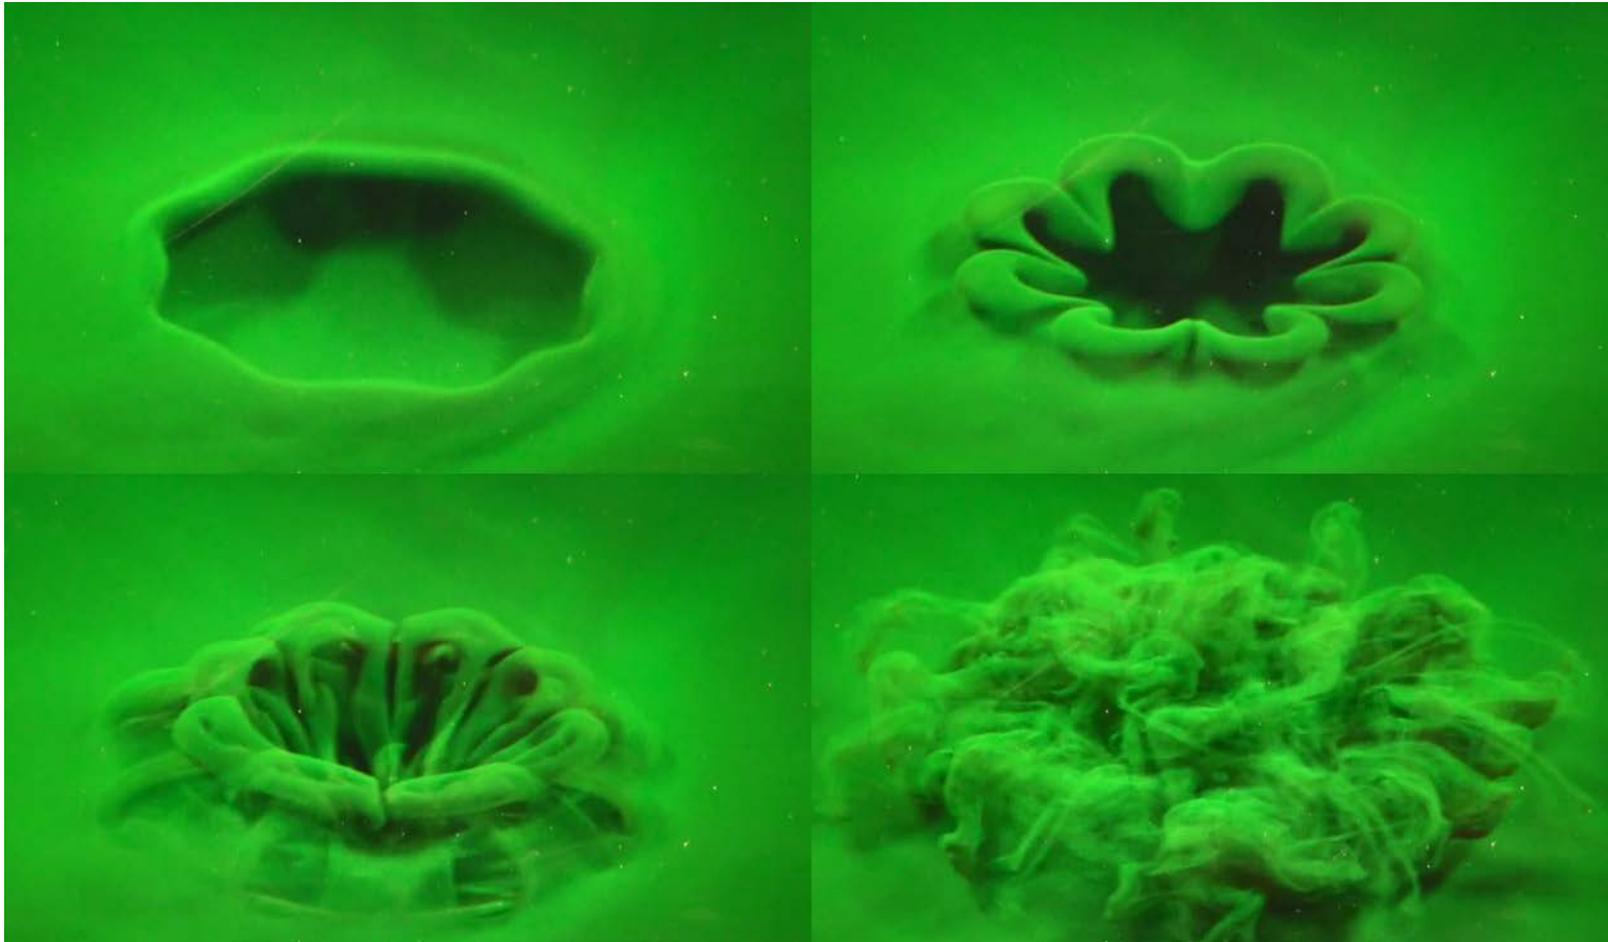


Fig. 2.13 Selected photographs of the vortex ring–interface interaction for $(Re_0, Ri_0, \theta_0) = (2070, 4.6, 0^\circ)$, taken from a dye visualisation experiment where the lower layer is dyed green to visualise the distortion of the density interface. Here, the ring diameter $a = 4.9$ cm and the ring propagation speed $U = 42.2$ cm s $^{-1}$. Small perturbations to the flow in the upper layer were found to heavily compromise the symmetry of the instability. The time interval over which the four photos were taken is approximately 2.5 seconds.

Numerical simulations of the ring impact for $\theta_0 = 45^\circ$ were presented in the PhD dissertation of Stock (2006) for $AR = 0.1, 0.3$ and 1 , where A and R are defined by Equation (2.48). The interaction was considered as one of a number of case studies to test the regularised inviscid vortex sheet method Stock developed, the numerical method being the primary focus of Stock's thesis. The analysis given for these simulations are brief, but interesting features are observed.

For $AR = 0.3$, as the ring approaches the interface, vorticity is generated baroclinically first by the portion of the ring closer to the interface. The strength of this secondary vorticity relative to the secondary vorticity produced along the rest of the deformed interface is responsible for the deformation of the ring, causing it to bend. The simulation was not run much further than this point as it became difficult to discern the actual ring dynamics from numerical instabilities.

For $AR = 1$, the strength of the stratification is much greater, with the ring only weakly distorting the density interface. A secondary vortex ring forms from the baroclinically produced vorticity and is swept around the outside of the primary ring. This leads to a three-dimensional instability exhibiting bi-helical vortex lines, similar to those observed by Lim (1989) for the case of an oblique ring impact at a no-slip wall. Stock attributes a hastening in the turbulent breakdown of the ring to this instability.

Dye experiments were conducted by Kuehn et al. (2010) to visualise the trajectory of the vortex ring, so its change in propagation angle as it interacted with the stratification could be measured. A discussion was given as to whether a generalised version of Snell's law could be made to predict the evolution of the vortex ring trajectory. It was concluded that such a law would need to account for the structural evolution of the vortex ring as the flow evolves. No further attempt was made to describe the internal dynamics of the vortex ring and the experimental technique used provides insufficient resolution to discern any particular features of the ring's internal structure.

2.3.7 Vortex rings propagating through a weak, linear stratification

In this section, we briefly cover vortex rings propagating through a weak, linear stratification, to support some of the assumptions made in Section 6.3.2 for our periodic-mixing experiments. Previous studies of this interaction include Maxworthy (1977), van Atta (1989), Orlandi et al. (1998), and Scase and Dalziel (2006). Here, we restrict our attention to the observations of Maxworthy (1977), for vortex rings propagating vertically downwards through the stratification.

Figure 2.14 (a) illustrates the experimental setup used by Maxworthy (1977) for observing vortex rings propagating through a linear stratification. Vortex rings first propagate through

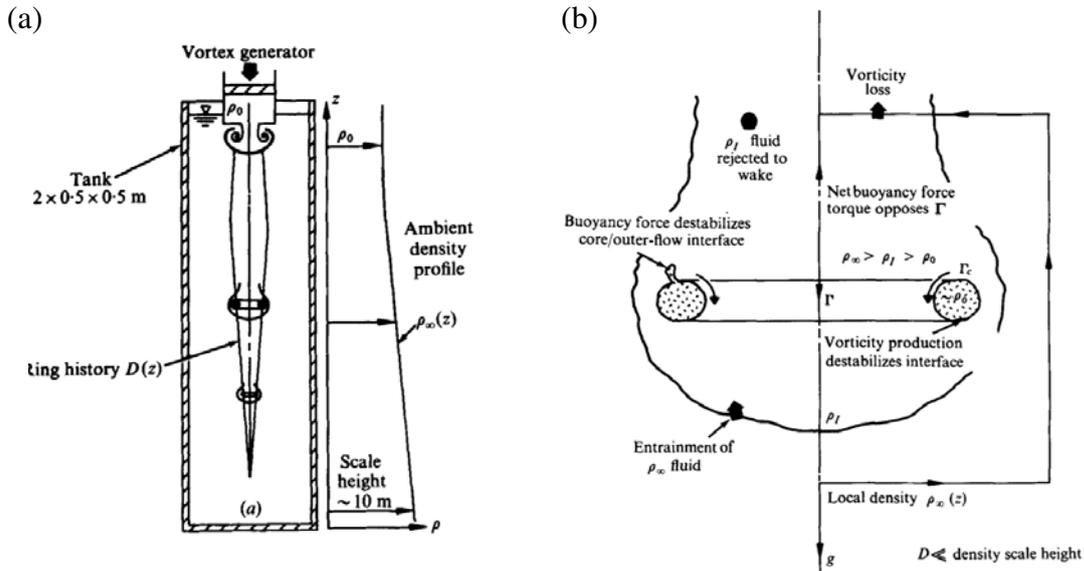


Fig. 2.14 (a) Illustration of the experimental setup used by Maxworthy (1977) to observe vortex rings propagating through a linear stratification. A plot of the ambient density profile $\rho_\infty(z)$ and a sketch of the ring diameter $D(z)$ as it propagates through the stratification are both provided. (b) Sketch of the vortex ring after it has propagated well into the linearly-stratified region. Three distinct regions emerge: the core of the ring, a mixed intermediate region, and the ambient. Vorticity is produced baroclinically at the interfaces between adjacent regions. Figures reproduced from Maxworthy (1977).

an unstratified region with density ρ_0 and depth 50 cm, enabling them to fully form before then entering a region of linearly stratified fluid, of depth 150 cm and density increasing from 1.00 g/cm^3 to 1.15 g/cm^3 . In Figure 2.14 (a), Maxworthy sketches the evolution of the ring diameter as it propagates through the tank, first expanding in the unstratified region (consistent with the model of Maxworthy (1972) described in Section 2.2.1), before then diminishing in size as it propagates through the linearly stratified region.

Figure (b) provides a detailed sketch of the vortex ring after it has propagated well into the linearly-stratified region, illustrating the dynamical behaviour of the flow observed by Maxworthy. When the ring first enters the linearly-stratified region, all fluid moving with the ring has density ρ_0 . As the ring continues to propagate, the portion of the ring external to the core is first to mix with the surrounding ambient, creating mixed fluid with an intermediate density $\rho_I(x, t)$. It is this intermediate fluid that is rejected into the wake of the ring. Due to there being very little entrainment across the interface between the core of the ring and the intermediate fluid, the ring retains a self-similar form, and the core of the ring retains its initial density to a good approximation. Over time, baroclinic vorticity production at the interface between the core of the ring and the surrounding fluid leads to this interface

destabilising, enabling some of the core fluid to be entrained into the intermediate fluid and exposing more of the ring core. This process leads to a gradual reduction in the size and circulation of the ring until its ultimate collapse.

2.4 Aspects of stratified mixing

The last section of this chapter covers aspects of stratified mixing that are relevant to our analysis of a periodic sequence of vortex rings mixing an initially two-layered stratification. In particular, we restrict our attention to mixing and how mixing efficiency is quantified, and to revisiting the ring–eddy analogy following its introduction in Section 1.2.

2.4.1 The energetics of mixing

Scalar mixing consists of two physical processes that occur simultaneously: the stirring of scalar iso-surfaces, and the enhancement of diffusion as a result of the stirring-induced intensification of the scalar gradients and the associated temporary increase in interfacial area over which diffusion can occur. The ring–interface interactions considered in this thesis can be thought of as events that result in mixing of the density field. For salt solutions, the fluid density takes the form

$$\rho = \rho(S, T, P), \quad (2.50)$$

where S_k is the concentration of solute s_k , T is the temperature and P is the pressure. Throughout our experiments, as the fluid depth was kept at less than 60 cm, variations in pressure had a negligible effect on the fluid density. Also, as the fluid was kept to being close to thermal equilibrium with the laboratory environment, variations in fluid temperature were sufficiently small that ring-induced mixing of the temperature field had a negligible effect on the density field when compared with the mixing of the solute-concentration fields. We therefore restrict our definition of mixing in this thesis to mixing of the solute-concentration fields. This can effectively be thought of as mixing of the density field, assuming a linear equation of state for the fluid density. This assumption is valid for our purposes given the solutes and range of solute concentrations we use to stratify the fluid.

Mixing can be quantified as the irreversible increase in gravitational potential energy (PE) that results from a given energy input that stirs the fluid (Winters et al. (1995)). To distinguish between reversible and irreversible changes in PE, Winters et al. (1995) partitioned PE into background and available potential energy (BPE and APE). The BPE of a system is the minimum potential energy realisable following an adiabatic rearrangement of the density field $\rho(\mathbf{x}, t)$ to a sorted density field $\rho_s(z, t)$, and the APE is the excess potential energy in

addition to the BPE that makes up the total PE. For a closed system, the BPE and APE can be written as

$$\text{BPE} = \int g\rho_s z dV, \quad \text{APE} = \text{PE} - \text{BPE} = \int g(\rho - \rho_s)z dV. \quad (2.51)$$

The total amount of energy available in the system for mixing (AE) is the sum of kinetic energy (KE) and APE, which can be written as

$$\text{AE} = \text{KE} + \text{APE} = \int g(\rho - \rho_s)z + \frac{1}{2}\rho|\mathbf{u}|^2 dV. \quad (2.52)$$

A system can only be in a steady state of rest when $\text{AE} = 0$, in which case $\text{KE} = 0$, $\text{PE} = \text{BPE}$ and the BPE increases only through heat diffusion and molecular diffusion of any solutes present in the fluid.

An irreversible mixing rate M was introduced by Peltier and Caulfield (2003) to distinguish between increases in BPE associated with an initial kinetic energy reservoir and increases associated with molecular diffusion of the mean density profile. This leads to the energetic pathways between KE, APE and BPE to be written as

$$\begin{aligned} \frac{d}{dt}\text{KE} &= H - \varepsilon, \\ \frac{d}{dt}\text{APE} &= -H - M, \\ \frac{d}{dt}\text{BPE} &= M + D_\rho. \end{aligned} \quad (2.53)$$

Here, H is the buoyancy flux, allowing for the direct transfer from KE to APE and vice-versa; $\varepsilon > 0$ is the kinetic-energy dissipation rate; and D_ρ is the rate at which the potential energy would increase in a stable stratification at rest, encapsulating both heat and molecular diffusive processes. Given an initial source of KE that leads to a temporary increase in APE, Peltier and Caulfield (2003) associate irreversible mixing with an imperfect return of APE back into the KE reservoir, with the fraction of energy not transferring back to KE instead transferring to the BPE reservoir via diffusive processes.

An important property of turbulent mixing processes that is useful to characterise is the ‘mixing efficiency’, that is the increase in BPE of a system as a proportion of the energy expended in carrying out the mixing process. A number of measures have been introduced to estimate this quantity, the most common of which are listed in Table 1 of Gregg (2018). From the energetic pathways given by Equations (2.53), Peltier and Caulfield (2003) introduce

instantaneous and cumulative measures of the mixing efficiency, η_I and η_C , given by

$$\eta_I = \frac{M}{M + \varepsilon} \quad \text{and} \quad \eta_C = \frac{\int_0^T M(\tau) d\tau}{\int_0^T M(\tau) d\tau + \int_0^T \varepsilon(\tau) d\tau}. \quad (2.54)$$

For η_C , time T is selected such that the mixing and dissipation associated with the mixing process occurs over the time interval $[0, T]$. For our periodic mixing experiments presented in Chapter 6, we introduce and make use of a modified form of η_C to estimate the mixing efficiency of the ring–interface interaction, taking into account the increases in BPE associated with heat flux through the tank walls.

As discussed by Tailleux (2009), while the APE framework introduced by Winters et al. (1995) and developed further by Peltier and Caulfield (2003) is valid for Boussinesq flows with a linear equation of state, their results do not generalise to fluids with non-linear equations of state, or compressible flows. For our experiments, we are able to assume the flow is Boussinesq and the equation of state is linear, allowing us to adopt the APE framework.

2.4.2 Turner’s mixing box and the ring–eddy analogy

This last section complements our discussion of Turner (1968), Maxworthy (1972), Linden (1973) and Olsthoorn and Dalziel (2015, 2017, 2018) in Section 1.2 on the history of the ring–eddy analogy by discussing in more detail the findings of Turner (1968) and Linden (1973). The relevance of eddy-induced mixing to turbulent mixing in zero-mean shear flows was first demonstrated by Turner (1968). Turner performed experiments in a two-layer density stratification, using a vertically oscillating grid, in either one or both of the fluid layers, to mix the stratification. Experiments were conducted for $\text{Ri} \gtrsim 0.1$, where the Richardson number is defined by

$$\text{Ri} = g \frac{\Delta\rho}{\rho} \frac{1}{ln^2}, \quad (2.55)$$

using the oscillation frequency n of the grid and a fixed lengthscale l imposed by the geometry of the grid.

The purpose of Turner’s experiments was to demonstrate the effects of molecular diffusivity on turbulent entrainment at a density interface. This was achieved by using salt and heat in separate experiments as the single stratifying agent. For the case where only one layer is stirred, Figure 2.15 (a) shows the measured rate of mixing as a function of the Richardson number. For the entrainment velocity u_e , Turner obtains the following functional form for $\text{Ri} \gtrsim 1$,

$$\frac{u_e}{ln} \propto \text{Ri}^{-m}, \quad (2.56)$$

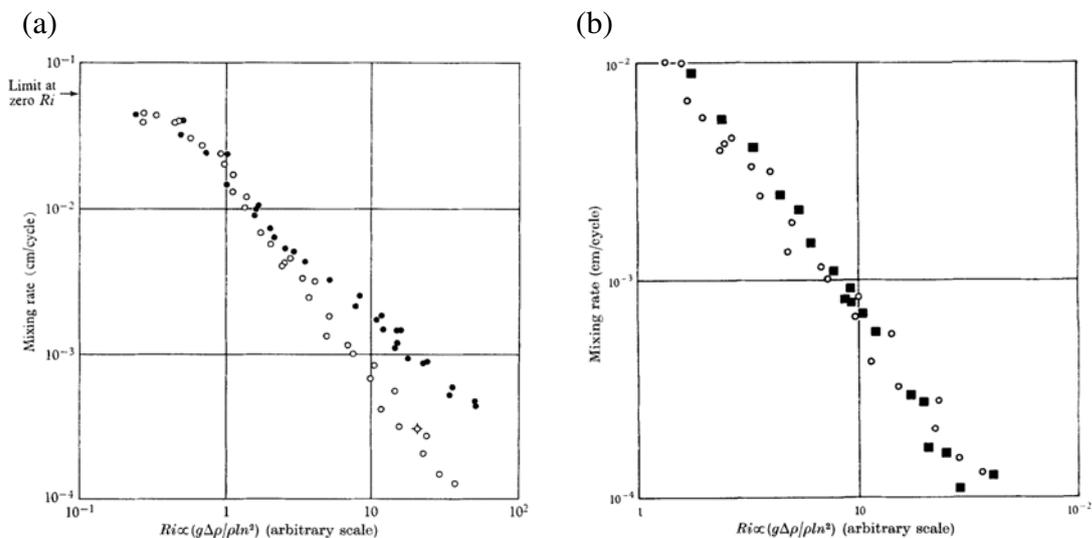


Fig. 2.15 Plots of the measured mixing rate as a function of Richardson number. (a) Experimental results when a single layer stirred. Filled black circles and empty black circles correspond to experiments where heat and salt are the stratifying agents respectively. (b) Experimental results when salt is the stratifying agent. White empty dots and black filled squares correspond to experiments where one layer is stirred and both layers are stirred respectively. Graphs reproduced from Turner (1968).

where $m = 1$ and $m = 3/2$ for temperature and salinity stratifications respectively.

The enhanced entrainment velocity associated with using heat instead of salt as the stratifying agent was attributed to the thermal diffusivity of water being significantly larger (by a factor of 700) than the molecular diffusivity of salt, effectively resulting in a higher mass diffusivity for the heat experiments. At higher mass diffusivities, a greater amount of diffusion can occur when fluid elements are displaced from their neutral-buoyancy level, thus leading to a larger proportion of the work done to initially displace the fluid being converted into background potential energy. The difference in functional forms for u_e leads Turner to conclude that the structure of the density interface depends on the interplay between mechanical stirring, which sharpens the interface from the side that is being stirred, and molecular diffusion, which tends to spread it out.

For each stratifying agent, Turner also compares the mixing rates for experiments when either one layer is stirred or both layers are stirred. Figure 2.15 (b) presents the results for when salt is the stratifying agent. For both salinity and temperature stratifications, no significant change is observed in the mixing rate whether one or both layers are being stirred. Turner attributed this indifference to the nature of the mixing at the density interface. In particular, it was observed that the largest eddies that interact with the interface appear to be responsible for the majority of mixing, through a process likened to the breaking of steep,

forced internal waves. Turner argued that these mixing events were sufficiently uncommon that both layers could be regarded as statistically independent from each other, as the mixing sites cover only a small fraction of the interfacial area at any one time. This observation motivated the study of eddy-induced mixing at a sharp density interface to better understand turbulent mixing in zero-mean shear flows.

The suggestion that vortex rings could be used as idealised eddies to understand the entrainment mechanism better at a turbulent density interface was made by Maxworthy (1972). The first study utilising this analogy was presented shortly after by Linden (1973), who conducted experiments where vortex rings vertically impacted a sharp density interface. Experiments were carried out for a range of Reynolds numbers Re and Froude numbers Fr , given by $360 \leq Re \leq 1080$, $0.14 \leq Fr \leq 0.80$. The Froude number here was defined previously in Equation (2.47) and is related to the Richardson number defined in Equation (2.49) by $Ri_0 = Fr^{-2}$, which lies in the range $1.5 \lesssim Ri_0 \lesssim 51$. Results were obtained using a shadowgraph technique.

Linden proposes a model to predict the entrainment velocity at a turbulent density interface by approximating the flow as a discrete number of energy-containing eddies, where each eddy interacts with the interface in a similar way to the vortex ring. By equating the rate per unit area at which kinetic energy is made available for entrainment at the mixing sites to the rate of increase in potential energy of the system, Linden obtains an estimate for the entrainment velocity,

$$\frac{\bar{\rho} D^3 U^2}{\tau l^2} \sim g \Delta \rho d u_e \implies u_e \sim \frac{1}{\tau} \frac{\bar{\rho}}{g \Delta \rho} \frac{D^3 U^2}{l^2 d} . \quad (2.57)$$

Here, τ is defined as the timescale at which energy is made available for mixing, $\bar{\rho}$ is the mean density, d is the height at which fluid is raised by the recoil, and l is the lengthscale of the interfacial distortion created by a single eddy interaction. The other parameters have been defined previously by Equation (2.47).

Three assumptions are made to relate the entrainment velocity to the Froude number. Firstly, the height d of the interfacial recoil is assumed to scale as $d \sim D$. This was justified by observing that the lengthscale over which mixing occurred in the vortex ring experiments appeared to scale with the size of the ring, and that interfaces were observed to remain sharp in the experiments of Turner (1968). Secondly, for the ring–interface interaction, it is assumed that an energy balance exists between the initial bulk kinetic energy of the vortex ring and the potential energy increase associated with the distortion of the interface. This assumption yields the scaling

$$\frac{l}{D} \sim Fr^{-1}, \quad (2.58)$$

which is shown to agree with experimental results for $Fr \lesssim 0.3$, but is seen to break down for larger Froude numbers. Note also that the energy balance assumed here neglects contributions to the kinetic energy from the rotational component of the ring motion. Thirdly, the timescale τ relating to the rate at which energy is made available for mixing is assumed to scale with the buoyancy response timescale of the perturbed interface, giving

$$\tau \sim \left(\frac{\bar{\rho} l}{g \Delta \rho} \right)^{1/2}. \quad (2.59)$$

Making these three assumptions, Linden obtains the following non-dimensional form of the entrainment velocity,

$$\frac{u_e}{U} \sim Fr^3 = Ri^{-3/2}. \quad (2.60)$$

This scaling agrees with the experimental results of Turner (1968) for salinity stratifications for the range of Froude numbers considered, which is expected as the vortex-ring interactions entrain fluid in an analogous way to the dominant entrainment mechanism observed in Turner's grid-mixing experiments.

Chapter 3

Experimental Apparatus, Setup and Methodology

In this chapter we discuss our methodology for carrying out the experiments conducted for this thesis. In Section 3.1 we describe our procedure for creating a two-layer stratification in the two experimental tanks we used. This includes a discussion on the design of a new piece of apparatus that was essential for being able to conduct experiments in the larger of the tanks we used. In Section 3.2, we talk through our setup and procedure for generating reproducible vortex rings. Then, in Section 3.3 we describe the procedure for setting up our PIV/LIF experiments, after which we describe our methodology for making PIV and LIF measurements. Finally, in Section 3.4 we introduce the conductivity probe and thermistor used to take vertical profiles for our periodic mixing experiments, as well as describing our procedure for calibrating these two apparatus. The experimental setup and procedure for our periodic mixing experiments is discussed in Chapter 6.

3.1 Creating a two-layer stratification

In this thesis, we investigate the influence that the propagation angle of a vortex ring has on its interaction with a two-layer density stratification. Experiments were conducted in two tanks: a smaller tank of size $500 \times 200 \times 500 \text{ mm}^3$ and a larger tank of size $1200 \times 400 \times 750 \text{ mm}^3$ (dimensions in Length x Width x Height). This section describes how we create a two-layer stratification in each of these tanks.

Note that preliminary experiments (including those from which the photographs presented in Figure 2.13 were taken) were conducted in the ‘Rayleigh-Taylor’ tank, developed by Dalziel (1993). The filling methodology we employed for the Rayleigh-Taylor tank is similar

to that described by Davies Wykes and Dalziel (2014), as well as the dye and lighting arrangement used to produce Figure 5 of Davies Wykes and Dalziel (2014). Importantly, the separating barrier between the two layers was removed slowly to minimise the amount of mixing induced by the barrier removal. This ensured a sharp density (and dye concentration) jump across the interface, enabling a high clarity in the distortion of the interface made by impinging vortex rings (as showcased in Figure 2.13).

3.1.1 Filling the tank

We start by prescribing the height of the two fluid layers. This determines the fluid volume required for each layer. We then prescribe the density of the two fluid layers. For each layer, the desired density is achieved by mixing the appropriate mass of salt crystals per litre to the required volume of fresh water. The fresh water is drawn from a break tank feeding the lab that is, in turn, supplied by the water mains. The density of the solution is then verified using a density meter (Anton Paar DMA 5000), which measures the fluid density at 20°C to a precision of 10^{-6} g/cm³. In the experiments conducted for this thesis, the upper-layer fluid is a NaCl (sodium chloride) solution. When the refractive indices of the two fluid layers must be matched, the lower layer is a NaNO₃ (sodium nitrate) solution (this is discussed in Section 3.3.1. Otherwise, the lower layer is a NaCl solution.

The upper-layer fluid is prepared in the tank being used for the experiment. The lower-layer fluid is prepared in an external reservoir. The two solutions are left overnight to allow any excess of dissolved gases to come out of solution from each of the fluid layers, after which any air bubbles that had nucleated on the surface of the tank were removed using a squeegee. Leaving the solutions overnight was necessary for our PIV/LIF experiments to reduce the risk of bubbles blocking the cameras' view of the flow, and was necessary for our periodic mixing experiments to significantly reduce the risk of bubbles forming at the tip of the conductivity probe, which could compromise the measurements made by the probe. Leaving the solutions overnight was also helpful in allowing them to converge towards thermal equilibrium with the laboratory environment, which we discuss further in Section 3.1.4. After the two solutions have been left overnight, the tank filling process can be completed.

Aside from the apparatus used to fill each tank, the tank-filling methodology is the same for the two tanks that we use. Using a pump system with connected tubing, the lower-layer fluid is injected into the bottom of the tank beneath the upper layer. The injection flow rate is manually controlled using a digital interface on the pump system, taking care to ensure that negligible mixing occurs between the two solutions in the tank. Initially, the injection flow rate is kept low to prevent mixing between the two fluids. Once the density interface is

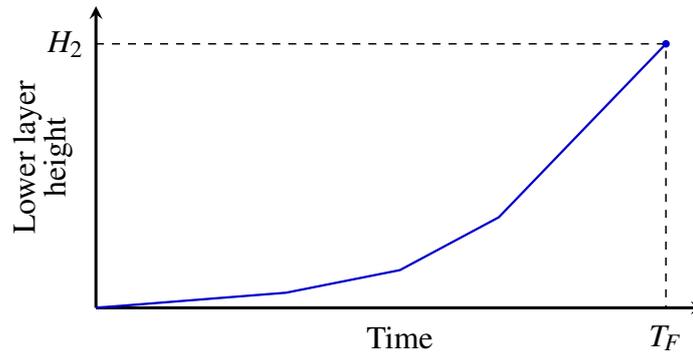


Fig. 3.1 Sketch of the lower-layer height as a function of time during the lower-layer filling process.

sufficiently far above the injection site, the flow rate is gradually increased up to a maximum rate, after which the flow rate is kept constant. A sketch of the lower-layer height as a function of time during the lower-layer filling process is given by Figure 3.1. Once the lower layer has reached the prescribed height, the fluid injection is stopped and the tank filling process is completed.

3.1.2 Lower-layer filling apparatus

Different digital pump systems were used for the two tanks during the lower-layer filling process. For the smaller tank, fluid was driven through a Masterflex Easy-Load peristaltic pump head (model no. 77202-60), which was fitted to a Masterflex L/S Digital Drive (model no. 7551-00). For the larger tank, fluid was driven through a Micropump (model no. GJ-N25.PF1S-A), which was fitted to a Masterflex Digital Gear Pump drive (model no. 75211-70). For each tank, the associated digital pump system was capable of delivering fluid to the bottom of the tank at the required range of flow rates for the lower-layer filling process.

For each tank, different apparatus were connected to the pump tubing to transport lower-layer fluid to the bottom of the tank. For the smaller tank, the tubing is connected to the top end of a vertical cylindrical hollow rod of diameter 9 mm, which is made to stand in one corner of the tank. Lower-layer fluid is injected into the tank from the opposite end of the rod. A small sponge is fitted at the bottom end of the rod to diffuse the kinetic energy of the injected fluid. This was sufficient to prevent significant mixing from occurring during the filling process. This filling setup is identical to that used by Davies Wykes and Dalziel (2014), and allows the filling process to take less than two hours to complete.

For the larger tank, it was necessary to replace the rod with a different apparatus that could supply lower-layer fluid to the tank at higher flow rates without significant mixing

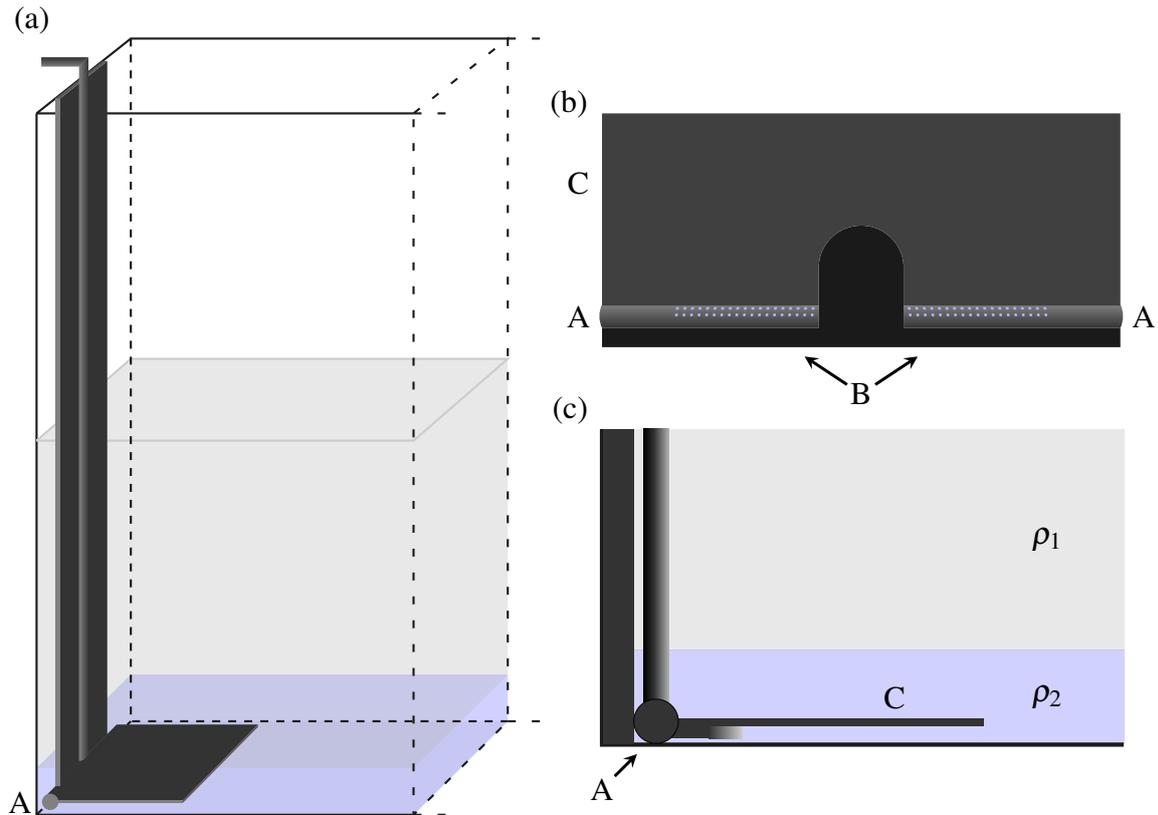


Fig. 3.2 Sketch of the apparatus used to fill the larger tank. (a) View of the apparatus positioned in the tank. (b) View of the underside of the apparatus from directly beneath it. (c) Side view of the apparatus at the bottom of the tank.

of the two fluids. A sketch of the apparatus developed for this purpose is given in Figure 3.2. The apparatus is positioned against one of the narrow walls of the tank, as shown in Figure 3.2 (a). Lower-layer fluid is pumped into the apparatus at the top of the tank, before travelling to the bottom of the tank through a closed pipe of diameter 15 mm. The fluid then enters and fills a horizontal cylindrical pipe of diameter 15 mm, labelled A, that spans the width of the tank.

Figure 3.2 (b) shows a view of the underside of the apparatus, from directly beneath it. Fluid enters the tank through a series of 80 small holes, labelled B, on the surface of pipe A. Each hole has a diameter of $D_h = 3$ mm, and was cut to face radially from the centre of the pipe. When the apparatus is positioned in the tank, half of the holes are centred at 3 mm above the tank base, and the other half are centred at 6 mm above the tank base.

For the range of flow rates used, the volume flux Q of lower-layer fluid supplied to the apparatus ranged from $4 \text{ ml/s} \leq Q \leq 40 \text{ ml/s}$. Therefore, assuming the volume flux through

each of the holes is equal, fluid would enter the tank at a velocity U_h ranging from 0.7 cm/s $\leq U_h \leq 7$ cm/s. Based on D_h and U_h , we can define a Reynolds number Re_h for the flow passing through each hole and a Richardson number Ri_h quantifying the relative strength of the potential energy of the buoyancy difference of the two fluids compared to the kinetic energy of the inflow. Writing $U_{h,\min} = 0.7$ cm/s with $r_h = U_h/U_{h,\min}$ ranging from 1 to 10, for the range of flow rates used we have

$$(Re_h, Ri_h) = \left(\frac{U_h D_h}{\nu}, \frac{g \Delta \rho D_h}{\rho_1 U_h^2} \right) = (20r_h, 6r_h^{-2}), \quad (3.1)$$

where we have taken $\Delta \rho / \rho_1 = 10^{-2}$ as this corresponds to the lowest density difference between the two layers that we used.

For the lowest flow rates (corresponding to $Ri_h = 6 \gtrsim O(1)$), the buoyancy force in combination with the orientation of the pipe-exit holes causes the injected lower-layer fluid to fall to the base of the tank before a significant amount of mixing can occur. Subsequently, the lower-layer fluid spreads along the base of the tank as a laminar gravity current. As $Re_h = 20 \sim O(1)$, we were confident that negligible mixing would occur as the lower-layer fluid settles from the pipe-exit holes to the base of the tank.

From the above analysis, we were confident that, using this tank-filling apparatus, we could inject the lower layer into the tank whilst inducing negligible mixing between the upper and lower-layer fluids. To keep the tank-filling time T_F sufficiently low for our experiments, it was necessary to resort to higher flow rates. At the highest flow rates used, from Equation (3.1) we have $(Re_h, Ri_h) = (200, 0.06)$. In this regime, the inflow consists of a series of turbulent jets coming out of each of the holes, capable of generating significant mixing in the presence of a stratification.

Figure 3.2 (c) shows a side view of the apparatus at the injection site. To prevent the turbulent energy of the inflow from generating mixing at the interface at higher flow rates, a flat horizontal plate is used to vertically confine the turbulent motion. The plate, labelled C, is fixed to the apparatus at $H_I = 7.5$ mm above the base of the tank. The plate spans the width of the tank and protrudes $L_I = 100$ mm into the tank.

Lower-layer fluid injected into the tank first enters the confined region between the tank base and the plate. As the ratio of length to height in this confined region is about 13 to 14, then towards the exit of the confined region, by the incompressibility condition we could expect the vertical velocity component to be $O(1)$ mm/s at the highest injection flow rates. Once the injected fluid exits the confined region, due to an adverse pressure gradient caused by the wall at the opposite end of the tank, the departing fluid would be pushed up by a small amount as it expanded into the rest of the lower-layer fluid. Though we did not investigate

the effect of this adverse pressure gradient in detail, by being aware of it, we knew to take care in keeping the injection flow rates at an intermediate level while the density interface was less than 2-3 cm above the confining plate to avoid inducing mixing.

By keeping the flow rate low initially to allow the density interface to pass above the confining plate, and only gradually increasing the flow rate to its maximum after this point, we were able to fill the lower layer in less than 90 minutes without inducing mixing at the density interface. The filling time was short enough that, for our PIV/LIF experiments presented in Chapter 5, the view of the interface through the cameras remained clear as sufficiently few PIV particles had settled at the interface during the tank-filling process. This was essential in allowing us to carry out our PIV/LIF experiments in the larger tank. Note that as PIV/LIF experiments were not conducted in the smaller tank, a similar confining plate was not required for the smaller tank as we could keep injection rates low to prevent mixing, while not having to worry about the settling issue associated with the use of PIV particles.

3.1.3 Molecular diffusion and interfacial thickness

The sharpness of the density interfaces produced for our experiments was limited by molecular diffusion, which acted to smooth the density field. Table 3.1 lists the diffusivities of the various species used in our experiments. For an initially sharp and deep two-layer stratification with $\rho_1 < \rho_2$ and evolving only via molecular diffusion with diffusivity κ , it is simple to show that the evolution of the density field $\rho(z, t)$ is given by

$$\rho(z, t) = \frac{1}{2}(\rho_1 + \rho_2) + \frac{\Delta\rho}{2} \operatorname{erf}\left(\frac{z}{2\sqrt{\kappa t}}\right), \quad (3.2)$$

where $-1 < \operatorname{erf}(x) < 1$ is the error function.

Defining the thickness of the density interface as $\Delta H = 2Z(t)$ where Z is such that $\operatorname{erf}(Z/(2\sqrt{\kappa t})) = 0.9$, we can use the diffusivities given in Table 3.1 to estimate the thickness of the density interface after the tank-filling process. Taking the tank-filling time T_F to be 90 minutes ($t = T_F = 5400$ s), we have interface thickness $\Delta H_\rho \sim 12.5 - 14.4$ mm for the ionic species present in our experiments, and $\Delta H_{\text{LIF}} \sim 6.8$ mm for rhodamine 6G. These interface-thickness estimates are consistent with our density measurements using the conductivity probe and thermistor (introduced in Section 3.4) and with our LIF images.

Writing $\Delta h_\rho = \Delta H_\rho/a$ for the relative-interface-thickness of the density field where $a \approx 49$ mm is the vortex-ring diameter, we have $\Delta h_\rho \sim 0.25 - 0.3$ at the end of the tank filling process. The range of Δh_ρ was deemed to be sufficiently small so that the interfacial thickness could be treated as constant across all PIV/LIF experiments. For our periodic-ring-mixing

Ion/dye	Cl ⁻	Na ⁺	NO ₃ ⁻	R6G
Diffusivity [10^{-9} m ² /s]	2.032	1.334	1.902	0.4

Table 3.1 Table presenting the molecular diffusivities of chloride (Cl⁻), sodium (Na⁺) and nitrate ions (NO₃⁻), and the molecular diffusivity of rhodamine 6G dye (R6G). Data obtained from Weast et al. (1984) and Gendron et al. (2008).

experiments presented in Chapter 6, the role of molecular diffusion on the evolving density field is discussed in Section 6.1.1.

It is worth noting that, for our periodic-mixing experiments, the mixing produced by vortex rings in the early stages of the initial adjustment phase will have a significant dependence on the prescribed initial value for Δh_ρ . This is because the interfacial thickness is a control parameter for the interfacial gradient Richardson number, which will control the amount of mixing produced both by the vortex ring, and by interfacial waves generated by the ring's impingement onto the interface. As we do not have access to apparatus with sufficient precision to examine closely the relationship between the interfacial thickness and the mixing produced by a single vortex ring, we do not explore this relationship further. For our purposes, the important point is that the initial interfacial thickness is kept fixed to a good approximation throughout our experiments.

3.1.4 Thermal effects

Using the methodology detailed above, we were able to generate two-layer stratifications with moderately sharp density interfaces in both of the tanks that were used. Variations in fluid temperature led to small differences between the prescribed stratification and the final stratification that was produced. Here, we go through the procedures followed in our methodology that impacted the temperature of each solution before the lower-layer fluid was added to the bottom of the tank, and comment on the typical final temperature stratifications that formed before the beginning of an experiment.

When the fresh water used for each solution was initially drawn from the break tank, it was typically 2 – 3°C cooler than the temperature of the lab, with the precise temperature difference depending on the weather and the time of year. Solutions were then prepared by dissolving salt crystals into the fresh water. For both salts used (NaCl and NaNO₃), the dissolution reaction is endothermic, meaning the reactions result in a cooling of the solution, thus further increasing the temperature difference between the solutions and the lab.

When the solutions were left overnight, they warmed up as they converged towards thermal equilibrium with the lab. This led to a reduction in the solubility of air in the

solutions, leading to an excess of dissolved air which came out of solution. As mentioned in Section 3.1.1, this was responsible for the nucleation and growth of air bubbles that formed on the surfaces in contact with the solution. By leaving the solutions overnight, a substantial majority of the initial excess of dissolved gases came out of solution, with any bubbles that had formed being easily removable. This significantly reduced the risk of bubbles forming during an experiment which could have impacted our ability to take measurements.

Two mechanisms primarily affected the temperature of the solutions as they were left overnight: evaporative cooling at the free surface, and a heat flux from the laboratory environment to the solutions, across the reservoir and tank walls. The effects of evaporative cooling were mostly eliminated by closing the top of the tank and the reservoir, isolating the solutions from direct exposure to the lab environment. Due to the reservoir walls being thinner than the tank walls and having a similar thermal conductivity, the heat flux across the reservoir walls was greater than the heat flux across the tank walls. This resulted in a faster warming of the lower-layer fluid in the reservoir than the upper-layer fluid in the tank.

Figure 3.3 (a) shows a vertical profile of the fluid temperature as a function of height for a typical experiment, measured using a thermistor a few minutes after the lower-layer filling process had been completed. The positive temperature gradients in each layer indicate that, within the measured region, each layer is thermally stably stratified. For this example, the upper layer is cooler, which we attribute to evaporative cooling at the free surface and convection mixing the cooled fluid throughout the upper layer. The lower layer is not cooled in this way, allowing it to become warmer than the upper layer. As a result, the temperature profile is unstably stratified in the vicinity of the density interface, allowing for double-diffusive behaviour to develop in the flow. As the average temperature between the two fluid layers was consistently less than 1°C at the start of an experiment, any double-diffusive effects in our experiments were negligible, so no attempt was made to mitigate them.

Figure 3.3 (b) shows the temperature of the laboratory over the course of a week in September 2020, measured at a fixed position near the tanks. The blue lines correspond to 8am for each day. Temperature fluctuations occurred over an hourly timescale due to the air ventilation system, which was on between 8am and 8pm on weekdays. These fluctuations could cause the air temperature to vary by as much as 1.5°C over a 24 hour period. For our periodic mixing experiments, which lasted about 18 hours, the temperature of the stratification could vary by as much as 1°C . For experiments where external heat fluxes played a significant role in the energetics of the system, this is quantified and discussed. The average temperature in the lab also changed from day to day and on a seasonal timescale due to the varying weather, but these temperature changes had a small effect over the course of a single experiment relative to the effect of the lab ventilation system.

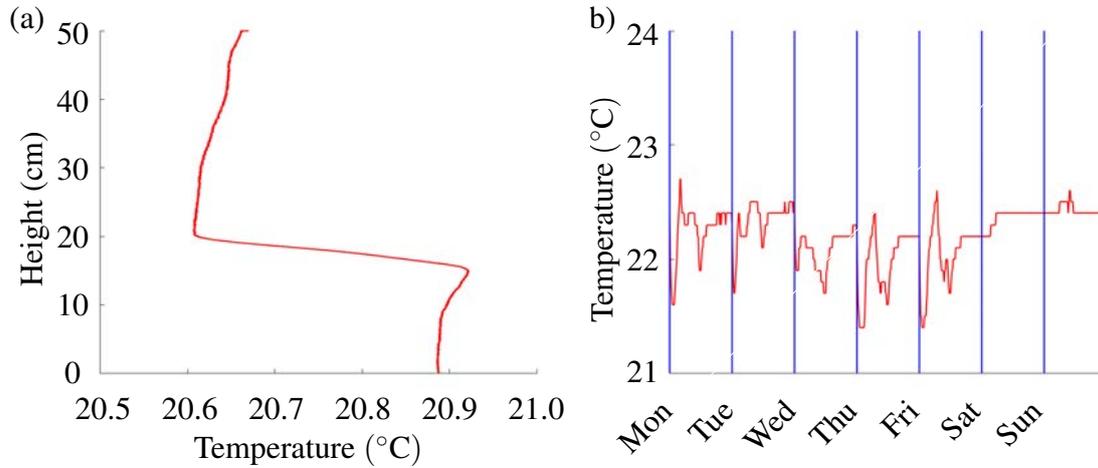


Fig. 3.3 (a) Example of a typical vertical profile of the temperature stratification, taken a few minutes after the lower-layer filling process had been completed. (b) The temperature of the laboratory measured using an RC-5 data logger (Elitech) for a week in September 2020. The blue lines correspond to 8am for each day.

3.2 Vortex rings

In the experiments conducted for this thesis, a method was required for generating reproducible vortex rings. The methodology we adopt is to generate vortex rings at the outlet of a cylindrical tube with circular cross-section, by making use of a piston-like mechanism. For details of the dynamics associated with the ring formation process, the reader is referred to our discussion in Section 2.2.2. Here, we restrict our discussion to the apparatus and setup used to generate vortex rings, and the controllable parameters that allowed us to generate reproducible rings.

The methodology we used for generating vortex rings is similar to that used in many previous studies (see Scase and Dalziel (2006), Bethke and Dalziel (2012), Olsthoorn and Dalziel (2017) for example). A hollow, piece-wise cylindrical tube with circular cross-section (hereafter referred to as a ‘ring tube’) was positioned in the tank, with one open end fully submerged in the stratification and the other open end above the free surface of the stratification. A closed section of tubing, filled with air, connected the end of the ring tube above the free surface to a pair of bike pumps, with internal diameter $D_{BP} = 26$ mm. Actuation of the bike pumps delivers an impulse to the fluid contained within the ring tube, causing some of the fluid to be ejected from the tube at its submerged end, forming a vortex ring in the process.

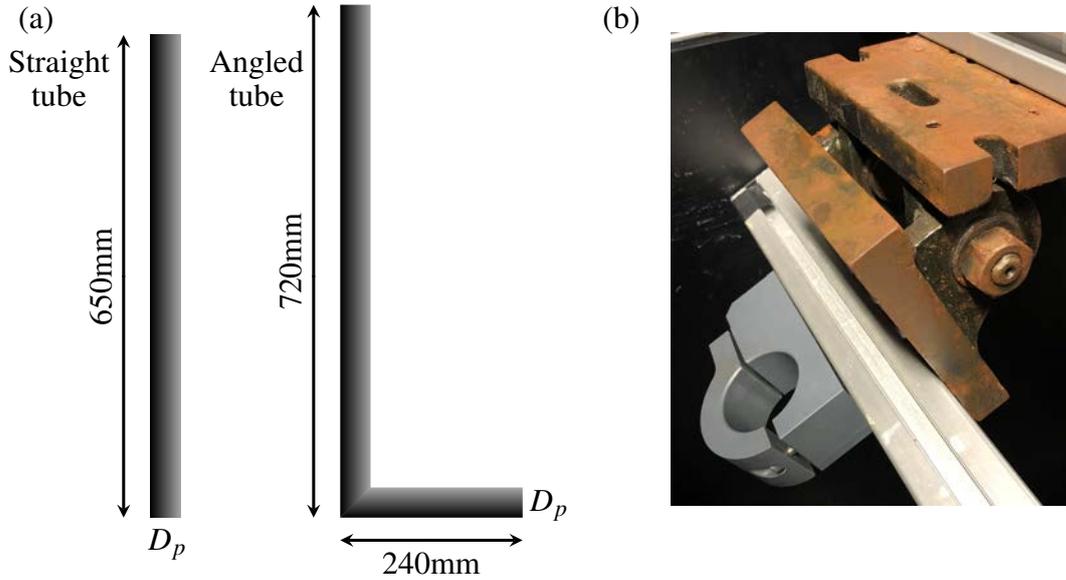


Fig. 3.4 (a) Simple illustration of the two ring tubes used, which both have an inner diameter $D_p = 39\text{mm}$. Shading is used to indicate the cylindrical geometry of the tubes. (b) A photograph of the angle setter. The ring tube in use is secured to the angle setter, which we use to set the ring propagation angle θ_0 .

Two ring tubes were used in our experiments: a straight tube, and an angled tube. A simple illustration of the ring tubes is given in Figure 3.4 (a). The straight tube is a hollow cylindrical tube of length 650 mm and inner diameter $D_p = 39\text{ mm}$. The angled tube, which also has an inner diameter of $D_p = 39\text{ mm}$, is a hollow, piece-wise cylindrical L-shaped tube, consisting of two straight sections with a right angle at the connecting bend. The longer and shorter section have lengths 720 mm and 240 mm respectively. For the angled tube, vortex rings are generated at the outlet of the shorter section of tubing.

For a given experiment, the ring tube in use was fixed to the ‘angle setter’, an apparatus allowing us to vary θ_0 , the propagation angle of the vortex ring with respect to the vertical. A photograph of the angle setter is given in Figure 3.4 (b). Due to geometrical constraints of the tanks, the use of two different tubes was required so we could investigate the full range of propagation angles $0^\circ \leq \theta_0 \leq 90^\circ$, with the straight tube used for $\theta_0 \leq 30^\circ$ and the angled tube for $\theta_0 > 30^\circ$.

As with Olsthoorn and Dalziel (2017), the bike pumps were attached to a digitally controlled traverse, which was actuated by a stepper motor. The stepper motor was operated via a multi-functional instruments card, as were all the other digitally controlled apparatus used for the experiments presented in this thesis. Two different cards were used in our experiments: a DAQ device (National Instruments) for the PIV/LIF experiments, and a UEI

card (PDL-MF, United Electronic Industries) for the periodic mixing experiments. Both devices provided us with the same functionality for operating the stepper motor, allowing us to follow the same ring generation procedure across all our experiments.

Control over parameters governing the ring formation process, in particular the ring formation number F_p (defined in Section 2.2.2) and propagation speed U , was made through control of the stepper motor operating the bike pump traverse. The stepper motor moves in discrete steps, in response to discrete pulses sent via the instruments card. For each pulse, the stepper motor rotates by a fixed angle, corresponding to a fixed linear displacement $\Delta z_s = 0.33$ mm of the bike pump traverse. For each vortex ring generation, we were able to prescribe F_p and U through prescribing the traverse displacement profile, $Z_s(t)$.

The traverse displacement profile we used for our experiments is given by

$$Z_s(t) = \Delta z_s \left\lfloor \frac{N_s}{2} \left(1 + \tanh \left[\frac{N_s}{w_s} \left(\frac{-T_s + 2t}{2T_s} \right) \right] \right) \right\rfloor, \quad 0 \leq t \leq T_s, \quad (3.3)$$

where the step total N_s is the total number of pulses sent to the stepper motor, the traverse time T_s is the total time taken for the full displacement of the bike pump traverse, and w_s is a non-dimensionalised width. The notation $\lfloor \cdot \rfloor$ represents the floor function, outputting the largest integer less than or equal to the input value. Figure 3.5 (a) plots the non-dimensionalised displacement profile against time. The functional form of $Z_s(t)$ was motivated by allowing the traverse to have a smooth acceleration and deceleration phase, increasing the repeatability of the traverse displacement profiles between different ring generations of the same input parameters.

The pulse total was kept fixed at $N_s = 300$ for all our experiments, which set the vertical displacement of the bike-pump handles to $L_B = 99$ m. Relating the total volume of air displaced in the two bike pumps to the geometry of the ring tube, it follows that our prescription of N_s sets the formation number at $F_p = 2.3$ for the ring formation process. At $F_p = 2.3$, from Figures 2.5 (a) and 2.3 (a) we can see that vortex rings will have a moderate core thickness, with a ratio of core to ring diameter $b/a \approx 0.2$. The dimensionless width was kept fixed at $w_s = 0.06N_s$ for simplicity; the impact of varying w_s has on the ring formation process was not investigated. The traverse time T_s is the only parameter in Equation (3.3) that we varied, allowing us to vary the propagation speed of the vortex ring without changing the functional form of $Z_s(t)$.

For all experiments, the ring tube was positioned at height H_B such that

$$\frac{(H_B - H_2) \cos \theta_0}{a} = 4,$$

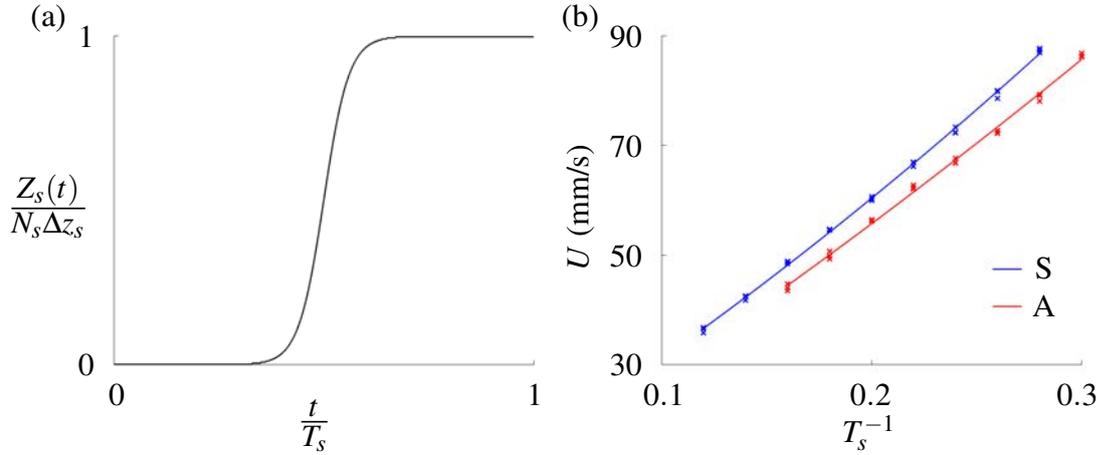


Fig. 3.5 (a) The non-dimensionalised displacement profile of the bike pump traverse used to generate vortex rings. (b) Calibration curves for the straight (S) and angled (A) ring tubes, where PIV measurements were used to relate the ring propagation speed U to the traverse time T_s . Across all calibration measurements, precision errors for U were less than ± 0.6 mm/s.

where H_2 is the height of the density interface. In this way, the distance along the initial propagation direction of the ring between the centre of the tube outlet and the interface is $4a$ for all θ_0 . Therefore, assuming our vortex rings are reproducible for given T_s , the amount of ring-transported kinetic energy delivered to the interface will be independent of the prescription of θ_0 .

Figure 3.5 (b) plots the calibration data obtained from PIV to relate T_s to U for each ring tube, as well as the corresponding fitted quadratic calibration curves. The calibration curves were used to prescribe U in our experiments. To calibrate our bike pumps, PIV was used to relate T_s to U , using the methodology detailed in Section 3.3. For each T_s , three experiments were conducted, where a vortex ring was generated and propagated through fresh water. For the j th experiment, the positions $\mathbf{x}_{i,j}(t) = (x_{i,j}, z_{i,j})$ ($i = 1, 2; j = 1, 2, 3$) of the two vortex-core cross-sections of the ring visible in the light sheet were measured using the methodology detailed in Section 5.1. By making 120 measurements of $\mathbf{x}_{i,j}(t)$ when the center of the ring was between two and four ring diameters from the tube outlet, a velocity measurement $U_j(T_s)$ ($j = 1, 2, 3$) was obtained by taking the gradient of a line of best fit for $\mathbf{x}_{i,j}(t)$ against t , obtained using a least squares procedure. The measurement U_j was taken to be the propagation speed U of the vortex ring when at a distance $3a$ from the tube outlet.

We attribute the reduction in U for the angled tube relative to the straight tube at a given T_s to energy losses at the bend of the angled tube. Here, a separating flow will develop as fluid travels around the corner of the tube, resulting in an increase of the work required by

the bike pumps to generate a vortex ring at a given velocity. For our PIV/LIF experiments, we were able to measure U directly, which we always found to be in good agreement with the calibration data. These calibration validations further increased our confidence in our procedure for generating reproducible vortex rings.

3.3 Planar data acquisition

Implementing the methodology outlined in sections 3.1 and 3.2, we have a robust procedure for producing a sharp two-layer density stratification and generating reproducible vortex rings. Using Particle Image Velocimetry (PIV) and Planar Laser Induced Fluorescence (PLIF), this section describes how, using two cameras, two-dimensional planar velocity and density data was obtained in the larger tank that we used. First, we discuss our choice of solutes for the upper and lower fluid layers to eliminate the effects of refractive index variations in the stratification.

3.3.1 Refractive index matching

When light travels through a boundary between two optically isotropic media, the light path is deflected according to Snell's law:

$$n_1 \sin \theta_1 = n_2 \sin \theta_2, \quad (3.4)$$

where θ_1 is the angle of incidence, θ_2 is the angle of refraction and n_1, n_2 are the refractive indices of the respective media. The principle is illustrated in Figure 3.6. A change in light refraction can lead to an inability to focus on regions in the field of view of a camera. As PIV particles needed to be kept in focus so accurate velocity measurements could be made, it was necessary to minimise refractive index variations in our experiments.

The refractive index n of a medium is dependent on several intrinsic properties of the medium. For salt solutions, we assume n takes the form

$$n = n(\mathbf{S}, \lambda, T, P), \quad (3.5)$$

where S_k is the concentration of solute s_k , λ is the wavelength of the light travelling through the medium, T is the temperature and P is the pressure. To obtain a uniform refractive index field $n = n_0$, each of these parameters are considered.

A monochromatic light source of wavelength $\lambda_0 = 532$ nm was used to illuminate the flow, ensuring there was no dispersion of light due to n varying with λ . As the range of

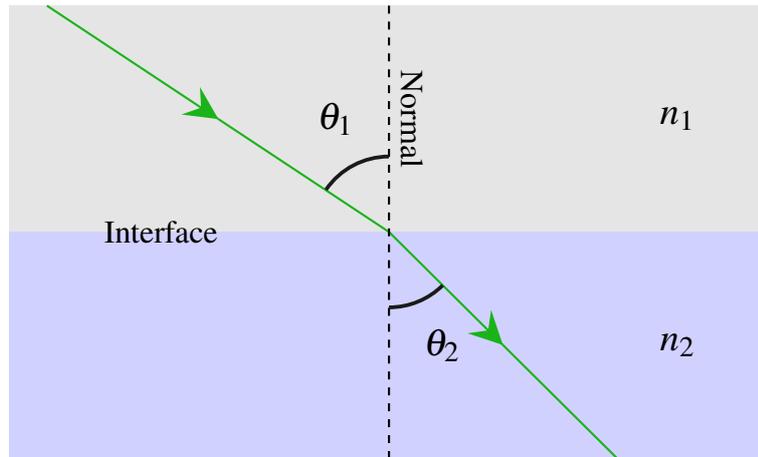


Fig. 3.6 Diagram illustrating the deflection of a light ray at the boundary between two isotropic media, as described by Snell's law given by Equation (3.4).

temperatures in the stratification was less than 1°C and the fluid depth was typically less than 60 cm, refractive index variations due to temperature and pressure were found to be negligible. Thus, for our purposes, we were able to simplify Equation (3.5) and treat n as a function of solute concentrations only, i.e. $n = n(\mathcal{S})$.

Suppose we created a two-layer stratification with density difference $\Delta\rho = 10^{-2} \text{ g/cm}^3$, using fresh water to form the upper layer and sodium chloride solution to form the lower layer. From Weast et al. (1984), the refractive index difference between the two layers would be $\delta n \approx 2.4 \times 10^{-2}$. Using other solutes readily available in place of sodium chloride, we similarly found that $\delta n \approx O(10^{-2})$. For a refractive-index difference at this order of magnitude, we would have had an unacceptable loss of focus of the PIV particles in the flow. Thus using a single solute to form our stratification was not feasible.

To minimise variations in n , we used the method of refractive index matching, which makes use of two different solutes to produce solutions with different densities and matching refractive indices. The technique was first introduced by McDougall (1979) to study plumes, using magnesium sulfate and sugar as the two solutes. As discussed in Section 3.1.1, we used NaCl and NaNO₃ to create the solutions forming the upper and lower layers respectively.

The calibration we used between refractive index and fluid density for each solute was carried out by Lefauve (2018), who used a handheld refractometer illuminated under a green light source in an attempt to match the light wavelength λ_0 . The precision of this calibration was found to be sufficient for our purposes. Density plots for each calibration measurement of NaNO₃ and NaCl against n are given in Figure 3.7.

At this stage it is unclear whether the mixing of index matched NaCl and NaNO₃ solutions will cause the refractive index to vary. Fortunately, there were no observable variations in

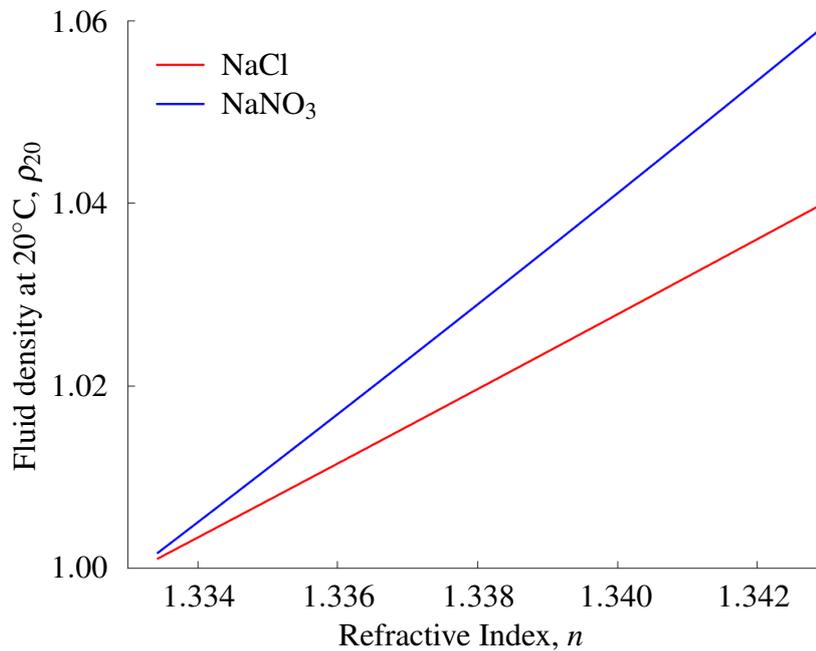


Fig. 3.7 Plots of the density of sodium chloride (NaCl) and sodium nitrate (NaNO₃) solutions at 20°C against the refractive index, n under green light. Data obtained by Lefauve (2018).

refractive index due to mixing in our experiments, so we did not investigate this possibility. It is also worth noting that the molecular diffusivities of the two salts are sufficiently similar (see Table 3.1) that we do not introduce any identifiable double-diffusive behaviour at the interface, making NaCl and NaNO₃ an ideal choice of solutes in our experiments.

3.3.2 Plane illumination

In this section, we describe how, using a laser system, we illuminated the plane coinciding with the central axis of the vortex ring over its full trajectory. A Nano L 50-100 PIV pulsed laser (Litron Lasers), located on the floor beneath the tank base, was used to illuminate the flow. The laser system produces two 532 nm coaxial laser beams that can each be fired independently at a repetition rate of up to 100 Hz. Each laser pulse carries 50 mJ of energy and illuminates the flow for 8 ns. Optical lenses are attached to the laser system. Each laser beam first travels through a concave lens that expands the beam, before travelling through the subsequent lenses which ultimately leads to the formation of a light sheet. The resulting light sheet passes through the base of the tank, illuminating a thin, near two-dimensional region of the tank interior. A schematic of the experimental setup is given in Figure 3.8.

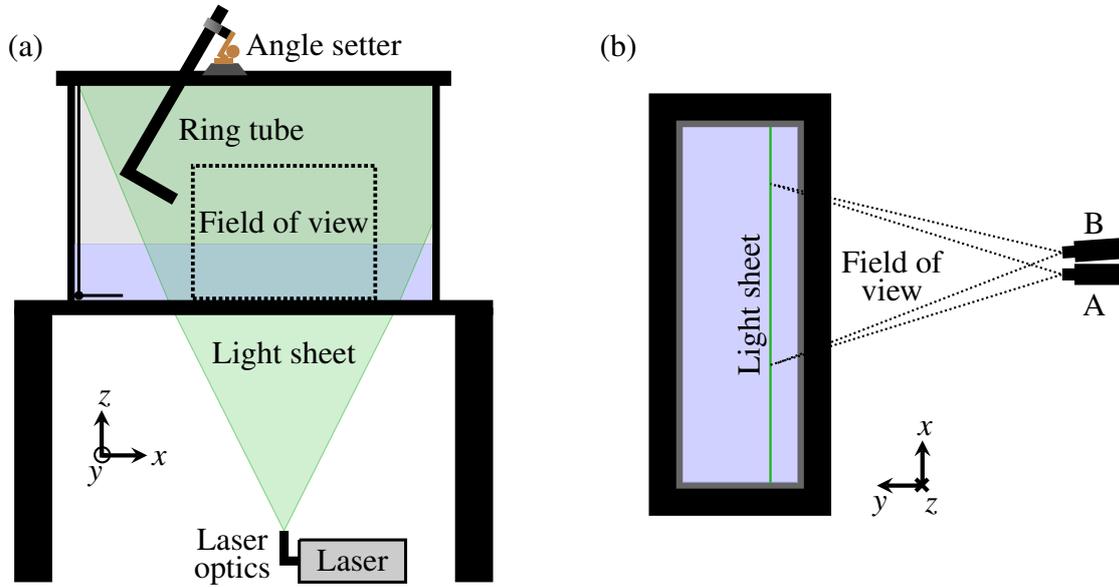


Fig. 3.8 Schematic of the planar PIV/LIF setup. (a) Front view; the central plane of the light sheet is aligned with the axis of the ring tube. (b) Plan view; two cameras (labelled ‘A’ and ‘B’) used to take measurements are positioned to have overlapping fields of view of the light sheet. For ease of interpretation, we have omitted illustrations of the ring tube, angle setter and tank filling apparatus in the plan view.

We use the subscript L to denote quantities associated with the properties of the light sheet. Writing $\mathbf{x} = (x, y, z)$ for the position vector of the world coordinate system, the light sheet has thickness $\Delta y_L(x, z)$, light intensity $I_L(\mathbf{x})$ and an illumination period of $t_L = 8$ ns corresponding to each laser firing. The plane $y = 0$ is set to coincide with the central plane of the light sheet, which is set up to be parallel with the direction of gravity through careful alignment of the laser system.

The precise alignment of the light sheet was chosen to optimise the cameras’ view of the illuminated flow. The positioning of the cameras with respect to the light sheet is discussed in Section 3.3.3. After the light sheet had been correctly aligned, the ring tube was aligned with the light sheet so that the symmetry plane of the ring–interface interaction during its laminar phase would coincide with the central plane of the light sheet. The alignment of the light sheet and the tube is discussed below. The laser was set to reduced power and safety goggles were used during the alignment process to minimise the risk of eye damage.

There are six degrees of freedom associated with the precise alignment of the laser system (and, in turn, the light sheet). Three of these are the translational positioning of the laser system in the x , y and z directions, and the other three are related to the orientation of the light sheet, by rotation of the laser system through the x , y and z axes. Translation in the y

direction and rotation through the x and z axes required a much higher level of precision for alignment of the laser system than the other degrees of freedom, and so are discussed separately. By initially placing the laser system on the floor beneath the centre of the tank (with the laser optics pointing upwards), there was no need for further alignment of the laser system in the other three degrees of freedom.

Before orienting the light sheet, we first repositioned the laser system by translation in the y direction, fixing the distance between the central plane of the light sheet and the tank wall nearest the cameras (hereafter referred to as the ‘near-camera wall’). As the fluid in the tank was seeded with tracer particles (as discussed in Section 3.3.4), the image quality could be compromised if too many particles were obscuring the view of the light sheet. It was therefore desirable for the light sheet to be as close as possible to the near-camera wall. Also, it was necessary for the tank walls to have a negligible effect on the vortex ring dynamics. As we choose to illuminate the central axis of the vortex ring, we positioned the laser so that the central plane of the light sheet was two and a half vortex ring diameters (~ 12.4 cm) from the near-camera wall.

Once the laser system had been positioned, we oriented the light sheet to be parallel with the near-camera wall. This requires rotational alignments of the light sheet, both through the x and z axes. Rotational alignment through the z axis was achieved by using two white 1 m rulers, made to stand in the tank with the same orientation at different ends of the light sheet, approximately 60 cm apart. The laser system was then rotated about the z axis so that the light sheet hit both metre rulers at the same position to the nearest millimetre.

Rotational alignment through the x axis was achieved by observing the reflection of the light sheet from the base of the tank back onto the laser optical lenses. A small target for the reflection of the light sheet was made using yellow post-it note paper, which fluoresces slightly under the light of the laser, making the reflection visible through the laser-safety goggles. The target was securely fixed to the optical lenses. Using the target, the laser was then levelled so that the reflection of the light sheet coincided with the centreline of the optical lenses. With this methodology, the light sheet could be aligned with sufficient precision for our purposes. Though the rotational alignment procedure could perturb the initial y -position of the light sheet by a few millimetres, it was not necessary to correct for this.

After aligning the laser system, we refocused the light sheet optical lenses to adjust the thickness profile of the light sheet. In the y - z plane, the light sheet has an hourglass shape, contracting to a minimum thickness of $\Delta y_{L\min}$ before expanding again, as illustrated in Figure 3.9. The minimum thickness $\Delta y_{L\min} \approx 1$ mm, corresponding to approximately 2% of the vortex ring diameter. If the light sheet had been too thick, the ability to extract in-plane

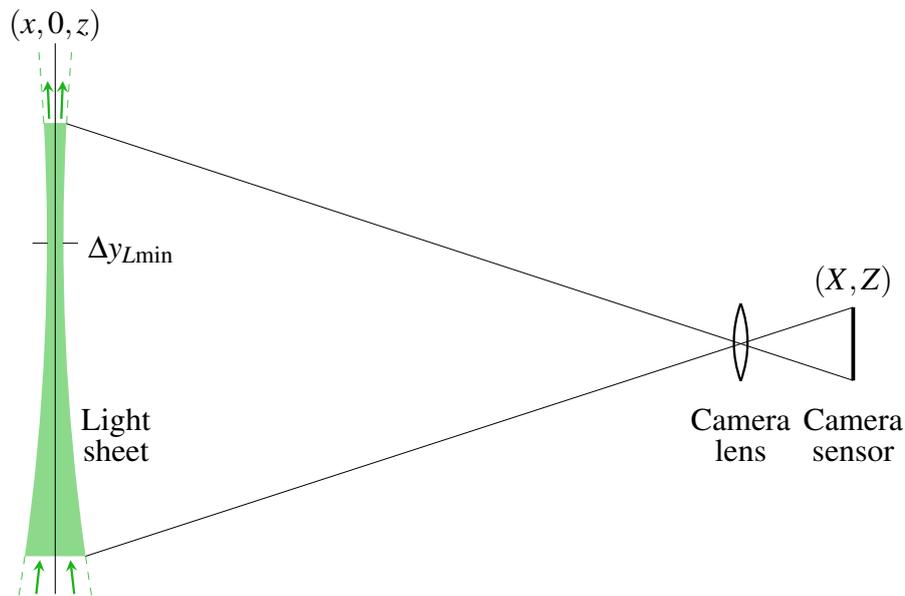


Fig. 3.9 Side-view illustration of the three-dimensional light sheet projected onto a camera sensor and quantised onto the sensor array (defined in Section 3.3.3). The light sheet thickness Δy_L has an hour-glass like shape. The z -position of the minimum thickness Δy_{Lmin} is adjustable by refocusing the laser optics.

velocity information from our PIV images would have been hindered, as would have the ability to discern structural features of the ring–interface interaction in the central plane of the light sheet from our LIF images. To reduce the light sheet thickness across the field of view of the cameras to a suitable level, we refocused the light-sheet optics to move the Δy_{Lmin} z -position to 2–3 cm above the prescribed height of the density interface.

After refocusing the laser optical lenses, we then aligned the cameras so they had clear, overlapping views of the light sheet (as illustrated in Figure 3.8). The camera alignment is discussed separately in Section 3.3.3. After aligning the cameras, we then aligned the ring tube with the light sheet and the field of view of the cameras. Due to the considerable difference in refractive index between air and the salt solutions used for our PIV/LIF experiments, the tube alignment procedure was carried out while the tank was partially filled with a NaCl solution, of the same refractive index as the solutions used in our experiments. This was done to reduce the risk of a slight misalignment of the ring tube with respect to the light sheet passing through the two-layer stratification.

The ring tube was securely held in position by the angle setter (introduced in Section 3.2 and illustrated in Figure 3.8 (a)). Rotational alignment of the tube through the y axis was made first, by using the angle setter to fix the propagation angle θ_0 . After this, we made an approximate repositioning of the ring tube along the y axis, to a position where the tube

would coincide with the light sheet when the laser is fired. The position where the tube is held by the angle setter can be adjusted. By making use of this ability for adjustment, along with the ability to move the angle setter along the x axis, we then moved the tube outlet in the x - z plane to a suitable position given the field of view of the cameras.

The ring-tube-alignment procedure up to this stage was made by hand while the lasers are off. Targets cut into the tube exterior were then used to facilitate the precise translational and rotational alignment of the tube with the light sheet. The targets were positioned on the tube so that they were fully illuminated by the light sheet only when the tube axis coincided with the central plane of the light sheet. By using these targets while firing the laser at reduced power, we aligned the ring tube by hand to the correct position.

Provided the ring tube had been correctly aligned with the light sheet, the central axis of each vortex ring would remain in the central plane of the light sheet over the full ring trajectory, assuming the trajectory was not perturbed by any residual motion in the tank. Two processes that impart unwanted motion into the tank were addressed.

Firstly, the lower-layer-filling process introduces highly energetic fluid into the tank. After the completion of the lower-layer filling process, we waited for the typical rms velocity in the region of the fluid that the vortex ring would pass through to fall beneath 0.5 mm/s. This took about five minutes. Though waiting longer would have allowed the typical residual velocities to fall further, it would have also given more time for more PIV seeding particles to settle at the density interface, eventually completely obscuring the cameras view of the illuminated region of the density interface. We found that waiting five minutes was long enough for the typical rms velocities to be sufficiently small for our purposes, and short enough that the effect of additional particles settling during this time was not significant.

Secondly, the air flow due to the ventilation system in the lab was capable of imparting significant kinetic energy into the fluid at the free surface. This effect was removed by using black corrugated card to provide roofing at the top of the tank, sheltering the upper-layer fluid from the air flow. By making these adaptations to our experimental procedure, we improved our ability to ensure that the central axis of each vortex ring would remain in the central plane of the light sheet over the full ring trajectory.

Through adjusting the focus on the laser optics and the power of the lasers, there was some control over the thickness and intensity profiles of the light sheet, which could be finely tuned with the camera optics to optimise the image quality for PIV and LIF data processing. These considerations are discussed in more detail in Sections 3.3.4 and 3.3.5 respectively. We now go on to describe our camera setup and calibration methodology.

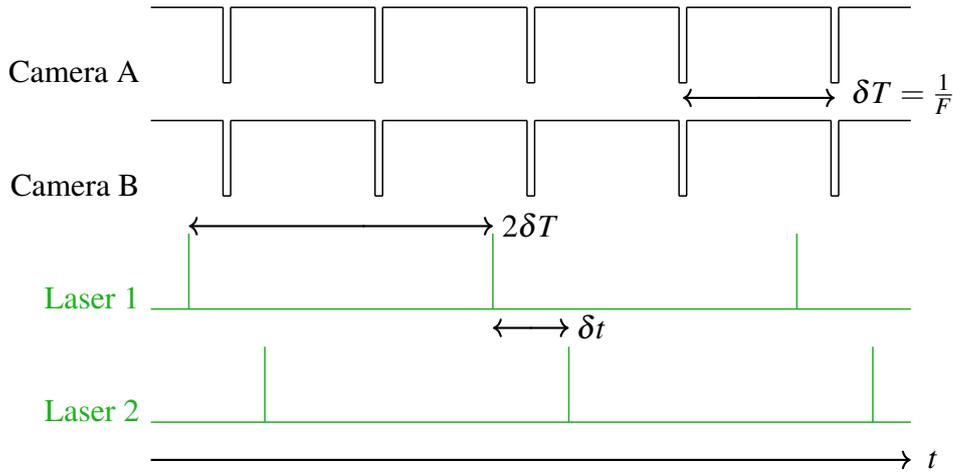


Fig. 3.10 Schematic of the timing sequence between the camera shutters opening and closing for frame capture and the laser pulses, where the laser-pulse offset δt is the time interval between consecutive laser pulses and δT is the inverse sampling frequency of the camera.

3.3.3 Camera setup

Two cameras (Allied Vision Technologies, Bonito CMC-4000 4MP) were used to record the evolution of the flow in the light sheet, from the instant the ring begins to distort the interface, to the complete turbulent breakdown of the ring and the interface restoring from the impact. Each camera has 4 Mpx resolution and was equipped with a 50 mm, f/1.4 Nikon lens. One of the cameras, which we refer to as ‘Camera A’, was used to obtain PIV measurements. The other camera, which we refer to as ‘Camera B’, was used to obtain LIF measurements. The cameras are illustrated in Figure 3.8 (b).

Figure 3.10 illustrates the relationship between the frame capture of the cameras and the pulsing of the lasers. The cameras and lasers were controlled through a DAQ Device (National Instruments) that was programmed using DigiFlow, allowing for precise synchronisation between the laser pulses and frame capture of the two cameras. The recording interval of each camera frame overlapped with a single laser pulse, so each frame provided a near-instantaneous snapshot of the flow. The offset δt between laser 1 and laser 2 pulsing could be made smaller than the inverse frame rate $\delta T = F^{-1}$, allowing us to decrease the minimum time interval between two snapshots. For our experiments, our precise choice of δt was motivated by our particular implementation of a PIV algorithm we developed, which we discuss in Section 4.4.

Each camera was set with a frame rate of $F = 160$ fps, twice the repetition rate set for each of the lasers. Based on the maximum repetition rates of each of the lasers, in theory it was possible to set the frame rate at $F = 200$ fps, with the recording of each camera frame

overlapping with a single laser pulse. However, due to technological issues that emerged when attempting to control the cameras and lasers with the lasers firing towards the highest possible repetition rates, we found that the highest repetition rate where we were able to reliably acquire image sequences was at 80 Hz.

We refer to the photographs taken by Cameras A and B as ‘PIV images’ and ‘LIF images’, respectively. A bandpass optical filter centred around 532 nm was attached to Camera A, so each PIV image captured only the light scattered from the illuminated particles in the camera’s field of view. A 570 nm long-pass optical filter was attached to Camera B, so for each LIF image, only the fluorescing dye in the camera’s field of view was visible. Each image recorded by one of the cameras is a two-dimensional projection onto the camera sensor of the filtered light from the three-dimensional volume illuminated by a single laser pulse. A schematic of this projection is given in Figure 3.9.

Each projection onto a camera was quantised onto the ‘sensor array’: a matrix with the same size as the pixel resolution of the camera. The sensor arrays associated with Camera A and Camera B are denoted A and B , and have two-dimensional pixel coordinate systems (X_A, Z_A) and (X_B, Z_B) , respectively. Each time the cameras (or light sheet) were repositioned, a calibration was required allowing the measurements made in each sensor array to be mapped back onto the region of the central plane of the light sheet in the field of view of both cameras, which we refer to as the ‘image plane’. As discussed in Section 3.3.2, the world coordinate system (x, y, z) is defined so the image plane is coincident with the plane $y = 0$.

This calibration was achieved using a grid, consisting of a rectangular array of dots equally spaced 20mm apart, allowing a mapping between the pixel and world coordinate systems to be constructed. After filling the tank with an NaCl solution of equal refractive index to the salt solutions used in experiments, an image was taken with each camera of the grid standing in the tank, coinciding with the central plane of the light sheet. Using these calibration images, projection mappings $\mathcal{P}_A : \mathbf{x} \mapsto \mathbf{X}_A$ and $\mathcal{P}_B : \mathbf{x} \mapsto \mathbf{X}_B$ and inverse mappings \mathcal{P}_A^{-1} , \mathcal{P}_B^{-1} were constructed by making use of the known dot locations in the world coordinate system. The inverse mappings allow measurements to be mapped from the pixel coordinate systems to $(x, 0, z)$ in the world coordinate system. The Jacobian of each inverse function was used to convert differential quantities, such as the velocity field, from pixel units to world units.

Bilinear mappings were constructed for \mathcal{P}_A , \mathcal{P}_B by using a least squares fitting procedure. The rms error of \mathcal{P}_A , \mathcal{P}_B across our different camera calibrations was consistently around 1 pixel, corresponding to 0.2 mm in the world coordinate system. This error corresponds to less than 0.5% of the vortex ring diameter, and so was deemed negligible for our purposes.

Once the calibration between the pixel and world coordinate systems had been made, we were ready to take quantitative measurements of the flow in the world coordinate system. The next two subsections discuss the experimental considerations incorporated into our procedure and methodology for obtaining PIV and LIF measurements.

3.3.4 Particle image velocimetry

The basic principle behind our PIV methodology to calculate velocity fields is as follows. Tracer particles were seeded in the fluid to visualise the flow. As discussed in Section 3.3.3, the laser and cameras were synchronised so each frame capture coincided with a single light sheet illumination. A bandpass optical filter centred around 532nm was attached to Camera A so each PIV image captured only the light scattered from the illuminated particles in the camera's field of view. To calculate displacement vectors in the sensor array of Camera A, pairs of PIV images were compared using a pattern matching algorithm available with DigiFlow. This algorithm obtains displacement measurements by interrogating the movement of groups of particle projections in small regions of the sensor array. A second algorithm, which makes use of the displacement measurements made from the pattern matching algorithm, was then used in conjunction with the camera calibration described above to construct velocity fields in the world coordinate system.

The dynamical and optical properties of the particles used, the particle-seeding density, the camera setup and the properties of the light sheet all play a crucial role in taking accurate, high-resolution PIV measurements. The pattern matching algorithm and the algorithm used for constructing the final velocity fields are also important; these are discussed separately in Chapter 4. In this subsection, we restrict our attention to optimising the experimental setup for taking PIV measurements. First, we discuss our choice of tracer particles. We then describe our methodology for seeding the two fluid layers with particles. Finally, we discuss adjustments made to camera A to produce high quality PIV images.

Tracer particles

To make accurate PIV measurements, the particles used must act as passive tracers in the flow. The ability for particles to closely follow the fluid streamlines is characterised by the Stokes number, the appropriate form of which is dependent on the particle size (Xu and Bodenschatz (2008)). For particles smaller than the Kolmogorov lengthscale η , the Stokes number St_K is given by

$$St_K = \frac{1}{18} \left(\frac{\rho_p}{\rho_f} \right) \left(\frac{d_p}{\eta} \right)^2, \quad (3.6)$$

where d_p is the particle diameter, ρ_p is the particle density and ρ_f is the density of the ambient fluid. For $St_K \ll 1$, particles will faithfully follow the flow. It was important for us to verify that the Stokes number is small for the tracer particles we used.

The settling velocity was also an important consideration for our choice of tracer particles. During the lower-layer filling process, particles remain suspended in the fluid for approximately one hour. If the settling velocity is too great, particles will migrate to either the top or bottom of their respective layers in this time, prohibiting the use of PIV. For spherical particles at low Reynolds numbers, the settling velocity U_g is given by

$$U_g = d_p^2 \frac{(\rho_p - \rho_f)}{18\mu} \mathbf{g}, \quad (3.7)$$

where μ is the dynamic viscosity of the fluid and \mathbf{g} is the acceleration due to gravity (Stokes (1851)).

Our third consideration was the particle image diameter, d_A , i.e. the diameter of the particle projections in the sensor array. Previous studies have demonstrated that the particle image diameter strongly influences the accuracy of PIV measurements, with different optimal values reported in the broad range of 2.0 to 3.0 pixels (Prasad et al. (1992), Westerweel (1997), Wilson and Smith (2013)). It was important that our choice of tracer particles would allow us to produce particle images of a suitable size with our camera and laser setup.

In our experiments, we used spheroidal polyamide particles (Orgasol 2002 ES5 NAT 3) as tracer particles, with a distribution of diameters centred at $d_p = 50 \mu\text{m}$ and mostly lying in the range $47 \mu\text{m}$ to $53 \mu\text{m}$ (information provided by the product data sheet). To determine the particle density distribution, we created a linear stratification in a measuring jug and added a small sample of particles. Figure 3.11 is a close-up photograph of the particle sample after they have settled to their neutral-buoyancy level, where it is seen that the particles have a bimodal density distribution. By taking samples of the fluid, the two bands were found to be centred at $\rho_{p-} = 1.023 \text{ g/cm}^3$ and $\rho_{p+} = 1.025 \text{ g/cm}^3$ respectively, with each band exhibiting a range of approximately 10^{-3} g/cm^3 .

Calculations for St_K and U_g are given in Appendix B. The order of magnitude estimate for St_K confirmed that the particles we opted to use would closely follow the streamlines of the flow. Using the average of the density peaks ρ_{p-}, ρ_{p+} and Equation (3.7), we obtain an average value for the particle-settling velocity,

$$U_g(\rho_f) = 4.9(\rho_f - 1.024) \times 10^3 \text{ mm/h}, \quad (3.8)$$

where ρ_f is measured in g/cm^3 . We return to the settling velocity shortly in our discussion on seeding the two fluid layers with particles.

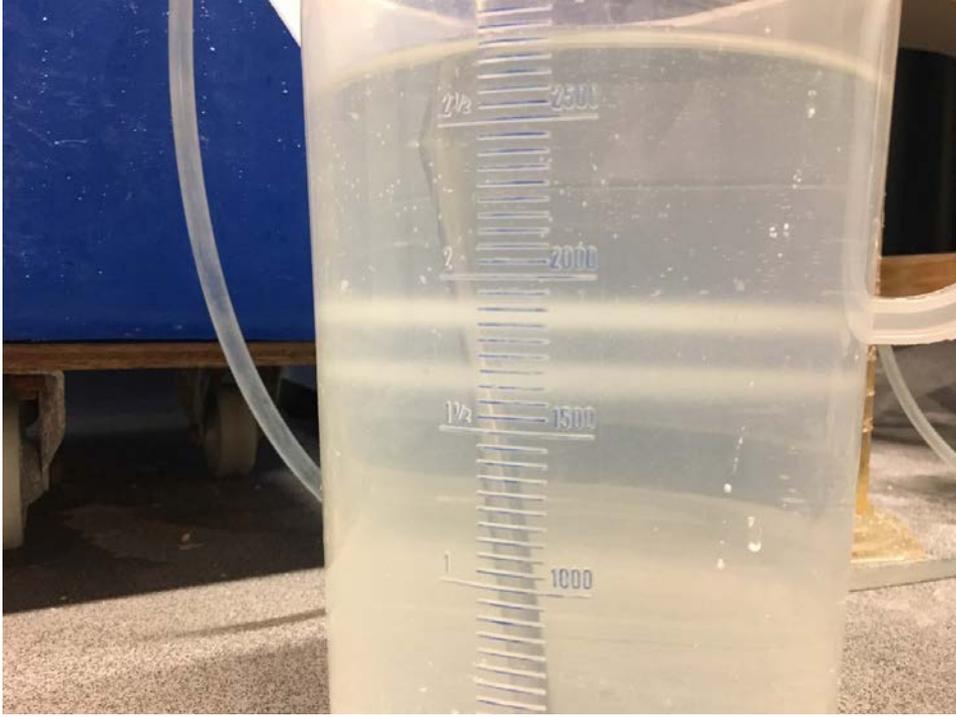


Fig. 3.11 Close-up photograph of tracer particles that have settled to their neutral-buoyancy level in a linear stratification. Particles are seen to settle in two bands, with peaks at $\rho_{p-} = 1.023 \text{ g/cm}^3$ and $\rho_{p+} = 1.025 \text{ g/cm}^3$ and each with an approximate range of 10^{-3} g/cm^3 .

Particle-Seeding

Before injecting the lower-layer fluid into the tank, we seeded each fluid reservoir with tracer particles. For each reservoir, we used a 5 ml scoop to add a precise volume of particles to a small beaker. A few millilitres of dishwasher rinse aid and some fluid from the reservoir were then added to the beaker and the mixture was well mixed, creating a suspension in which the particles did not clump together. This suspension was then poured into the reservoir and mixed thoroughly, ensuring a uniform particle-seeding density distribution in the reservoir.

Over time, the particle-settling velocity will impact the initial uniformity of the particle-seeding density distribution. For refractive-index-matched NaCl and NaNO₃ solutions of respective densities ρ_1, ρ_2 , using Equation (3.8) and the refractive-index-calibration measurements plotted in Figure 3.7, a minimum bound on the larger settling velocity in the two fluid layers can be obtained,

$$\max_{\rho_f = \rho_1, \rho_2} |U_g(\rho_f)| \gtrsim 22.5 \text{ mm/h}. \quad (3.9)$$

To retain a near-uniform particle-seeding density distribution in each refractive-index-matched layer, we prescribed the upper and lower-layer fluid densities such that $\rho_1 <$

$\rho_{p-}, \rho_{p+} < \rho_2$. In doing so, particles from both the upper and lower layer would slowly migrate and settle at the density interface. Eventually, the particle settling will lead to a dense cloud of particles at the interface, obscuring the illuminated region of the interface from the view of the cameras. This would prevent us from taking measurements at the interface and so had to be avoided.

The effect of particles settling at the interface was primarily exacerbated by the filling time required for the lower layer. We minimised the depth required of the lower layer (and in doing so, the required filling time) by positioning the cameras so the bottom of the field of view coincided with the base of the tank. In doing so, the lower-layer filling time could be reduced to less than 90 minutes (as mentioned in Section 3.1.2), which helped to mitigate the effect of particles settling at the interface.

We prescribed the upper and lower-layer fluid densities to be $\rho_1 \approx 1.0185 \text{ g/cm}^3$ and $\rho_2 \approx 1.0273 \text{ g/cm}^3$ for all our PIV/LIF experiments, corresponding to a refractive index $n = 1.3378$. This choice of fluid densities was empirically determined to be suitable with regards to the particle-settling issue. We opted to prescribe ρ_2 to be closer to the particle densities than ρ_1 as the entire vertical extent of the lower layer was in the field of view of the camera, whereas the free surface was typically located $\sim 20 \text{ cm}$ above the top of the field of view of the cameras. Though this preference increased the accumulation of particles in the upper layer settling at the interface, this was not found to be a significant problem.

It should be noted that the particle-settling velocity could be reduced by choosing a smaller size of tracer particle. The other appropriate tracer particles we had at our disposal have diameter $d_p = 20 \mu\text{m}$ and a similar density to the larger $50 \mu\text{m}$ particles. Using the smaller particles would have reduced $|U_g|$ by over a factor of 6, effectively eliminating any concerns associated with the particle-settling velocity. However, these smaller particles had an image diameter $d_A \approx 1$ pixel with our camera and laser setup, smaller than the optimal range of diameters for accuracy of PIV measurements (see the discussion above under the heading ‘Tracer Particles’). As we were able to sufficiently mitigate the effects of particles settling at the interface with the measures described above, we opted to use the larger 50 micron particles.

In addition to the tracer particles and the layer densities and heights, the average particle-seeding density in the stratification was also an important consideration for the number of particles that would settle at the density interface by the end of the lower-layer filling process. The more tracer particles that are added to the stratification, the shorter the time taken before the interface becomes unacceptably clouded with particles. However, if too few tracer particles are added, the spatial resolution at which accurate PIV measurements could be made would have been unacceptably compromised for our purposes. A suitable

average particle-seeding density was determined through trial and error. Note that no rigorous attempt was made to optimise the average particle-seeding density. However, the average particle-seeding density we prescribed was deemed as being close to optimal, as we found increasing the seeding density further lead to the interfacial region beginning to look blurry in the view of the camera by the end of the tank-filling process.

Image quality

To improve the quality of the PIV images obtained, we made two adjustments to camera A. Firstly, to avoid pixel saturation, we adjusted the camera aperture so that the brightest light scattered by the particles corresponded to a pixel response that was just below saturation. By using the camera aperture to make full use of the 8-bit range of the camera while avoiding pixel saturation, we made the maximal increase possible for the signal to noise ratio of the scattered light that is quantised onto the sensor array. This in turn improved the performance of the pattern matching algorithms that were used for the PIV image-interrogation process.

Secondly, after focusing the camera onto the central plane of the light sheet (so that the particle images are sharp and not blurred), the particle image diameter d_A could be increased by slightly defocusing the camera. As mentioned previously in our discussion on ‘Tracer Particles’, optimal values for d_A have been reported in the range of 2 to 3 pixels for the improved performance of commonly used PIV pattern matching algorithms. We chose to slightly defocus camera A so that the typical particle image diameters observed would fall in this optimal range.

3.3.5 Laser-induced fluorescence

To track the distortion of the density interface and the development of any secondary structures during the ring–interface interaction, we used the LIF imaging technique. This technique requires the use of a fluorescent dye. When such a dye absorbs a photon at a given wavelength, it will emit a photon a short time later at a longer wavelength (corresponding to a lower energy). The dye we used is Rhodamine-6G (R6G), which emits photons at 570 nm under the 532 nm light sheet. As discussed in Section 3.3.3, a 570 nm low-pass filter was attached to Camera B, so for each LIF image, only the fluorescing dye in the camera’s field of view was visible. As the diffusivity D_{R6G} of R6G is similar to the diffusivities of the salts used for the upper and lower layers (see Table 3.1), the dye could be treated as a passive scalar in the flow. By adding R6G to one of the fluid layers in the two-layer stratification, using camera B we were able to track the movement of the dyed fluid during the ring–interface interaction.

An example LIF image taken using camera B is shown in Figure 3.12. Before starting the lower-layer filling process, we added a few millilitres of a 205 mg/l R6G solution into the upper-layer fluid and mixed it thoroughly, so that the R6G concentration in the upper layer was about 1.6×10^{-2} mg/l. Despite the uniform concentration of R6G in the upper layer, the initial light intensity profile of the upper layer observed through camera B will have spatial variations due to the thickness profile of the light sheet and imperfections at the base of the tank. In a similar way as with camera A, we adjusted the aperture of camera B so that the brightest regions of fluorescing light would correspond to just below pixel saturation, allowing us to make full use of the 8-bit digital range of the camera.

During the ring–interface interaction, as the dyed and undyed layers mix, the intensity of the fluorescing light decreases proportionally, allowing us to locate regions of mixed fluid easily. Using commonly known post-processing imaging techniques (Shan et al., 2004), it was possible to perform a calibration to associate the intensity of the fluorescing light to the fluid density at any given location on the image plane of camera B. For our purposes, the use of any such post-processing techniques was not necessary as we only needed to be able to differentiate between the upper and lower-layer fluids.

For each experiment at fixed (Ri, θ_0) , the laser was turned on 30 seconds before we started recording, allowing it to warm up, and turned off immediately after recording a sequence of LIF images. By having the laser turned on only when we were conducting experiments, we were able to avoid possible difficulties associated with LIF, such as photobleaching or quenching.

3.4 Vertical profiles

For the periodic mixing experiments presented in Chapter 6, a conductivity probe and a thermistor were used to obtain vertical density and temperature profiles of the evolving stratification. In this section, we go through the design and setup of the apparatus used to make these measurements, and explain our methodology for calibrating the probe and thermistor.

3.4.1 Apparatus and setup

For the experiments where a conductivity probe and thermistor were used, NaCl solutions were used to form the upper and lower layers of the stratification. As sodium chloride is an electrolyte, when it is dissolved into fresh water it produces an electrically-conducting solution. This is because the sodium and chloride ions in the fluid, which carry positive and negative charge respectively, are disassociated, allowing for their free movement from



Fig. 3.12 LIF snapshot of the ring–interface interaction for $(\text{Ri}_0, \theta_0) = (1.18, 15^\circ)$ during the transition to turbulence. The fluorescing dye clearly marks large disconnected regions of upper and lower layer fluid in the lower layer indicate significant out-of-plane motion. The instability structures in convectively unstable regions of the flow are also clearly visible, illustrating the presence of small length scales and high in-plane velocity gradients.

one another. For the range of salt concentrations used in our experiments, the electrical conductivity of the solution is a strictly increasing function of the salt concentration, for the full range of temperatures seen in the two-layer stratification. Therefore, by measuring the electrical conductivity and temperature of the fluid, we could uniquely determine its density.

The electrical conductivity of the fluid was measured using a conductivity probe. The probe used for this thesis was designed by Stuart Dalziel in 1995, and has since been used for a number of publications (e.g. Holford and Linden (1999), Davies Wykes and Dalziel (2014), Olsthoorn and Dalziel (2015)). Figure 3.13 shows a schematic of the conductivity probe. The probe consists of two concentric cylindrical electrodes that are separated by an electrically insulating material. The inner electrode has an internal diameter $D_I = 2.3$ mm, and the outer electrode has an external diameter $D_O = 8.0$ mm. Attached to the bottom of the probe is a tip, made out of an acetal resin. At the end of the tip is a very narrow opening, of diameter 0.3mm, allowing fluid to enter the interior of the probe.

The top of each electrode is connected to a bridge circuit. When an electrically-conducting solution passes through the probe tip opening into the interior of the probe, a complete electrical circuit is established between the two electrodes. Due to the narrow size of the opening, the resistivity of the circuit is dominated by the fluid passing through the probe tip opening. Therefore, as the electrical conductivity of the fluid is equal to the inverse of its resistivity, the conductivity of the circuit, which is measured from the bridge circuit, is approximately equal to the conductivity of the fluid passing through the probe tip opening.

The temperature of the fluid was measured using a P25 NTC thermistor (General Electric). We chose to use this thermistor as its response time when plunged into water was sufficiently short (~ 23 ms) that we could treat the temperature measurements as instantaneously taken. Another advantage of using this thermistor was that, as it is hermitically sealed, its calibration remained stable over a long period of time. The thermistor was fixed to the narrow end of an 88 cm lightweight tapered carbon boom (Hyperflight), varying from 5 mm to 10 mm across its length. The carbon boom was chosen to carry the thermistor as it produced negligible mixing when translated through the fluid stratification, because of its slender geometry. Note that the probe similarly produced negligible mixing when translated through the fluid stratification.

The probe and carbon boom were secured to a carriage that held them vertically, at a horizontal distance of 7 cm apart. The probe and boom were held such that the tip of the conductivity probe and the thermistor were at the same vertical position, so that they were sampling fluid from the same horizontal plane at any instance in time. The carriage was attached to a traverse that was powered by a stepper motor, allowing the probe and thermistor to be moved vertically through the stratification. The stepper motor was controlled in a

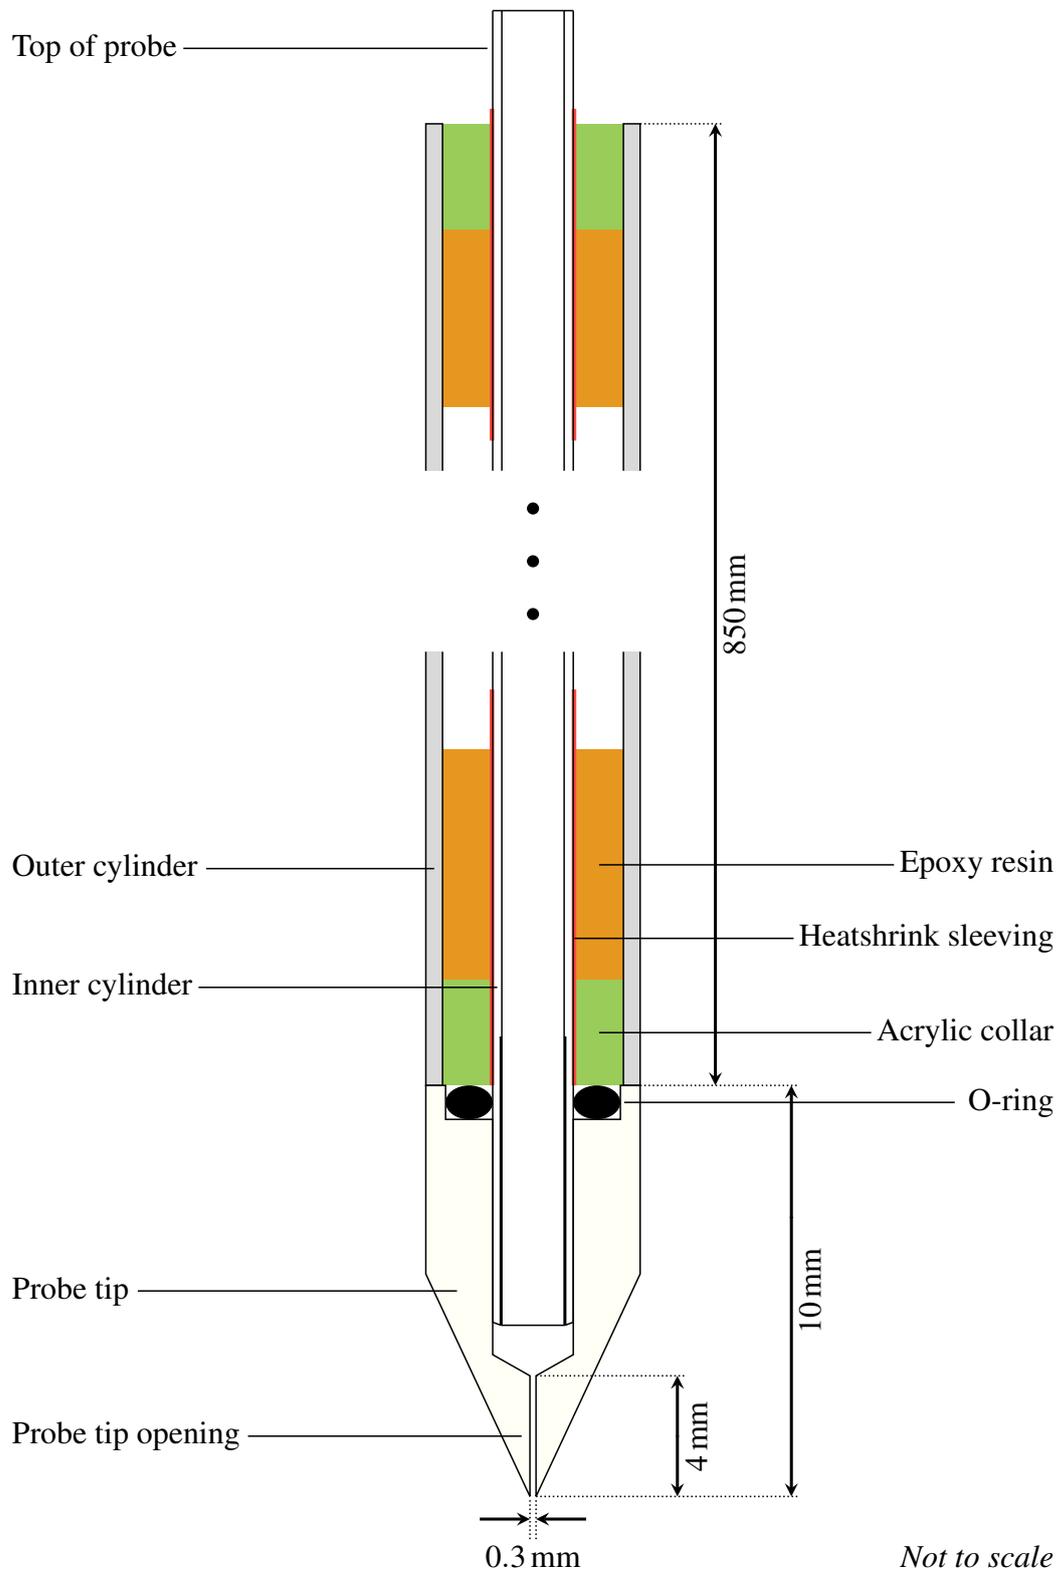


Fig. 3.13 Schematic of the conductivity probe used in this thesis.

similar way to the stepper motor used for the bike-pump traverse (discussed in Section 3.2). This allowed for precise movement of the probe and thermistor through the stratification.

As mentioned earlier, we could treat the temperature measurements from the thermistor as being instantaneously taken as it moved through the stratification. This was also possible for the conductivity measurements from the probe. The probe samples the stratification by continuously drawing fluid through the probe tip as the probe traverses. Fluid is drawn through the probe by a hydrostatic pressure head. To create the pressure head, flexible tubing was attached to the top of the probe and a syringe was used at the tube outlet to draw out the air inside the probe and tubing. After removing the syringe, provided the tube outlet was sufficiently below the free surface of the stratification, the resulting pressure head would allow the probe to draw in and sample the fluid continuously.

A portion of the tubing near the tube outlet was fed through a digitally controlled pinch valve, so siphoning of the fluid by the probe could effectively be turned on and off. During experiments, the pinch valve was positioned roughly 1.5 m below the free surface. This corresponded to a flow rate of about 0.1 – 0.2 ml/s through the probe, and a fluid velocity of about 2 m/s through the probe tip opening, which was large enough for measurements from the conductivity probe to be treated as instantaneously taken.

Note that additional details regarding the use of the conductivity probe and thermistor for our periodic mixing experiment are discussed in Section 6.1. Additionally, a detailed analysis of the conductivity probe is given in Appendix C.

3.4.2 Calibration

Calibration of the conductivity probe and thermistor was required so we could obtain density measurements from our experiments. This involved two separate calibration procedures. Firstly, a calibration of the thermistor so that the thermistor readout voltage, T_V , could be related to the fluid temperature T . Secondly, a joint calibration of the probe and thermistor so that the probe readout voltage C_V , together with T_V , could be related to the fluid density when measured at 20°C (which we denote as ρ_{20}). By carrying out these two calibration procedures, we were able to relate the readout voltages C_V and T_V to the fluid density ρ when sampling the stratification.

The methodology for the two calibration procedures were similar. A 1600 ml fluid sample was prepared in a beaker and left in the refrigerator until it reached the lowest temperature that was required for the given calibration being carried out. After this, the sample was placed on a magnetic plate (Variomag Maxi Series, Thermo Scientific) that was positioned beneath the probe and thermistor, and a magnetic puck was put in the fluid sample. The

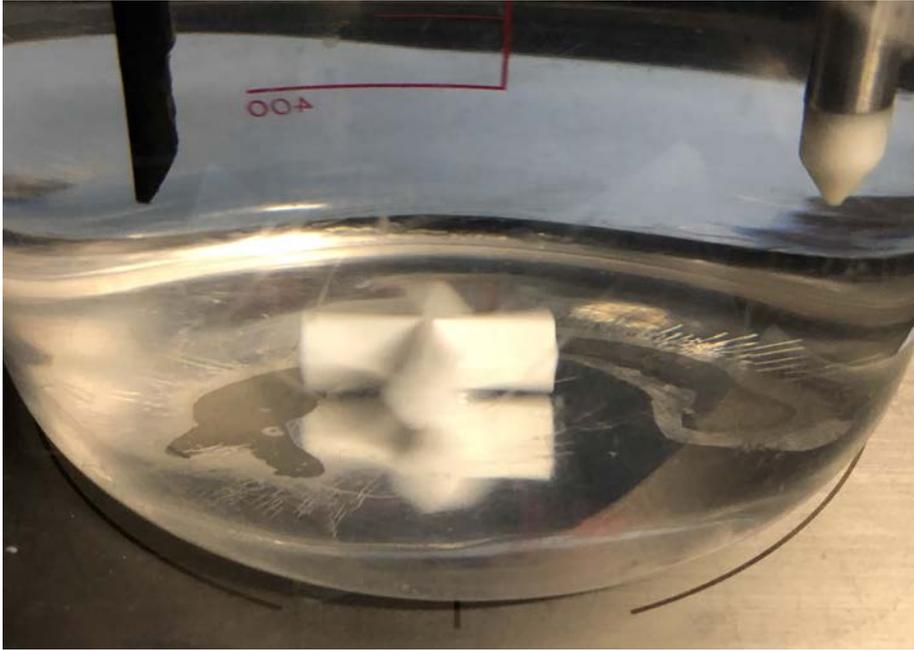


Fig. 3.14 Close-up photograph of the probe tip and thermistor submerged in a fluid sample with the magnetic puck stirring the fluid.

probe and thermistor were then lowered into the fluid sample and the magnetic plate was turned on. This caused the magnetic puck to spin, keeping the fluid sample well mixed.

Figure 3.14 shows a close up photograph of the probe tip and thermistor submerged in a fluid sample with the magnetic puck stirring the fluid. After the magnetic plate was turned on, we began sampling the fluid. The fluid was sampled at a rate of 2048 Hz for 3 seconds, every 45 seconds. For each 3 second sampling period, the mean voltage from the probe and thermistor were recorded (these were the values we denote as C_V , T_V respectively). The sample was slowly heated as the motor actuating the puck warmed up. As this happened, the puck ensured the temperature of the sample remained homogeneous. The motor provided enough heat to the fluid sample for it to go through the full range of temperatures we required for our calibrations.

To calibrate the thermistor, a pair of pre-calibrated thermocouples were used to measure the fluid temperature T during each sampling period. Figure 3.15 (a) plots the thermocouple measurements against the corresponding thermistor voltage. The thermistor voltage ranged from 0 V to 10 V, which was set to correspond to temperatures ranging between 17°C and 25°C, covering the full range of temperatures recorded during our experiments. A quadratic curve was fitted to the data using a least squares procedure, which was subsequently used to calculate the temperature from the thermistor voltage, $T = T(T_V)$.

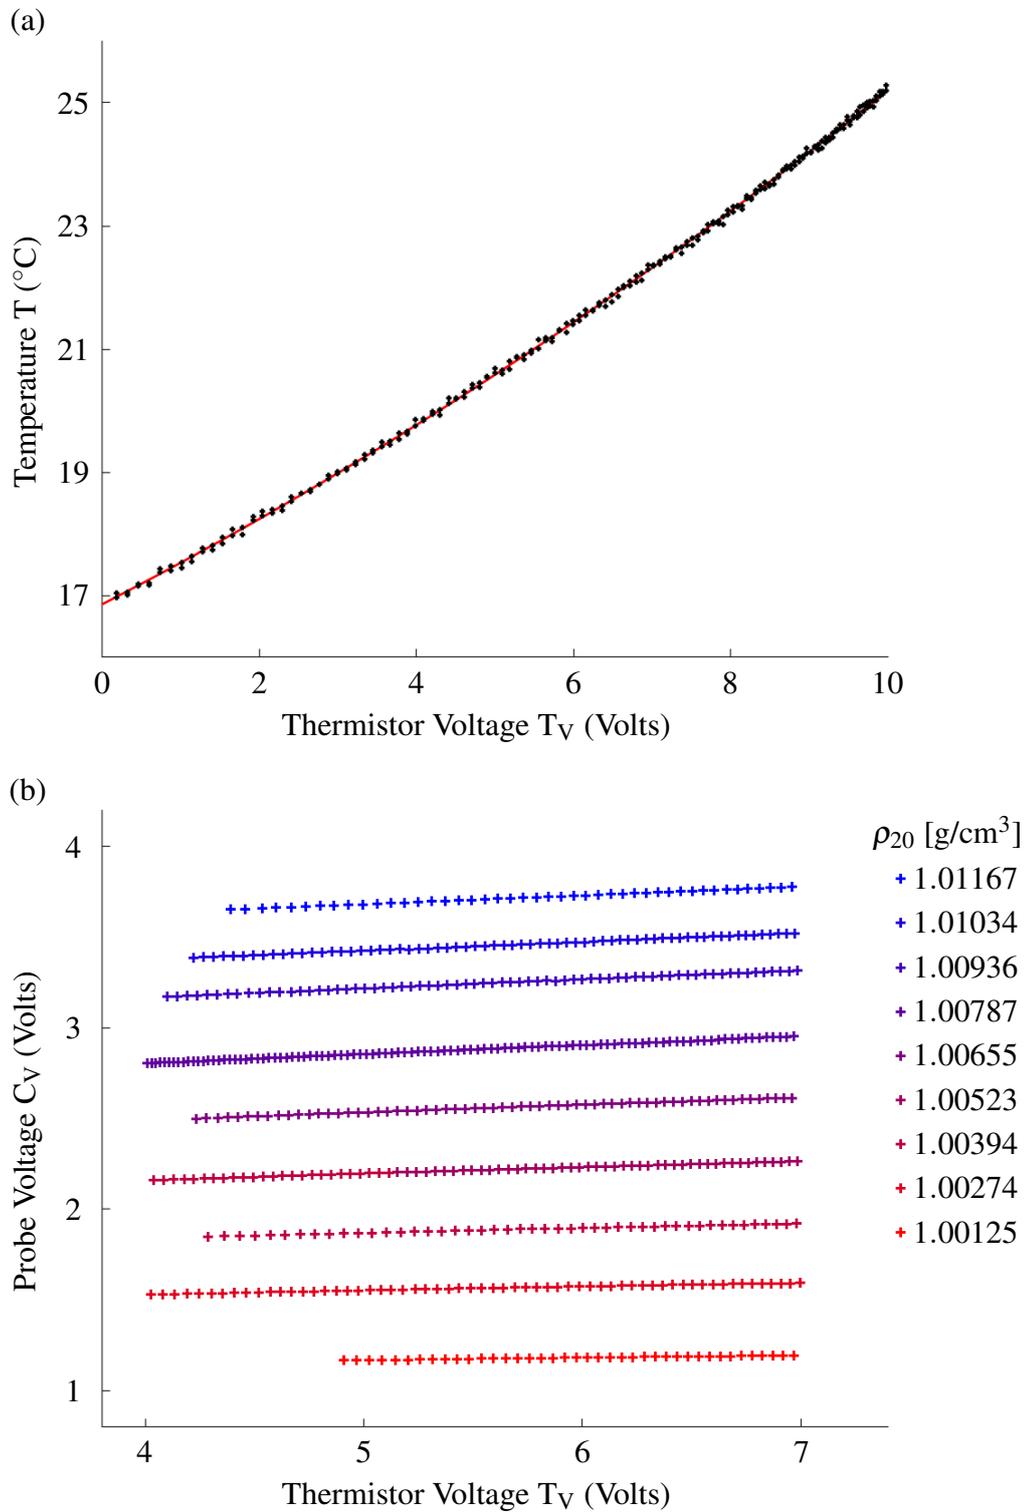


Fig. 3.15 (a) Calibration curve relating the thermistor voltage output to the temperature of the sampled fluid. (b) Example set of the data points taken for a probe/thermistor calibration.

For each set of experiments that were conducted, a joint calibration of the probe and thermistor was required to relate C_V and T_V to ρ_{20} . Associated with each set of experiments was a fixed prescription of the upper and lower-layer fluid densities. As the probe voltage was quantised and limited to a maximum value (after which the voltage would start to clip), it was desirable for the range of measured probe voltages to span as large a range of measurable voltages as possible, thereby maximising the signal to noise ratio of the probe measurements. This was achieved by making use of a rotary switch on the bridge circuit box, which provided a choice of different resistors to connect to the electrical circuit of the probe. For each set of experiments, the resistor was chosen which maximised the signal to noise ratio of the probe measurements.

Another consideration before beginning the probe/thermistor calibration was to ensure the probe tip material was fully saturated with water. When dry and first submerged in water, the probe tip would start absorbing water, leading to a slight change in its geometry. This was capable of having a significant effect on readings from the conductivity probe, because the conductivity of the circuit is sensitive to changes in the geometry of the probe tip opening. By ensuring the probe tip was already saturated with fresh water before calibration, we could be confident that the effect of water absorption by the probe tip material would not impact our probe measurements. Moreover, after calibration, it was necessary to keep the probe tip constantly submerged in water so the calibration would remain valid for the experiments to be carried out. When it was necessary to remove the probe tip from being submerged in water, this was done for no longer than a minute, which was a sufficiently short amount of time that the probe tip remained fully saturated.

For each probe/thermistor calibration, nine fluid samples of different densities were used, following the procedure outlined above using the magnetic stirrer. Figure 3.15 (b) plots the probe voltage against the thermistor voltage for the full set of data points taken from one of our calibrations. The range of densities of the samples was chosen to cover the prescribed upper and lower-layer densities for the associated set of experiments. The range of thermistor voltages calibrated for covered the range of temperatures measured during experiments.

Similarly for the thermistor calibration, a polynomial surface was fitted to the data to calculate the fluid density at 20°C from the probe and thermistor voltages, C_V , T_V . The polynomial fitting surfaces obtained from our calibrations were of the form

$$T(T_V) = \sum_{k=0}^2 a_k T_V^k, \quad \rho_{20}(C_V, T_V) = \sum_{k=0}^1 \sum_{m=0}^4 b_{k,m} T_V^k C_V^m. \quad (3.10)$$

The order of the fitting polynomials were chosen as they were of low enough order so as to not introduce high wavelength deviations in the fitting surfaces from the calibration data,

while providing a significantly smaller rms error compared to fitting surfaces of lower orders. For example, we obtained an rms error of $2.2 \times 10^{-5} \text{ g/cm}^3$ for the calibration corresponding to Figure 3.15 (b), which was sufficiently small for our purposes. By determining T and ρ_{20} for a given fluid sample, it was then possible to determine the density of the fluid. Across all probe/thermistors calibrations carried out for our experiments, density measurements were made with a precision error of less than $1 \times 10^{-4} \text{ g/cm}^3$, which was sufficiently small for our purposes.

3.5 Summary

The experimental apparatus, setup and methodology presented in this chapter was used to carry out the experiments presented in Chapters 5 and 6. We discuss the experimental setup and procedure for our periodic-mixing experiments separately in Section 6.1.1. For our PIV/LIF experiments, an algorithm was developed to increase the dynamic range over which accurate velocity measurements could be made. We now move on to introduce this algorithm in Chapter 4 and explain how we implement the algorithm to process our PIV data.

Chapter 4

A Variation of Multi-Frame PIV for Localised Turbulence

As discussed in Section 3.3, Particle Image Velocimetry (PIV) was used to obtain near-instantaneous velocity data as the vortex ring interacts with the density stratification. In keeping with the discussion of our camera setup in Section 3.3.3, we use the term ‘PIV images’ to refer to the photographs taken by a camera that captures the light scattered by particles in the illuminated volume. We also use the terms ‘particle projection’, for the light scattered by a particle onto the camera sensor, and ‘particle image’, for the individual digitised and quantised particle projections on the sensor array (defined in Section 3.3.3).

The processing of a sequence of PIV images can be split into two components. Firstly, a pattern-matching algorithm is used to compare pairs of PIV images to produce pixel-displacement measurements. Note that in this chapter we use the term ‘image-pair interrogation’ to refer to this algorithm being applied to a given pair of PIV images, and the compared pair of PIV images is referred to as being ‘interrogated’. Secondly, an algorithm that makes use of the pixel-displacement measurements is used to construct velocity fields of the flow. For our PIV processing, we make use of the pattern-matching algorithms available with DigiFlow, which are discussed in detail of Section 5.6.4.3 of Dalziel (2017).

In this chapter, we introduce the multi-frame PIV algorithm we have developed, which makes use of the pattern-matching capabilities available with DigiFlow to increase the dynamic velocity range in which a low relative measurement error can be obtained for the constructed velocity fields in all regions of the flow. In Section 4.1, we introduce the concept of a ‘multi-frame’ PIV algorithm, provide a short review of previously published algorithms that make use of this idea and motivate the need for our particular algorithm. Then, in Section 4.2 we introduce the input parameters associated with each image-pair interrogation that we vary when using our multi-frame algorithm. We introduce the multi-frame algorithm we

have developed in Section 4.3. Then, in Section 4.4 we detail our particular implementation of our multi-frame algorithm used for our experiments. Finally, in Section 4.5 we discuss the advantage made through our multi-frame algorithm relative to the in-built functionality available within DigiFlow, and discuss possible ways our algorithm could be developed further.

4.1 Introduction

We define the ‘image-pair time step’, ΔT , as the separation interval between an interrogated pair of images. For ‘multi-frame’ PIV algorithms such as the one presented in this chapter, the principal idea is to make use of measurements made from interrogating pairs of PIV images separated by several different image-pair time steps $\Delta T = \Delta T^{(k)}$ apart ($1 \leq k \leq K$, $\Delta T^{(m)} > \Delta T^{(m+1)}$) when constructing velocity fields of the flow. At each instance in time $T^{(n)}$ where a velocity field is to be constructed, calculations are performed on each interrogation window to optimise $k = k_{ij}$ locally for the local flow and particle imaging conditions (i and j are subscripts associating the image-pair-time-step index k_{ij} to a particular interrogation window). In doing so, measurement errors associated with tracking the trajectory of particle images can be reduced significantly, thus allowing velocity measurements to be made that more accurately represent the velocity field of the imaged flow.

Though the concept of a multi-frame PIV algorithm is not new, it has received relatively little attention in the literature. The earliest published study on the subject (Pereira et al. (2004)) proposes an algorithm that uses local values for the signal to noise ratio $SNR_{ij}^{(n-1)}$ and optimal $k_{ij}^{(n-1)}$ from the velocity field constructed at time $T^{(n-1)}$ to predict $SNR_{ij}^{(n)}$ and the optimal $k_{ij}^{(n)}$ for each interrogation window at time $T^{(n)}$. Local predictions for $k_{ij}^{(n)}$ are then adjusted using a set of validation criteria based on the predicted values for $SNR_{ij}^{(n)}$ and $k_{ij}^{(n)}$, before obtaining the final values for $k_{ij}^{(n)}$ that are used to construct the velocity field at time $T^{(n)}$. The last validation criterion applied requires a linear variation between the displacement over the initially predicted time step and the criteria-adjusted time step. Implicit in this criterion is the assumption that the motion of particles is not significantly influenced by flow dynamics of second or higher orders in time, over the range of possible optimal time steps being considered. The viability of this assumption is dependent on the frame rate of the camera as well as the dynamics of the flow.

A later study by Hain and Kähler (2007) argues that, to obtain a minimum measurement error, the validation criteria should consider the curvature and acceleration of the particle

images in addition to the correlation peak height and signal to noise ratio, so that second order dynamics of the flow can be taken into account. To construct a velocity field at time $T^{(n)}$ using the images taken at times $T^{(n\pm X)}$ (where the $T^{(i)}$ are equidistantly spaced apart in time), the objective function

$$\left(\frac{\sin(X\alpha)}{X\alpha} - 1\right)^2 + \left(\frac{\sigma_x}{2r_\alpha \sin(X\alpha)}\right)^2,$$

is minimised for each interrogation window to determine the local optimum choice of X . Here, α and r_α are the degree and radius of curvature of the trajectory of the group of particle images local to the interrogation window, and σ_x is the measurement error associated with the precise positioning of the particle images. Determination of optimal X is made iteratively. Local determination of α and r_α at time T_n is made by using the longest valid displacement vector (made interrogating images at times $T_{n\pm M}$ for maximum valid M), and the vector measured interrogating images at T_{n-M} and T_n , assuming the trajectory of the particle images is circular between T_{n-M} and T_{n+M} . The inaccuracy in determining α and r_α is high, so a 3×3 median filter is applied to these parameters to reduce the level of noise.

Similar to the multi-frame algorithm concept, Persoons and O'Donovan (2011) propose a 'multi-pulse' algorithm, where a series of double exposure images are recorded at N different pulse separations, with the successive images recorded cycling through the pulse separations. Then, conventional multi-grid algorithms are used on each image to process the velocity fields, after which composite velocity fields are constructed with a locally optimised pulse separation. When the camera frame rate is significantly limited compared to the repetition rate of the lasers, the highest speeds at which a multi-pulse algorithm could be used to obtain accurate velocity measurements is significantly higher than the highest resolvable speeds if a multi-frame algorithm were used. However, if there are no significant limitations in the camera and laser technology used, multi-frame algorithms are arguably more flexible than multi-pulse algorithms, as multi-frame algorithms are able to draw from several different image-pair interrogations to construct a single velocity field, whereas multi-pulse algorithms are more restrictive as measurements can only be made by interrogating images against themselves.

For our PIV/LIF experiments, ideally we would have been able to obtain fully three-dimensional near-instantaneous velocity and density data of the ring-interface interaction. To do this, the recent plane-scanning technology developed for 3D PIV/LIF by Partridge et al. (2019) would have been suitable. However, at present, the technology available in our laboratory is incapable of scanning volumes at a fast enough rate to be able to reconstruct near-instantaneous three-dimensional velocity and density fields of the ring-

interface interaction. Moreover, no such technology has been developed anywhere else, to the best of our knowledge. We therefore opted to take two-dimensional planar PIV and LIF measurements of the flow instead.

Due to the limited repetition rate of the lasers, to obtain low measurement errors in the fastest regions of the flow and regions of high vorticity, it was necessary to introduce a time offset δt between the i th pulses of the first and second laser that was smaller than the inverse frame rate of the camera. When interrogating pairs of images separated in time by δt , the measurement error relative to the particle image displacement can be higher in slower-moving regions of the flow, where the displacement is comparable to the precision error of the sub-pixel and window-deformation techniques used in the pattern-matching algorithm. Reducing the relative measurement error in these slower moving regions is desirable so that the PIV measurements obtained in these regions perform better at characterising the flow.

To obtain a low relative measurement error in all regions of the flow, we have developed a multi-frame PIV algorithm to process our PIV images. The algorithms developed by Pereira et al. (2004) and Hain and Kähler (2007) we deemed as being unsuitable for our purposes as they implicitly assume a range of suitable image-pair time steps exist for all regions of the flow. In our case, we are restricted in the choice of image pairs we interrogate in the fastest moving regions of the flow, and regions of high vorticity. Associated with this restriction are additional considerations for constructing the final velocity fields, which are not accounted for by Pereira et al. (2004) and Hain and Kähler (2007) as they do not pose a problem for their cases. Also, the algorithm by Persoons and O'Donovan (2011) is not suitable as it requires the timescale over which the Eulerian velocities vary to be sufficiently smaller than the image-pair time steps, which is not the case for our experiments.

The algorithm we have developed assumes the use of pattern-matching algorithms capable of making high-precision velocity measurements, such as those available with DigiFlow. The algorithm is designed for flows where there are localised regions of turbulence in an otherwise laminar flow, such as the ring-interface interactions considered in this thesis. After introducing the notation and definitions required to explain our multi-frame algorithm, separate discussions are given on the generalised algorithm and our particular implementation of it. We conclude this chapter with a further discussion on how our multi-frame algorithm could be extended further to increase its functionality.

4.2 Image pair interrogation

Displacement measurements of the particle projections on the plane of the camera sensor are made by interrogating pairs of PIV images. We denote the ordered sequence of PIV

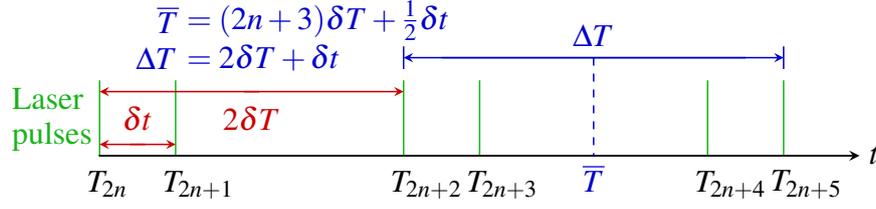


Fig. 4.1 Diagram illustrating how the parameters \bar{T} , ΔT associated with an image pair $\{\mathcal{I}_i, \mathcal{I}_j\}$ are defined, where image \mathcal{I}_k is a snapshot of the flow at time T_k when the laser pulses. In the illustrated case, $(i, j) = (2n + 2, 2n + 5)$.

images taken during an experiment as

$$\Phi(\delta T, \delta t) = \{\mathcal{I}_0, \mathcal{I}_1, \dots\},$$

where δT , δt are associated with the laser pulse timings initially defined in Section 3.3.3. The k th PIV image \mathcal{I}_k captures a snapshot of the projected positions of particles in the light sheet, illuminated by the k th laser pulse, which fires at time

$$T_k = 2 \lfloor k/2 \rfloor \delta T + (k \bmod 2) \delta t, \quad (4.1)$$

where $\lfloor \cdot \rfloor$ denotes the floor function, outputting the largest integer less than or equal to the input value.

Pixel-displacement measurements are made in the \mathbf{X}_A coordinate system of camera A by interrogating the movement of particle images between pairs of PIV images $\{\mathcal{I}_i, \mathcal{I}_j\}$. The parameters \bar{T} and ΔT , defined by

$$\bar{T} = \frac{1}{2} (T_i + T_j), \quad \Delta T = T_j - T_i,$$

uniquely identify the pair of PIV images that are being interrogated. The relationship of \bar{T} and ΔT with respect to the laser pulse timings δT and δt is illustrated in Figure 4.1.

4.2.1 Measurement error of a single particle trajectory

To develop an understanding for how the choice of ΔT influences the accuracy and precision of velocity measurements, consider first how the velocity of a single particle can be estimated by sampling its position at two instances in time. For a particle with a trajectory $\mathbf{r}_p(t)$ that is sampled at times $T_{\pm} = \bar{T} \pm \Delta T/2$, its actual displacement $\delta \mathbf{r}_p$ is related to the

velocity $d\mathbf{r}_p/dt$ at time \bar{T} , which is made explicit using a central difference approximation,

$$\begin{aligned}
\delta\mathbf{r}_p &= \mathbf{r}_p(T_+) - \mathbf{r}_p(T_-) \\
&= \sum_{n \text{ odd}} \frac{(\Delta T)^n}{2^{n-1}n!} \left. \frac{d^n \mathbf{r}_p}{dt^n} \right|_{t=\bar{T}} \\
&= \Delta T \left. \frac{d\mathbf{r}_p}{dt} \right|_{t=\bar{T}} + O(\Delta T^3) \\
\Rightarrow \frac{\delta\mathbf{r}_p}{\Delta T} &= \left. \frac{d\mathbf{r}_p}{dt} \right|_{t=\bar{T}} + \boldsymbol{\varepsilon} \Delta T^2, \tag{4.2}
\end{aligned}$$

where $\boldsymbol{\varepsilon} = O(1)$. Equation (4.2) demonstrates that, provided $|\boldsymbol{\varepsilon} \Delta T^2| \ll |d\mathbf{r}_p/dt|$ at time \bar{T} , an accurate estimate of the particle velocity can be obtained using this sampling strategy. Clearly, as $\Delta T \rightarrow 0$, Equation (4.2) becomes an equivalence, thus choosing a sufficiently small sampling interval ΔT will guarantee an accurate estimate of the particle velocity.

For PIV, velocity measurements of particles illuminated by the light sheet are made by interrogating the movement of groups of particle images on the camera A sensor. The coordinate system calibration described in Section 3.3.3 associates the particle projection positions on the sensor to the particle positions in the light sheet. The light scattered from an illuminated particle at position $\mathbf{r}_p(t)$ is projected onto the sensor, centred at a position $\mathbf{X}_p(t)$ with light intensity profile $I_p(|\mathbf{X} - \mathbf{X}_p(t)|)$. Each projection is then spatially discretised onto the sensor array and the light intensity is quantised by the bit depth of the camera, giving the particle images we see in the PIV images obtained. This discretisation and quantisation limits the highest precision attainable for the velocity measurements obtained due to the uncertainty in locating the precise positions of particle images on the camera sensor. This precision is limited further both by the fill factor of the camera sensor (the fraction of each pixel that is photosensitive), which is commonly around 50% for many cameras, and random electrical noise associated with thermal fluctuations and variability in light sensitivity between pixels.

Consider again the trajectory $\mathbf{r}_p(t)$ of a single particle in the flow. Writing $\mathbf{R}_p(t)$ for the estimate of $\mathbf{X}_p(t)$ made using the discretised and quantised information available on the sensor array, we can write the trajectory of the particle projection as

$$\mathbf{X}_p(t) = \mathbf{R}_p(t) + \boldsymbol{\eta}_p, \tag{4.3}$$

where $\boldsymbol{\eta}_p$ is the error associated with precisely locating the particle projection, arising due to the discretisation and quantisation of the particle projection onto the sensor array. We assume that $\mathbf{R}_p(t)$ is accurate to the nearest pixel, ensuring $|\boldsymbol{\eta}_p| < 1$ px, and that $\boldsymbol{\eta}_p$ is a random variable, with mean $\overline{\boldsymbol{\eta}_p} = \mathbf{0}$ and variance $\boldsymbol{\sigma}_{\boldsymbol{\eta}_p}^2 = \overline{\boldsymbol{\eta}_p^2}$.

Suppose we estimate the velocity of the projection of a single particle on the camera sensor by sampling its position at times $t = T_{\pm}$ as before. Writing $\delta\mathbf{R}_p(\Delta T)$ for the pixel-displacement measurement and \mathbf{U}_p for the particle projection velocity, using Equation (4.3) we have

$$\begin{aligned}\delta\mathbf{R}_p(\Delta T) &= \mathbf{R}_p(T_+) - \mathbf{R}_p(T_-) \\ &= \mathbf{X}_p(T_+) - \mathbf{X}_p(T_-) + \sqrt{2}\boldsymbol{\sigma}_{\eta_p} \\ &= \Delta T \left. \frac{d\mathbf{X}_p}{dt} \right|_{t=\bar{T}} + \sqrt{2}\boldsymbol{\sigma}_{\eta_p} + O(\Delta T^3) \\ \Rightarrow \frac{\delta\mathbf{R}_p(\Delta T)}{\Delta T} &= \mathbf{U}_p|_{t=\bar{T}} + \frac{\sqrt{2}\boldsymbol{\sigma}_{\eta_p}}{\Delta T} + \boldsymbol{\epsilon}\Delta T^2 + O(\Delta T^3),\end{aligned}\quad (4.4)$$

where $\boldsymbol{\epsilon} = O(1)$ is defined in a similar way as for Equation (4.2). The left hand side of Equation (4.4) is the measurement made for the particle projection velocity, which is given in precise form by the first term on the right hand side of Equation (4.4). The second term on the right hand side of Equation (4.4) represents the measurement error associated with the precision limits of the camera, and the last two terms are associated with the nonlinear behaviour of the flow.

From Equation (4.4), the relative measurement error (RME) for estimating the velocity of the projection of a single particle on the camera sensor is given by

$$\text{RME} = \frac{\sqrt{2}|\boldsymbol{\sigma}_{\eta_p}|/\Delta T + |\boldsymbol{\epsilon}|\Delta T^2}{|\mathbf{U}_p|_{t=\bar{T}}} + O\left(\frac{\Delta T^3}{|\mathbf{U}_p|_{t=\bar{T}}}\right). \quad (4.5)$$

The expression given for RME clearly demonstrates the importance in choice of ΔT for obtaining a small relative measurement error. As $\Delta T \rightarrow 0$, the precision error dominates the measurement made for the particle projection displacement. For larger ΔT , the truncation error $\boldsymbol{\epsilon}$ may become significant relative to the projection displacement measurement. Additionally, $\boldsymbol{\sigma}_{\eta_p}$ might be expected to increase with ΔT , especially if out-of-plane velocity or in-plane divergence is large. The optimal choice of ΔT for minimising RME will depend on the local flow and optical conditions, and will therefore vary both spatially and temporally.

4.2.2 Interrogation windows

Rather than measuring the displacement of individual particle projections, PIV uses pattern-matching algorithms to produce displacement measurements by interrogating the movement of groups of particle projections in small regions on the camera sensor. For each image-pair interrogation, displacement measurements are made by first subdividing

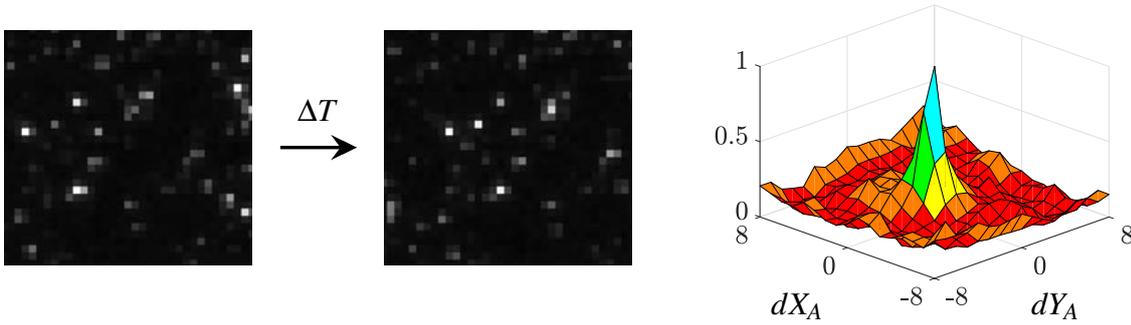


Fig. 4.2 Particle pattern-matching in two interrogation windows, separated by a time ΔT . In this example, the measured pixel displacement is $(\delta X_A, \delta Y_A) = (6, 0)$. The normalised absolute difference signal is shown on the right, where (dX_A, dY_A) is the pixel deviation from the measured pixel displacement.

the images into small, possibly overlapping regions, commonly referred to as ‘interrogation windows’. For each interrogation window, a single displacement measurement is made by using a pattern-matching function locally to determine the pixel displacement shift of the group of particle images within the interrogation window. Figure 4.2 illustrates the pattern-matching process.

Later stages of the pattern-matching algorithm determine the displacement to sub-pixel accuracy, first by spatially interpolating the PIV images to allow smaller displacement shifts to be probed, then by distorting the shapes of the interrogation windows to conform with the in-plane velocity gradients in the flow. The details of these stages are discussed in Dalziel (2017). In this chapter, our discussion on the interrogation windows is restricted to the size and distribution of windows prescribed before the pattern-matching algorithms are initiated. For our implementation of the pattern-matching algorithms, the interrogation windows are initially square, equally spaced apart and are all of the same size (with the exception of windows nearest the edge of the camera sensor), though in general this does not need to be the case. For parameters and variables associated with the interrogation windows, we use the subscript ‘ I ’ to avoid ambiguity.

Each interrogation window has size $n_I \times n_I$ in pixel units, where n_I is chosen to be an odd integer, so that each interrogation window is centred at an integer pixel position \mathbf{X}_I , hereafter referred to as an ‘interrogation point’. If an interrogation point is close to the edge of the sensor array, its corresponding interrogation window may be reduced in size if its area overlaps with the edge of the sensor array. For such cases, the interrogation point is positioned as though the interrogation window had not been reduced in size. By adopting this strategy, adjacent interrogation points are spaced equally a fixed pixel distance δ_I apart. The choice of

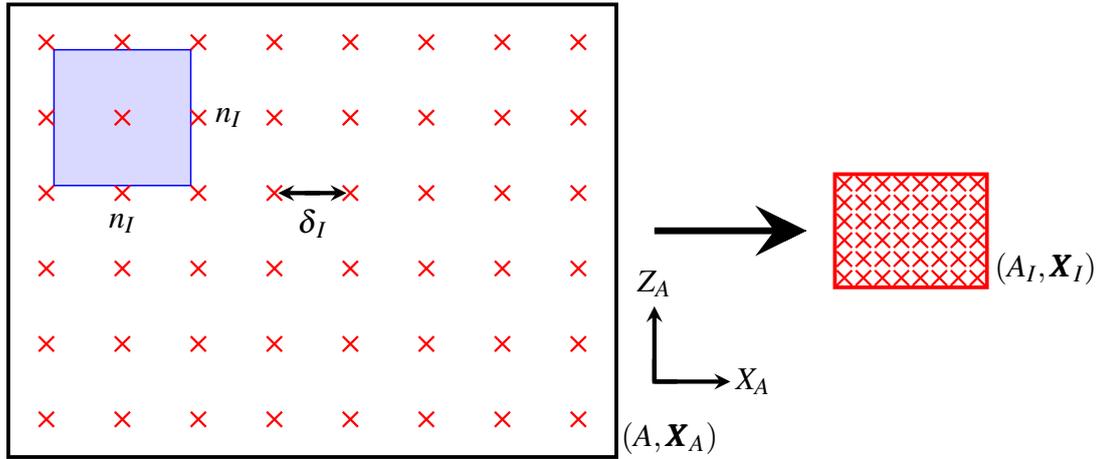


Fig. 4.3 Diagram illustrating the relationship between the sensor array A and the interrogation array A_I . Interrogation points \mathbf{X}_I , which are spaced a pixel distance δ_I apart, are marked with a red cross. A single interrogation window of size $n_I \times n_I$ pixels is shaded blue. The displacement measurements $\delta \mathbf{X}_I$ obtained from interrogating a pair of PIV images are taken to be evaluated on the interrogation array.

n_I plays a role in setting the finest spatial resolution at which accurate velocity measurements can be made, and the choice of δ_I plays a role in setting the amount of sampling of the PIV images.

The full set of interrogation points on the sensor array A form a rectangular grid of $a_{I,X} \times a_{I,Y}$ points. We define the ‘interrogation array’, A_I , as an $a_{I,X} \times a_{I,Y}$ matrix for which variables evaluated on A_I are taken to be evaluated at the pixel-valued positions of the interrogation points. In keeping with our notation convention for interrogation windows, variables evaluated on the interrogation array are labelled with a subscript I . The relationship between the sensor array of Camera A, the interrogation windows, the interrogation points and the interrogation array is illustrated in Figure 4.3.

4.2.3 Measurement error of a group of particle projections within an interrogation window

The purpose of the multi-frame algorithm we have developed is to reduce the RME of displacement measurements to an acceptably low level across all regions of the flow. By using interrogation windows to interrogate groups of particle projections rather than interrogating the movement of single particle projections, additional considerations arise that should be made when aiming to reduce the RME of a given measurement. Note that the functional form of the RME given by Equation (4.5) remains valid to a good approximation when making

displacement measurements using interrogation windows, though additional dependencies arise relating to the parameters associated with the interrogation windows.

Some of these considerations have been discussed previously in Section 3.3, which we will briefly re-iterate. As discussed in Section 3.3.2, the refractive indices of the upper and lower-layer fluids are matched to avoid losing focus of particles in the light sheet. As discussed in Section 3.3.4, the camera aperture is adjusted both to avoid pixel saturation and to make full use of the 8-bit digitised signal, maximising the signal-to-noise ratio of the light scattered by illuminated particles against any background electrical noise impacting the PIV images produced. The camera is also slightly defocused to increase the particle-image diameter d_A to improve the performance of the pattern-matching process.

As discussed in Section 3.3.4, tracer particles are prepared so that there is an initially uniform particle-seeding density distribution (though particles would gradually settle towards the density interface). Writing \mathbb{N}_p for the mean number of particles per mm^3 in the tank and N_p for the number of particle images within an interrogation window, the mean of N_p can be written in terms of the thickness of the light sheet $\Delta y_L(x, z)$ and the size n_I of interrogation windows as

$$\overline{N_p} = n_I^2 \zeta^2 \Delta y_L(x, z) \mathbb{N}_p, \quad (4.6)$$

where ζ is the mm/px scale associated with the area of the central plane of the light sheet projected onto a single camera pixel.

Assuming a uniform velocity within an interrogation window, the standard deviation σ_η of the sample mean of precision errors η for measurements made in the window will scale as

$$|\sigma_\eta| \sim N_p^{-1/2} \sim n_I^{-1}. \quad (4.7)$$

This motivates the use of larger n_I where possible. Alternative strategies for decreasing $|\sigma_\eta|$ would be to increase $\Delta y_L(x, z)$ or \mathbb{N}_p . For our experiments, increasing the light-sheet thickness was not possible due to significant out-of-plane motion existing outside the central plane of the ring-interface interaction. This would have impacted the ability to match patterns in the interrogation windows both due to particles entering and leaving the light sheet, and an increased disparity in particle trajectories projected onto a given interrogation window.

As discussed in Section 3.3.4, increasing \mathbb{N}_p significantly using $50\mu\text{m}$ particles would reduce the visibility of the light sheet at the density interface (due to an increased number of particles settling there). Alternatively, it would have been possible to increase \mathbb{N}_p by using the smaller $20\mu\text{m}$ particles available to us. However, this would have led to a reduction in the particle-image diameter d_A , which would have also impacted the ability to match patterns in the interrogation windows. Based on this consideration we opted to use $50\mu\text{m}$ particles.

The spatial velocity gradients of the particle projections on the camera sensor affect both the range of suitable n_I and δ_I . If significant spatial velocity gradients are present within a single interrogation window, then, as only a single displacement measurement is made within the window, it is not possible for the information extracted from interrogating that window alone to deduce the velocity gradients present in that window. This can considerably affect the accuracy of derived gradient properties such as the vorticity. Moreover, the ability to match patterns in the interrogation window would be reduced due to differences in particle trajectories within the same window. Later stages of the pattern-matching algorithms we implement that are part of the in-built functionality of DigiFlow can take into account in-plane velocity gradients to an extent by distorting the images to conform with the gradients of the flow. Ideally, the optical conditions of the flow will allow a choice of n_I and δ_I that are smaller than the length scale over which the curvature of particle trajectories vary significantly.

For fixed n_I , the choice of δ_I effectively sets the average number of interrogation windows covering each pixel. If $\delta_I > n_I$ then it is possible for regions of the PIV images not to be compared by the pattern-matching algorithms, leading to an unnecessary loss of information for constructing velocity fields of the flow. Conversely, taking $\delta_I \ll n_I$ can result in significant oversampling of the PIV images, which can lead to high levels of noise particularly in the gradients of the displacement measurements obtained.

Similarly to ΔT , the optimal choice of both n_I and δ_I will vary locally, both spatially and temporally. The algorithm presented in this chapter focuses only on optimising ΔT spatially and temporally, and assumes fixed values for n_I and δ_I . For all our experiments, we take $(n_I, \delta_I) = (27, 8)$, with the choice of $\delta_I = 8$ made as it was observed to be not too small so as to introduce high levels of noise in the gradient of displacement measurements, while being small enough to obtain an acceptably high spatial resolution in our measurements. The distorted passes of the pattern-matching algorithm (part of the in-built functionality of DigiFlow) carry out a local optimisation of n_I based on a measure of how well particle patterns match, with the size and the shape of the interrogation window changing on each distorted pass. Note that a natural extension of our algorithm would be to incorporate an ability to optimise locally the initial values of n_I and δ_I . Such an extension is left for future work.

4.3 Multi-frame PIV algorithm

The multi-frame algorithm we have developed makes use of varying the interrogation time step ΔT both spatially and temporally, increasing the dynamic velocity range at which

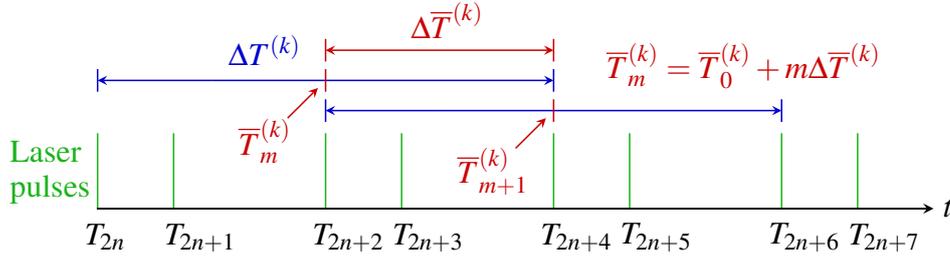


Fig. 4.4 Diagram illustrating how, for a PIV image sequence $\Phi(\delta T, \delta t)$, the parameters $\bar{T}_m^{(k)}$ and $\Delta \bar{T}^{(k)}$ are defined for $k < K$. In the illustrated case, $\Delta T^{(k)} = 4\delta T$ and $\Delta \bar{T}^{(k)} = 2\delta T$.

a low RME can be obtained. For a PIV image sequence $\Phi(\delta T, \delta t)$, ΔT is restricted to taking a discrete set of values, set by the laser pulse timings δT and δt . For our algorithm, image pairs are interrogated at K different time steps, $\Delta T^{(k)}$, where $\Delta T^{(k)} > \Delta T^{(k+1)}$ for $1 \leq k < K$. We take $\Delta T^{(K)} = \delta t$, corresponding to the smallest utilisable time step, and $\Delta T^{(1)}$ is chosen to obtain a low RME for the slowest resolvable regions of the flow. For the ordered sequence of $N^{(k)}$ image pairs $P(\bar{T}_m^{(k)}, \Delta T^{(k)})$ ($0 \leq m < N^{(k)}$) interrogated at a given time step $\Delta T^{(k)}$, we denote the time step between successive image-pair mean times as $\Delta \bar{T}^{(k)}$, so $\bar{T}_m^{(k)} = \bar{T}_0^{(k)} + m\Delta \bar{T}^{(k)}$.

Figure 4.4 illustrates the definitions of $\Delta T^{(k)}$, $\bar{T}_m^{(k)}$ and $\Delta \bar{T}^{(k)}$ for $k < K$. Apart from the image pairs interrogated utilising the smallest image-pair time step $\Delta T^{(K)}$, for each image pair, we specify that the two images correspond to laser pulses generated by the same laser (in particular, the laser that generates pulses at times T_{2m}). This decision is made to avoid possible increases in σ_η associated with a slight misalignment, or difference in light intensity, between the light sheets generated by the two lasers we use. By restricting our choice of image pairs in this way, for all $k < K$ we have $\Delta T^{(k)} = 2a_k\delta T$ and $\Delta \bar{T}^{(k)} = 2b_k\delta T$ for integers a_k, b_k .

A pseudocode for the algorithm we have developed is given by Figure 4.5. The algorithm starts by using function **A** to interrogate all specified pairs of images separated by the largest time step $\Delta T^{(1)}$, based on the input values $N^{(1)}$, $\bar{T}_0^{(1)}$ and $\Delta \bar{T}^{(1)}$. Subsequently, for each $\Delta T^{(k)}$ used, the displacement measurements $(\delta \mathbf{X}_I)_n^{(k)} = (\delta X_I, \delta Z_I)_{n,ij}^{(k)}$ are evaluated on the interrogation points specified by $J_n^{(k)}$ and at time $\bar{T}_n^{(k)}$. Here, the input parameter $J_n^{(k)}$ is an array of indicator functions of the same size as the interrogation array, specifying for function **A** the set of interrogation points where pixel-displacement measurements $(\delta \mathbf{X}_I)_n^{(k)}$ are to be produced. For the largest time step, $J_n^{(1)}$ specifies that displacement measurements are made at all interrogation points for the given n_I and δ_I to be used, where both n_I and δ_I are both fixed input parameters.

```

function MULTI-FRAME ALGORITHM( $\Phi(\delta T, \delta t), \{N^{(k)}, \Delta T^{(k)}, \bar{T}_0^{(k)}, \Delta \bar{T}^{(k)}\}_{k=1}^K$ )
  for  $k = 1$  to  $K$ 
    for  $n = 0$  to  $N^{(k)}$ 
      A :  $(\delta \mathbf{X}_I)_n^{(k)} \leftarrow \left[ P \left( \bar{T}_0^{(k)} + n \Delta \bar{T}^{(k)}, \Delta T^{(k)} \right), J_n^{(k)} \right]$ 
    end for
    if  $k < K$ 
      B :  $\{J_m^{(k+1)}\}_{m=0}^{N^{(k+1)}} \leftarrow \{(\delta \mathbf{X}_I)_m^{(k)}\}_{m=0}^{N^{(k)}}$ 
    end if
    C :  $\{(\Delta \mathbf{X}_I)_m^{(k)}\}_{m=0}^{N^{(K)}} \leftarrow \{(\delta \mathbf{X}_I)_m^{(k)}\}_{m=0}^{N^{(k)}}$ 
  end for
  for  $n = 0$  to  $N^{(K)}$ 
    D :  $\mathbf{U}_I(\bar{T}_n^{(K)}) \leftarrow \{(\Delta \mathbf{X}_I)_n^{(k)} / \Delta T^{(k)}\}_{k=1}^K$ 
    E :  $\mathbf{U}(\bar{T}_n^{(K)}) \leftarrow \mathbf{U}_I(\bar{T}_n^{(K)})$ 
    F :  $\mathbf{u}(\bar{T}_n^{(K)}) \leftarrow \mathbf{U}(\bar{T}_n^{(K)})$ 
  end for
end function

```

Fig. 4.5 A pseudocode for the multi-frame algorithm discussed in this chapter.

After producing displacement measurements $(\delta \mathbf{X}_I)_n^{(1)}$ using $\Delta T^{(1)}$, function **B** is used to determine which regions of the flow require further interrogation using the next smallest image-pair time step $\Delta T^{(2)}$. To do this, a set of criteria are used on the displacement measurements $\{(\delta \mathbf{X}_I)_m^{(1)}\}_{m=0}^{N^{(1)}}$ to determine the set of indicator arrays $\{J_m^{(2)}\}_{m=0}^{N^{(2)}}$ to be applied by function **A** when producing the displacement measurements $(\delta \mathbf{X}_I)_n^{(2)}$. Suitable interrogation criteria for determining the indicator arrays will depend on the dynamics of the imaged flow. The interrogation criteria we use for our experiments are described in Section 4.4. Also, using the displacement measurements $(\delta \mathbf{X}_I)_n^{(1)}$, function **C** uses a temporal interpolation scheme to calculate displacement measurements $\{(\Delta \mathbf{X}_I)_n^{(1)}\}_{n=0}^{N^{(K)}} = \{(\Delta \mathbf{X}_I, \Delta \mathbf{Z}_I)_n^{(1)}\}_{n=0}^{N^{(K)}}$, which are evaluated at the image-pair mean times $\bar{T}_n^{(K)}$ prescribed for interrogating images at the smallest prescribed time step $\Delta T^{(K)}$.

After using functions **A-C** with the input parameters associated with using the largest image-pair time step, this process is then repeated for the next smallest time step $\Delta T^{(2)}$, and so on until this process has been carried out for all prescribed $\Delta T^{(k)}$. Then, for each image-pair mean time $\bar{T}_n^{(K)}$, function **D** uses a set of validation criteria on the pixel velocity measurements $\{(\Delta \mathbf{X}_I)_n^{(k)} / \Delta T^{(k)}\}_{k=1}^K$ to construct pixel velocity fields $\mathbf{U}_I(\bar{T}_n^{(K)})$ defined on the interrogation array, with px/s units. These velocity fields are then spatially interpolated using function **E** to construct pixel velocity fields $\mathbf{U}(\bar{T}_n^{(K)})$ defined on the sensor array. Finally, the function **F** uses the inverse mapping \mathcal{P}_A^{-1} constructed from the coordinate system calibration (discussed in Section 3.3.3) to map the pixel velocity fields $\mathbf{U}(\bar{T}_n^{(K)})$ to

fluid velocity measurements $\mathbf{u}(\bar{T}_n^{(K)})$ defined on the central plane of the light sheet, in m/s units.

Implicit in the pseudocode given by Figure 4.5 is the choice of interrogation criteria associated with function **B**, the temporal interpolation scheme associated with function **C**, the validation criteria associated with function **D** and the spatial interpolation scheme associated with function **E**. Our implementation of each of these for our experiments will be discussed in Section 4.4.

4.4 Implementation of algorithm

For each PIV image sequence $\Phi(\delta T, \delta t)$ recorded in our experiments, the frame rate of the camera is set at 160fps (corresponding to $\delta T = (1/160)\text{s}$), and the laser-pulse offset is set at $\delta t = (1/320)\text{s}$. Then, from Equation (4.1), the laser pulse timings in seconds are given by

$$\{T_k\}_{k \geq 0} = \{0, 1/320, 4/320, 5/320, \dots\}$$

and so on. Our choice of δt is small enough to make accurate pixel-displacement measurements in regions of high vorticity, including the primary vortex cores and any secondary baroclinically generated structures. Also, our choice of δt is large enough to allow image-pair time steps to be employed that can ensure a low RME is attainable for all velocity scales exhibited in the flow.

Input parameters we prescribe for the multi-frame algorithm are given by Table 4.1. We write $(A_n, B_n)^{(k)}$ for the subscripts of the n th image pair $\{\mathcal{I}_{A_n}, \mathcal{I}_{B_n}\}^{(k)}$ interrogated at the k th image-pair time step $\Delta T^{(k)}$, taken at times T_{A_n} and T_{B_n} . Six different $\Delta T^{(k)}$ were used to select image pairs to interrogate. As explained in Section 4.3, the n th constructed velocity field is taken to be evaluated at time $\bar{T}_n^{(6)}$. For each k , we set $A_0^{(k)} = 2m$, where m is a non-negative integer determined by choosing maximal $\bar{T}_0^{(k)}$ such that $\bar{T}_0^{(k)} \leq \bar{T}_0^{(6)}$. For $k \leq 4$, our choice of $\Delta \bar{T}^{(k)}$ is made so that the time interval between two images comprising an image pair has a 50% overlap with the time interval associated with the next image pair. For $k = 5, 6$, all image pairs of the associated time step where $A_n^{(k)} = 2m$ are interrogated.

As discussed in Section 4.3, the algorithm begins by interrogating all specified image pairs separated by the largest employed time step, $\Delta T^{(1)}$, producing pixel-displacement measurements $(\delta \mathbf{X}_I)_n^{(1)}$. When using DigiFlow, for regions of the flow where velocity scales are too fast to resolve accurately using a given image-pair time step, spurious displacement vectors may be outputted. To illustrate this by means of an example, Figure 4.6 (a) plots a typical displacement array $(\delta \mathbf{X}_I)_m^{(1)}$ from one of our experiments, with Figure 4.6 (b) plotting

k	$\Delta T^{(k)}$	$\bar{T}_0^{(k)}$	$\Delta \bar{T}^{(k)}$	$(A_n, B_n)^{(k)}$	$s^{(k)}$
1	$32\delta T$	$16\delta T$	$16\delta T$	$(0, 32), (16, 48), \dots$	4
2	$16\delta T$	$16\delta T$	$8\delta T$	$(8, 24), (16, 32), \dots$	4
3	$8\delta T$	$16\delta T$	$4\delta T$	$(12, 20), (16, 24), \dots$	4
4	$4\delta T$	$16\delta T$	$2\delta T$	$(14, 18), (16, 20), \dots$	4
5	$2\delta T$	$15\delta T$	$2\delta T$	$(14, 16), (16, 18), \dots$	4
6	δt	$16\delta T + \frac{1}{2}\delta t$	$2\delta T$	$(16, 17), (18, 19), \dots$	N/A

Table 4.1 Table of the input parameters used to determine which image pairs are interrogated in the implementation of our multi-frame algorithm. The n th image pair $\{\mathcal{I}_{A_n}, \mathcal{I}_{B_n}\}^{(k)}$ interrogated at the k th image-pair time step $\Delta T^{(k)}$ is taken to be evaluated at the time $\bar{T}_n^{(k)}$. Interrogation points with pixel-distance measurements $|(\delta X_I, \delta Z_I)_{n,ij}^{(k)}| > s^{(k)}$ are interrogated further, utilising the time step $\Delta T^{(k)}$.

the associated LIF field illustrating the ring-interface interaction. Regions with spurious vectors are easily identified through the seemingly random curl of the pixel-displacement field, coinciding with the location of the ring-interface interaction.

Using the displacement measurements $(\delta \mathbf{X}_I)_n^{(1)}$, robust selection criteria are necessary to determine the full set of interrogation points that require further interrogation at smaller time steps. The simplest solution would be to interrogate the entire image plane for all time steps to be employed. However, this strategy can be computationally very expensive, both in time and memory. For our experiments, for each $\bar{T}_n^{(1)}$, interrogation points (i, j) are identified with pixel-distance measurements $|(\delta X_I, \delta Z_I)_{n,ij}^{(1)}| > s^{(1)}$, where $s^{(k)}$ is the maximum distance threshold associated with the time step $\Delta T^{(k)}$ (for simplicity, we take $s^{(k)} = 4$ for $k \leq 5$). The smallest rectangle covering these interrogation points is determined. Then, this rectangle is expanded by three interrogation points in each direction, and only the interrogation points in the resulting rectangle are interrogated using smaller time steps. This interrogation strategy ensures that slow moving regions outside the bounding rectangle are only interrogated at the largest time step, thus greatly reducing the computational time associated with interrogating image pairs at the smaller time steps.

Included in Figure 4.6 (a) is a drawing of the algorithm-produced rectangle bounding the interrogation points identified for further interrogation using smaller time steps. The minimum and maximum X_A and Z_A coordinates for each displacement array $(\delta \mathbf{X}_I)_n^{(1)}$ determine the indicator arrays $J_n^{(2)}$ identifying the regions to be interrogated using the time step $\Delta T^{(2)}$. The process of interrogating image pairs at a given time step, identifying regions for further

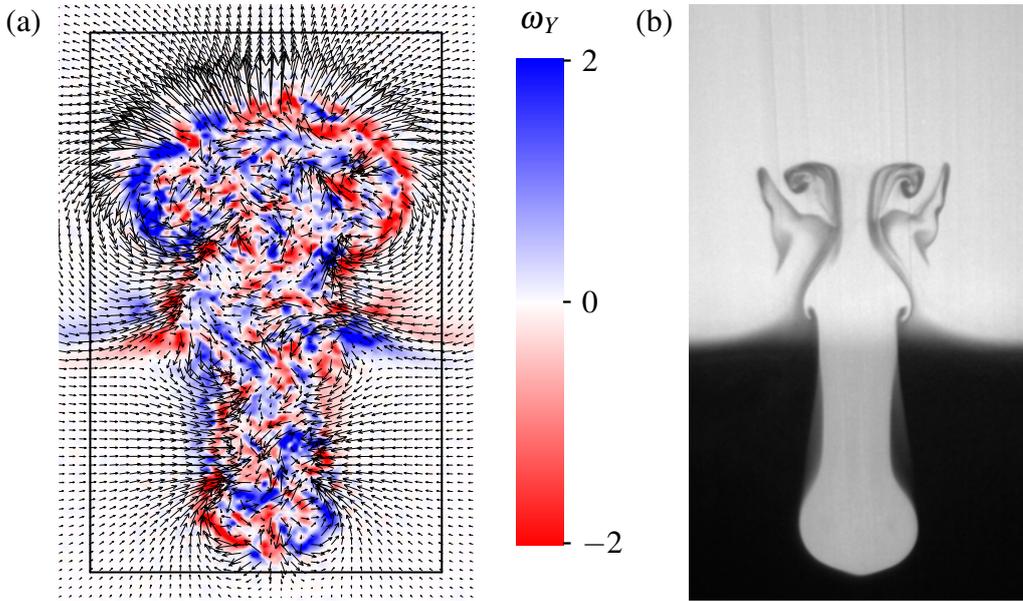


Fig. 4.6 (a) Example of a displacement array obtained using the largest image-pair time step we employ using our PIV algorithm. The rectangle drawn encloses the region to be interrogated at the next smallest time step. The scalar field is the curl of the pixel-displacement field, denoted ω_Y . (b) LIF image of the ring-interface interaction at the instance in time corresponding to the measured displacement array in Figure 4.6 (a).

interrogation, interrogating those regions at the next smallest time step, and so on, continues until all image-pair time steps have been utilised.

The pixel-distance criterion described above does not guarantee the use of the smallest time step $\Delta T^{(6)}$ for interrogating regions where the trajectory of particle images is highly curved, such as the core of the vortex ring, or regions where baroclinically-produced vorticity rolls up into vortical structures. This is because the pixel distance measurements in these regions are often small, but the curvature of the particle-image trajectories are sufficiently large that measurements made utilising larger time steps might inaccurately represent the flow. For our choice of $\Delta T^{(k)}$, this consideration is particularly important for measurements at the smallest time step because $\Delta T^{(k)}/\Delta T^{(k-1)} = 4$ for $k = 6$, and is equal to 2 for all smaller k . To account for this, an additional selection criterion based on the Q-criterion (Hunt et al. (1988)) is used on the measurements made utilising the smallest time step to identify these regions.

The Q-criterion is a method for vortex identification based on the pressure field p and the second invariant Q of the velocity tensor $\nabla \mathbf{u}$, defined by

$$Q = \frac{1}{2} (\|\boldsymbol{\Omega}\|^2 - \|\mathbf{S}\|^2), \quad (4.8)$$

where \mathbf{S} and $\mathbf{\Omega}$ are the symmetric and anti-symmetric components of $\nabla \mathbf{u}$. The Q -criterion identifies vortices as connected regions where the vorticity magnitude exceeds the strain-rate magnitude (corresponding to $Q > 0$), and the pressure converges to a local minimum somewhere interior to the connected region.

For two-dimensional PIV, out-of-plane velocities cannot be measured, so Q cannot be calculated directly. However, as our measurements are made in the central plane of the ring-interface interaction, we assume that out-of-plane velocity gradients are small relative to in-plane velocity gradients during the laminar stage of the flow. We therefore use a two-dimensional analogue Q_{2D} of Q to identify rotational regions, given by

$$Q_{2D}(u, w) = - (u_x^2 + w_z^2 + 2u_z w_x) , \quad (4.9)$$

where this formulation makes use of incompressibility to substitute $v_y = -(u_x + w_z)$.

To determine $J_n^{(6)}$, interrogation points are identified with either $|(\delta X_I, \delta Z_I)_{n,ij}^{(5)}| > s^{(5)}$, or $(Q_{2D})_{n,ij}^{(5)} > Q_{\min}$, where Q_{\min} is a prescribed minimum threshold. For each $|(\delta X_I, \delta Z_I)_{n,ij}^{(5)}|$, we take Q_{\min} to be 5% of the maximal value exhibited by $(Q_{2D})_{n,ij}^{(5)}$. In the same way as previously described, the smallest rectangle covering these interrogation points is determined, then expanded by three interrogation points in each direction. Only the interrogation points in the resulting rectangle are interrogated using $\Delta T^{(6)}$. Figure 4.7 exhibits subsequent displacement measurements made after producing the displacement measurements shown in Figure 4.6 (a), for the image-pair time steps corresponding to $k = 2, 4$ and 6.

After producing all displacement measurements $(\delta \mathbf{X}_I)_n^{(k)}$, measurements are linearly interpolated in time onto the times $\bar{T}_n^{(6)} = (16 + 2n)\delta T + \delta t/2$. Defining $r_m^{(k)}$ as

$$r_m^{(k)} = \frac{\bar{T}_n^{(6)} - \bar{T}_m^{(k)}}{\bar{T}_{m+1}^{(k)} - \bar{T}_m^{(k)}} ,$$

for m such that $\bar{T}_m^{(k)} < \bar{T}_n^{(6)} < \bar{T}_{m+1}^{(k)}$, we define the time-interpolated displacement measurements $(\Delta \mathbf{X}_I)_n^{(k)}$ as

$$(\Delta \mathbf{X}_I)_n^{(k)} = (1 - r_m^{(k)})(\delta \mathbf{X}_I)_m^{(k)} + r_m^{(k)}(\delta \mathbf{X}_I)_{m+1}^{(k)} . \quad (4.10)$$

As displacement measurements are interpolated onto times $\bar{T}_n^{(6)}$, no interpolation is required of the displacement measurements made utilising the smallest time step, and so these are used directly in constructing the pixel velocity fields $\mathbf{U}_I(\bar{T}_n^{(K)})$.

A heuristic approach is used to construct the pixel velocity fields $\mathbf{U}_I(\bar{T}_n^{(6)})$ with the aim of obtaining a low RME in all regions of the flow. For each $\bar{T}_n^{(6)}$, pairs of interpolated displace-

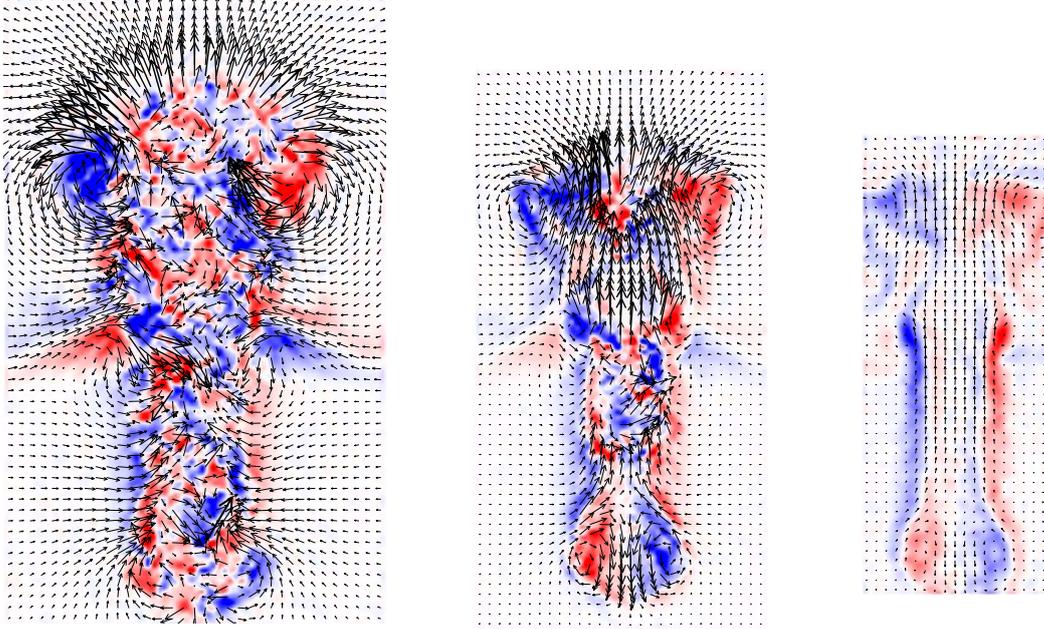


Fig. 4.7 Following the displacement measurements shown in Figure 4.6 (a), from left to right, displacement measurements produced utilising the image-pair time steps $\Delta T^{(2)}$, $\Delta T^{(4)}$ and $\Delta T^{(6)}$. For each set of displacement measurements, the scalar background is the curl of the pixel displacement field.

ment measurements $P_n^{(k)} = \left((\Delta \mathbf{X}_I)_n^{(k)}, (\Delta \mathbf{X}_I)_n^{(k-1)} \right)$ are considered sequentially, beginning with the measurements made utilising the two smallest time steps, then those made utilising the second and third smallest time steps, and so on until all such pairs have been considered. By approaching pairs of displacement measurements in order of increasing image-pair time step, spurious vectors such as those exhibited in Figure 4.6 (a) could be easily avoided when constructing $\mathbf{U}_I(\bar{T}_n^{(6)})$ by keeping track of the interrogation points where suitable pixel velocity vectors had already been determined.

When considering pair $P_n^{(k)}$, two types of measurement-validation criteria are used on the interrogation points (i, j) where a pixel velocity vector $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(k)})$ is yet to be constructed, with an additional criterion applied to the measurements made utilising the smallest time step. These criteria decide whether the velocity vector at (i, j) should be constructed using either the pair $P_n^{(k)}$, one of the pairs $P_n^{(k+1)}$ or $P_n^{(k-1)}$, or a later pair that is yet to be considered. By adopting this approach, each $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)})$ is constructed from at most two interpolated displacement measurements.

For each $\bar{T}_n^{(6)}$, the pair $P_n^{(6)}$ is first considered. The first validation criteria applied are based on pixel distance. Two maximum-pixel-distance thresholds, S_0 and S_1 , are used. The threshold S_0 is applied so that, if $|(\Delta X_I, \Delta Z_I)_{n,ij}^{(k)}| > (\Delta T^{(k)} / \Delta T^{(k-1)}) S_0$ or $|(\Delta X_I, \Delta Z_I)_{n,ij}^{(k-1)}| >$

S_0 , then

$$(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)}) = \alpha_{n,ij} \frac{(\Delta X_I, \Delta Z_I)_{n,ij}^{(k)}}{\Delta T^{(k)}} + (1 - \alpha_{n,ij}) \frac{(\Delta X_I, \Delta Z_I)_{n,ij}^{(k-1)}}{\Delta T^{(k-1)}}, \quad (4.11)$$

where $\alpha_{n,ij}$ is a weighting constant to be determined. This criterion places a maximal limit on the pixel distance of the measurements that are used to construct $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(k)})$ for $k < 6$, helping to avoid using spurious measurements that can be made when interrogating image pairs separated by larger time steps. The threshold S_1 is applied so that, if $S_1 < |(\Delta X_I, \Delta Z_I)_{n,ij}^{(k-1)}| \leq S_0$, then the velocity vector $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)})$ is to be constructed in the same way as for Equation (4.11) by either using $P_n^{(k)}$ or $P_n^{(k-1)}$. In our implementation of these criteria, we take $S_0 = 5\text{px}$ and $S_1 = 2.5\text{px}$, and apply these thresholds to all $P_n^{(k)}$ considered. After applying the criteria for $\{S_0, S_1\}$, all remaining measurements that have not satisfied any of the $\{S_0, S_1\}$ criteria correspond to very slow-moving regions of the flow, and so are constructed using $P_n^{(1)}$.

Note that our prescription of $S_0 = 5\text{px}$ was chosen from an intermediate range of suitable values for our purposes. While prescribing significantly smaller values of S_0 would lead to a larger RME in the constructed $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(k)})$, prescribing significantly larger values of S_0 would risk spurious velocity measurements being accepted as valid. We prescribed S_1 so that $S_0/S_1 = 2$, equalling the smallest ratio of time steps $\Delta T^{(k)}/\Delta T^{(k-1)}$.

As previously discussed, a pixel-distance criterion alone cannot not guarantee the construction of valid velocity vectors in regions where the trajectory of particles images are highly curved. To account for this in our algorithm, for each $\bar{T}_n^{(6)}$ we define a temporal average of the second invariant tensor, Q_n , as

$$Q_n = \frac{1}{5} \sum_{m=n-2}^{n+2} Q_{2D}(\Delta \mathbf{X}_I)_m^{(6)}, \quad (4.12)$$

where Q_{2D} is defined in Equation (4.9). We introduce a maximum threshold Q_{\max} such that, if $(Q_n)_{ij} > Q_{\max}$, then the velocity vector $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)})$ takes the form of Equation (4.11) with $\alpha_{n,ij} = 1$. Similarly to Q_{\min} , by taking Q_{\max} to be 5% of the maximal value exhibited by $(Q_n)_{ij}$ (across all n), we found that this criterion was sufficient in identifying regions where particle trajectories are highly curved, thus ensuring that velocity vectors in these regions would be constructed using displacement measurements made utilising the smallest time step.

The second validation criteria applied to all pairs $P_n^{(k)}$ compares the two interpolated-displacement measurements by making use of the variable $R_n^{(k)}$, defined by

$$(R_n^{(k)})_{ij} = \left\| \frac{\Delta T^{(k)}}{\Delta T^{(k-1)}} (\Delta X_I, \Delta Z_I)_{n,ij}^{(k-1)} - (\Delta X_I, \Delta Z_I)_{n,ij}^{(k)} \right\|. \quad (4.13)$$

For $(R_n^{(k)})_{ij} \lesssim O(0.1)\text{px}$, pixel-displacement measurements made at (i, j) using the time steps $\Delta T^{(k)}$ and $\Delta T^{(k-1)}$ are in close agreement, suggesting both measurements are suitable for constructing a valid velocity vector for $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)})$. Conversely, if $(R_n^{(k)})_{ij} \gtrsim O(1)\text{px}$, the measurements agree less well, possibly due to the curvature of particle trajectories, or to spurious measurements arising when utilising $\Delta T^{(k-1)}$, due to the pixel displacements of particle images being too large. If agreement between the measurements is poor, the measurement made utilising the smaller time step is preferred to avoid making use of an erroneous measurement.

In our algorithm, we introduce minimum and maximum thresholds $R_{\min} = 0.1 \text{ px}$ and $R_{\max} = 0.3 \text{ px}$ for $(R_n^{(k)})_{ij}$ that are used to set the weighting constant, $\alpha_{n,ij}$, introduced in Equation (4.11). These weighting constants are prescribed to be piecewise linear in $(R_n^{(k)})_{ij}$, given by

$$\alpha_{n,ij} = \begin{cases} 0.5 & (R_n^{(k)})_{ij} \leq R_{\min} \\ 0.5 \left(1 + \frac{(R_n^{(k)})_{ij} - R_{\min}}{R_{\max} - R_{\min}} \right) & R_{\min} \leq (R_n^{(k)})_{ij} \leq R_{\max} \\ 1 & (R_n^{(k)})_{ij} \geq R_{\max} \end{cases}. \quad (4.14)$$

As large differences in $\alpha_{n,ij}$ between adjacent interrogation points are susceptible to introducing high levels of noise in the velocity-gradient field, it is important that the profile prescribed for $\alpha_{n,ij}$ is continuous. The piecewise-linear profile given in Equation (4.14) was prescribed for simplicity, and proved to be sufficient for our purposes.

By applying the measurement-validation criteria outlined above, initial guesses are obtained for $k_{n,ij}$ and $\alpha_{n,ij}$, where $k_{n,ij}$ is defined as the value for k employed in Equation (4.11). Based on our construction of $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)})$ given in Equation (4.11), we can define the parameter $W_{n,ij} = k_{n,ij} + \alpha_{n,ij}$, from which $k_{n,ij}$ and $\alpha_{n,ij}$ can be extracted as $k_{n,ij}$ is taken to be the largest integer less than $W_{n,ij}$. Figure 4.8 (a) plots the initial guess for the W_n array associated with the flow also considered for Figures 4.6 and 4.7.

After obtaining initial guesses for $k_{n,ij}$ and $\alpha_{n,ij}$, three filters are applied to smooth the W_n array, removing large discontinuities between adjacent interrogation points. Values for $k_{n,ij}$ and $\alpha_{n,ij}$ are then extracted from the smoothed $W_{n,ij}$ array, then employed using Equation (4.11) to construct the $(\mathbf{U}_I)(\bar{T}_n^{(6)})$ array. First, a 3×3 median filter is applied to improve

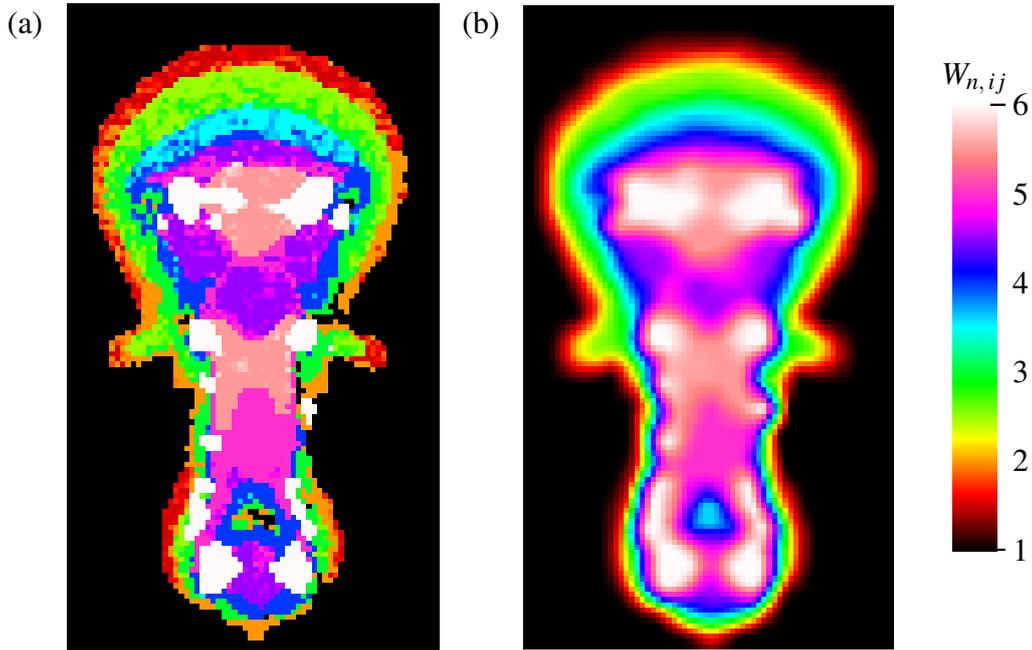


Fig. 4.8 Plots of $W_{n,ij} = k_{n,ij} + \alpha_{n,ij}$, indicating both the image-pair time steps utilised at each interrogation point to construct the PIV velocity field given in (b), and the weighting constants $\alpha = \alpha_{n,ij}$, defined in Equation (4.11). Plots are (a) before, and (b) after, three filters are applied to the initial guess of the W_n array in the following order: a 3×3 median filter, a 3×3 maximum filter, and a 5×5 low-pass convolution filter.

the level of similarity between $W_{n,ij}$ and adjacent neighbours on the W_n array. Secondly, a 3×3 maximum filter is applied to replace the choice of $k_{n,ij}$ with that of the highest neighbour. Finally, a 5×5 low-pass convolution filter is applied to smooth discontinuities in $W_{n,ij}$ between adjacent interrogation points. The filtered array is plotted in Figure 4.8 (b).

The filtered array W_n is then used in conjunction with Equation (4.11) to construct velocity vectors $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)})$. The constructed velocity vectors are then spatially interpolated onto the sensor array using a bi-cubic interpolation scheme (available using DigiFlow), then mapped onto the central plane of the light sheet to produce fluid velocity measurements $\mathbf{u}(\bar{T}_n^{(6)})$.

It is important that quantities derived using velocity gradients, such as the vorticity, divergence and the Q-field, are derived using the velocity vectors $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)})$ defined on the interrogation array, rather than the velocity measurements $\mathbf{u}(\bar{T}_n^{(6)})$. If such quantities are derived making use of spatially interpolated measurements, spurious features can arise in these quantities that are not representative of the flow, and only arise as a result of the particular spatial-interpolation scheme being used. We remove this possibility by deriving vector-gradient quantities using the vectors $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)})$. After deriving these quantities

on the interrogation array, in the same way as for the velocity vectors $(\mathbf{U}_I)_{ij}(\bar{T}_n^{(6)})$, these quantities are spatially interpolated then mapped onto the central plane of the light sheet.

4.5 Discussion

We have constructed a heuristic PIV algorithm that makes use of different image-pair time steps to interrogate sequences of PIV images, with the aim of obtaining a low RME in all regions of the flow. The algorithm we have developed makes use of pre-existing pattern-matching functionality available within DigiFlow. We use our algorithm to process all PIV images obtained from the experiments presented in Chapter 5, using the threshold values specified in Section 4.4 for the parameters related to implementing the algorithm.

Figure 4.9 (a) plots the algorithm-constructed velocity measurements at the time coinciding with the LIF image given in Figure 4.6 (b). For comparison, Figure 4.9 (b) plots the velocity measurements constructed using the same parameters prescribed for the pattern-matching functions used to construct the velocity field presented in Figure 4.9 (a), but utilising only the smallest time step when comparing pairs of PIV images in all regions of the flow. Figures 4.9 (c) and (d) re-plot the velocity measurements shown in Figures 4.9 (a) and (b), with the vorticity colour map rescaled to visualise slower-moving regions of the flow more clearly.

Away from the vortex ring, baroclinically-produced structures and the deformed regions of the density interface, the vorticity field is smooth and the flow is irrotational to a high approximation. A qualitative comparison between the measurements presented in Figure 4.9 demonstrates that our PIV algorithm is capable of significantly reducing the RME in slow-moving regions of the flow when compared to the pre-existing PIV functionality available in DigiFlow. This can be seen both in the improved resolution of the shear-generated vorticity at the interface (labelled A) and the reduction of noise in the vorticity field regions external to the vortex ring and any baroclinically-produced structures (for example, the region labelled B). Our algorithm also appears to accurately resolve the baroclinically-produced vorticity (such as the regions labelled C). As we do not have access to the full three-dimensional velocity field at this time, it is difficult to say which velocity field more accurately represents the baroclinically-produced vorticity. However, the level of agreement in the regions of high vorticity between Figures 4.9 (a) and (b) supports the notion that both PIV approaches are suitable for the interrogation of sequences of particle images.

To test our PIV algorithm against the pre-existing PIV functionality available in DigiFlow in a rigorous way would be to generate a synthetic sequence of particle images, advect the synthetic particle images using a known velocity field (such as a Rankine vortex), then test

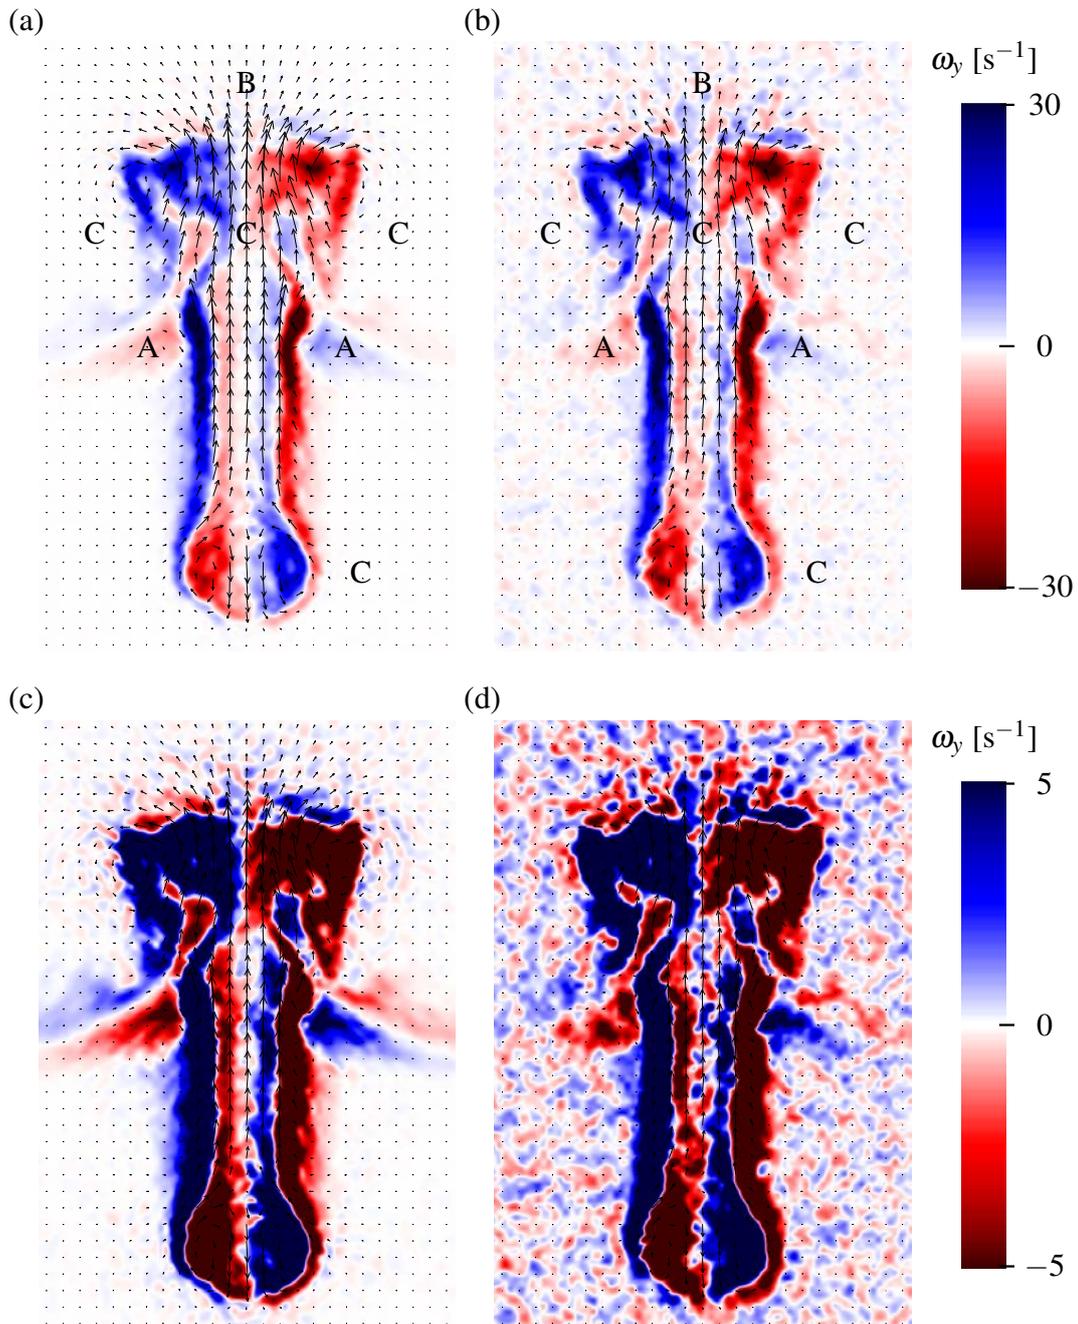


Fig. 4.9 (a) Example of a velocity field constructed using the algorithm presented in this Chapter, at the time coinciding with the LIF image given in Figure 4.6 (b). (b) For comparison, a velocity field evaluated at the same instance in time, utilising only the smallest available time step in all regions of the flow. Labels correspond to features identified in the text. Plots (c) and (d) correspond to (a) and (b) respectively, using a reduced scale in ω_y to illustrate the reduction in RME achieved using our algorithm in slower-moving regions of the flow. For all plots, a subsample of the constructed velocity vectors are printed, with the largest velocity vectors corresponding to approximately 11.5 cm/s.

both algorithms on the synthetic PIV images and compare the velocity measurements against the known velocity field. Such tests are not presented in this chapter, but should be carried out using different known velocity fields to assess in a precise way the advantages associated with using either PIV-processing approach. Without conducting such tests, based on the use of our algorithm in our PIV/LIF experiments presented in Chapter 5, at the very least it can be said with confidence that our algorithm significantly reduces the RME in near-irrotational regions where the measured particle-projection pixel distance would be less than $O(0.1)$ px if utilising the smallest image-pair time step.

A natural extension of our PIV algorithm would be to incorporate a functionality that allows for the interrogation window size and spacing (n_I and δ_I) to vary locally depending on the conditions of the flow. For instance, a criterion based on the Q-criterion could be used to identify regions where particle trajectories are highly curved, motivating the use of smaller n_I and δ_I to improve the resolution of such regions. The investigation of such possible extensions to our algorithm is left for future work.

Chapter 5

Influence of propagation angle on the ring–interface interaction

Following our review of vortex-ring interactions in Section 2.3, this chapter investigates the dynamics of a vortex ring obliquely impacting a density interface. Figure 5.1 sketches a vortex ring obliquely propagating at an angle θ_0 to the vertical before impacting a density interface, with the physical parameters annotated that are used to define Re_0 and Ri_0 . As there has been no previously published systematic study on this type of vortex interaction, the parameter space characterising such interactions remains almost entirely unexplored. To the best of our knowledge, this experimental investigation is the first systematic study on vortex rings obliquely impacting a density interface.

In this thesis we regard the evolution of the ring–interface interaction as being characterised by the three-dimensional parameter space (Re_0, Ri_0, θ_0) . As discussed previously in Section 1.3, the formation number F_p , the Schmidt number Sc , the thickness $\delta z/a$ of the density interface non-dimensionalised by the ring diameter, and the non-dimensionalised height $(H_B - H_2)/a$ between the ring-generation site and the density interface, will also impact the evolution of the ring–interface interaction. However, in our experimental setup we keep these four parameters fixed, with $F_p = 2.3$, $Sc = 500 - 700$, $\delta z/a \approx 0.3$ and $(H_B - H_2)/a \approx 4$. Note that taking $(H_B - H_2)/a \approx 4$ was chosen as it is large enough that vortex rings are given sufficient distance to fully form before interacting with the interface, while also being small enough that vortex rings are unable to dissipate a significant proportion of its kinetic energy before interacting with the interface.

The case of vertically propagating vortex rings impacting a density interface has been studied previously (Dahm et al. (1989), Marcus and Bell (1992), Stock et al. (2008), Olsthoorn and Dalziel (2017), Yeo et al. (2020)), and was discussed in Section 2.3.5. They concluded that both Re_0 and Ri_0 play an important role in the evolution of the ring–interface interaction,

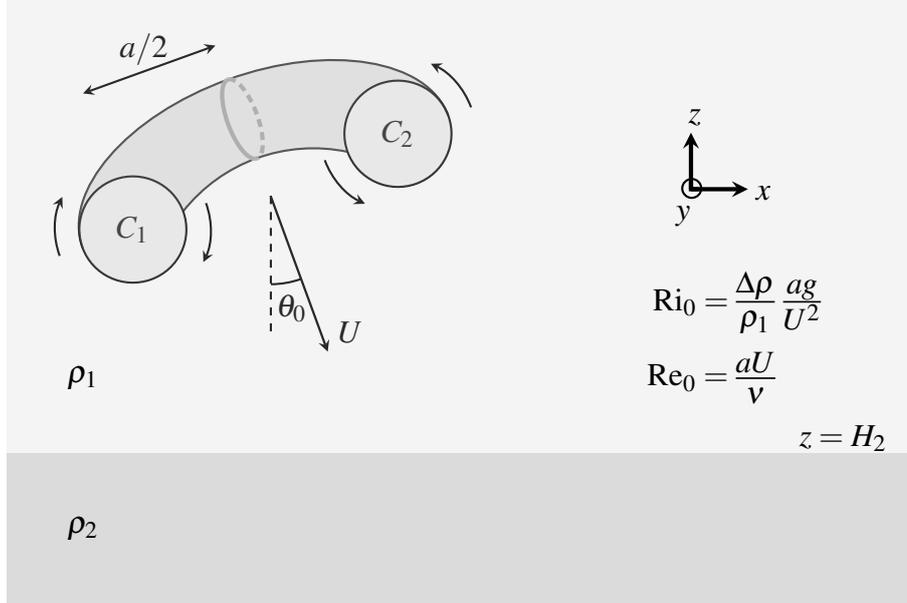


Fig. 5.1 Diagram of a vortex ring of diameter a propagating obliquely towards a sharply stratified density interface (at height $z = H_2$) at propagation speed U and initial propagation angle θ_0 . Here, the vortex-core cross-sections C_1 and C_2 coincide with the central plane of the flow. Richardson and Reynolds numbers Ri_0 , Re_0 are defined based on bulk properties of the flow and the fluid density of each layer. Note that the tube height $z = H_B$ is out of view in this diagram.

with Re_0 controlling the mode number of the instability structure that develops as the ring interacts with the stratification, and Ri_0 controlling the depth the vortex ring is able to penetrate into the lower layer. To study the influence of θ_0 on the ring–interface interaction, we compare cases with equal Re_0 and Ri_0 .

Using the methodology detailed in Chapter 3, this chapter presents two-dimensional PIV data and LIF observations obtained from probing the dynamics in the central plane of the ring–interface interaction. Ideally, we would have been able to obtain fully three-dimensional near-instantaneous velocity and density data of the ring–interface interaction. As discussed in Section 4.1, though the plane-scanning technology developed for 3D PIV/LIF by Partridge et al. (2019) would have been suitable for this purpose, at present this technology is incapable of scanning volumes at a fast enough rate to make use of in this investigation. Although we were unable to probe the full three-dimensional evolution of the ring–interface interaction in this way, our PIV and LIF observations provide useful insight into the three-dimensional evolution of the flow.

As there has been no previously published systematic study of vortex rings obliquely impacting a density interface, there are an abundance of questions that could be asked in

investigating these interactions. The primary motivation of this experimental investigation is to understand how deviating from $\theta_0 = 0^\circ$ in the parameter space characterising the ring–interface interaction alters the dynamics associated with the interaction, and the implications this may have for the ring-induced mixing and the ring-eddy analogy more broadly. As such, we restrict our attention in this chapter to vortex-ring interactions with initial propagation angles $\theta_0 \leq 20^\circ$.

Our decision to restrict our attention to $\theta_0 \leq 20^\circ$ was made for two reasons. First, for experiments using the angled tube (corresponding to $\theta_0 \geq 30^\circ$, it proved incredibly difficult to keep the central axis of the vortex ring aligned with the central plane of the light sheet throughout the interaction. We attributed this to the vortex ring interacting with the interface for a relatively long period of time (compared to smaller angles of θ_0), resulting in the ring deviating from the central plane of the light sheet before impacting the interface. Secondly, the dynamics of the ring–interface interaction varies significantly over the range $0^\circ \leq \theta_0 \leq 20^\circ$, deserving in itself a study with undivided attention. As a result, we opted to prioritise studying interactions with $\theta_0 \leq 20^\circ$, leaving interactions at higher θ_0 to be studied at a later time, when a three-dimensional PIV/LIF scanning system (such as that introduced by Partridge et al. (2019)) becomes available that can scan the flow at a sufficiently high rate to capture the full three-dimensional dynamics of the ring–interface interaction.

The coordinate system (x, y, z) we make use of in this chapter was shown in Figure 3.8, with \hat{x} and \hat{z} directions corresponding to the horizontal and vertical vector components parallel with the central plane of the vortex-ring trajectory, and \hat{y} corresponding to the out-of-plane vector component, pointing away from the cameras. For each experiment, we take the plane $z = 0$ to be at the density interface, by determining the position of maximal vertical gradient in the horizontally averaged pixel intensity from one of our LIF snapshots, taken before the ring perturbs the density interface. We take $t = 0$ to be the time when the ‘ring centre’ (defined below in Section 5.1) is one ring diameter above the density interface, i.e. at $z = a$. We take $x = 0$ to be the x component of the ring centre at $t = 0$.

Table 5.1 presents the parameters characterising the ring–interface interactions investigated in this chapter. As discussed in Section 3.2, the traverse time T_s for the bike pumps is prescribed, which sets the ring propagation velocity U . The range of θ_0 considered corresponds to that investigated in our periodic-mixing experiments, which are presented in Chapter 6. In Section 5.1 we introduce the quantities we make use of to infer properties of the evolution of the vortex ring during its interaction with the density interface, as well as describing how we use the PIV and LIF data to infer other aspects of the flow. In Section 5.2, we analyse the evolution of the tabulated cases corresponding to $(Re_0, Ri_0) = (1810, 3.09)$, discussing how varying θ_0 impacts both the evolution of the primary ring and the emer-

T_s [s]	U [mm/s]	a/U [s]	Re_0	Ri_0	θ_0
8.33	36.95 ± 0.47	1.33 ± 0.05	1810	3.09	$0^\circ, 5^\circ, 10^\circ, 20^\circ$
5.45	54.55 ± 0.10	0.90 ± 0.02	2680	1.48	$0^\circ, 5^\circ, 10^\circ, 20^\circ$
4.56	66.54 ± 0.42	0.74 ± 0.02	3270	0.98	$0^\circ, 5^\circ, 10^\circ, 20^\circ$

Table 5.1 Table of parameters for the experiments investigated in this chapter. The values and precision errors for U and a/U are determined from our bike pump calibration, discussed in Section 3.2, noting that measurements of the ring diameter were in the range $a = 49.2 \pm 1.0$ mm.

gence of secondary structures in the flow. We carry out a similar analysis in Section 5.3, restricting our attention to the lower Ri_0 (and higher Re_0) cases $(Re_0, Ri_0) = (2680, 1.48)$ and $(3270, 0.98)$. We conclude this chapter in Section 5.4 with a summary of our findings and a final discussion.

5.1 Observed measures

Before investigating the cases of the ring–interface interaction listed in Table 5.1, in this section we lay out the measures we make use of to infer aspects of the evolution of the flow. As previously discussed, we use PIV and LIF to probe the flow in the central plane of the ring–interface interaction. By probing the flow in this way, two vortex-core cross-sections, C_1 and C_2 (labelled in Figure 5.1) of the primary vortex ring are observable in our PIV measurements while the primary ring remains intact.

To identify and track $C_1(t)$ and $C_2(t)$, a 5% maximum-threshold is first used on the Q-field, $Q = Q_{2D}$ (introduced in Equation (4.9)), calculated from the earliest PIV velocity field obtained at time T_0 . Using this threshold calculation, the area $A_i(T_0)$ and the centroid $\mathbf{x}_i(T_0)$ of $C_i(T_0)$ ($i = 1, 2$) are determined. Subsequently, given $A_i(T_m)$ and $\mathbf{x}_i(T_m)$, the vortex core C_i at time T_{m+1} is identified by searching in the square region of size $(\sqrt{A_i(T_m)} + 4\zeta)$ centred on $\mathbf{x}_i(T_m)$, after which $A_i(T_{m+1})$ and $\mathbf{x}_i(T_{m+1})$ are determined (here, ζ is the mm/px scale). This strategy allows us to track C_1 and C_2 as the ring interacts with the density interface.

Using the centroids $\mathbf{x}_1, \mathbf{x}_2$, we define the ‘ring centre’, $\bar{\mathbf{x}} = (\bar{x}, \bar{z})$, of the vortex ring as

$$\bar{\mathbf{x}}(T_k) = \frac{1}{2} (\mathbf{x}_1(T_k) + \mathbf{x}_2(T_k)) . \quad (5.1)$$

Using the ring centre, we fix the time origin at time T_K , where K minimises $|\bar{z}(T_K) - a|$. We also define the ‘plane diameter’, $D(T_m)$, given by

$$D(T_m) = |\mathbf{x}_1(T_m) - \mathbf{x}_2(T_m)|. \quad (5.2)$$

In the case $\theta_0 = 0^\circ$, the plane diameter can be considered as equivalent to the ring diameter (modulo any instabilities developing in the flow). For $\theta_0 > 0^\circ$, the structural evolution of the vortex ring is partly influenced by the variation in buoyancy forces acting on different portions of the ring, which can potentially lead to the ring deforming asymmetrically (such as the vortex ring ‘bending’ identified by Stock (2006) and described in Section 2.3.6). As such, the plane diameter should not be interpreted as equivalent to the ring diameter, but only as a measure of the size of the ring.

By making use of \mathbf{x}_1 and \mathbf{x}_2 , two different angles can be calculated, which we use to infer aspects of the evolution of the ring–interface interaction. Firstly, the propagation angle, $\theta_p(t)$, which we define as the angle the trajectory of the ring centre $\bar{\mathbf{x}}$ makes with the vertical. The angle $\theta_p(T_m)$ is given by

$$\theta_p(T_m) = \tan^{-1} \left(\frac{d\bar{x}}{d\bar{z}} \right) \Big|_{t=T_m}. \quad (5.3)$$

Here, the gradient $d\bar{x}/d\bar{z}$ at $t = T_M$ is taken to be the gradient of the line calculated using the method of least squares on the points $\bar{\mathbf{x}}(T_k)$ for $k = M - 4, M - 3, \dots, M + 4$. Secondly, the core angle, $\theta_c(t)$, defined as the angle that the line intersecting \mathbf{x}_1 and \mathbf{x}_2 makes with the horizontal. This is given by

$$\theta_c(T_m) = \tan^{-1} \left(\frac{z_2 - z_1}{x_2 - x_1} \right) \Big|_{t=T_m}. \quad (5.4)$$

Figure 5.2 illustrates the angles $\theta_p(t)$ and $\theta_c(t)$. Before the ring interacts with the interface, the two angles are equal, with $\theta_p(t) = \theta_c(t) = \theta_0$. Changes in $\theta_p(t)$ are associated with changes in direction of the bulk motion of the vortex ring, whereas changes in $\theta_c(t)$ are associated with differences in the forces acting on C_1 and C_2 . The differences in the inferences that can be made between probing the evolutions of $\theta_p(t)$ and $\theta_c(t)$ motivates the use of both measures in our investigations of the ring–interface interaction.

The velocity-gradient fields we make use of in this chapter are the vorticity, ω_y , the Q-field, Q_{2D} , and the two-dimensional divergence, $\nabla \cdot \mathbf{u}_{2D} = \partial u / \partial x + \partial w / \partial z$. These fields are used in conjunction with the pixel-intensity fields (obtained using LIF), which allow us to track the movement of the dyed upper-layer fluid and non-dyed lower-layer fluid in the

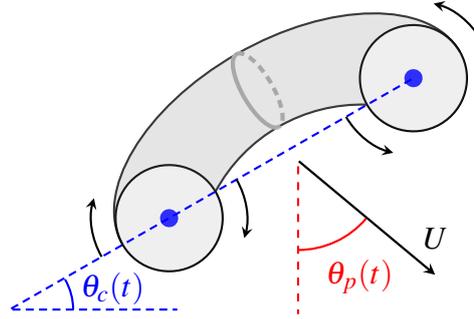


Fig. 5.2 Diagrams illustrating the definitions of the propagation angle $\theta_p(t)$ and core angle $\theta_c(t)$. Note that before the ring interacts with the density interface, the two angles are equal, with $\theta_p(t) = \theta_c(t) = \theta_0$.

central plane of the flow. By making use of these fields and the measures outlined above, we infer aspects of the three-dimensional evolution of the flow, both in the primary ring and secondary structures that emerge in the central plane of the flow.

It should be noted that the $\nabla \cdot \mathbf{u}_{2D}$ fields derived from our PIV measurements were found to have a higher level of noise relative to the ω_y and Q_{2D} fields. The noise in the $\nabla \cdot \mathbf{u}_{2D}$ field does not appear to be a consequence of employing the PIV algorithm introduced in Chapter 4. We return to this point in Section 7.3. For the purposes of our analysis in this chapter, the important point is that the $\nabla \cdot \mathbf{u}_{2D}$ contours do a sufficiently good job of indicating the regions of convergence on to or away from the central plane of the flow.

Using our LIF measurements, we define the maximum penetration depth $Z(t)$ of the vortex ring at a given time as the vertical distance between the resting position of the density interface, and the z position of the most deeply penetrating upper-layer fluid into the lower layer that is made visible by the light sheet. For a given experiment, the resting position of the interface is determined by vertically averaging an LIF image where the interface is unperturbed, then identifying the position of maximal gradient. For each LIF image, the z position of the most deeply penetrating upper-layer fluid is determined by applying a maximum filter to each row of pixels, then finding the lowest-lying row of pixels with a pixel intensity greater than 0.3. In doing so, we are then able to determine $Z(t)$ for each LIF image.

5.2 Non-penetrating rings at small propagation angles

The experiments presented in this section are those listed in the first row of Table 5.1, for which $(Re_0, Ri_0) = (1810, 3.09)$ and $\theta_0 \leq 20^\circ$. We describe these ring–interface interactions as ‘non-penetrative’, as the maximum-penetration depth of the vortex ring into the lower layer is less than one ring diameter, and the entirety of the mixing appears to occur only

in the upper-layer region. Additionally, the $\theta_0 = 0^\circ$ case corresponds closely to the ‘non-penetrative’ case $(\text{Re}_0, \text{Ri}_0) = (2000, 3.0)$ investigated by Olsthoorn and Dalziel (2017), who present three-dimensional reconstructions of the velocity field.

In Section 5.2.1, we describe the evolution for the case $\theta_0 = 0^\circ$. Then, in Section 5.2.2 we discuss the influence of θ_0 on the evolution of the primary vortex ring. Finally, in Section 5.2.3 we demonstrate for the cases $\theta_0 = 5^\circ$ and 10° that significant azimuthal flow arises during the later stages of the flow, which we attribute to azimuthal pressure gradients developing as a result of the azimuthally asymmetric production of baroclinic vorticity.

5.2.1 Vertical propagation

LIF snapshots of the central plane of the flow are presented in Figure 5.3 for the case $\theta_0 = 0^\circ$. For each sequence of snapshots, the first snapshot is taken at $t = 1$ s and consecutive snapshots are separated in time by 0.5 seconds (or equivalently, starting at $0.75a/U$ and separated in time by $0.375a/U$). The presented snapshots showcase the full evolution of the flow, from the ring beginning to impinge on the interface, up to the stage where the flow is in a dissipative state of relaxation.

As the ring impinges onto and deforms the density interface, the deformation leads to an azimuthally symmetric production of baroclinic vorticity on the deformed interface. To help visualise this stage of the flow, Figure 5.4 (a) presents the LIF snapshot at $t = 2.0$ s, overlaid with velocity vectors and ω_y contours, with the blue and red contours respectively representing anti-clockwise and clockwise in the central plane of the flow. The baroclinically-generated vorticity is swept around the outside of the vortex ring and accumulates at the upper rim of the deformed interface. Subsequently, the accumulated vorticity separates, leading to the emergence of fine-scale coherent structures that develop as the flow continues to evolve (exhibited from $t = 2.0$ s to $t = 4.0$ s in Figure 5.3).

As the interface begins to recoil after the ring reaches its maximum-penetration depth into the lower layer (around $t = 2.5$ s), the upward motion of the deformed interface in combination with the continuing production of baroclinic vorticity along the deformed interface leads to the formation of a secondary vortex ring of larger diameter than the primary ring. This secondary ring (hereafter referred to as the ‘recoil ring’) can be seen clearly at $t = 3.5$ s, both in Figure 5.3 and Figure 5.4 (b). As the flow continues to evolve, the recoil ring subsequently pinches off from the lower layer and self-propagates into the upper layer. In doing so, the recoil ring enhances the ‘splashing’ mechanism first observed by Linden (1973), which describes the transient enhancement of vertical buoyancy flux associated with the ejection of lower-layer fluid into the upper layer.

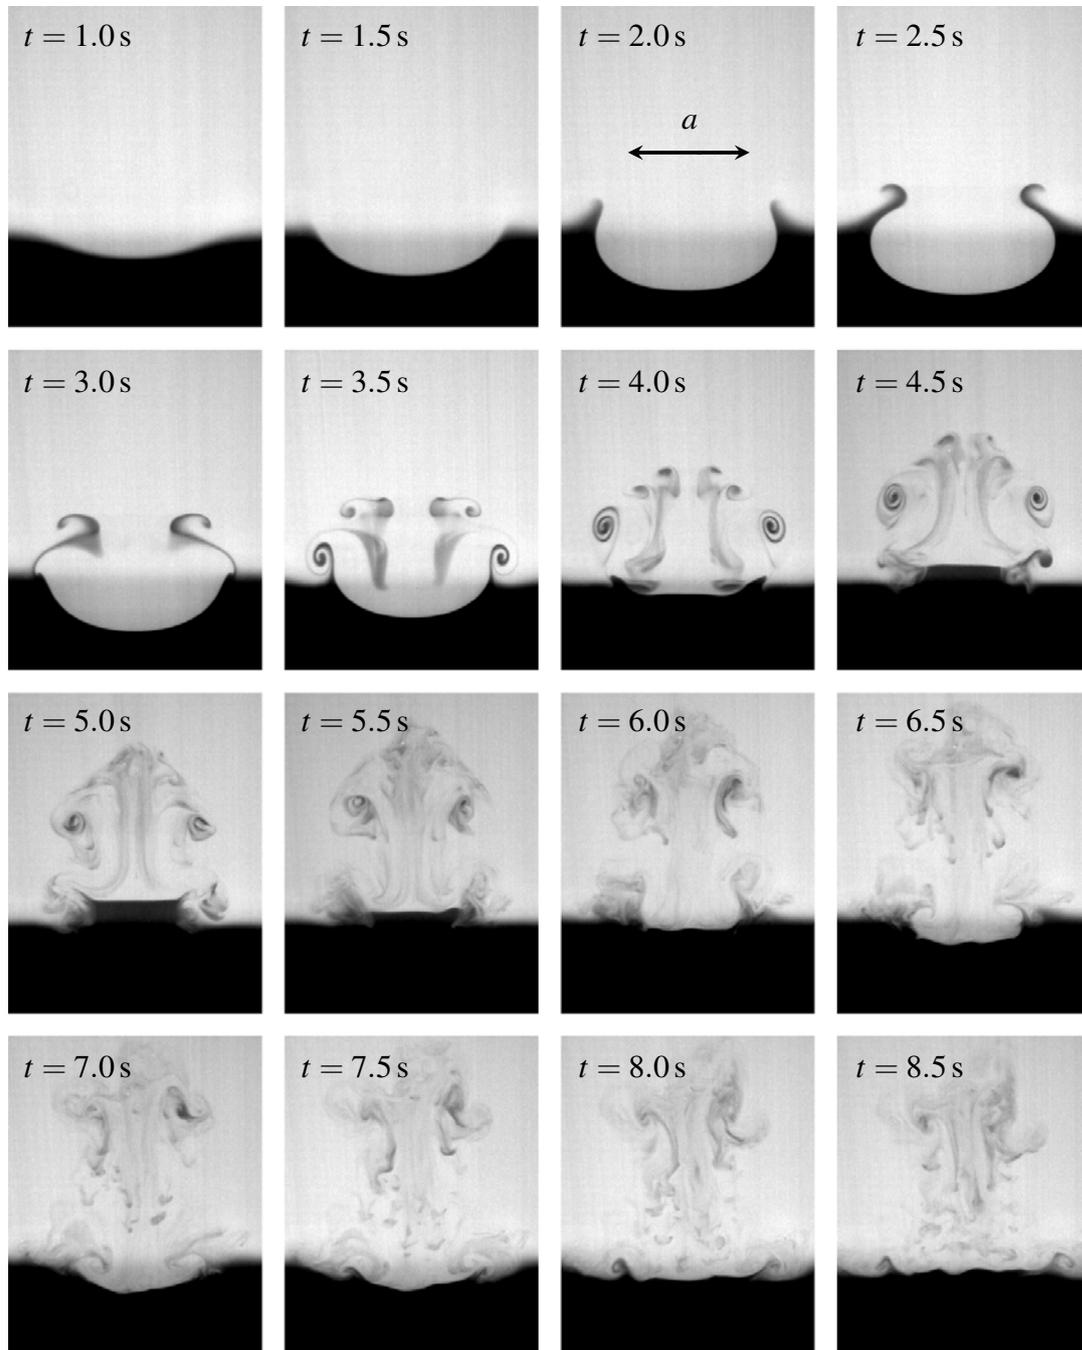


Fig. 5.3 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (1810, 3.09, 0^\circ)$. The time separation between consecutive snapshots is 0.5 seconds, with the first frame taken at time $t = 1\text{ s}$. Note that the advective timescale $a/U = 1.33\text{ s}$.

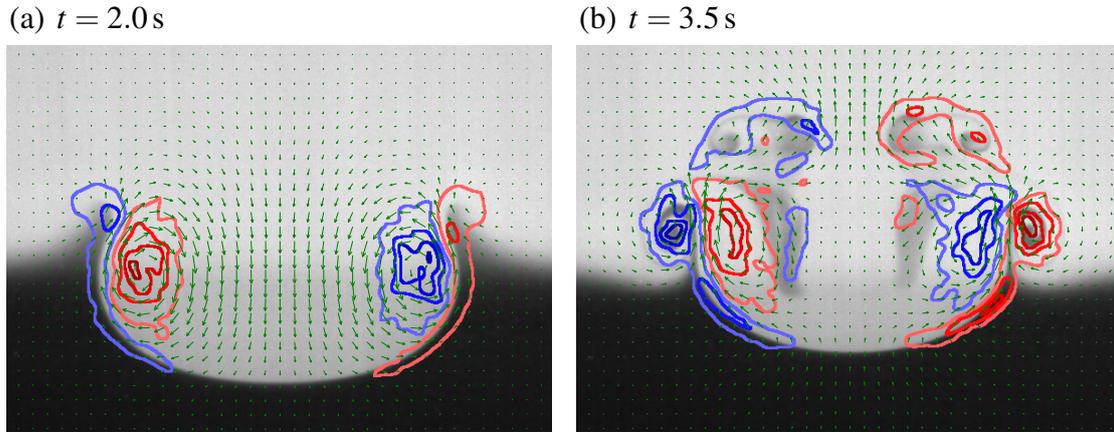


Fig. 5.4 LIF images at (a) $t = 2.0$ s and (b) $t = 3.5$ s, overlaid with a subsample of 50% the PIV algorithm-constructed velocity vectors, and ω_y contours, with the blue and red contours respectively representing anti-clockwise and clockwise in the central plane of the flow. The in-plane vorticity contours are drawn for $|\omega_y| = 5, 15, 25$ and 35 s^{-1} . Note that the largest velocity vectors in (a) correspond to a speed of 59.2 mm/s .

Flow in the azimuthal direction of the primary vortex ring appears to be small throughout the ring–interface interaction, as there is no strong indication in Figure 5.3 of fluid structures emerging in the central plane of the flow due to out-of-plane motion. We attribute this to the azimuthal symmetry of the ring–interface interaction before the turbulent breakdown of the flow.

5.2.2 Oblique propagation: evolution of the primary ring

In a similar way to Figure 5.3 for $\theta_0 = 0^\circ$, Figures 5.5, 5.6 and 5.7 present LIF snapshots of the ring–interface interaction, exhibiting the evolution of the flow for the cases $\theta_0 = 5^\circ$, 10° and 20° respectively. Clear symmetry breaking is observed for $\theta_0 = 5^\circ$ and 10° in the separated vortical structures and the recoil ring in the laminar stage of the flow, and for $\theta_0 = 20^\circ$, the vertical buoyancy flux associated with the splashing mechanism appears to be significantly reduced. These observations support the notion that the mixing energetics of the ring–interface interaction may be significantly sensitive to θ_0 over the range of angles $\theta_0 \leq 20^\circ$. Before returning to analyse these secondary structures in more detail, we first consider the evolution of the primary ring as it interacts with the stratification.

Figures 5.8 (a) to (f) plot the observed measures of the vortex ring introduced in Section 5.1, in addition to the non-dimensionalised penetration depth $Z(t)/a$ into the lower layer. For each figure, data are plotted between $t = 0$ and $t = 4$ seconds. Vertical dashed lines are drawn in plots (a) to (d), marking the instances in time that correspond to the first seven LIF

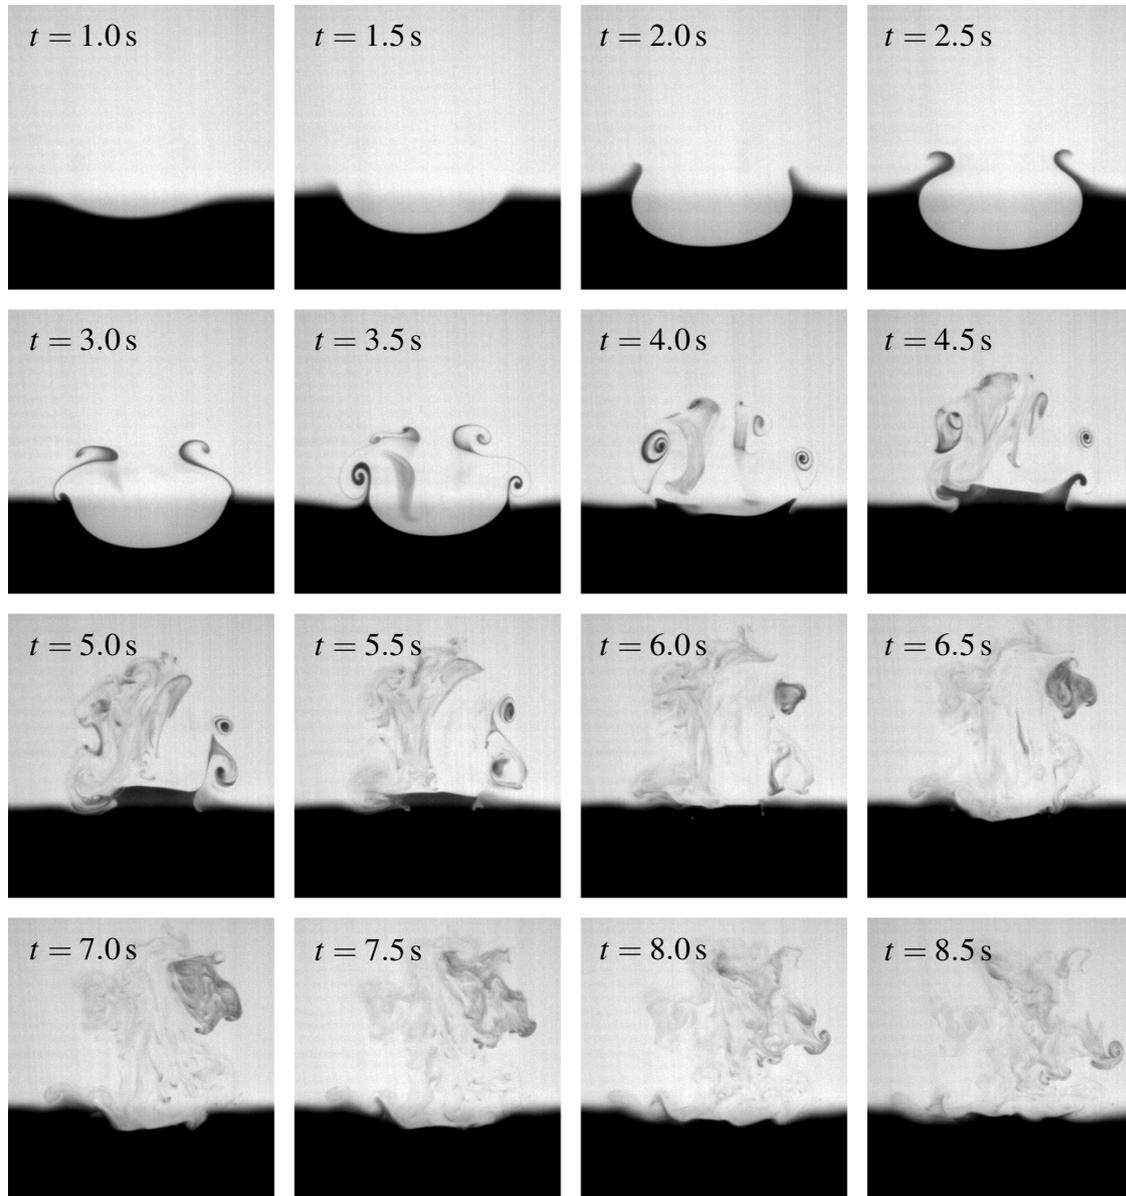


Fig. 5.5 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (1810, 3.09, 5^\circ)$. The time separation between consecutive snapshots is 0.5 seconds, with the first frame taken at time $t = 1$ s. Note that the advective timescale $a/U = 1.33$ s.

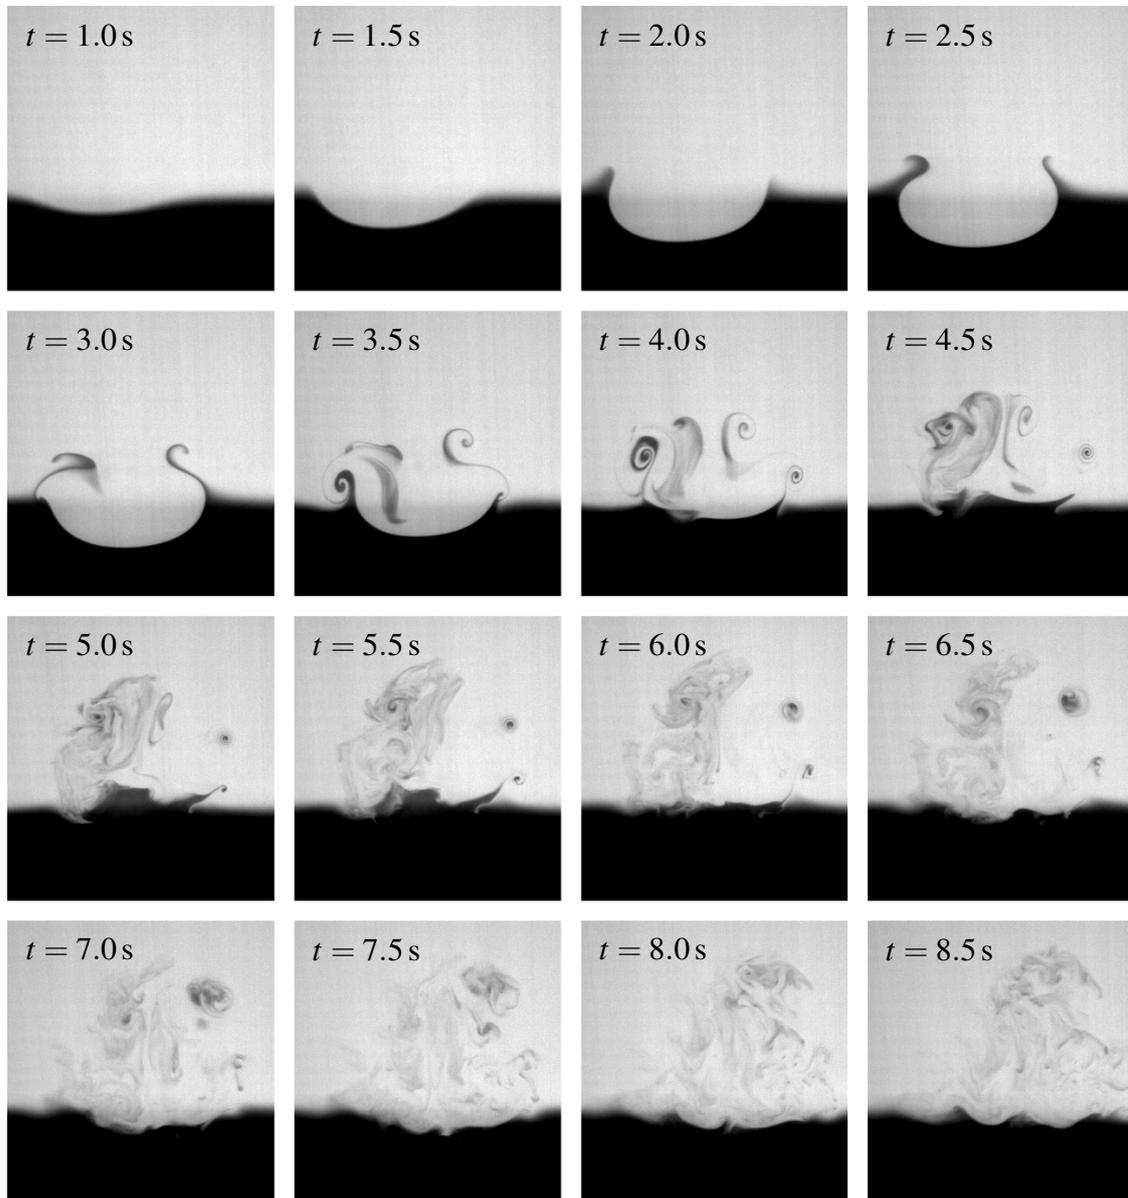


Fig. 5.6 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (1810, 3.09, 10^\circ)$. The time separation between consecutive snapshots is 0.5 seconds, with the first frame taken at time $t = 1$ s. Note that the advective timescale $a/U = 1.33$ s.

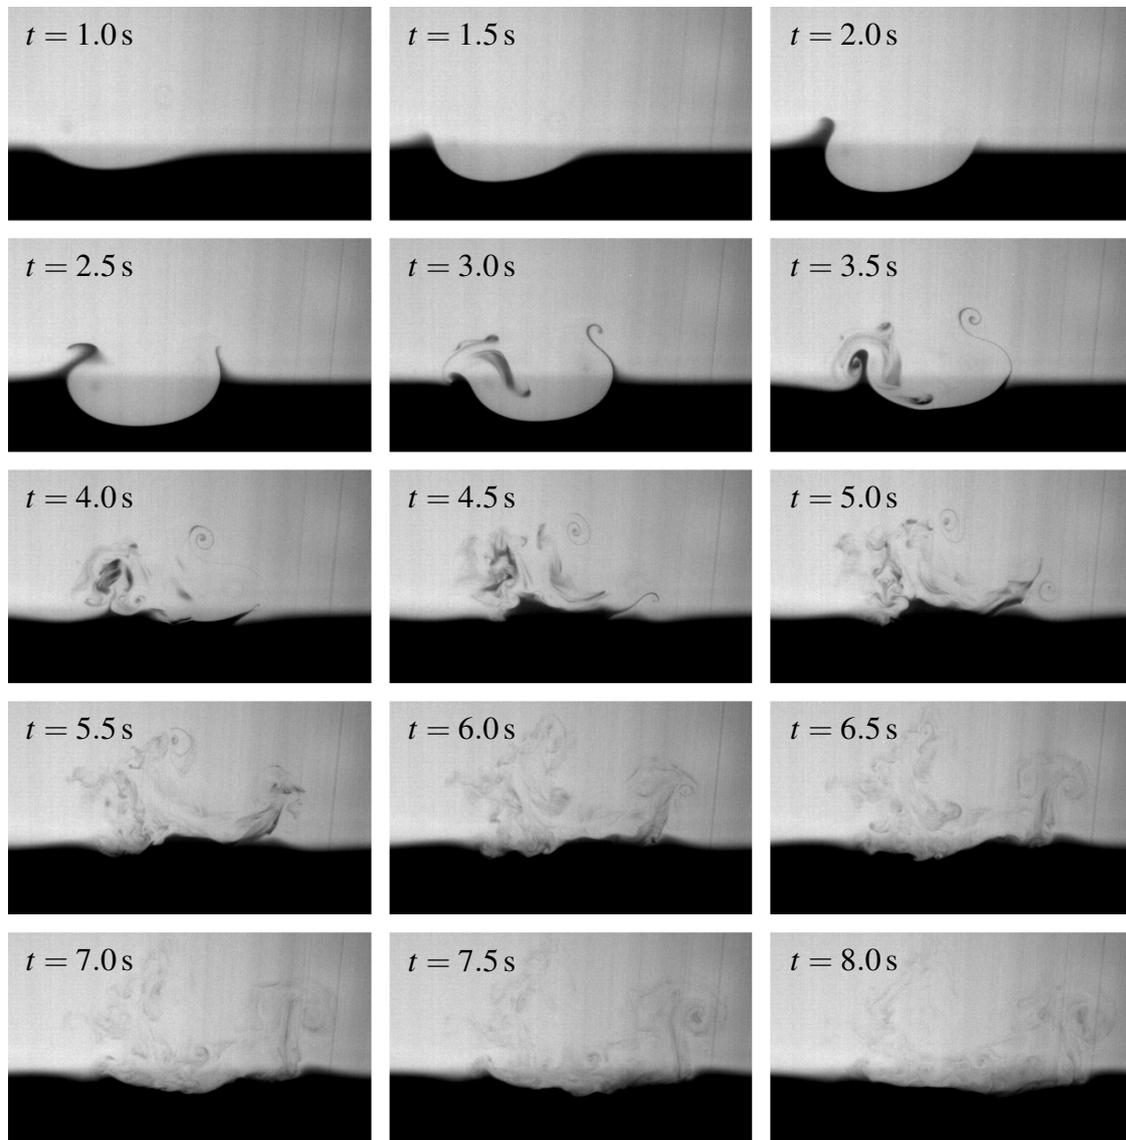


Fig. 5.7 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(Re_0, Ri_0, \theta_0) = (1810, 3.09, 20^\circ)$. The time separation between consecutive snapshots is 0.5 seconds, with the first frame taken at time $t = 1$ s. Note that the advective timescale $a/U = 1.33$ s.

snapshots presented in each of Figures 5.3 to 5.7. In plots (e) and (f), filled circles mark the position of the ring centroids and ring centre every 0.5 seconds.

As the primary ring begins to feel the presence of the density interface, the plane diameter $D(t)$ increases due to the interaction of the ring with the baroclinic vorticity generated on the deformed regions of the interface. For $\theta_0 > 0^\circ$, the portion of the vortex ring centred at core C_1 is in closer proximity with the interface than that corresponding to core C_2 . For $t < 1$, the effect of this on the structural evolution of the ring appears to be small for the range of θ_0 considered, though slight reductions in $\theta_c(t)$ are observed for $\theta_0 = 10^\circ, 20^\circ$.

For $1 < t < 2$, $D(t)$ continues to increase, reaching its maximum in all cases between $1.8 < t < 2$. The vertical component of the buoyancy force acting on the ring leads to a reduction in its bulk vertical momentum, the effect of which can be seen between $1 < t < 1.5$ by the inflection points in $Z(t)$ and the increase in $\theta_p(t)$ for $\theta_0 > 0^\circ$. For the obliquely propagating cases, as C_1 impinges on the interface before C_2 , there is azimuthal variation in the total time-integrated vertical buoyancy force that has acted on the ring, which is maximal at C_1 and minimal at C_2 . Over $1 < t < 2$, this leads to a reduction in $\theta_c(t)$ of about 60%. Though we are unable to directly observe the structural evolution of the ring outside of the central plane of the flow, the reduction in $\theta_c(t)$ in combination with the azimuthal variation in time-integrated vertical buoyancy force (described above) suggest that the ring deforms as it impinges onto the interface. Such deformations were observed in Lim (1989) and Stock (2006), for the cases $\theta_0 = 38.5^\circ$ impacting a no-slip wall, and $\theta_0 = 45^\circ$ impacting a density interface respectively.

For all θ_0 considered, the ring reaches its maximum-penetration depth into the lower layer around $2.3 < t < 2.5$. Around this time, the propagation angle passes through $\theta_p(t) = 90^\circ$, indicating that the recoiling of the density interface is beginning to eject the vortex-ring fluid back into the upper layer. For the case $\theta_0 = 0^\circ$, the vortex-core centroids move slightly in the positive x -direction, leading to $\theta_p(t)$ increasing from 0° to 180° as the interface begins to eject the ring fluid back into the upper layer, rather than $\theta_p(t)$ decreasing from 0° to -180° . Using least squares to fit a line through $\theta_c(t)$ for $\theta_0 = 0^\circ$ between $t = 0$ and $t = 2$ gives

$$\theta_c(\theta_0 = 0^\circ, t) = 0.18 - 0.15t,$$

indicating that the sign of $d\theta_p/dt$ as the interface recoils may be related to the precision limit of 0.1 degrees for setting θ_0 in our experiments. The subtle asymmetry of the 5th LIF snapshot of Figure 5.3 at $t = 3.0$ s suggests that the precise value of θ_0 for this experiment may differ slightly from its nominal value. As the precision limit of θ_0 is small compared to the non-zero values of θ_0 considered in this chapter, we do not pursue this further.

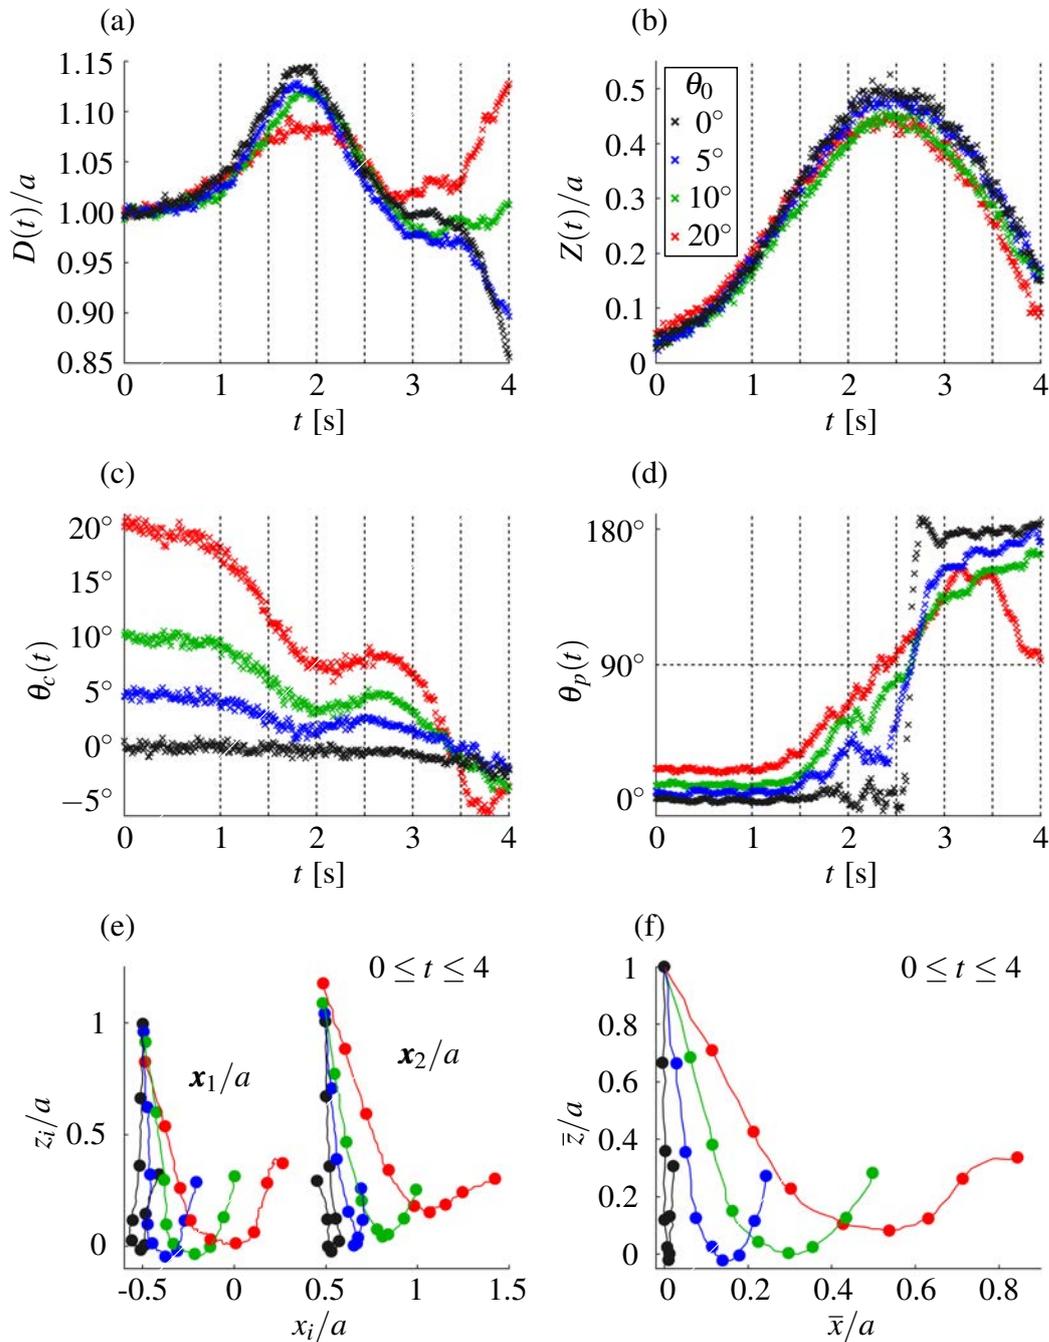


Fig. 5.8 Plots of the observed measures of the vortex ring introduced in Section 5.1, in addition to the non-dimensionalised penetration depth $Z(t)/a$ into the lower layer. For each figure, data are plotted between $t = 0$ and $t = 4$ seconds. Vertical dashed lines are drawn in plots (a) to (d), marking the instances in time that correspond to the first seven LIF snapshots presented in each of Figures 5.3, 5.5, 5.6 and 5.7. In plots (e) and (f), filled circles mark the position of the ring centroids and ring centre every 0.5 seconds. The legend corresponding to all plots is given in (b).

The plane diameter $D(t)$ decreases over the period $2 < t < 3$, due to horizontal buoyancy forces acting to compress the ring while it is surrounded radially by lower-layer fluid. For $2 < t < 2.5$, a small increase is observed in $\theta_c(t)$ due to the horizontal buoyancy forces pushing core C_1 to the right more rapidly than the shallower C_2 , while the difference in vertical positions of the vortex core is approximately maintained. This can also be seen in the vortex-centre trajectories plotted in Figure 5.8 (e).

For $3 < t < 4$, the density interface recoils, ejecting the primary ring back into the upper layer. At this stage, secondary structures are starting to become the more prominent features in the flow, and the flow appears to be transitioning into turbulence. For the cases $\theta_0 = 5^\circ$ and 10° , the recoil ring (analogous to that described in Section 5.2.1) effectively guides the primary ring back into the upper layer, causing $\theta_p(t)$ to further increase. To visualise this stage of the flow, Figure 5.9 (a) presents the LIF snapshot at $t = 4$ s for the case $\theta_0 = 10^\circ$, overlaid with velocity vectors and contours of Q_{2D} to identify vortex cores C_1 and C_2 and the cross-sections of the recoil ring (denoted S_1 and S_2) in the central plane of the flow.

Figure 5.9 (b) presents the corresponding LIF snapshot for $\theta_0 = 20^\circ$ at $t = 4$ s, with overlaid velocity vectors and Q_{2D} contours. From the plots presented in Figure 5.8, for $3.5 < t < 4$ we see an increase in $D(t)$, a decrease in $\theta_p(t)$ and a significant rightward drift in x_2 , all in contrast to the evolution of these measures for smaller θ_0 . Comparing Figures 5.9 (a) and (b), these differences appear to be attributable to the recoil ring (labelled S) failing to pass around the outside of the primary ring on the side of C_2 . In particular, the portion of the recoil ring centred at S_2 acts to advect the portion of the primary ring centred at C_2 mostly toward the right, as opposed to mostly vertically for the cases corresponding to smaller θ_0 . As S_2 is unable to pass around the outside of C_2 , the vertical buoyancy flux associated with the splashing mechanism described in Section 5.2.1 appears to be significantly reduced.

5.2.3 Oblique propagation: azimuthal flow

For the cases $\theta_0 = 5^\circ, 10^\circ$, significant flow converging onto the central plane of the ring–interface interaction is observed in the LIF snapshots presented in Figures 5.5 and 5.6, emerging around $t = 5$ seconds for both cases. The flow converges onto the plane at two locations, coinciding with the position of vortex-core cross sections C_2 and S_2 . Here, we argue that these flows are associated with azimuthal pressure gradients in the primary and recoil rings, where a pressure minimum within each ring is located on the central plane of the flow.

To help visualise the convergence of fluid at the location of C_2 , Figures 5.10 (a) and (b) each present an LIF snapshot, for the cases $\theta_0 = 5^\circ$ at $t = 5.5$ s and $\theta_0 = 10^\circ$ at $t = 6.5$ s respectively. Each snapshot is overlaid with velocity vectors, and contours where $\nabla \cdot \mathbf{u}_{2D} > 0$,

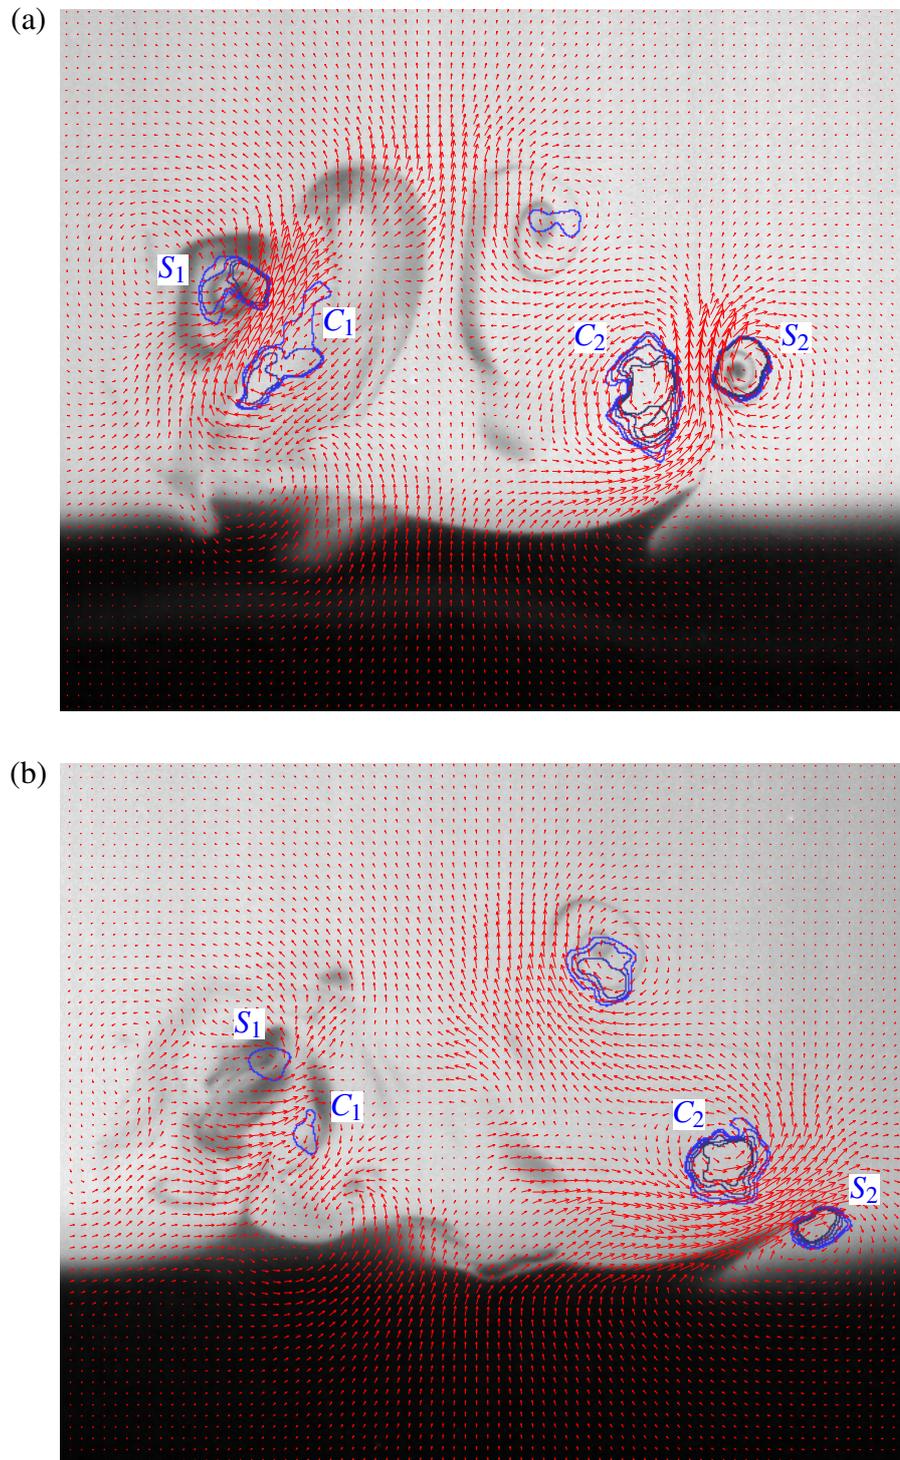


Fig. 5.9 LIF snapshots at $t = 4$ s for the cases (a) $\theta_0 = 10^\circ$ and (b) $\theta_0 = 20^\circ$, overlaid with velocity vectors and contours of $Q_{2D} = 1, 2, 3, 4$ and 5 s^{-2} . Vortex S and cores C_1 and C_2 are labelled in both figures where they are made reference to in the text.

associated with the flow converging onto the central plane at C_2 . Note that $\nabla \cdot \mathbf{u}_{2D} > 0$ corresponds to a source of fluid on the central plane (and a divergence of in-plane velocity vectors), whereas $\nabla \cdot \mathbf{u}_{2D} < 0$ corresponds to a sink on the central plane (and convergence of in-plane velocity vectors). By relating these figures to their corresponding LIF-snapshot sequences (Figures 5.5 and 5.6), it can be seen that upper-layer fluid emerges into the central plane around the location of C_2 , pushing outward the surrounding fluid in the central plane.

The in-plane divergence associated with S_2 (identified in Figures 5.10 (a) and (b) by the drawn circles) can be seen in Figures 5.5 and 5.6 by the emergence of lower-layer fluid at the location of S_2 from $t = 5.5$ s onwards, with the converging flow particularly pronounced for the 5° case. Subsequently, for the $\theta_0 = 5^\circ$ case, the emerging lower-layer fluid appears to undergo extensional straining along an axis at 45° to the vertical. This does not appear to occur for the $\theta_0 = 10^\circ$ case, where instead the flow appears to be break down and enter a state of dissipative relaxation.

Due to the oblique impingement of the primary ring on the interface, baroclinic vorticity is produced azimuthally asymmetrically on the deformed interface, with a greater volume of baroclinic-vorticity-carrying fluid swept around the side of C_1 and entrained back through the centre of the ring relative to C_2 . To understand the effect this will have on the primary ring, consider the Biot-Savart law, which relates the vorticity field to the rotational component of the velocity field by

$$\mathbf{u}_{\text{rot}}(\mathbf{x}, t) = \frac{1}{4\pi} \int_V \frac{\boldsymbol{\omega}(\mathbf{x}', t) \times (\mathbf{x} - \mathbf{x}') dV'}{|\mathbf{x} - \mathbf{x}'|^3}. \quad (5.5)$$

As the baroclinic vorticity is of opposite sign to the vorticity of the primary ring, the rotational velocity component associated with C_1 reduces at a greater rate compared to C_2 , due to the cancellation of vorticity in the integral of Equation (5.5) and the $|\mathbf{x} - \mathbf{x}'|^{-2}$ scaling of the integrand. The effect of this can be seen for example by the Q_{2D} contours in Figure 5.9 (a), where higher Q_{2D} contours are observed on C_2 and S_2 relative to C_1 and S_1 .

The decrease in \mathbf{u}_{rot} at C_1 leads to a local increase in pressure at C_1 , consequently introducing a non-zero azimuthal pressure gradient in the primary ring. Assuming the pressure along the primary ring decreases with increasing distance from C_1 as a result of reduced vorticity cancellation in Equation (5.5), it follows that a pressure minimum along the ring exists at C_2 . It can similarly be argued that a pressure minimum along the recoil ring exists at S_2 . The local pressure minima associated with each ring would explain the observed flow at C_2 and S_2 converging onto the central plane.

Figure 5.11 presents LIF snapshots of the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2140, 2.24, 5^\circ)$, to illustrate the evolution of the recoil ring as the flow transitions to turbulence for a case with similar

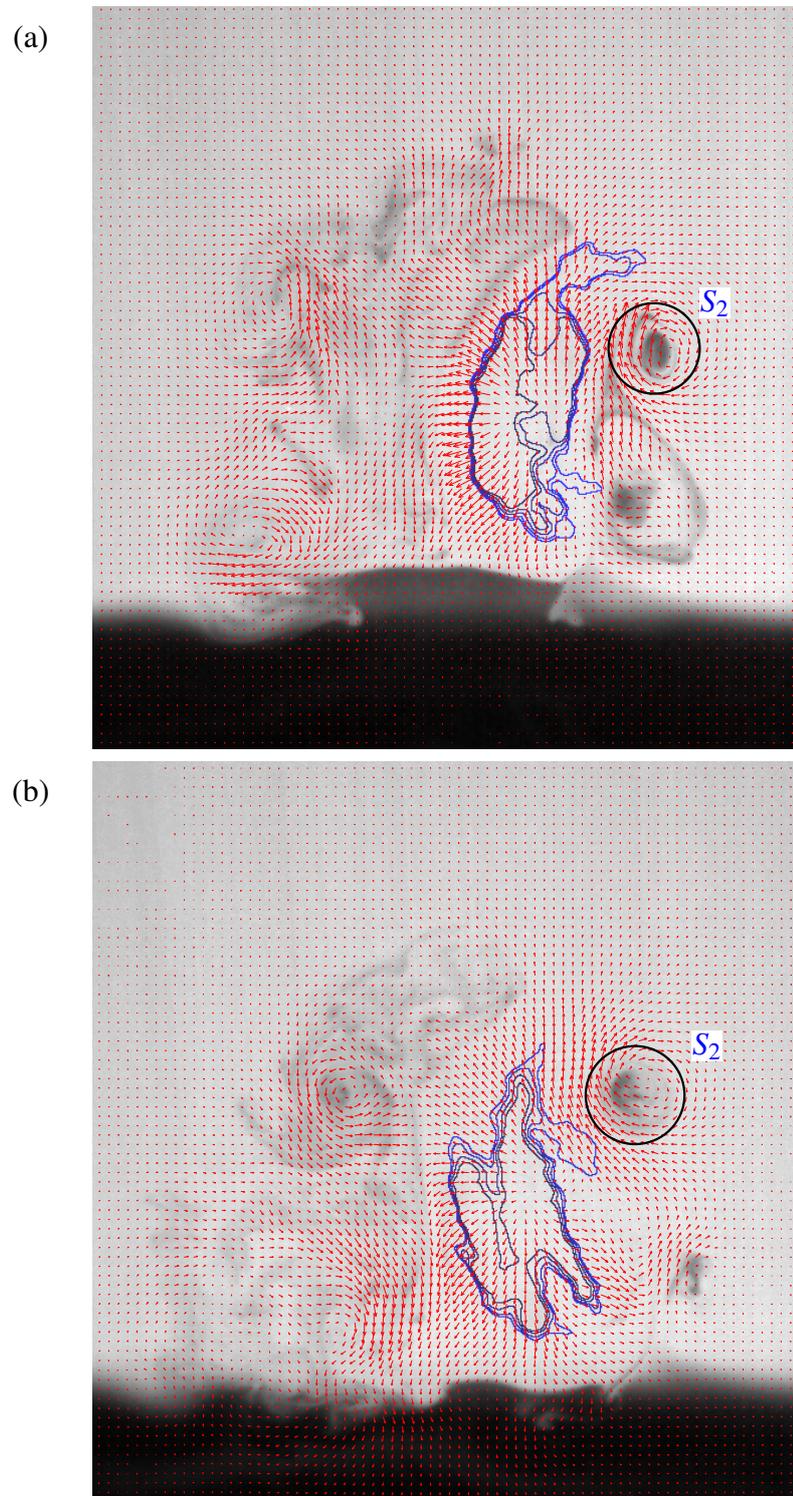


Fig. 5.10 LIF snapshots at (a) $t = 5.5$ s for $\theta_0 = 5^\circ$, and (b) $t = 6.5$ for $\theta_0 = 10^\circ$. Snapshots are overlaid with velocity vectors and contours of $\nabla \cdot \mathbf{u}_{2D} = 5, 10, 15, 20$ and 25 s^{-1} . Circles are drawn to clearly identify the vortex core S_2 .

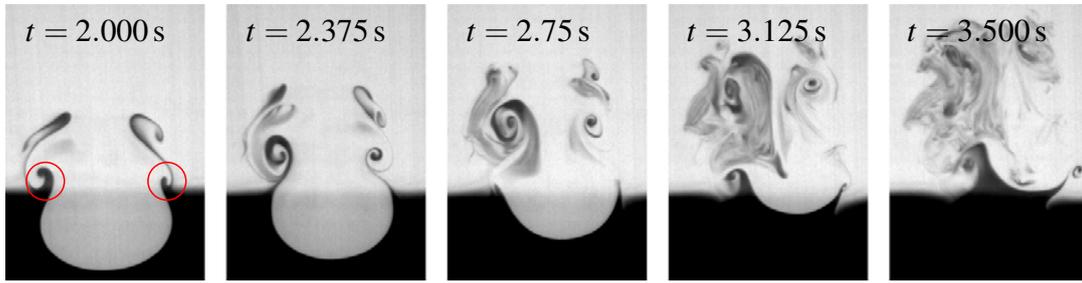


Fig. 5.11 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2140, 2.24, 5^\circ)$. Each snapshot is labelled with its associated instant in time. Red circles are drawn around the recoil ring cores exhibited in the first snapshot for clarity.

Re_0 but significantly lower Ri_0 . Red circles are drawn around the vortex cores associated with the recoil ring in the first snapshot. Our observations indicate that the azimuthal flow associated with the recoil ring for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (1860, 3.09, 5^\circ)$ does not manifest itself in a similar way below some critical Richardson number Ri_c , where $2.24 < \text{Ri}_c < 3.09$. In an analogous way to the Crow-like instability identified by Olsthoorn and Dalziel (2017), it is plausible that the disappearance of the azimuthal-flow structure is related to a transition in the timescale dominating the evolution of the recoil ring, where for $\text{Ri}_0 < \text{Ri}_c$ the instability leading to strong azimuthal flow is unable to develop before the turbulent breakdown of the ring. We do not investigate this conjecture in any detail here, but such flow transitions imply a sensitive dependence of the evolution of the ring–interface interaction on Ri_0 for small θ_0 .

5.3 Penetrating rings at small propagation angles

The experiments presented in this section are those listed in the second and third rows of Table 5.1, corresponding to $(\text{Re}_0, \text{Ri}_0) = (2680, 1.48)$ and $(3270, 0.98)$ respectively and $\theta_0 \leq 20^\circ$. The cases $\text{Ri}_0 = 1.48$ and 0.98 are respectively above and below the case $\text{Ri}_0 = 1.24$ described as ‘penetrative’ by Olsthoorn and Dalziel (2017), who show at this value of Ri_0 for $\theta_0 = 0^\circ$ that the Crow-like instability they identify is unable to develop before the ring–interface interaction loses coherence and transitions into turbulence. The cases considered in this section also correspond to two of the three penetrative cases of $(\text{Re}_0, \text{Ri}_0)$ investigated in our periodic-mixing experiments presented in Chapter 6.

The layout of this section is similar to that of Section 5.2. In Section 5.3.1, we describe the evolution of the two cases corresponding to $\theta_0 = 0^\circ$. Then, in Section 5.3.2, we introduce LIF snapshots for the cases corresponding to $\theta_0 > 0^\circ$, identifying features in the flow that

are not present in the $\theta_0 = 0^\circ$ case. In Section 5.3.3, we discuss the influence of θ_0 on the evolution of the primary vortex ring. Finally, in Section 5.3.4, we describe the secondary structures that emerge in the flow when $\theta_0 > 0^\circ$ and offer explanations for their origins.

5.3.1 Vertical propagation

We first consider the case $(Re_0, Ri_0) = (2680, 1.48)$, for which selected LIF snapshots of the ring–interface interaction are presented in Figure 5.12. Note that the time separation between consecutive snapshots is not the same for all consecutive-snapshot pairs. Instead, snapshots have been chosen to make the evolution of (visible) secondary structures in the flow as clear as possible. To complement these snapshots, Figure 5.13 presents LIF snapshots at $t = 2.00$ s, $t = 2.75$ s and $t = 3.50$ s, overlaid with velocity vectors and ω_y contours.

For $t \lesssim 1$ s, as the ring impinges onto the interface, baroclinic vorticity is produced on the deformed interface and swept outwardly by the ring, accumulating on the rim of the deformed area of the interface. The evolution of the flow up to this stage is similar to the non-penetrative case investigated in Section 5.2.1. Shortly after the ring penetrates into the lower layer (corresponding to $\bar{z} < 0$), an in-plane stagnation point develops immediately above the vortex ring, indicated first by the cross drawn in Figure 5.13 (a) at $t = 2.00$ s. This stagnation point can be interpreted as the point along the central axis of the ring where the ring-induced entrainment of fluid to pass through its centre is balanced by the flow induced by the baroclinic vorticity, which attempts to eject the ring-entrained upper-layer fluid in the lower layer back into the upper layer.

Two features of the flow are exhibited in all ring–interface interactions considered in this section. Firstly, the ejection of ring-entrained upper-layer fluid from the lower layer back into the upper layer by the surrounding baroclinic vorticity. Hereafter, we refer to the ejected ring fluid in the lower-layer region as the ‘ejected fluid column’, or EFC. Secondly, the baroclinic vorticity generated on the deformed interface between the ring-entrained upper-layer fluid and surrounding lower-layer fluid in the lower layer. We refer to this interface as the ‘baroclinic-vorticity column’, or BVC.

As the ring penetrates the lower layer and reaches its maximum-penetration depth (at around $t = 4$ s), the interaction of the ring with the surrounding baroclinic vorticity results in the outermost regions of the vortex-ring core being peeled away, leading to a reduction in the ring diameter and the ejection of ring-entrained fluid into the EFC. The associated reduction in circulation of the ring is made apparent by the $|\omega_y| = 2 \text{ s}^{-1}$ contours in Figures 5.13 (b) and (c) that extend from the ring into the EFC and above. As fluid leaves the EFC and returns to the upper-layer region, the previously-described baroclinic vorticity (of opposite circulation to the ring) accumulating immediately above the resting position of the

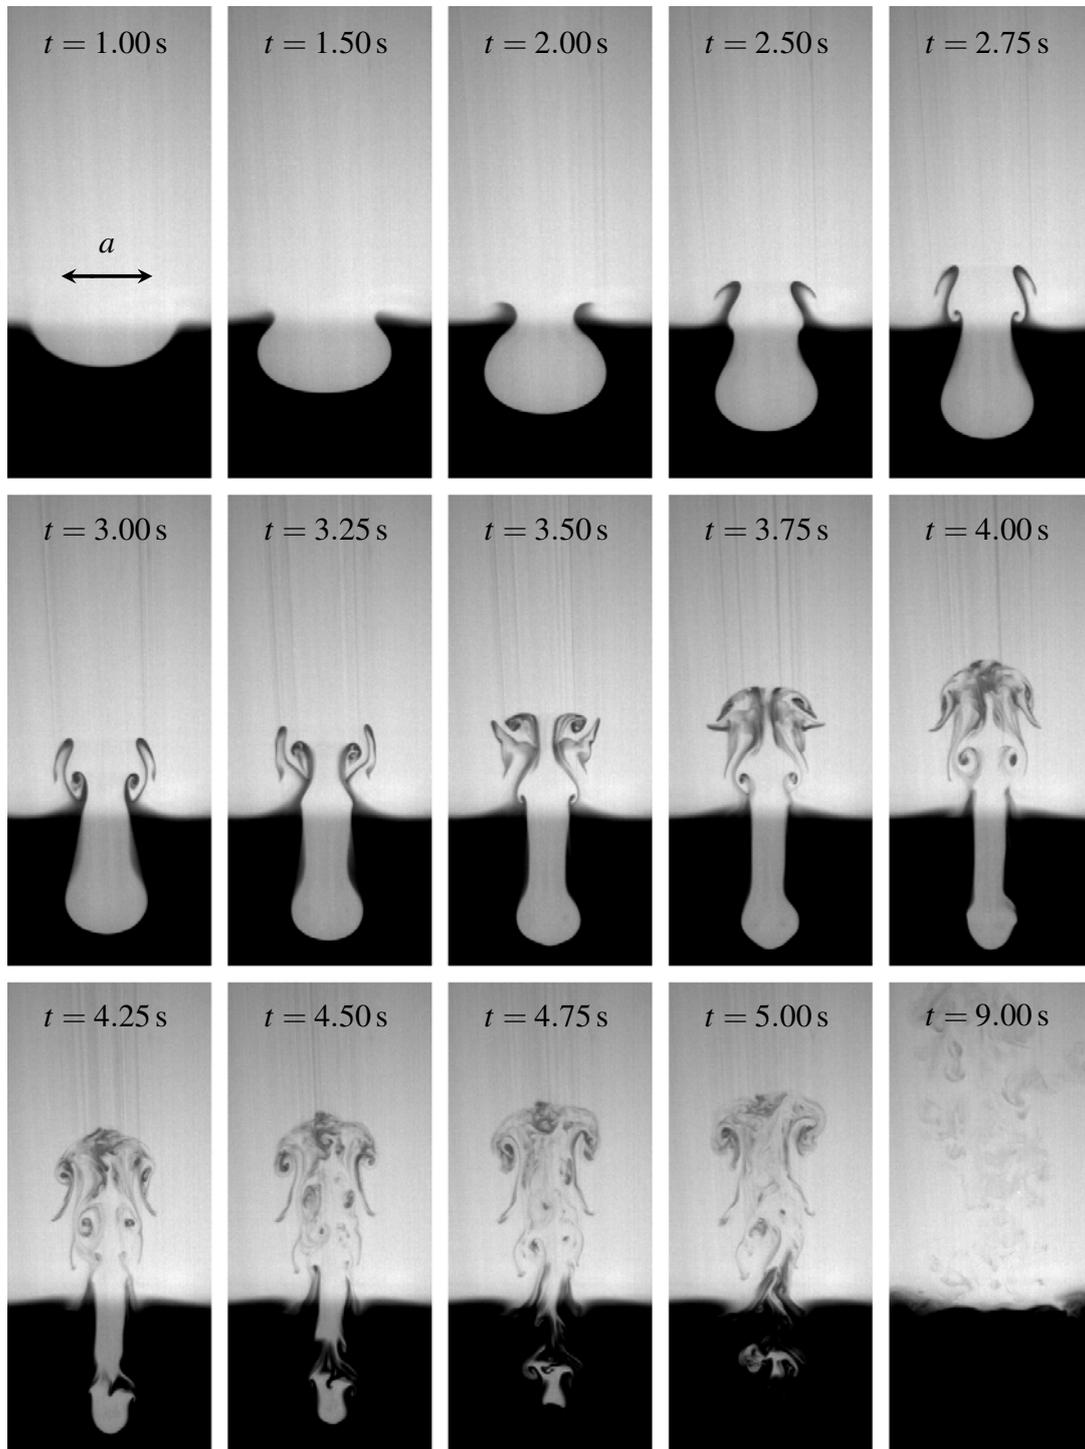


Fig. 5.12 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(Re_0, Ri_0, \theta_0) = (2680, 1.48, 0^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.90$ s.

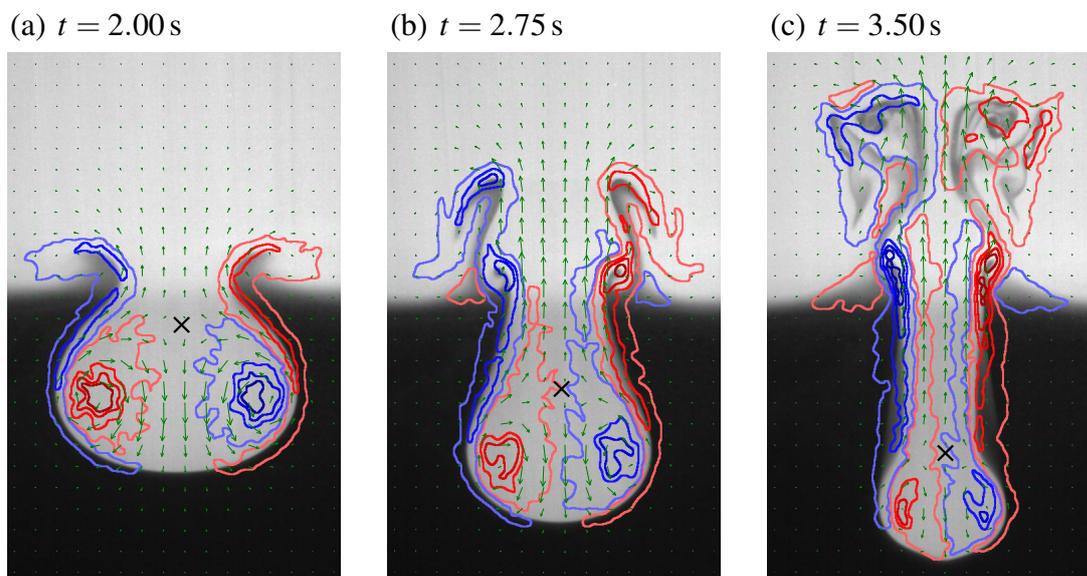


Fig. 5.13 LIF images at (a) $t = 2.00$ s, (b) $t = 2.75$ s and (c) $t = 3.50$ s, overlaid with (sparsely drawn) velocity vectors and ω_y contours, with the blue and red contours respectively representing anti-clockwise and clockwise in the central plane of the flow. The in-plane vorticity contours are drawn for $|\omega_y| = 2, 15, 25$ and 35 s^{-1} . Black crosses are drawn to identify the in-plane stagnation point separating the ring-entrained fluid and ring fluid ejected by the surrounding baroclinic vorticity.

interface separates in discrete events, leading to the formation of secondary vortex rings. The accumulated baroclinic vorticity can be seen in Figures 5.13 (b) and (c) shortly before two distinct separation events, and the propagation of the secondary rings these events generate into the upper layer can be seen in the subsequent LIF snapshots in Figure 5.12, made clear by the lower-layer fluid entrained into the secondary-ring cores.

The BVC appears to remain stable for most of the ring–interface interaction, preventing significant mixing from occurring between the lower-layer fluid and the EFC. As the interface recoils during $4 < t \leq 5$ seconds, the BVC begins to destabilise, indicated by the disconnected patches of upper-layer fluid in the lower layer exhibited in Figure 5.12. These patches appear to remain unmixed as they rise towards the upper layer, suggesting no significant mixing occurs in the lower layer in this stage of the flow. The secondary rings that propagate into the upper layer increase the vertical buoyancy flux associated with the splashing mechanism. Consequently, the vast majority of mixing appears to occur in the upper layer.

The evolution of the ring–interface interaction is considerably different for the case $(\text{Re}_0, \text{Ri}_0) = (3270, 0.98)$, as shown in Figure 5.14. The BVC is seen to destabilise as early as $t = 2.5$ s while the vortex ring is still penetrating into the lower layer, leading to a rapid loss of coherence in the flow as it transitions to turbulence. This appears to result in a substantial

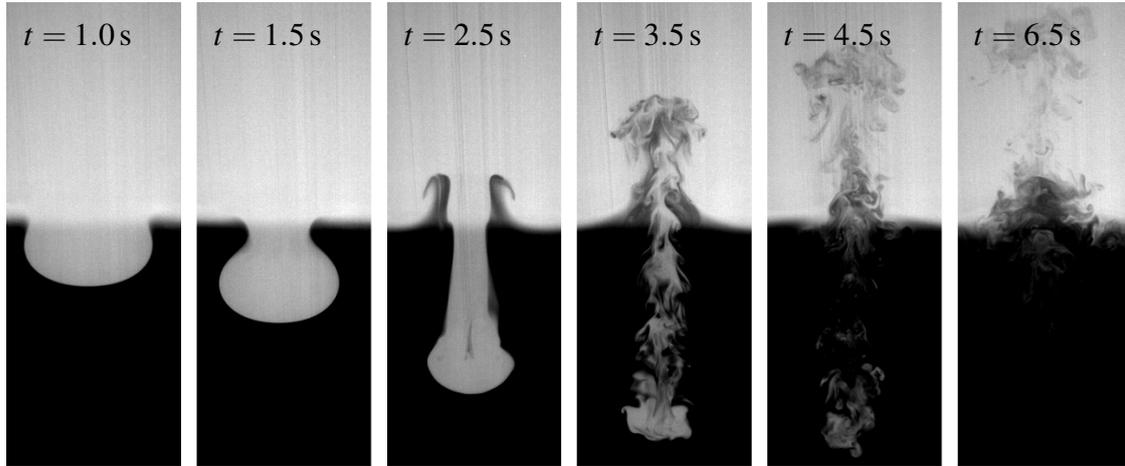


Fig. 5.14 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(Re_0, Ri_0, \theta_0) = (3270, 0.98, 0^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.74$ s.

increase in the total interfacial area between the lower-layer fluid and ring-entrained upper-layer fluid relative to the case $(Re_0, Ri_0) = (2680, 1.48)$, indicative of a substantial increase in the amount of upper-layer fluid mixed into the lower layer. Lower-layer fluid is ejected into the upper layer, though unlike the $(Re_0, Ri_0) = (2680, 1.48)$ case, the associated vertical buoyancy flux does not appear to be enhanced by the development of large-scale secondary vortex rings. Nonetheless, the ejection of lower-layer fluid into the upper layer is suggestive that significant mixing occurs in the upper layer as well as the lower layer.

5.3.2 Oblique propagation: a first glance

Before discussing the evolution of the primary ring and secondary structures in the flow for the oblique cases in more detail, we briefly point out the similarities and differences in the evolution observed in the central plane of the flow for given (Re_0, Ri_0) as θ_0 increases. We follow the order of our discussion given in Section 5.3.1, first considering the cases corresponding to $(Re_0, Ri_0) = (2680, 1.48)$, then the cases corresponding to $(Re_0, Ri_0) = (3270, 0.98)$.

For $(Re_0, Ri_0) = (2680, 1.48)$, Figures 5.15, 5.16 and 5.17 present selected LIF snapshots of the ring–interface interaction for the cases $\theta_0 = 5^\circ$, 10° and 20° respectively. For $\theta_0 = 5^\circ$, azimuthal asymmetries are apparent in the flow as early as $t = 2.0$ s, indicating that different dynamics will be at play along the circumference of the BVC as the flow evolves, both in the $\theta_0 = 5^\circ$ case and at larger values of θ_0 . In the view of the camera, the interface between the right-hand side of the EFC and the surrounding lower-layer fluid will be more susceptible to

having convectively unstable regions relative to the interface along the left-hand side, due to the orientation of the EFC relative to the vertical. For clarity, hereafter we refer to the right-hand and left-hand sides of the EFC as the ‘front’ and ‘back’ of the EFC respectively, and the interfacial regions outside the view of the camera as the ‘sides’ of the EFC. The terms front, back and sides are similarly applied to the BVC.

For the $\theta_0 = 5^\circ$ case, the structure of the lower-layer fluid ejected into the upper layer between $t = 2.50$ s and $t = 3.50$ s has some resemblance to the analogous structure for the $\theta_0 = 0^\circ$ case over the same interval in time (shown in Figure 5.12), though much of the coherence is lost. Additionally, a substantial majority of the ejected lower-layer fluid visible in the upper layer is above the back of the EFC; a feature also seen for the $\theta_0 = 10^\circ$ and 20° cases. We propose in Section 5.3.4 that this transport of lower-layer fluid is attributable to the generation of axial vorticity, leading to a dipole that extends through much of the EFC.

For the $\theta_0 = 5^\circ$ and 10° cases, between $t = 2.50$ s and $t = 3.50$ s, lower-layer fluid gradually enters the central plane of the flow immediately above the vortex ring in the region occupied by the EFC, entering at the front and back of the EFC for the cases $\theta_0 = 5^\circ$ and 10° respectively (see Figures 5.15 and 5.16). Without being able to probe the full three-dimensional evolution of the flow, it is difficult to identify the mechanisms responsible for these out-of-plane flows. These flows may be attributable to the axial dipole described above, to a slight misalignment between the light sheet and the central plane of the flow, or to something else. As this flow does not appear to have significant consequences for the subsequent evolution of the flow from our observations of the central plane, we do not pursue this further.

For the case $\theta_0 = 20^\circ$ (with LIF snapshots presented in Figure 5.17), lower-layer fluid is seen to roll up along the front of and immediately above the EFC from $t = 3.0$ s to $t = 3.8$ s. Additionally, from $t = 3.8$ s to $t = 4.6$ s, upper-layer fluid is seen to converge onto the central plane of the flow above the back of the EFC, in the central region of the ejected lower-layer fluid residing in the upper layer. We return to these features in Section 5.3.4. For all cases corresponding to $(Re_0, Ri_0) = (2680, 1.48)$, a substantial majority of the mixing appears to occur in the upper layer.

For $(Re_0, Ri_0) = (3270, 0.98)$, Figures 5.18, 5.19 and 5.20 present selected LIF snapshots of the ring–interface interaction for the cases $\theta_0 = 5^\circ$, 10° and 20° respectively. The cases corresponding to $\theta_0 = 5^\circ$ and 10° are very similar to the $\theta_0 = 0^\circ$ case (for which LIF snapshots are presented in Figure 5.14), insofar as the primary ring and the BVC rapidly destabilise, and a substantial fraction of the total mixing appears to occur in the lower layer. For $\theta_0 = 20^\circ$, the cross section of the EFC in the central plane of the flow thins from $t = 2.0$ s

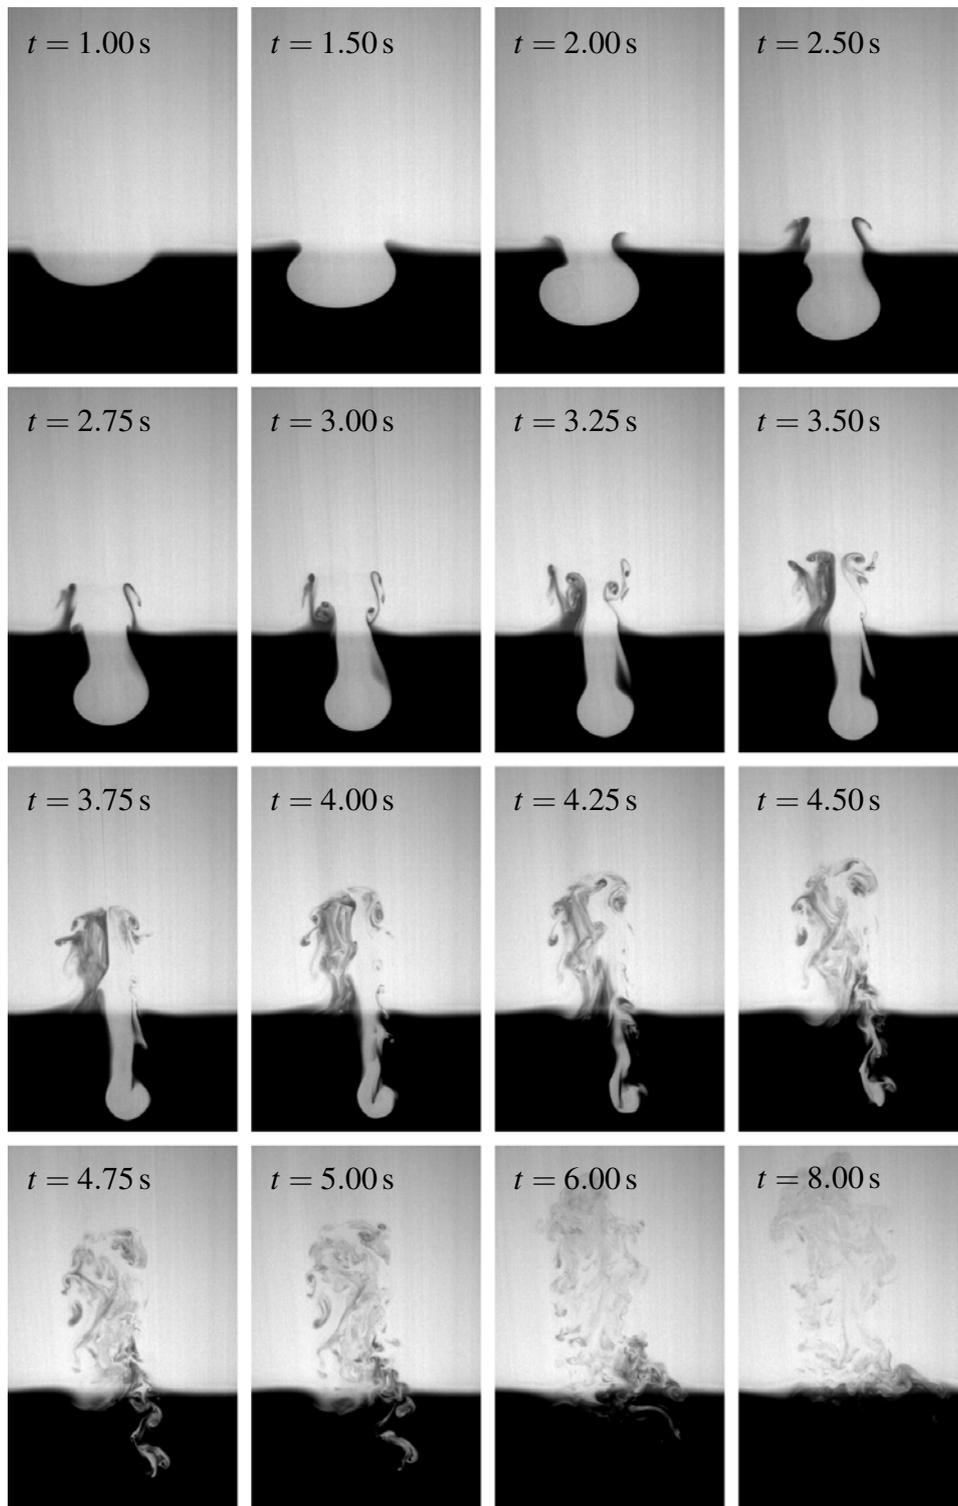


Fig. 5.15 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 5^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.90$ s.

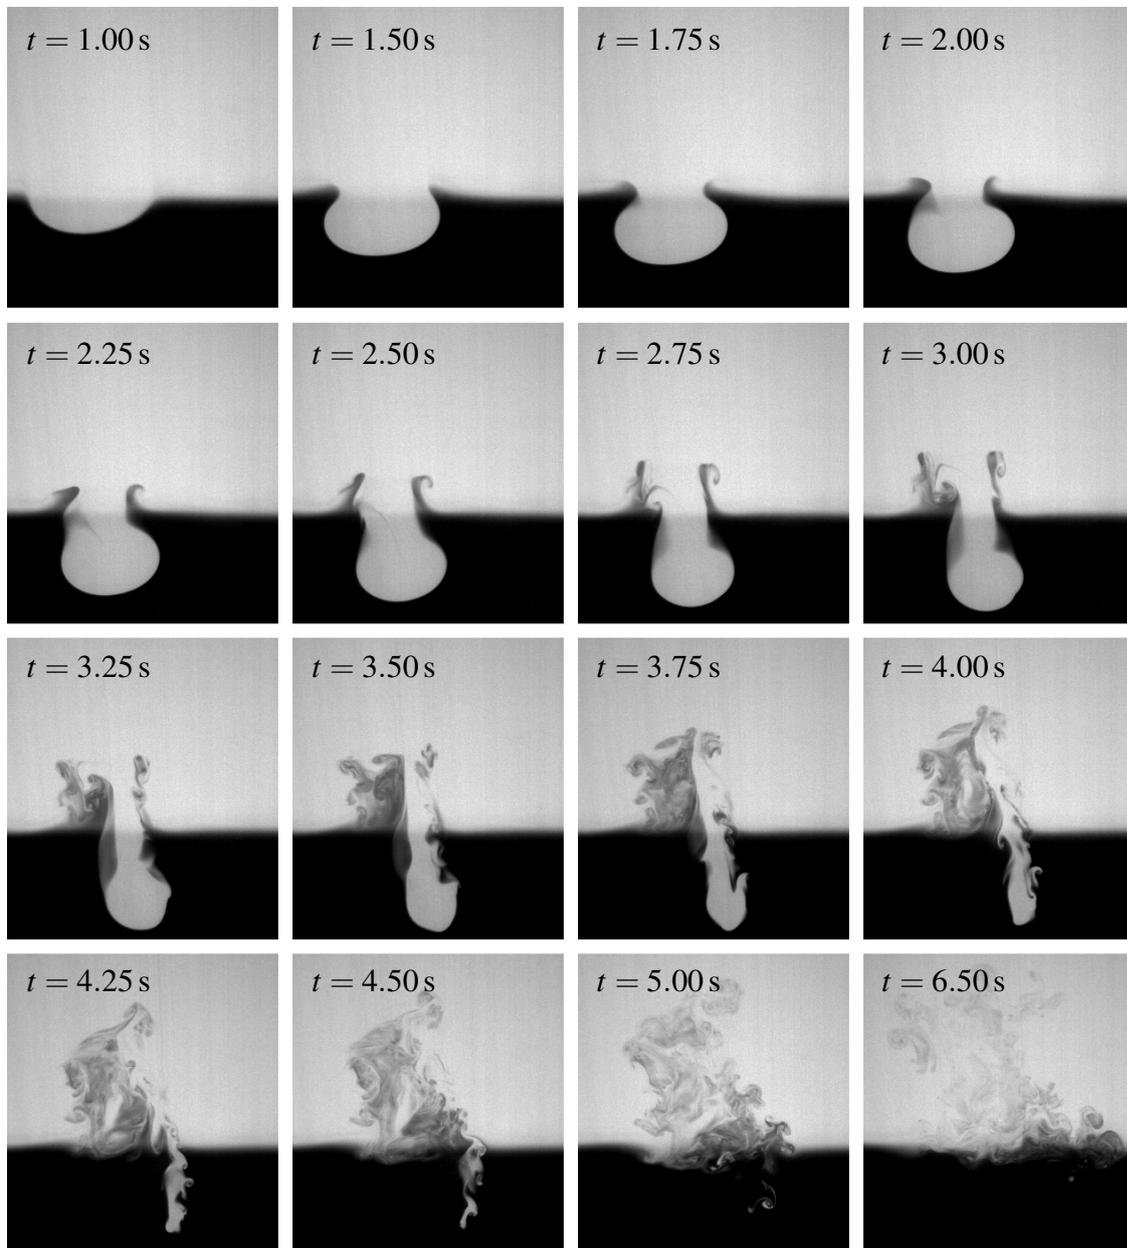


Fig. 5.16 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 10^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.90$ s.

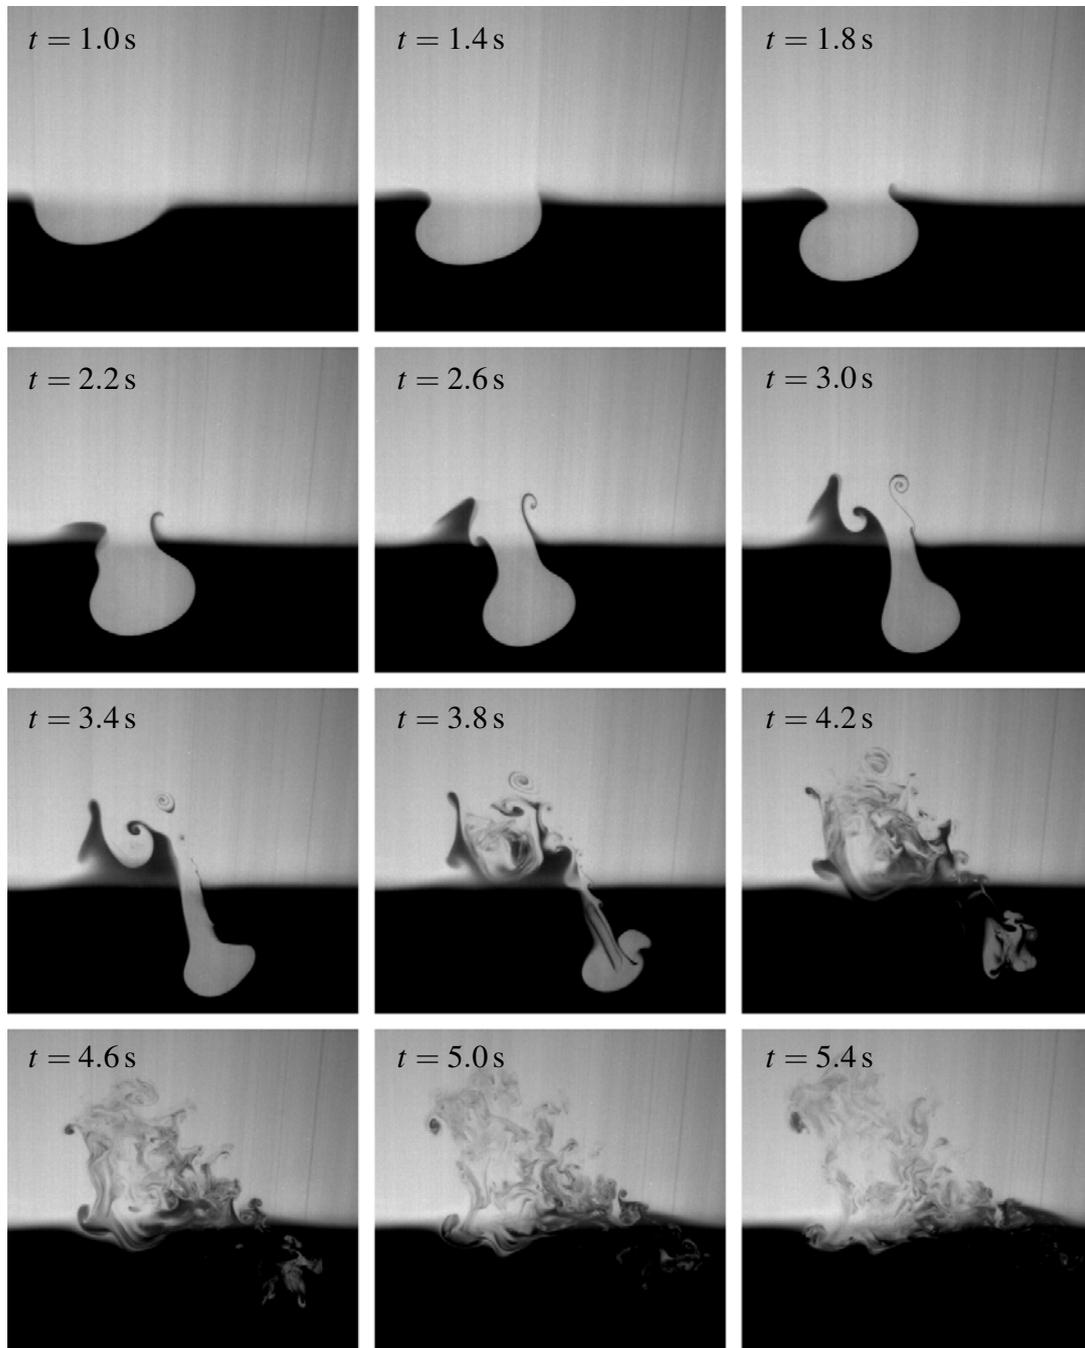


Fig. 5.17 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 20^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.90$ s.

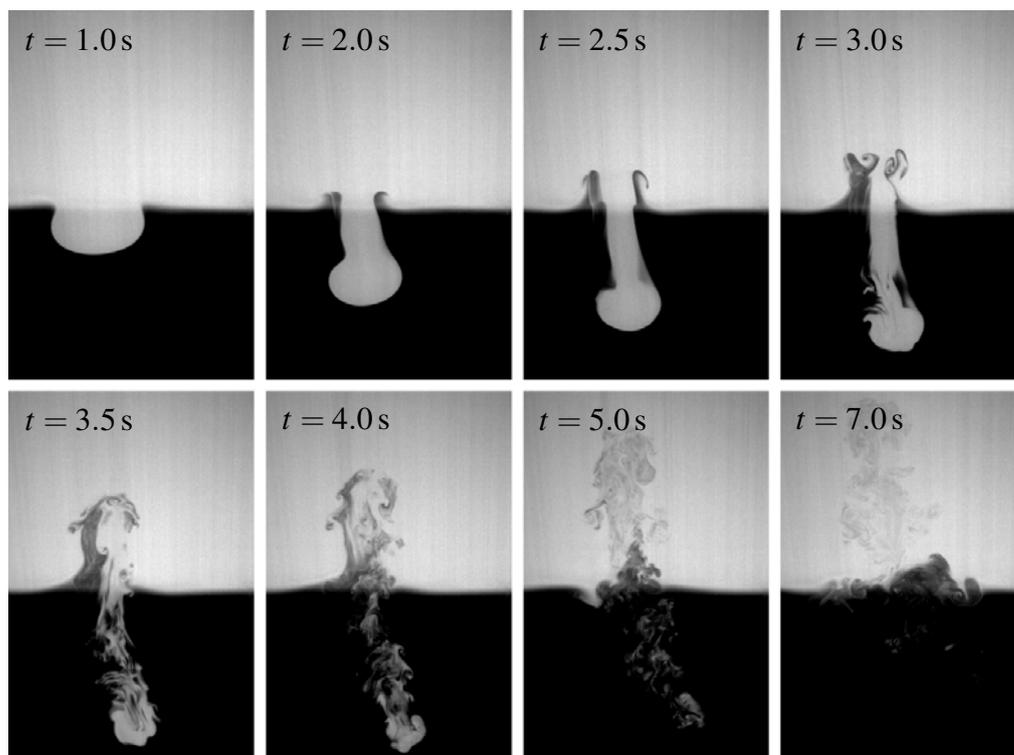


Fig. 5.18 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (3270, 0.98, 5^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.74$ s.

to $t = 3.0$ s, after which disconnected patches of upper-layer fluid emerge in the vicinity of the EFC from $t = 3.0$ s to $t = 4.0$ s. We return to these structural features in Section 5.3.4.

5.3.3 Oblique propagation: evolution of the primary ring

Before considering in greater detail the secondary structures that develop in the flow, we discuss the evolution of the primary ring for the case $(\text{Re}_0, \text{Ri}_0) = (2680, 1.48)$. Figures 5.21 (a) to (f) plot the observed measures of the vortex ring introduced in Section 5.1, in addition to the non-dimensionalised penetration depth $Z(t)/a$ into the lower layer. For the measures derived from \mathbf{x}_1 and \mathbf{x}_2 , data in Figures 5.21 (a), (c) and (d) are plotted up to the stage where either C_1 or C_2 can no longer be reliably identified in the central plane of the flow using the tracking method described in Section 5.1.

As the ring approaches the density interface, the plane diameter $D(t)$ increases slightly due to the interaction of the ring with the baroclinic vorticity developing on the deformed regions of the interface, reaching a maximum in all cases at $t \approx 1.2$ s. Similarly to the

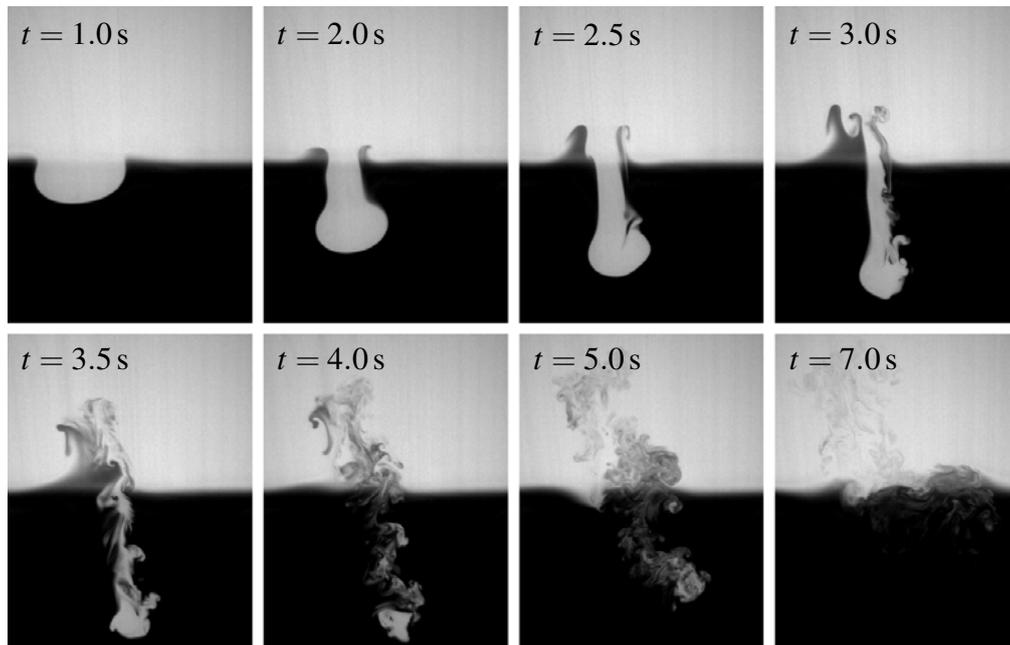


Fig. 5.19 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (3270, 0.98, 10^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.74$ s.

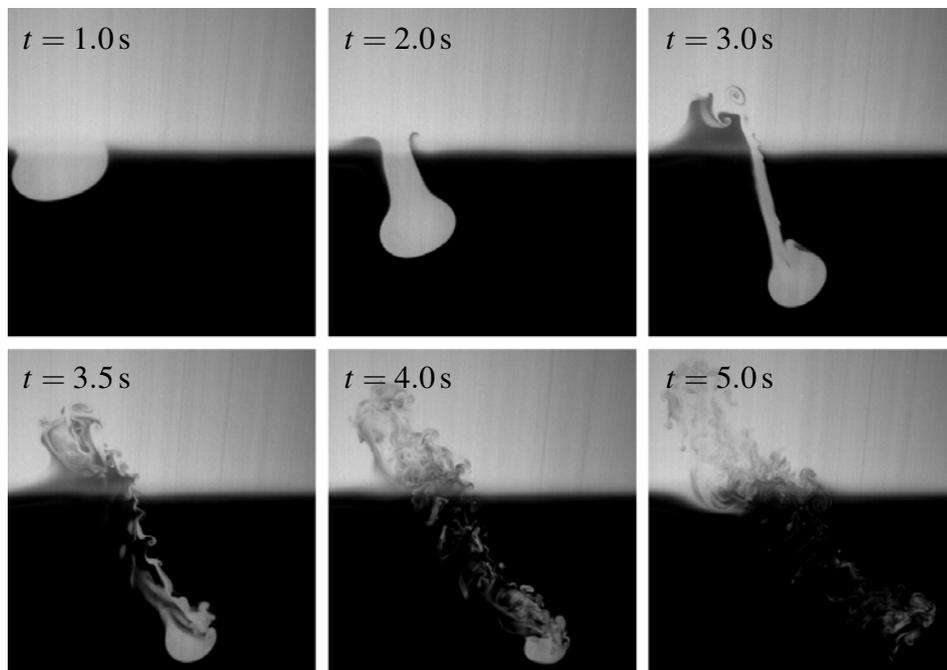


Fig. 5.20 LIF snapshots of the evolution of the ring–interface interaction in the central plane of the flow, for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (3270, 0.98, 20^\circ)$. Each snapshot is labelled with its associated instant in time. Note that the advective timescale $a/U = 0.74$ s.

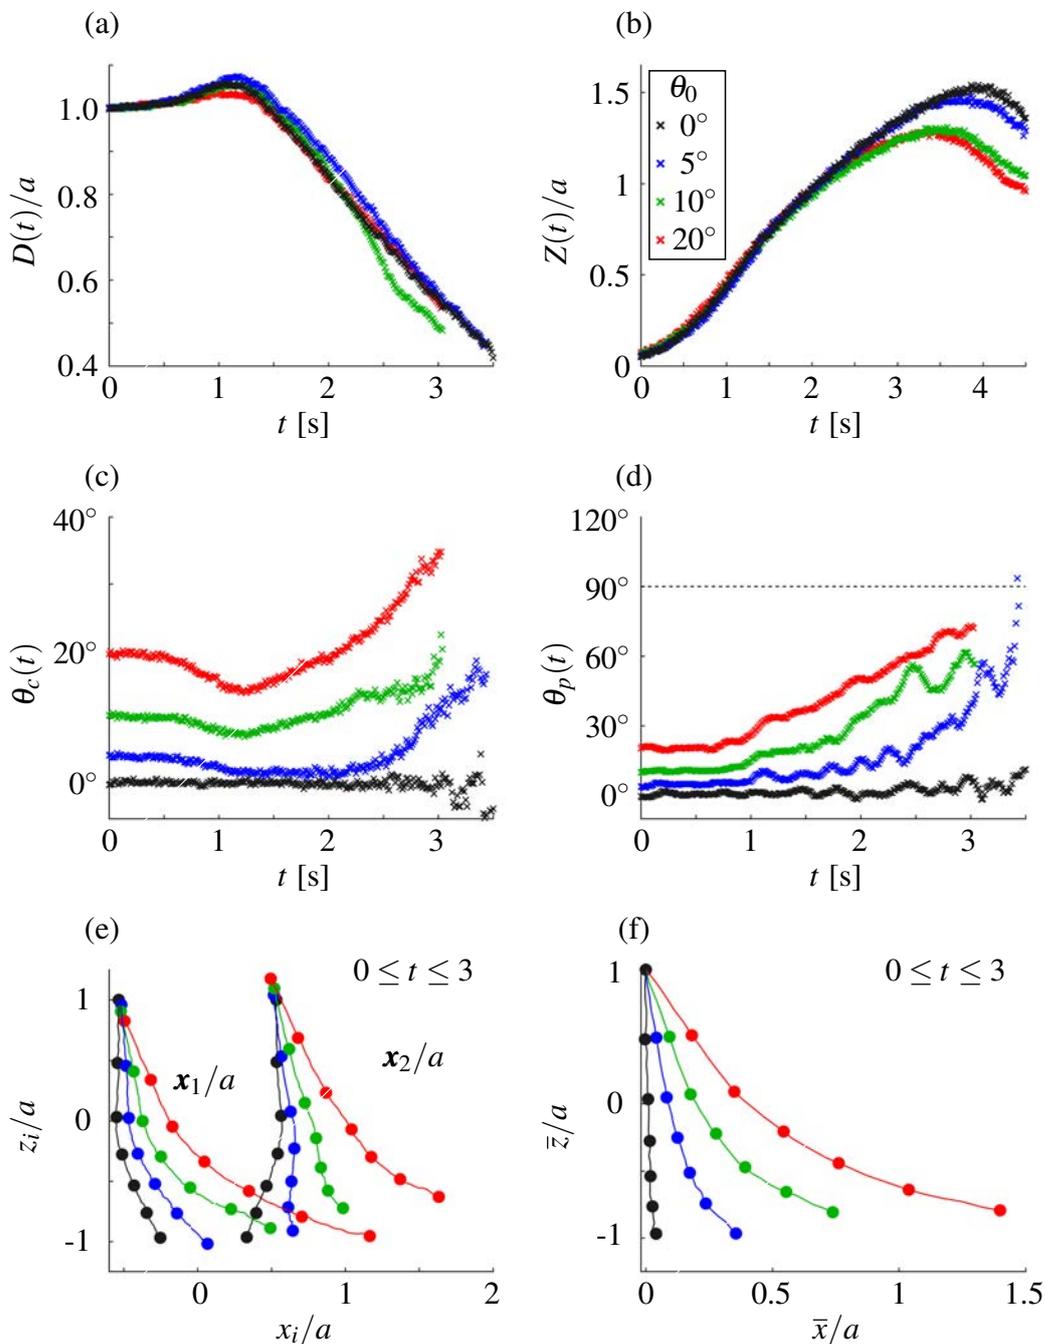


Fig. 5.21 Plots of the observed measures of the vortex ring for the cases corresponding to $(Re_0, Ri_0) = (2680, 1.48)$, in addition to the non-dimensionalised penetration depth $Z(t)/a$ into the lower layer. For plots (a), (c) and (d), data are plotted up to the stage where either C_1 or C_2 can no longer be reliably identified in the central plane of the flow. In plots (e) and (f), filled circles mark the position of the ring centroids and ring centre every 0.5 seconds. The legend corresponding to all plots is given in (b).

non-penetrative cases presented in Section 5.2, this is accompanied with a reduction in the core angle $\theta_c(t)$ due to the interface interacting more strongly with ring-core C_1 relative to C_2 . These features are evident in Figures 5.21 (a) and (c).

Similarly to the $\theta_0 = 0^\circ$ case, the ring is gradually eroded by the surrounding baroclinic vorticity. The associated reduction in $D(t)$ for $t \gtrsim 1.2$ s is demonstrated in Figure 5.21 (a). The observed linear reduction in $D(t)$ can be predicted using the following argument. Assuming the ring-volume $V(t) \sim D^3$ and the peeling of the ring is reducing its volume at a rate proportional to the interfacial area A between the ring and the BVC (which can be assumed to scale as $A \sim D^2$), then one obtains the following relation,

$$\frac{dV}{dt} \sim -A(t) \quad \Rightarrow \quad \frac{dD}{dt} \sim -1.$$

The high level of agreement in the gradient of dD/dt suggests that over the range of θ_0 considered, the propagation angle does not effect the rate at which fluid is ejected by the ring into the EFC in a significant way.

As the ring propagates through the lower-layer region, the buoyancy forces acting on the ring push C_1 towards C_2 , leading to the observed increase in $\theta_c(t)$. This can also be seen in the trajectories of \mathbf{x}_1 and \mathbf{x}_2 plotted in Figure 5.21 (e). Shortly before the ring reaches its maximum penetration depth, the primary ring begins to destabilise as the straining induced by the BVC begins to overcome the rotational motion of the ring. This transition can be attributed to the reduction in the circulation of the ring as it is eroded by the surrounding baroclinic vorticity.

Note that the evolution of $Z(t)$ in Figure 5.21 (b) suggest that a change in regime is observed between $\theta_0 = 5^\circ$ and $\theta_0 = 10^\circ$. While we currently have no explanation why such a change in regime would occur, it is conceivable that such a transition may be the result of a significant change in the turbulent breakdown of the ring between the $\theta_0 = 5^\circ$ and $\theta_0 = 10^\circ$ cases. Further work would be needed to explore this.

After the vortex ring breaks down significantly, the remaining core-entrained upper-layer fluid begins to return to its neutral buoyancy level. Comparing the snapshots for $4 \lesssim t \lesssim 5$ seconds between Figures 5.12 and 5.17 (corresponding to $\theta_0 = 0^\circ$ and 20° respectively), it appears that the remaining core-entrained fluid is more effectively mixed into the lower layer for the case $\theta_0 = 20^\circ$, despite penetrating less deeply into the lower layer compared to the $\theta_0 = 0^\circ$ case. This is suggested by both the smaller structures of upper-layer fluid and the lower light intensity of these structures in the LIF images for the $\theta_0 = 20^\circ$ case (compared to the $\theta_0 = 0^\circ$ case), implying a greater amount of mixing has occurred. However, confirming this would require an understanding of the dynamics outside of the central plane

of the flow. Observations such as this raise the question of whether the mixing efficiency of the ring–interface interaction for given $(\text{Re}_0, \text{Ri}_0)$ might be maximised by prescribing a non-zero value of θ_0 . This was one of our initial motivating questions for our periodic-mixing experiments, introduced in Chapter 6.

Figures 5.18, 5.19 and 5.20 present LIF snapshots for the case $(\text{Re}_0, \text{Ri}_0) = (3270, 0.98)$, for $\theta_0 = 5^\circ$, 10° and 20° respectively. In a similar way as for the $\theta_0 = 0^\circ$ case, the primary ring and BVC rapidly destabilise. For the cases $\theta_0 = 10^\circ$ and $\theta_0 = 20^\circ$, the ring appears to first destabilise at the side of the ring corresponding to C_2 , which we attribute to the interface between this portion of the vortex ring and the lower-layer fluid directly above it being Rayleigh-Taylor unstable.

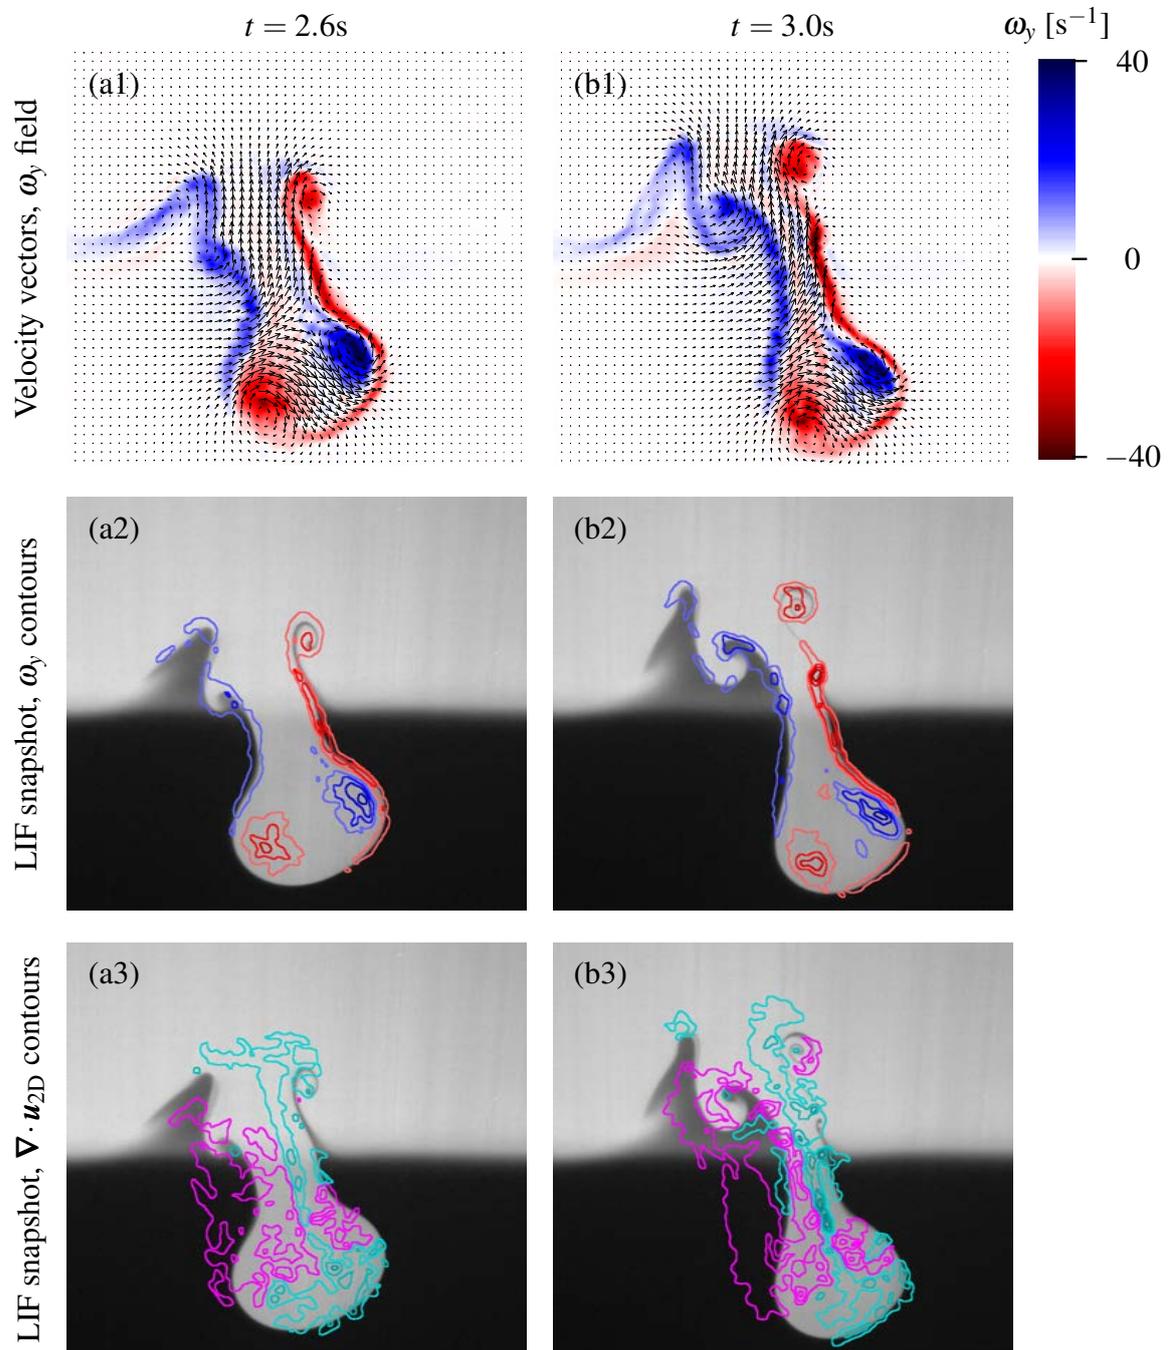
5.3.4 Oblique propagation: secondary structures

Owing to the complex three-dimensional flow that develops on the BVC as it begins to destabilise, here we restrict our attention to the coherent structures in the flow that can be easily discerned from our observations of the central plane of the flow. In particular, we discuss the evolution of the BVC and the structures that emerge as it destabilises, as well as discussing the development of significant azimuthal flows, both along the surface of the BVC and vortical structures that develop in the upper layer. Our discussion focuses primarily on the ring–interface interaction corresponding to $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 20^\circ)$, then relating the observed features to similar features that arise for other cases of $(\text{Re}_0, \text{Ri}_0, \theta_0)$.

To help visualise the three-dimensional evolution of the flow for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 20^\circ)$, Figure 5.22 present panels of the in-plane vorticity ω_y overlaid with velocity vectors, LIF snapshots overlaid with vorticity contours, and LIF snapshots overlaid with $\nabla \cdot \mathbf{u}_{2D}$ contours for the times corresponding to snapshots 5-8 of Figure 5.17. Vorticity contours are coloured in the same way as the vorticity fields, with blue and red respectively representing anti-clockwise and clockwise in the central plane of the flow. The magenta and cyan $\nabla \cdot \mathbf{u}_{2D}$ contours respectively represent flow converging on to and diverging away from the central plane.

First consider the panels in Figure 5.22 corresponding to $t = 2.6$ s. In panel (a1), the ω_y field clearly marks out C_1 and C_2 , in addition to the cross-section of the BVC coinciding with the central plane of the flow. This is confirmed by the overlap between the ω_y contours and the deformed density interface in panel (a2).

The bulk momentum of the ring and EFC has a non-zero horizontal component, due to the initial non-zero bulk horizontal momentum of the vortex ring. Lower-layer fluid converges onto the central plane in the wake region at the back of the ring and EFC, indicated by the large magenta $\nabla \cdot \mathbf{u}_{2D}$ contour in panel (a3). The velocity vectors in panel (a1) clearly show



(See next page for figure caption.)

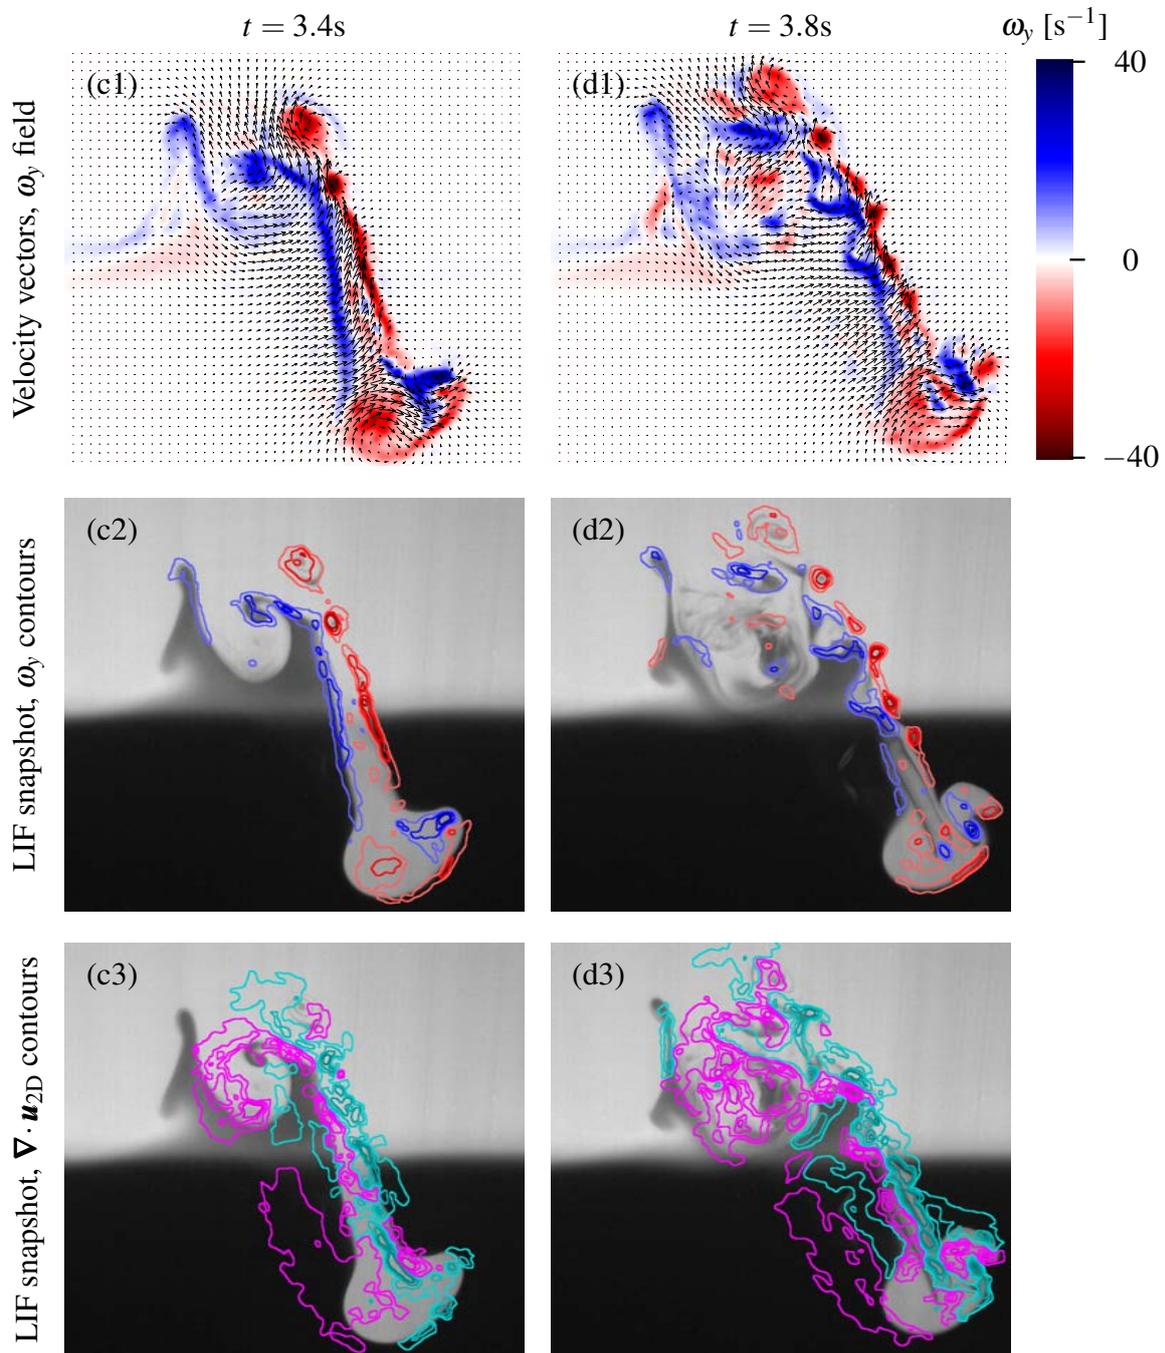


Fig. 5.22 Plots of the in-plane vorticity ω_y overlaid with velocity vectors, LIF snapshots overlaid with vorticity contours, and LIF snapshots overlaid with $\nabla \cdot \mathbf{u}_{2D}$ contours for the times $t = 2.6\text{ s}$, 3.0 s , 3.4 s and 3.8 s . The magenta and cyan $\nabla \cdot \mathbf{u}_{2D}$ contours respectively represent flow converging on to and diverging away from the central plane.

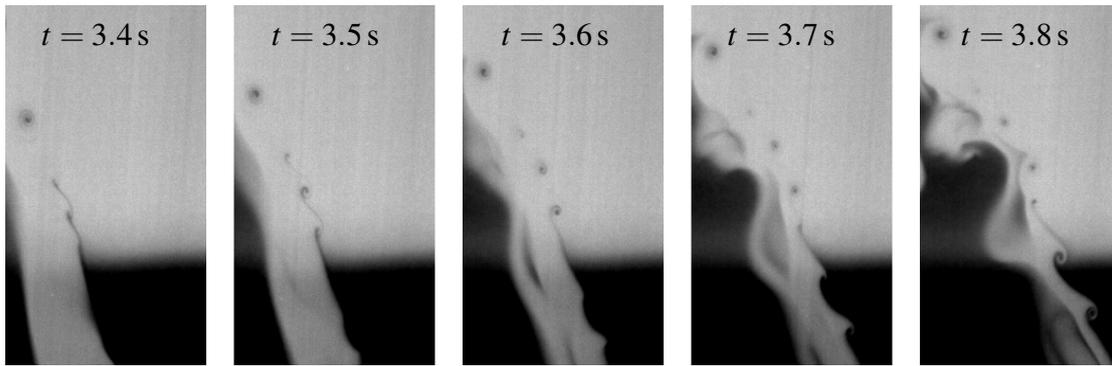


Fig. 5.23 Close-up LIF snapshots of the small secondary vortices that develop on the leading edge of the BVC as it begins to destabilise for the case $(Re_0, Ri_0, \theta_0) = (2680, 1.48, 20^\circ)$. The time interval over which the snapshots are taken correspond to the panels presented in Figure 5.22 (c-d).

the horizontal direction of the flow at the back of the EFC as fluid enters this region, as well as the near-stagnant flow at the front of the EFC. The (vertical) buoyancy forces acting on the front and back of the EFC result in the back of the EFC moving towards the front of the EFC (in the view of the camera, this appears as the left side of the dyed-undyed interface moving towards the right side). To satisfy three-dimensional incompressibility in the EFC, some of the EFC fluid must be displaced out of the central plane of the flow. This divergence from the central plane is indicated by the elongated cyan contour in panel (a3) that stretches across the entirety of the EFC.

As the flow evolves from $t = 2.6$ s to 3.4 s, the cross-section of the EFC in the central plane of the flow becomes progressively thinner as the back of the BVC approaches the front. While some of the EFC fluid is ejected into the upper layer, the strengthening of the cyan $\nabla \cdot \mathbf{u}_{2D}$ contours in the EFC shown in panels (b3) and (c3) indicates that EFC fluid continues to be displaced away from the central plane of the flow. As the deformed interface in the central plane remains largely laminar over this period of time, it is unclear whether the deformation of the BVC is leading to a significant increase in the amount of mixing occurring in the lower layer relative to the 0° case.

At $t = 3.8$ s, the destabilisation of the BVC becomes visible in the central plane of the flow, with the front of the BVC organising itself into a series of small coherent vortices. These vortices can be seen in panels (d1) and (d2) of Figure 5.22 by the concentrated regions of negative ω_y . Close-up LIF snapshots of the vortices are presented in Figure 5.23, made apparent by the observed roll-up of lower-layer fluid. Such vortices do not appear to develop for the $\theta_0 = 5^\circ$ and 10° cases, which is suggestive of a minimum- θ_0 criterion for the development of these vortices.

Around the same time, a strong convergent flow onto the central plane develops immediately above the density interface, indicated by the magenta $\nabla \cdot \mathbf{u}_{2D}$ contours that emerge at $t = 3.8$ s in Figure 5.22 (d3). By observing the subsequent evolution of the flow in the snapshots presented in Figure 5.17, it can be seen that there is a strong convergence of mostly upper-layer fluid into the central plane of the flow. This leads to both the pushing of lower-layer fluid higher into the upper layer and splashing onto the density interface. The associated splashing can be seen to occur for all cases with $\theta_0 > 0^\circ$ considered in this section, with the splashing visible in the presented LIF snapshots at $t = 4.5$ s and $t = 5$ s for $(\text{Re}_0, \text{Ri}_0) = (2680, 1.48)$ and $(3270, 0.98)$ respectively.

The particular physical mechanism leading to the strong in-flow is unclear. However, one plausible explanation for its origin is that the converging upper-layer fluid has been transported from the EFC by a vorticity dipole, that extends axially along the EFC. The regions of in-plane divergence and convergence support this explanation, as they appear consistent with the development of such a dipole, with two vorticity axes that are located either side and equidistant from the central plane of the flow. One can then imagine that, as upper-layer fluid is ejected from the vortex ring into the EFC, the ejected fluid is displaced away from the central plane by the axial dipole, and circulates around the dipole as the upper-layer fluid continues to rise towards the upper layer. If the strength of the dipole is sufficient, then this transported upper-layer fluid would re-emerge in the central plane of the flow towards the back of the EFC, possibly in a way similar to the convergent flow observed for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 20^\circ)$. Further work would be needed to confirm the existence of this transport mechanism.

Figure 5.24 presents LIF snapshots of the ring–interface interaction for the case corresponding to $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2960, 1.18, 15^\circ)$, with consecutive snapshots separated in time by 0.125 seconds. Note that this ring–interface interaction is an intermediate case in $(\text{Re}_0, \text{Ri}_0, \theta_0)$ -space of the ring–interface interactions considered in this section. While vortex roll-up structures emerge on the leading edge of the BVC in a similar way to the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2680, 1.48, 20^\circ)$, a striking difference is the emergence of disconnected regions of upper-layer fluid emerging in the lower layer. These can be seen from the 4th snapshot onwards in Figure 5.24, and are also seen to occur for the case $(\text{Re}_0, \text{Ri}_0, \theta_0) = (3270, 0.98, 20^\circ)$ from $t = 3.5$ s onwards in Figure 5.20.

The disconnected regions of upper-layer fluid that emerge in the lower layer can only be attributed to the structural evolution of the BVC as it destabilises, and is consistent with the development of an axial vorticity dipole (described above). The complex cross-sectional geometry of these disconnected regions indicates a significant increase in the interfacial area between the ring-entrained upper-layer fluid and surrounding lower-layer fluid as a result

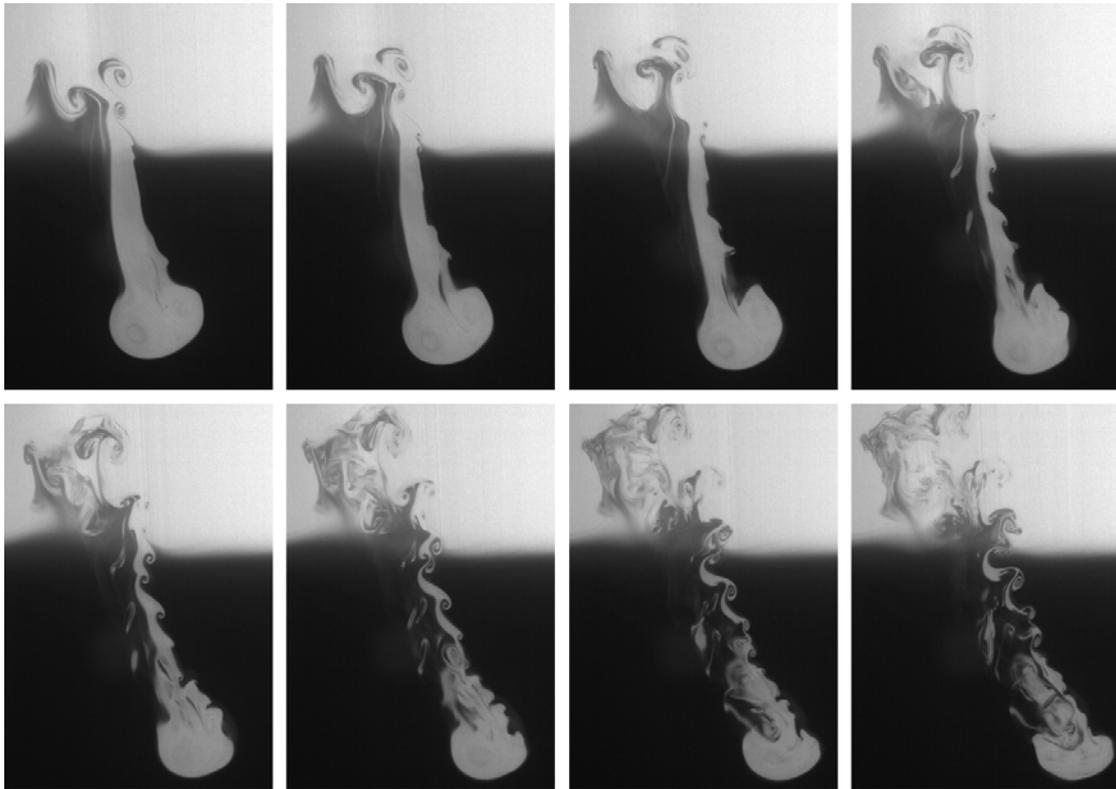


Fig. 5.24 Selected LIF snapshots from the case $(Re_0, Ri_0, \theta_0) = (2960, 1.18, 15^\circ)$, with consecutive snapshots separated in time by 0.125 seconds. In this case, there is a clear development of secondary vortices along the leading edge of the destabilising BVC, as well as disconnected regions of upper-layer fluid indicating significant out-of-plane flow.

of the BVC destabilisation. If there is significant variation in the increase of this interfacial area between different prescriptions of θ_0 , this would indicate that the amount of mixing that occurs during the ring–interface interaction and the location of the mixing may have a significant dependence on θ_0 over the range $0^\circ \leq \theta_0 \leq 20^\circ$.

5.4 Summary

The experiments presented in this chapter represent the first systematic study on the dynamics of vortex rings obliquely impacting a density interface. Two-dimensional PIV and LIF measurements were made in the central plane of the ring–interface interaction for a range of cases in the (Re_0, Ri_0, θ_0) parameter space, including cases both when the ring is able and when the ring is unable to penetrate the density interface. We have used our two-dimensional observations to infer features of the three-dimensional evolution of the flow,

identifying previously-undocumented dynamics of the ring–interface interaction for initial propagation angles in the range $0^\circ \leq \theta_0 \leq 20^\circ$.

For non-penetrative vortex rings at propagation angles $5^\circ \leq \theta_0 \leq 10^\circ$, although the ring partially reorients itself to impact more vertically on the density interface, the oblique impact causes an azimuthally asymmetric production of baroclinic vorticity on the deforming density interface. This results in the development of azimuthal pressure gradients both in the primary vortex ring and the largest secondary vortex ring, made apparent by the convergence of fluid onto the central plane in the regions coinciding with the cores C_2 and S_2 . For the case $\theta_0 = 20^\circ$, the secondary vortex core S_2 is unable to pass around the outside of the primary ring. This appears to reduce the vertical buoyancy flux associated with the splashing of the interface as the density interface recoils from the vortex-ring impact.

For penetrative vortex rings, our observations of the $\theta_0 = 0^\circ$ cases corresponding to $Ri_0 = 1.48$ and $Ri_0 = 0.98$ indicate that while the Crow-like instability identified by Olsthoorn and Dalziel (2017) increases the interfacial area between the upper and lower-layer fluid, it also acts to inhibit significant mixing from occurring in the lower layer. The most striking differences in the evolution of the ring–interface interaction between vertical-impact and oblique-impact cases were seen for $\theta_0 = 20^\circ$. In this case, a strong divergence from the central plane is seen in the EFC as the fluid column destabilises, with a corresponding convergence at the back of the fluid column. We propose that this is due to the formation of an axial vorticity dipole that extends along the EFC and shortly above it. Additionally, vortices are seen to emerge at the front of the fluid column, resulting from the development of a shear and convection-driven fluid instability. Both features act to increase the interfacial area between the ejected-fluid column and surrounding lower-layer fluid, enhancing the amount of irreversible mixing that can occur in the lower layer.

The observed differences in the dynamics of the ring–interface interaction indicate that, over the range of initial propagation angles $0^\circ \leq \theta_0 \leq 20^\circ$, the propagation angle may play a significant role in the mixing that occurs, both in the location of mixing events and change in vertical density profile that results from the ring–interface interaction. We investigate the mixing induced by vortex rings interacting with density interfaces in Chapter 6. Although the experiments presented in this section correspond only to a small region of the multidimensional parameter space governing the evolution of the ring–interface interaction, a wide range of dynamical features were observed, indicating a sensitive dependency of the control parameters on the evolution of the flow and the subsequent mixing that occurs. This suggests that careful interpretation is required when attempting to relate findings employing the ring-eddy analogy to properties of turbulent two-layer zero-mean-shear flows.

Chapter 6

Vortex-ring-induced mixing of an initially two-layered density stratification

In the previous chapter, we investigated the differences in the dynamical evolution of the ring–interface interaction between vertically and obliquely-propagating vortex rings over the range of initial propagation angles $0 \leq \theta_0 \leq 20^\circ$. Our observations revealed that breaking the azimuthal symmetry associated with the $\theta_0 = 0^\circ$ case introduces significant changes in the dynamical features of the flow, both for non-penetrative and penetrative types of ring–interface interaction. This chapter investigates the implications such dynamical changes have on the mixing properties associated with the interaction.

For the experiments presented in this chapter, we mix an initially two-layered density stratification by generating a periodic sequence of vortex rings that are made to propagate towards the density interface. Our experimental methodology is similar to that employed by Olsthoorn and Dalziel (2015), who demonstrated for the range of parameters $1700 \leq Re_0 \leq 2700$, $4 \leq Ri_0 \leq 12$, $\theta_0 = 0^\circ$ that the periodically-forced stratification tends to a self-similar form, corresponding to a Ri_0 -independent mixing-efficiency regime with mixing efficiency $\eta_c \sim 0.42$. Olsthoorn and Dalziel (2015) associate η_c to the mixing efficiency of the ring–interface interaction, by envisaging a banding process in which vortex rings mix with lower-layer fluid to form a new band of perfectly mixed fluid at the bottom of the upper layer, which develops and maintains a weak, constant density gradient. The envisaged process is described as analogous to the mixing induced by a buoyant plume in a filling box, for which an analytic expression for the density profile is obtained by Worster and Huppert (1983).

In keeping with the experiments presented in Chapter 5, we classify our experiments into two different mixing regimes, ‘non-penetrative’ and ‘penetrative’. Figures 6.1 (a) and (b) plot density profiles illustrating typical features of the evolving stratification in each

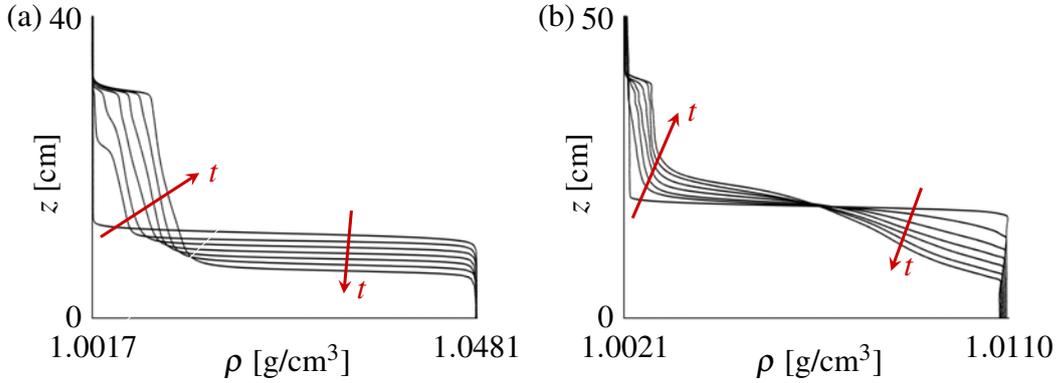


Fig. 6.1 Representative examples of the density profiles taken after every 100 ring generations for the cases in the (a) non-penetrative, and (b) penetrative, regimes respectively. Arrows are included to indicate the evolution of time. The selected experiments are labelled A0 and Q0 in Table 6.2, corresponding to $(\text{Re}_0, \text{Ri}_0, \theta_0) = (2960, 6.30, 0^\circ)$ and $(2960, 1.18, 0^\circ)$ respectively.

of the mixing regimes. In the non-penetrative regime (Figure 6.1 (a)), the initial density interface is maintained as fluid is entrained from the lower layer into the weakly-stratified upper layer. In this way, the experiments of Olsthoorn and Dalziel (2015) would be classified as non-penetrative. In the penetrative regime (Figure 6.1 (b)), the initial density interface is rapidly eroded and a partially-mixed middle layer is formed. This middle-layer stratification weakens over time by the continual ring-induced mixing.

The initial motivation for carrying out the experiments presented in this chapter was to investigate the functional dependencies of η on both a lower range of Ri_0 relative to Olsthoorn and Dalziel (2015), and for a range of non-zero propagation angles. While a Ri_0 -independent regime has been linked to η for $\text{Ri}_0 \geq 4$, we can expect that $\eta \rightarrow 0$ as $\text{Ri} \rightarrow 0$, as there would be almost no stratification to mix. Whether $\partial\eta/\partial\text{Ri}_0 \geq 0$ for $0 \leq \text{Ri}_0 \leq 4$ or $\eta(\text{Ri}_0)$ attains a maximum over this interval is one question these experiments aimed to address. We also aimed to investigate the behaviour of $\partial\eta/\partial\theta_0$ for a range of Ri_0 , to determine whether the dynamical differences between vertical and oblique cases of the ring–interface interaction have a significant impact on the mixing efficiency.

For non-penetrative mixing, insights made after the completion of the experiments presented in this chapter have led to a revised understanding of how the ring-induced mixing in the periodically-forced system should be interpreted. Crucially, we demonstrate that the ring-induced mixing should not be regarded as simply produced by a periodic sequence of ring–interface interaction events. We propose that the high mixing efficiency observed should instead be attributed to the ability of the ring to transport a large volume of fluid

from the top of the upper layer to the density interface, ring-induced upward entrainment across the interface (leading to some mixing), and the subsequent development of convective instabilities as the ring-mixed fluid finds its new neutral buoyancy level. Consequently, this reopens the question of what is the mixing efficiency of a ring–interface interaction event.

This chapter is laid out as follows. In Section 6.1 we outline our methodology for conducting the experiments presented in this Chapter. Then, in Section 6.2 we explain our method for quantifying mixing efficiency in the periodically-forced system, and discuss the measured mixing efficiencies from our experiments. In Section 6.3, we discuss in detail the evolution of the stratification for the non-penetrative regime. Our findings are summarised in Section 6.4.

6.1 Experimental methodology

6.1.1 Setup and procedure

The setup and procedure of our periodic-mixing experiments is similar to the experiments carried out by Olsthoorn and Dalziel (2015). A diagram of our experimental setup is given in Figure 6.2. These experiments were conducted in both the larger and smaller of our experimental tanks, described in Chapter 3. The tank-filling methodology outlined in Section 3.1 was used to prescribe the upper and lower-layer heights H_1 , H_2 and fluid densities ρ_1 , ρ_2 respectively. To generate reproducible vortex rings, we follow the methodology outlined in Section 3.2. For each of our experiments, we fixed the initial propagation angle θ_0 and the ring-tube height H_B , which we define as the vertical distance from the base of the tank to the upper edge of the ring tube outlet. Vertical profiles of the stratification were taken using the conductivity probe and thermistor introduced in Section 3.4. Measurements were made over a prescribed distance L , between the probe/thermistor resting position and the base of the tank. The resting position was chosen to be approximately a couple of ring diameters above the outlet of the ring tube to ensure measurements were made over the full extent of the mixed region of the stratification throughout the experiment. The upper-layer height H_1 was prescribed so the free surface is ~ 5 cm above the probe/thermistor resting position.

For each experiment, a periodic sequence of 600 vortex rings was generated to mix the stratification. At the start of each experiment and after every 10 vortex rings, the conductivity probe and thermistor were traversed through the stratification, sampling the fluid to obtain vertical profiles of electrical conductivity and temperature (and hence, density). Measurements were made only when the probe and thermistor were traversed downwards, before being withdrawn to their starting position. For each sequence of 10 ring generations,

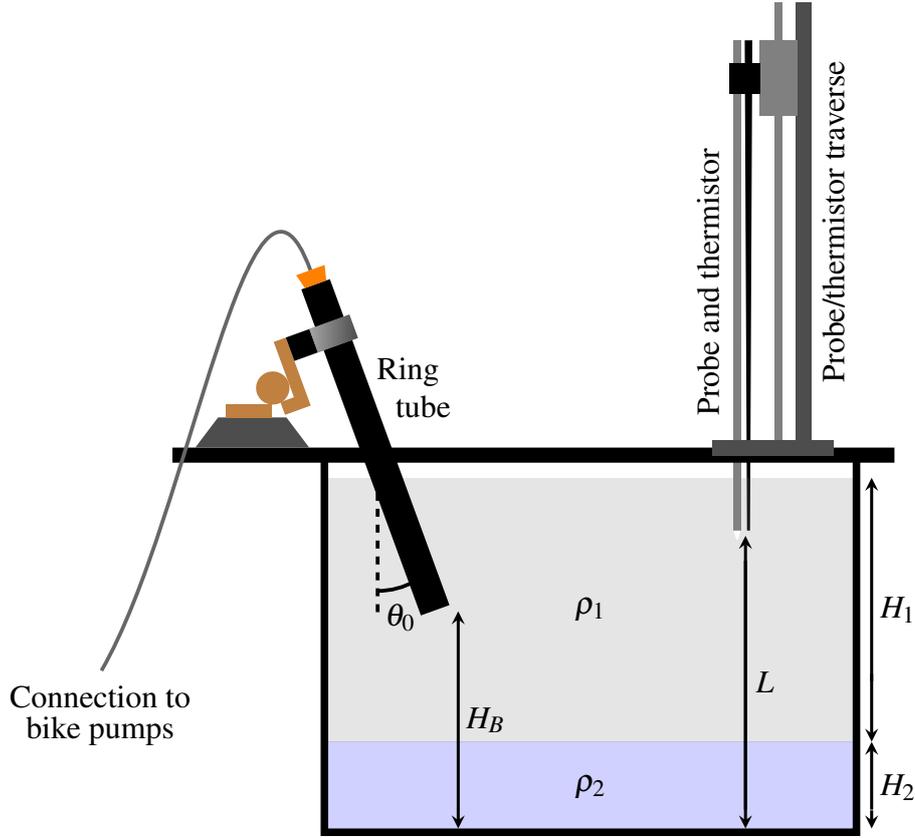


Fig. 6.2 Diagram of the experimental setup of a typical periodic-mixing experiment.

a fixed inter-ring period τ_R between the generation of each vortex ring was prescribed. A longer period τ_S between the last ring generation of one ring sequence and the first generation of the next ring sequence was prescribed to allow the probe and thermistor to move through the stratification sufficiently slowly so as to induce negligible mixing. Note that the effect the speed of the probe had on the measurements that were made is analysed in Appendix C. By prescribing the periods τ_R , τ_S , the time t_n of the n th vortex ring generation is then given by

$$t_n = n\tau_R + \left\lfloor \frac{n}{10} \right\rfloor (\tau_S - \tau_R). \quad (6.1)$$

For reasoning similar to that outlined in Olsthoorn and Dalziel (2015), the choice of τ_R was made to satisfy

$$\tau_{\text{dissip}} < \tau_R \ll \tau_{\text{diff}}, \quad (6.2)$$

where τ_{dissip} is the timescale for the macroscopic motion within the sampled fluid domain to dissipate, and τ_{diff} is the timescale for molecular diffusion of the density profile. From our PIV/LIF experiments, we found that τ_{dissip} increased for decreasing Ri_0 . For the lowest Ri_0

considered, the kinetic energy of the flow had significantly dissipated by ~ 45 seconds after the initial impact of the vortex ring with the density interface.

For a strongly-stratified region of vertical thickness l , the diffusion timescale is of the form $\tau_{\text{diff}} = l^2/\kappa$, where $\kappa \sim 1 \times 10^{-9} \text{ m}^2\text{s}^{-1}$ is the molecular diffusivity of salt. In our mixing experiments, stratified layers of approximately constant buoyancy gradients formed with thicknesses in the range $O(1) \text{ cm} < l < O(10) \text{ cm}$. Taking $l = 1 \text{ cm}$, we obtain $\tau_{\text{diff}} = 10^5 \text{ s}$. Alternatively, if we take $l = 1 \text{ mm}$, which is an order of magnitude smaller in thickness compared to the initial density interface, we obtain $\tau_{\text{diff}} = 10^3 \text{ s}$. Based on the estimates made for τ_{dissip} and τ_{diff} , we took $\tau_R = 75 \text{ s}$ for all our experiments.

The traverse time τ_S was prescribed by fixing the downward and upward probe/thermistor traverse speeds $U_S^\downarrow, U_S^\uparrow$. For all experiments, we fixed the traverse speeds at $U_S^\downarrow = 10 \text{ mm/s}$ and $U_S^\uparrow = 6 \text{ mm/s}$, unless otherwise stated. We chose to fix U_S^\uparrow at a lower speed than U_S^\downarrow . This is because when the probe and thermistor are being withdrawn from the base of the tank, they leave a wake behind them, enhancing their ability to mix the stratification. The prescribed speeds for $U_S^\downarrow, U_S^\uparrow$ are small enough that the probe and thermistor induce negligible mixing to the stratification relative to the ring-induced mixing.

Two additional steps were incorporated into our experimental procedure to improve the performance of the conductivity probe. Firstly, before each experiment, using the probe and thermistor we sampled the stratification five times over the period of one hour. This helped to improve the consistency between successive conductivity profiles obtained in the early stage of the experiment. We attributed this improvement in consistency to the probe electrodes converging towards an electrochemical equilibrium with being traversed through the density stratification, though we did not verify this hypothesis beyond demonstrating increased consistency and improved salinity conservation in the probe measurements.

Secondly, immediately before each probe/thermistor traverse, we opened the pinch valve for 90 seconds to flush the probe-interior with upper-layer fluid, removing the fluid siphoned by the probe when sampling the fluid during the previous probe/thermistor traverse. The fluid used to flush the probe was always taken from above the ring-tube outlet in a dynamically-isolated region of the stratification, thus had no impact on the ring-induced evolution of the stratification. Without flushing the probe, we found that the conductivity measurements behaved erratically in the uppermost 5 – 10 cm of the sampled stratification, consistently over-estimating the density of the upper layer fluid. We attributed this behaviour to lower-layer fluid held in the probe interior from the previous stratification sampling causing an increase in the electrical conductivity of the circuit established between the two probe electrodes and the bridge circuit. By flushing out lower-layer fluid held within the probe interior with

Ring case	1	2	3	4
T_s [s]	6.25	5.45	5.00	4.56
U [mm/s]	48.55 ± 0.22	54.55 ± 0.1	60.25 ± 0.32	66.54 ± 0.42
$Re_0 = aU/\nu$	2390	2680	2960	3270
δKE [mJ]	0.413	0.549	0.699	0.935

Table 6.1 Table presenting the ring propagation speeds U , Reynolds numbers Re_0 and kinetic energy δKE for the vortex rings generated in our periodic-mixing experiments. The values and precision errors for U are determined from our bike pump calibration, discussed in Section 3.2. We label each ring case from 1 to 4 for ease of reference.

upper-layer fluid, the calibration-calculated densities in the unmixed upper layer much more closely reflected the prescribed value of the upper-layer density.

It was also found that the conductivity profiles had a small but non-negligible dependency on the traverse speed and probe-siphoning volume flux. This prompted us to carry out a more thorough investigation of the conductivity probe, which is given in Appendix C. For our purposes here, we were able to demonstrate that these dependencies had a negligible effect on the conclusions we draw from our results.

6.1.2 Overview of input parameters

Vortex rings were generated using four different bike pump traverse times T_s , corresponding to four different values for propagation velocity U and kinetic energy δKE , given in Table 6.1. By varying T_s , we vary the bulk Reynolds number $Re_0 = aU/\nu$ between the four different vortex ring cases, with $2390 \leq Re_0 \leq 3270$.

We categorise the experimental results presented in this chapter into two groups: non-penetrative mixing, corresponding to $3.67 \leq Ri_0 \leq 6.30$, and penetrative mixing, corresponding to $0.98 \leq Ri_0 \leq 1.48$. Non-penetrative-mixing experiments were conducted in the smaller experimental tank of plan area $A = 0.5 \times 0.2 \text{ m}^2$, with probe-traverse height $L = 40 \text{ cm}$, tube height $H_B = 30 \text{ cm}$ and lower-layer height $H_2 = 12 \text{ cm}$. Penetrative-mixing experiments were conducted in the larger experimental tank of plan area $A = 1.2 \times 0.4 \text{ m}^2$, with probe-traverse height $L = 50 \text{ cm}$, tube height $H_B = 38.7 \text{ cm}$ and lower-layer height $H_2 = 18.3 \text{ cm}$. The input parameters for the experiments presented in this Chapter are given in Table 6.2, along with labels which are used in conjunction with θ_0 to uniquely identify each experiment (e.g. A10; C0; Q25).

Label	Ring case	ρ_1 (g/cm ³)	$\Delta\rho \times 10^2$ (g/cm ³)	Ri ₀	θ_0
A	3	1.0017	4.635 ± 0.025	6.30	0°, 10°, 20°
B	1	1.0021	2.738 ± 0.031	5.73	0°, 10°, 20°
C	2	1.0021	2.736 ± 0.017	4.56	0°, 10°, 20°
D	3	1.0021	2.735 ± 0.028	3.67	0°, 10°, 20°
P	2	1.0021	0.879 ± 0.025	1.48	0°, 5°, 15°, 25°
Q	3	1.0021	0.879 ± 0.025	1.18	0°, 5°, 15°, 25°
R	4	1.0021	0.879 ± 0.025	0.98	0°, 5°, 15°, 25°

Table 6.2 Table of parameters for the experiments presented in this chapter. A dividing line separates the experiments associated with non-penetrative and penetrative regimes (characterised by Ri₀). The ring cases correspond to vortex-ring parameters given in Table 6.1. The labels associated with each Ri₀ are used for reference throughout this chapter.

6.1.3 Data correction and validation

After carrying out a probe/thermistor calibration using the methodology outlined in Section 3.4.2, measurements from the conductivity probe were seen to drift by $O(10^{-2})$ V day⁻¹, leading to measurements for ρ_{20} drifting by $\Delta\rho_{20} \sim O(10^{-4})$ g cm⁻³ day⁻¹, with the magnitude of the error in ρ_{20} depending on the choice of resistor in the bridge circuit. Denoting the measured values and calibration-calculated values of the fluid density when measured at 20°C as $\rho_{20}^{(M)}$ and $\rho_{20}^{(C)}$ respectively, for a given experiment the drift $\delta\rho_{20} = \rho_{20}^{(M)} - \rho_{20}^{(C)}$ was seen to correspond anywhere from $\sim 0\%$ to 3% of the difference in ρ_{20} between the upper and lower layers of the initial stratification. As $\delta\rho_{20}$ also varied in ρ_{20} at a given time, the probe drift could in turn lead to measurement errors of up to 5-6% for the increase in potential energy of the system between two density profiles being taken.

To reduce the measurement error of changes in the potential energy of the system, the probe measurements were adjusted for each vertical profile to fit the corresponding probe/thermistor calibration. To do this, for each experiment we ensured that the traverse length was long enough to sample regions of the stratification in the upper and lower layer that were completely unmixed by the ring-induced mixing over the entirety of the experiment. In these unmixed regions, we have $\rho_{20} = \rho_{20,1}, \rho_{20,2}$, corresponding respectively to the prescribed values for the upper and lower layers of the initial two-layer stratification, and temperature measurements are made with the thermistor. Using these known values, a bi-iterative solver was used to invert the polynomial fitting surface for ρ_{20} defined in Equation (3.10) to calculate the calibration-expected probe voltages $C_{V,1,2}^{(C)}$ for the unmixed regions in the upper and lower layer.

For each vertical profile, probe measurement deviations $\Delta C_{V,1,2}$ for the upper and lower layers are given by

$$\Delta C_{V,1,2} = C_{V,1,2}^{(M)} - C_{V,1,2}^{(C)}, \quad (6.3)$$

where the probe measurements $C_{V,1,2}^{(M)}$ are calculated taking the mean probe measurement in the top and bottom centimetre of the measured region of the stratification. In these 1cm regions, we also calculate the mean of $\rho_{20}^{(M)}$, which we denote as $\rho_{20,1}^{(M)}$ and $\rho_{20,2}^{(M)}$ for the upper and lower layers respectively. Using the values calculated for $\Delta C_{V,1,2}$, $\rho_{20,1,2}^{(M)}$, we construct an adjusted probe voltage profile $C_V^{(A)}(z)$ in the following way,

$$C_V^{(A)}(z) = C_V^{(M)}(z) + \Delta C_{V,1} + \frac{\rho_{20}^{(M)}(z) - \rho_{20,1}^{(M)}}{\rho_{20,2}^{(M)} - \rho_{20,1}^{(M)}} (\Delta C_{V,2} - \Delta C_{V,1}), \quad (6.4)$$

where $C_V^{(M)}(z)$ is the measurement profile of the probe. This construction ensured that the adjusted measurements for ρ_{20} would correspond to the prescribed values for $\rho_{20,1}$, $\rho_{20,2}$ in the unmixed regions of the upper and lower layers. Using $C_V^{(A)}(z)$, we then used the probe/thermistor and thermistor calibrations to calculate vertical profiles for $\rho_{20}(z)$ and $\rho(z)$.

For the rest of this chapter, we drop the superscript (A) and all derived quantities will be calculated from voltage-adjusted measurements using Equation (6.4), unless otherwise stated. To validate our measurements, we make a salinity-conservation check. The conserved quantity in our periodic-mixing experiments is the mass of dissolved salt crystals in the stratification. As this is not directly measurable, we make a salinity-conservation check by using the following integral as a proxy,

$$S^{(k)} = \int_0^L (\rho_{20}^{(k)}(z) - \rho_{20,1}) dz, \quad (6.5)$$

for the k th probe/thermistor traverse. The integral $S^{(k)}$ can be interpreted as the salinity excess of the stratification relative to a homogeneous system of the same salinity as the upper-layer fluid. The integral was calculated for each probe/thermistor vertical profile, after adjusting the probe voltage using Equation (6.4). For each experiment, $S^{(k)}$ varied by less than $\pm 1\%$ from its mean value, which was sufficiently small to validate our probe-voltage adjustment methodology.

6.2 Energetics of sustained periodic forcing

Writing the increase in potential energy of the stratification due to the mixing induced by the n th vortex ring generated as $\delta\text{PE}_R^{(n)}$, we define the mixing efficiency $\eta^{(n)}$ of the n th vortex ring as

$$\eta^{(n)} = \frac{\delta\text{PE}_R^{(n)}}{\delta\text{KE}^{(n)}}, \quad (6.6)$$

where $\delta\text{KE}^{(n)}$ is the kinetic energy of the n th vortex ring. Our methods for quantifying $\delta\text{KE}^{(n)}$ and $\delta\text{PE}_R^{(n)}$ are discussed respectively in Sections 6.2.1 and 6.2.2. Mixing-efficiency measurements of the stratification as it evolves are then discussed in Section 6.2.3.

6.2.1 Quantifying kinetic energy

For simplicity, we assume in each experiment that all vortex rings have the same kinetic energy before interacting significantly with the stratification, giving

$$\delta\text{KE}^{(n)} = \delta\text{KE}. \quad (6.7)$$

The vortex-ring kinetic energy δKE determined for each value of T_s is given in Table 6.1, assuming the ring density $\rho_R = 1.002 \text{ g/cm}^3$.

By assuming axisymmetry of the vortex ring about its central axis, measurements of δKE for each value of T_s are calculated from a consecutive sequence of 100 PIV velocity fields, taken over 1.25 seconds. The i th measurement δKE_i is given by

$$\delta\text{KE}_i = 2\pi \int_z \int_{r=0}^{2a} \frac{1}{2} \rho_R (u_i^2 + v_i^2) H \left(2 - \frac{|\bar{\mathbf{x}}_i - \mathbf{x}|}{D_i} \right) r dr dz. \quad (6.8)$$

Here, r is a radial coordinate with $r = 0$ corresponding to the central axis of the ring, (u_i, v_i) are the velocity measurements from the i th velocity field made at each PIV interrogation point, H is a Heaviside step filter, and $\bar{\mathbf{x}}_i$ and D_i are the ring centre and plane diameter, defined in Section 5.1. The Heaviside filter is applied to ensure velocity measurements made in regions of the flow more than two plane diameters away from the centre of the ring do not contribute to the estimated value of δKE_i .

Figure 6.3 plots the measurements made for δKE_i for the ring case corresponding to $T_s = 6.25 \text{ s}$ against the non-dimensionalised distance $(\bar{z}_i - z_0)/a_0$ between the ring-tube outlet and the vortex ring, where the ring diameter $a_0 = 49.2 \text{ mm}$ is taken to be fixed. The observed linear reduction in δKE_i is a feature consistent with the measurements corresponding to the other cases, and can be attributed to both viscous dissipation and a loss of energy as

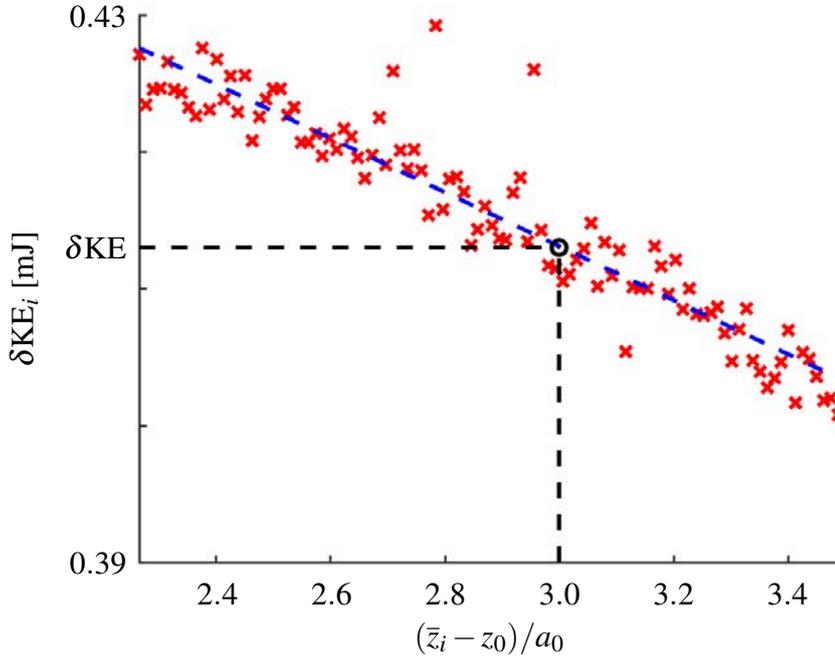


Fig. 6.3 Representative plot of the vortex-ring kinetic energy measurements δKE_i made over a sequence of 100 consecutive PIV velocity fields, plotted against the non-dimensionalised distance $(\bar{z}_i - z_0)/a_0$ between the vortex ring and the ring-tube outlet. A least squares approach is used to fit a line to the data, indicated by the blue dashed line. For each ring case, δKE is taken to be the value of the fitted line evaluated at $(\bar{z}_i - z_0)/a_0 = 3$.

fluid is rejected into the wake of the ring. For the representative case, an energy loss of approximately 0.02 mJ/s is observed, corresponding to about a 5% loss per second of the total ring energy.

In wanting to determine the mixing efficiency of the ring–interface interaction, it is desirable for our prescription of δKE to closely reflect the amount of kinetic energy delivered by the vortex ring to the initial density interface. We therefore choose to prescribe δKE as the kinetic energy of the ring when it is a distance a_0 above the density interface. Across all experiments, to the nearest ring diameter, the distance between the ring-tube outlet and the initial density interface is $4a_0$. Our prescription of δKE is therefore taken to be the value of δKE_i evaluated at $(\bar{z}_i - z_0)/a_0 = 3$. This is determined by using a least-squares approach to fit a line to the data, indicated by the blue dashed line in Figure 6.3.

6.2.2 Quantifying ring-induced increases in potential energy

Writing the vertical density profile after the n th ring has mixed the stratification as $\rho^{(n)}(z)$, we write the increase in potential energy of the stratification between before and after the n th ring-mixing event as

$$\delta\text{PE}^{(n)} = A \left(\int_0^L g\rho^{(n)}{}_z dz - \int_0^L g\rho^{(n-1)}{}_z dz \right). \quad (6.9)$$

Here, A is the area of the base of the tank, L is the length traversed by the probe and thermistor and $\rho^{(0)}(z)$ is the density profile of the stratification before any vortex rings are generated.

We decompose $\delta\text{PE}^{(n)}$ into contributions due to ring-induced mixing, $\delta\text{PE}_R^{(n)}$, and contributions due to temperature changes in the stratification, $\delta\text{PE}_T^{(n)}$, giving

$$\delta\text{PE}^{(n)} = \delta\text{PE}_R^{(n)} + \delta\text{PE}_T^{(n)}. \quad (6.10)$$

To calculate $\eta^{(n)}$ using Equation (6.6), it is necessary to be able to quantify $\delta\text{PE}_T^{(n)}$, so we can infer $\delta\text{PE}_R^{(n)}$. To do this, we first lay out all heat contributions to the energy budget, allowing us to understand the evolution of the temperature stratification better in our experiments.

Heat contributions to our experiments can be compartmentalised into two categories: ring-induced heat fluxes, and external heat fluxes. The ring-induced contributions to heat are made through heat of dissipation and heat of dilution. These heat contributions were found to be negligible, leading to temperature changes at most of $O(10^{-3})\text{K}$, which is demonstrated in Appendix D. Temperature changes in the fluid system due to external sources are made by the heat flux through the Perspex walls of the tank, and evaporative cooling at the free surface. Assuming a linear temperature gradient across the Perspex walls between the fluid system and lab environment, the heat flux H through each of the walls is given by

$$\frac{H}{A} = -k_p \frac{T_a - T}{w_p}, \quad (6.11)$$

where A is the area of the wall in contact with the fluid system, $k_p = 0.18\text{W/m/K}$ is the thermal conductivity of Perspex, w_p is the thickness of the walls and T_a , T are the air temperature and fluid temperature respectively.

For all experiments conducted in each of the two experimental tanks, using Equation (6.11) allowed us to verify that changes in the mean temperature of the stratification were mostly attributable to the heat flux through the tank walls. The heat flux calculation consistently over-estimated the observed increase in mean temperature of the stratification, which we expected as evaporative cooling would be working to reduce the mean stratification

temperature. No attempt was made to quantify the heat flux due to evaporative cooling as doing so would require us to measure a number of additional variables, including the temperature, velocity and relative humidity of the air directly above the free surface.

We estimate $\delta\text{PE}_T^{(n)}$ by calculating the potential energy of the stratification immediately before the n th ring is generated, and determining what the change in potential energy would be if the temperature field T is forward-advected to the time immediately after the n th ring has mixed the stratification, with the salinity field (or equivalently, ρ_{20}) remaining unchanged. Writing $\rho = \rho(\rho_{20}, T)$, our quantification for $\delta\text{PE}_T^{(n)}$ can be written as

$$\delta\text{PE}_T^{(n)} = A \left(\int_0^L g\rho(\rho_{20}^{(n-1)}, T^{(n)}) z dz - \int_0^L g\rho(\rho_{20}^{(n-1)}, T^{(n-1)}) z dz \right). \quad (6.12)$$

Incorporated into $\delta\text{PE}_T^{(n)}$ are changes in PE due to diffusion of the temperature field, and the ring-induced advection of the temperature field. These contributions to PE are small relative to the contribution made by the heat flux through the tank walls and so are neglected.

From Equations (6.9), (6.10) and (6.12), the potential energy contribution made by the mixing induced by the n th vortex ring is given by

$$\delta\text{PE}_R^{(n)} = A \left(\int_0^L g\rho(\rho_{20}^{(n)}, T^{(n)}) z dz - \int_0^L g\rho(\rho_{20}^{(n-1)}, T^{(n)}) z dz \right). \quad (6.13)$$

6.2.3 Mixing efficiency measurements

As conductivity and temperature profiles are only taken after every 10 vortex-ring generations, ring-induced increases in potential energy are determined between consecutive profiles being taken, then divided by the inter-measurement number of vortex rings to estimate the mixing efficiency $\eta^{(n)}$ (introduced in Equation (6.6)). We denote the ring-induced increases in potential energy made from the first n generated vortex rings as

$$\Sigma\delta\text{PE}_R^{(n)} = \sum_{k=0}^n \delta\text{PE}_R^{(k)}, \quad (6.14)$$

where $\Sigma\delta\text{PE}_R^{(0)} = 0$. Figures 6.4 (a) and (b) plot $\Sigma\delta\text{PE}_R^{(n)}$ against n for the experiments corresponding to $\text{Ri}_0 = 6.30$ and $\text{Ri}_0 = 0.98$ respectively. The observed evolution is representative of the other non-penetrative and penetrative mixing experiments listed in Table 6.2. For non-penetrative mixing, we see that $\delta\text{PE}_R^{(n)}$ appears to converge towards a constant for large n , corresponding to the constant mixing efficiency regime identified by Olsthoorn and Dalziel (2015). A similar regime is observed for penetrative mixing, though greater

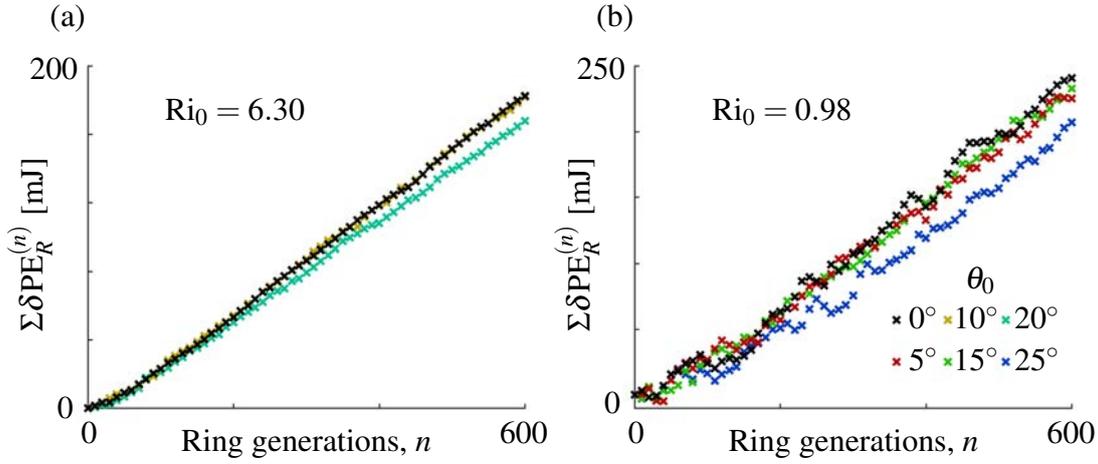


Fig. 6.4 Plots of $\Sigma\delta\text{PE}_R^{(n)}$ against the number of ring generations, n for the cases (a) $\text{Ri}_0 = 6.30$ and (b) $\text{Ri}_0 = 0.98$. Colours correspond to the legend given in (b).

variation in $\delta\text{PE}_R^{(n)}$ is observed, possibly attributable to the variation in density gradients as the stratification evolves.

Using Equation (6.14), we define the average mixing efficiency η_a^b of the k th vortex rings with $a < k \leq b$ as

$$\eta_a^b = \frac{\Sigma\delta\text{PE}_R^{(b)} - \Sigma\delta\text{PE}_R^{(a)}}{(b-a)\delta\text{KE}}. \quad (6.15)$$

Figure 6.5 (a) plots η_{200}^{600} against Ri_0 for each of the experiments listed in Table 6.2. The non-penetrative cases are broadly consistent with the mixing efficiency η_c determined by Olsthoorn and Dalziel (2015), with a mean average mixing efficiency of $1.05\eta_c$ obtained. Similar values for η_{200}^{600} are also obtained in the penetrative regime, with no indication of a drop in mixing efficiency as Ri_0 decreases. While we would expect $\eta \rightarrow 0$ as $\text{Ri}_0 \rightarrow 0$ for the ring–interface interaction, in the penetrative regime the initial interface erodes rapidly, shown for example in Figure 6.1 (b)). The associated changes in the ring-induced mixing of the stratification must be responsible for allowing η_{200}^{600} to be so high for low Ri_0 .

To relate η_0^n to η in the penetrative regime requires averaging over a sequence of vortex-ring generations small enough that the initial interface does not become significantly eroded. Figure 6.5 (b) plots η_0^{80} against Ri_0 for our penetrative-regime experiments. Due to both the large cross-sectional area of the tank the ring-mixed fluid redistributes itself over, and the measurement error of the conductivity probe as it moves through high conductivity gradients (which we discuss in Appendix C), there is a high level of uncertainty associated with our measurements of η_0^{80} . However, it can be seen that measured values of η_0^{80} are significantly

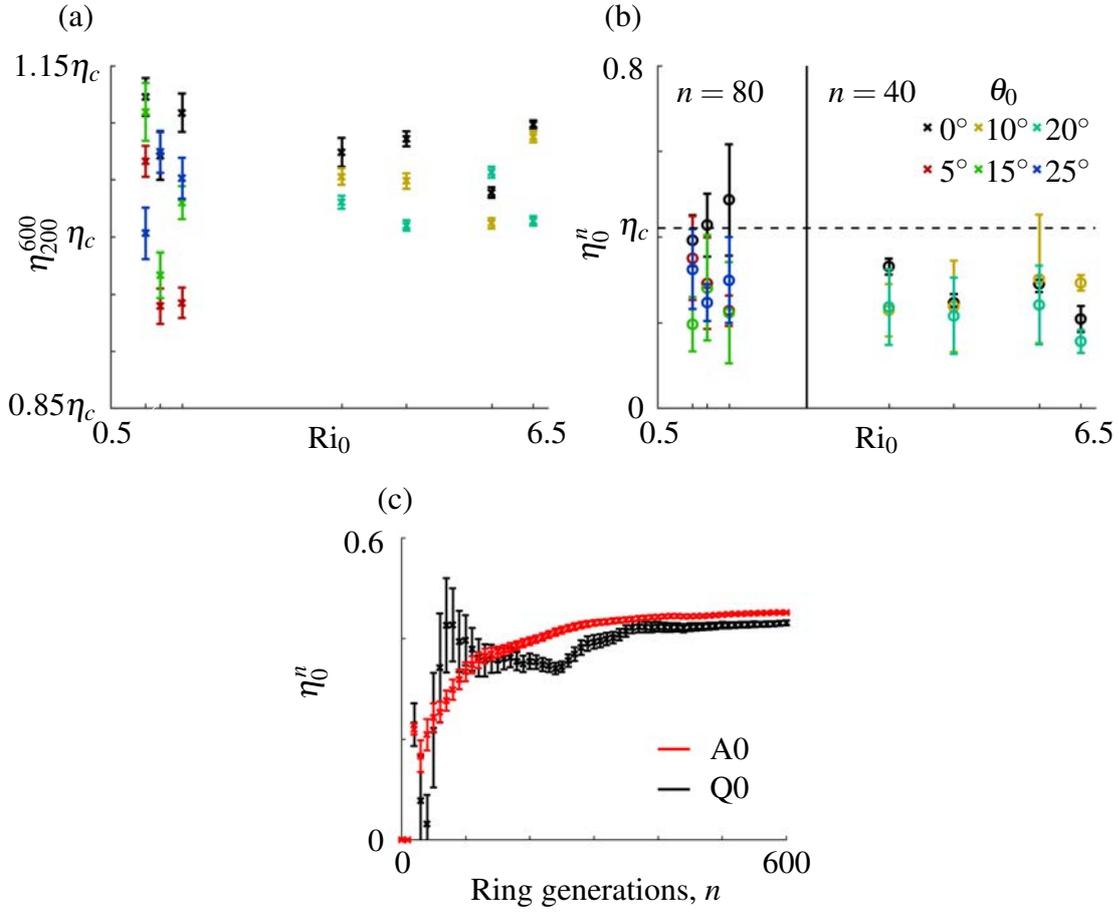


Fig. 6.5 Plots of averaged mixing efficiencies from each experiment listed in Table 6.2 against Ri_0 . (a) Plots of η_{200}^{600} against Ri_0 , where $\eta_c = 0.42$ is the mixing efficiency determined by Olsthoorn and Dalziel (2015). (b) Plots of η_0^n against Ri_0 , where $n = 80$ and $n = 40$ for the penetrative and non-penetrative regimes respectively. Colours for (a) and (b) correspond to the legend given in (b). (c) Plots of η_0^n against n for experiments A0 and Q0, corresponding to the non-penetrative and penetrative regimes respectively. Standard error bars are included in all plots to illustrate the uncertainty associated with our measurements.

lower than η_{200}^{600} . This supports our assertion that, in the penetrative regime, η_{200}^{600} is unsuitable as a proxy for evaluating the mixing efficiency of the ring–interface interaction.

For our non-penetrative-regime experiments, Figure 6.5 (b) plots η_0^{40} against Ri_0 . The level of disparity between η_0^{40} and the corresponding values of η_{200}^{600} is substantial, with $0.4 < \eta_0^{40}/\eta_c < 0.8$ and $1 < \eta_{200}^{600}/\eta_c < 1.1$. The data is suggestive that the ring-induced mixing of the stratification, once it has converged to its self-similar form, is significantly more efficient when compared to the mixing induced by the ring–interface interaction. To understand why this is the case, we now consider in detail the evolution of the stratification for non-penetrative mixing.

6.3 Evolution of stratification for non-penetrative mixing

We describe the evolution of the stratification in the non-penetrative regime as comprising of two stages. In keeping with Olsthoorn and Dalziel (2015), we refer to the first stage of the evolution as the ‘initial-adjustment phase’ (IAP). The second stage of evolution is referred to here as the ‘quasi-steady state’ (QSS). In Section 6.3.1 we describe the evolution of the IAP, restricting our discussion to the case $\theta_0 = 0^\circ$. Then, in Section 6.3.2, we discuss the evolution of vortex rings in the QSS and the associated mixing they induce. In Section 6.3.3, we describe the interplay between the mixing mechanisms in the QSS that appear to be responsible for locking the evolution of the stratification into a state of constant upward entrainment across the interface and constant mixing efficiency per ring generation. Finally, in Section 6.3.4 we discuss the influence of control parameters on the evolution of the QSS.

6.3.1 The initial-adjustment phase

To visualise the evolution of the stratification in the non-penetrative regime, Figure 6.6 plots selected density profiles for the case corresponding to $(\text{Re}_0, \text{Ri}_0, \theta_0) = (6.30, 2960, 0^\circ)$. At the start of the experiment, the stratification has two layers with a diffuse interface of thickness $\sim 2.5\text{cm}$. Vortex rings impact the density interface but are only able to deform the interface slightly due to the magnitude of the restoring buoyancy force relative to the inertial force of the ring (characterised by high Ri_0). The combined effect of the ring transporting lower-layer fluid around its exterior into the upper layer, and the ‘splashing mechanism’ described by Linden (1973), leads to the turbulent entrainment of lower-layer fluid into the upper layer. The direction of entrainment leads to the creation of mixed fluid, mostly with a low range of intermediate densities given by $0 < \rho' = \rho'_M \lesssim 0.1 - 0.2$, where $\rho' = (\rho - \rho_1)/(\rho_2 - \rho_1)$. Associated with this mixing is a decrease in the interfacial height H_2 and a slight sharpening of the density interface. The interface converges to an equilibrium thickness that is maintained throughout the rest of the experiment.

For $30 \lesssim n' \lesssim 120$, the stratification is essentially three-layered, with a ‘middle-mixed layer’ that forms and grows in thickness due to the ring-induced mixing. The interface between the upper and middle-mixed layers progressively steepens, which is clearly visible in Figure 6.6. In a similar way to Ri_0 , we can construct a bulk Richardson number Ri_M associated with the ring interacting with the upper-middle interface, which will scale as $\text{Ri}_M \sim 0.1\text{Ri}_0$ as $\rho' \sim 0.1$ in the middle-mixed layer. For the range of Ri_0 considered in this section, we have $\text{Ri}_M \sim 0.3 - 0.6$, indicating that the ring will be able to penetrate deeply into the middle layer and the splashing mechanism will be able to eject middle-layer fluid high into the upper layer due to the weak restoring buoyancy force. In a similar way as

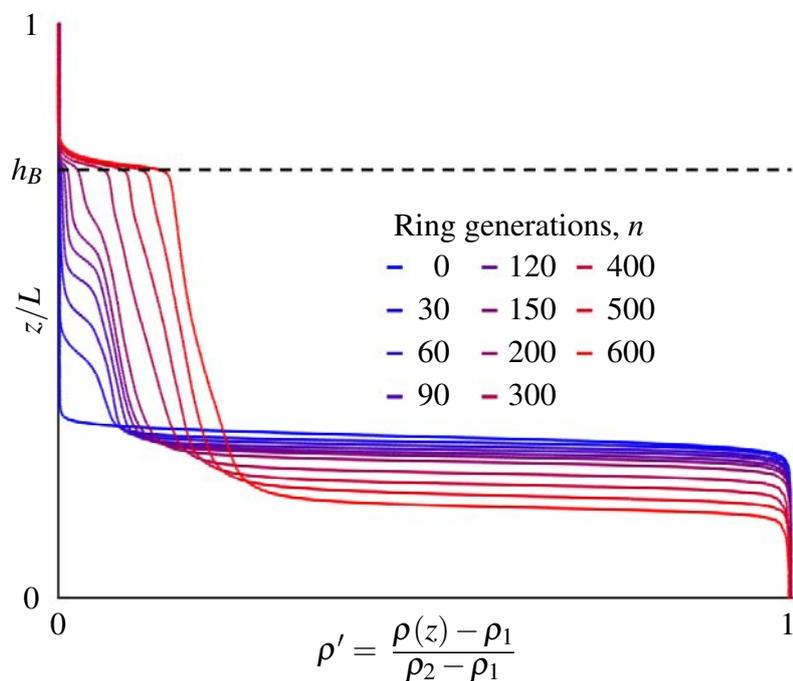


Fig. 6.6 Non-dimensionalised density profiles selected to visualise the evolution of the periodic-mixing system for case A0, corresponding to $(Re_0, Ri_0, \theta_0) = (2960, 6.30, 0^\circ)$. Here, the non-dimensionalised tube height $h_B = H_B/L = 0.75$, and unmixed upper and lower-layer densities are $\rho_1 = 1.0017\text{g/cm}^3$ and $\rho_2 = 1.0481\text{g/cm}^3$ respectively.

for $n' \lesssim 30$, the splashing mechanism leads to the formation of a mixed region above the middle-mixed layer, with $\rho' \sim 0.1\rho'_M \sim 0.01$. This mixed region grows in vertical extent due to the successive ring mixing events up to the height of the ring-tube, at $z = H_B$. The increase in density of the middle-mixed layer and the decrease in H_2 indicates that the generated rings continue to be able to entrain and mix lower-layer fluid into the middle-mixed layer.

For $120 \lesssim n' \lesssim 300$, the interface between the middle-mixed layer and the mixed region above is gradually eroded by the ring-induced mixing. The ring-induced splashing at the eroding interface leads to the weakly mixed layer growing to slightly above the height of the ring-tube opening, at $z = H_B + \delta H_B$ ($\delta H_B \sim 1\text{cm}$ in our experiments). After its erosion, the stratification is again three-layered. The unmixed upper layer occupying the region $z > H_B + \delta H_B$ is effectively dynamically isolated from the ring-induced mixing for the remainder of the experiment. The dynamically active region of the stratification can therefore be regarded as two-layered, and so after this evolutionary stage of the experiment we refer to the middle-mixed layer as the upper layer, and refer to the dynamically-isolated upper layer as the ‘isolated layer’.

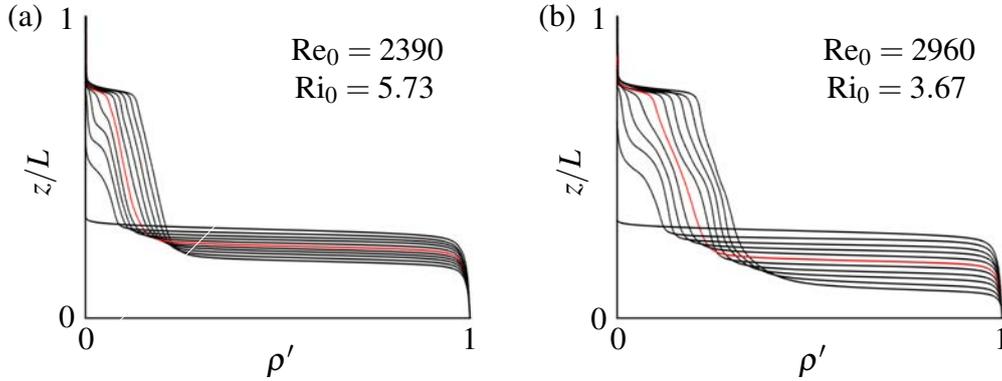


Fig. 6.7 Non-dimensionalised density profiles after every 60 ring generations for the cases (a) B0 and (b) D0. Red profiles mark the transition between the IAP and the quasi-steady state.

Figures 6.7 (a) and (b) respectively plot density profiles after every 60 ring generations for the cases $(Re_0, Ri_0, \theta_0) = (2390, 5.73, 0^\circ)$ and $(2960, 3.67, 0^\circ)$, with red profiles marking the end of the IAP. The IAP for these cases clearly follow a similar evolution to the case illustrated in Figure 6.6. Note that for the case corresponding to $Ri_0 = 3.67$, variations in the upper-layer density gradient develop after the IAP. We return to this point in Section 6.3.4.

The homogenisation of the upper-layer density gradient over the IAP clearly demonstrates that each ring induces mixing both in the upper layer and at the density interface. Additionally, the mixing induced by each ring is seen to produce fluid with a range of intermediate densities, as opposed to perfectly mixed fluid of a single density. These observations appear to be in contrast to the banding process described by Olsthoorn and Dalziel (2015) for the QSS, where each vortex ring is envisaged to produce a perfectly mixed layer of fluid at the bottom of the upper layer. However, the evolution of the stratification associated with the QSS is less clear, particularly the mechanisms in play that control the micro-structure of the upper layer. To understand these mechanisms, we now consider the evolution of the stratification in the QSS.

6.3.2 Ring evolution and mixing in the quasi-steady state

In the QSS, each ring generation can be thought of as a four-stage process: the formation of the vortex ring; its propagation through the stratified upper layer; the interaction of the ring with the density interface; and the relaxation of the system, where the fluid finds its new neutral buoyancy level and the remaining kinetic energy is eventually dissipated. We now

discuss each of these stages in turn, introducing the parameters that control the evolution of each ring generation after the IAP.

The gradual development of the interface between the isolated and upper layers is suggestive that the ring formation process may be playing a role in mixing the two layers. However, it was observed in real time that the gradual development of the isolated–upper interface is attributable to splashing that occurred at the interface shortly after each ring–interface interaction, due to the lightest ring-entrained fluid finding its new neutral buoyancy level and overshooting into the isolated layer under its own inertia. We therefore assume that a negligible amount of fluid initially above the height of the ring tube is entrained into the ring during the formation process. Additionally, assuming the entirety of the ring roll-up process will occur within one ring-diameter of the tube-height opening, for simplicity we take the ring density $\rho_T^{(n)}$ of the n th generated vortex ring to be

$$\rho_T^{(n)} = \rho^{(n)}(z = H_B - a/2). \quad (6.16)$$

We infer the dynamics of the ring propagating through the stratified upper layer from the observations of Maxworthy (1977) on vortex rings propagating vertically through a linear stratification. This comparison is justified for vertically propagating rings as the ring diameter is much smaller than the density scale height in the upper layer, i.e.

$$\frac{a}{\rho_T^{(n)}} \left| \frac{d\rho}{dz} \right| \ll 1. \quad (6.17)$$

We do not go into the detail of Maxworthy’s observations here, but importantly they justify two further assumptions. Firstly, that the ring will retain a similar form as it propagates through the upper layer, and secondly, that the dynamic pressure gradient induced by the ring (that is radial in the frame of the ring core) allows the vorticity-concentrated ring core to retain its initial density of $\rho_T^{(n)}$.

Conversely, as the ring propagates downward through the upper layer, the ring-entrained fluid external to the core (its ‘atmosphere’) will gradually increase in density. This is due to a proportion of the entrained fluid being continually rejected into the wake of the ring as it becomes lighter than the surrounding stratification, which in turn leads to the ring entraining heavier fluid from the ambient into the ring atmosphere. We can therefore approximate the ring as comprising of fluid with two densities: the core, with density $\rho_T^{(n)}$, and the ring atmosphere, with intermediate density $\rho_I^{(n)}(z)$ such that

$$\rho_T^{(n)} < \rho_I^{(n)}(z) < \rho^{(n)}(z). \quad (6.18)$$

As $\rho_I^{(n)}(z)$ is set by the dynamics of the ring and the background stratification in the upper layer, it follows that the evolution of the ring as it propagates through the upper layer is controlled by Re_0 , θ_0 and the gradient Richardson number $\text{Ri}_g^{(n)}(z)$, which we define as

$$\text{Ri}_g^{(n)}(z) = \frac{-g}{\rho_T^{(n)}} \frac{d\rho^{(n)}}{dz} \frac{a^2}{U^2}. \quad (6.19)$$

As the stratification evolves in the QSS, $\rho_T^{(n)}$ strictly increases. A natural way to characterise the interaction of the ring with the density interface would be to introduce a transitional Richardson number, $\text{Ri}_T^{(n)}$, given by

$$\text{Ri}_T^{(n)} = \frac{g(\rho_2 - \rho_T^{(n)})}{\rho_T^{(n)}} \frac{a}{U^2}. \quad (6.20)$$

Note that $\text{Ri}_T^{(0)} = \text{Ri}_0$, and $\text{Ri}_T^{(n)}$ only decreases as $\rho_T^{(n)}$ increases. Over the number of ring generations each experiment is run for, $\text{Ri}_T^{(n)}$ is seen to decrease linearly in the QSS to a good approximation, by a constant δRi_T per ring generation (corresponding to an approximately constant increase in $\rho_T^{(n)}$ of $\delta\rho_T$ per ring generation).

Figure 6.8 (a) plots δRi_T against θ_0 for our non-penetrative-regime experiments, where δRi_T is determined by fitting a line through $\text{Ri}_T^{(n)}$ from $n = 200$ to $n = 600$ using a least squares approach. Note that we return to the differences in the functional dependencies of δRi_T on θ_0 for the different cases of $(\text{Re}_0, \text{Ri}_0)$ in Section 6.3.4. Figure 6.8 (b) plots $\text{Ri}_T^{(n)}$ against n for the experiments corresponding to $\theta_0 = 0^\circ$. Across this sequence of ring generations, we see a total reduction in $\text{Ri}_T^{(n)}$ of $0.55 - 0.75$. If we assume that $\text{Ri}_T^{(n)}$ characterises the ring-induced mixing in precisely the same way as Ri_0 characterises the mixing associated with the ring-interface interaction, the observed reduction in $\text{Ri}_T^{(n)}$ would represent a significant change in the ring-induced mixing at the density interface.

As discussed in Olsthoorn and Dalziel (2015), in the QSS, the change in interfacial height δH_2 per ring generation is constant in a given experiment, and is seen to scale as $\delta H_2 \sim \text{Ri}_0^{-1}$. Figure 6.8 (c) plots the product of Ri_0 and the non-dimensionalised interfacial-height change $A\delta H_2/V_R$ in QSS per ring generation against θ_0 . Writing $H_2^{(n)}$ for the interfacial height immediately before the n th ring generation, $H_2^{(n)}$ is determined as the position at the interface with maximal density gradient $|\partial\rho/\partial z|$. Then, δH_2 is determined by fitting a line through $H_2^{(n)}$ from $n = 200$ to $n = 600$ using a least squares approach.

The ring-induced upward entrainment of fluid across the density interface is seen to be both independent to changes in $\text{Ri}_T^{(n)}$, and highly correlated with Ri_0^{-1} . Additionally, in the early stages of the IAP, the ring-induced upward entrainment is also seen to depend on

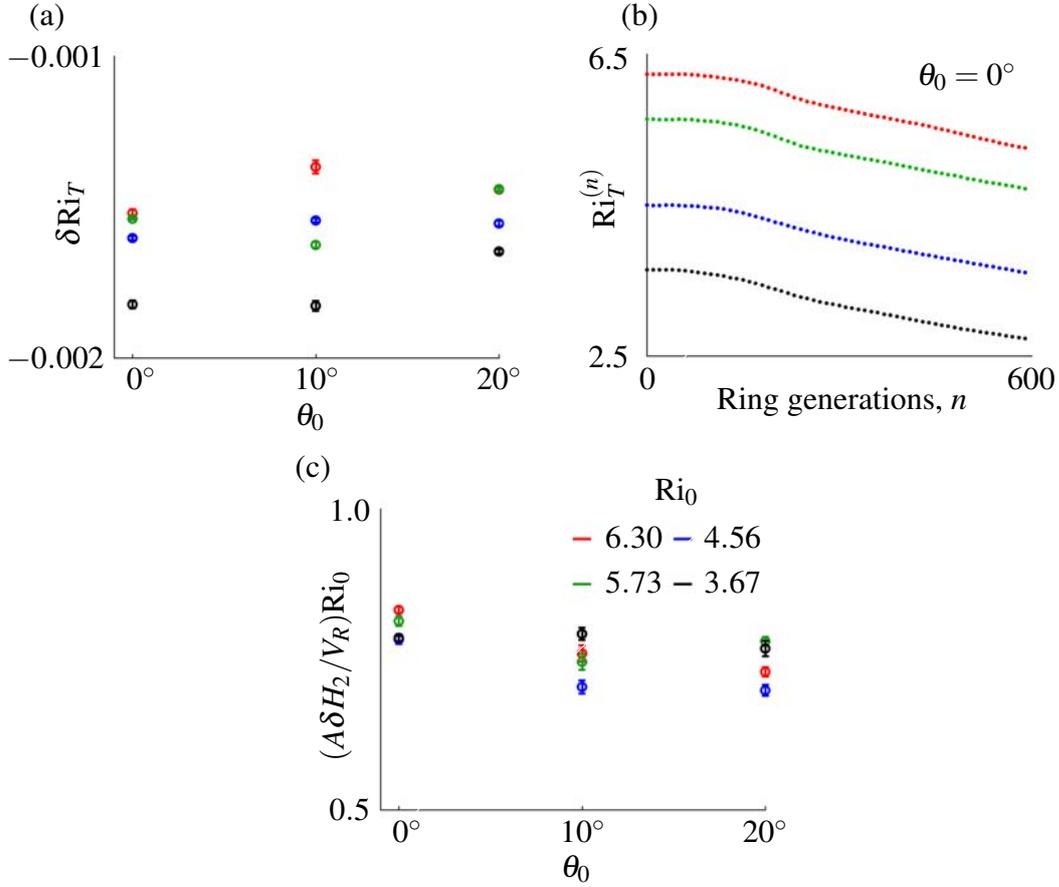


Fig. 6.8 (a) The average change in transitional Richardson number δRi_T per ring generation, against θ_0 . (b) Plot of $Ri_T^{(n)}$ against n for the cases corresponding to $\theta_0 = 0^\circ$, demonstrating the typical evolution of $Ri_T^{(n)}$ in the non-penetrative regime. (c) The product of Ri_0 and the non-dimensionalised interfacial height change $A\delta H_2/V_R$ in the quasi-steady state per ring generation, against θ_0 . Standard error bars for (a) and (c) are included.

Ri_0 . If $Ri_T^{(n)}$ characterised the interaction of the ring with the density interface during the QSS in the same way that Ri_0 characterises the ring–interface interaction, we would observe a dependency of δH_2 on $Ri_T^{(n)}$. As this dependency is not observed, we must conclude that the upper-layer $Ri_g^{(n)}(z)$ profile evolves in such a way that maintains a constant rate of upward-entrainment across the density interface. We return to this shortly in Section 6.3.3.

When the vortex ring reaches the density interface, the buoyancy force arrests the downward motion of the ring. The ring entrains fluid from the lower layer into the upper layer, leading to some irreversible mixing. As the vortex ring transports fluid of density $\rho_T^{(n)}$ to the interface, in an analogous way to the ring-mixed fluid produced by the first 30 ring generations in the IAP, we expect the ring-mixed fluid produced to have a distribution of

densities, given by $0 < \rho_M^{(n)} < 0.1 - 0.2$, where $\rho_M^{(n)} = (\rho - \rho_T^{(n)}) / (\rho_2 - \rho_T^{(n)})$. Subsequently, due to the weak stratification exhibited in the upper layer, the ring-mixed fluid redistributes itself throughout the upper layer as it finds its new neutral buoyancy level.

Additionally, as the vortex ring breaks down, the fluid transported within the ring core (of density $\rho_T^{(n)}$) begins to return to its neutral buoyancy level as the hydrostatic pressure gradient starts to dominate over the dynamic pressure gradient induced by the vortex ring. As the core fluid and ring-mixed fluid finds its new neutral buoyancy level, convective instabilities will develop, leading to further irreversible mixing. This mixing mechanism is not exhibited in the ring–interface interaction, and so would appear to explain the relatively high mixing efficiencies observed in the QSS compared to the IAP (shown in Figures 6.5 (a) and (b)).

6.3.3 Mixing equilibrium in the quasi-steady state

In the QSS, we observe over the range of parameters considered that the ring-induced upward entrainment of fluid from the lower layer across the density interface is controlled predominantly by two factors. Firstly, the increase in the ring density $\rho_T^{(n)}$. This is directly associated with both the vertical confinement of the evolving stratification, imposed by the prescription of the ring-tube height H_B , and the mixing associated with core and ring-mixed fluid finding its new neutral buoyancy level after each ring interaction with the density interface. Secondly, the increase in the vertically-integrated buoyancy $\langle B^{(n)} \rangle$ of the upper layer. By buoyancy conservation, $\langle B^{(n)} \rangle$ can be written as

$$\langle B^{(n)} \rangle = \int_{H_2^{(n)}}^{H_B} g (\rho^{(n)}(z) - \rho_1) dz = g(H_2 - H_2^{(n)}) (\rho_2 - \rho_1). \quad (6.21)$$

The increase in $\langle B^{(n)} \rangle$ is directly associated with the upward entrainment of lower-layer fluid across the density interface.

For a given experiment, in the QSS the stratification evolves so that $\rho_T^{(n)}$ and $\langle B^{(n)} \rangle$ follow an equilibrium curve in $(\rho_T, \langle B \rangle)$ -space, which we denote as $\rho_{T,e} = f(\langle B \rangle)$. A balance is reached through a competition between two mixing mechanisms. First, the upward entrainment of lower-layer fluid across the density interface, leading to an increase in $\langle B^{(n)} \rangle$. The associated strengthening of the stratification in the upper layer then leads to a decrease in the kinetic energy delivered by subsequent vortex rings to the interface (hence inhibiting their ability to entrain fluid upwardly across the interface). Second, the ring-induced transport of fluid from the top of the upper layer to the interface and the subsequent convective mixing that occurs, leading to increases in ρ_T . This reduces the buoyancy difference between the ring-transported fluid and the upper-layer stratification, allowing subsequent vortex rings to

deliver more kinetic energy to the interface (hence enhancing their ability to entrain fluid upwardly across the interface).

From this perspective, the IAP can be viewed as a state before the equilibrium $\rho_{T,e} = f(\langle B \rangle)$ is reached. If $\rho_T^{(N)} > f(\langle B^{(N)} \rangle)$ for given N (which is the case at the start of the IAP), the buoyancy upwardly entrained across the interface per ring generation decreases as the ring-entrained upper-layer buoyancy inhibits subsequent vortex rings from mixing at the density interface. This inevitably leads to $\rho_T^{(n)} \rightarrow f(\langle B^{(n)} \rangle)$. Conversely, if $\rho_T^{(N)} < f(\langle B^{(N)} \rangle)$ for given N , subsequent vortex rings would mix the upper-layer stratification, homogenising it sufficiently so that the ring-induced upward entrainment across the interface balanced the mixing of the upper-layer stratification, corresponding to $\rho_T^{(n)}$ converging onto $f(\langle B^{(n)} \rangle)$. Over the range of parameters considered, the competition between these two mixing mechanisms appears to be responsible for locking the evolution of the stratification into a state corresponding to constant upward entrainment across the interface, and constant mixing efficiency per ring generation.

6.3.4 Influence of control parameters on the QSS

The control parameters (Re_0, Ri_0, θ_0) , in addition to the vertical distance $(H_B - H_2)/a$, will predominantly control how the stratification evolves in the equilibrium state described in Section 6.3.3. To aid our discussion how the input parameters control the evolution of the flow, Figures 6.9 (a) to (d) plot overlaid density profiles, comparing the evolution of the stratification between $\theta_0 = 0^\circ$ and 20° for each of the cases of (Re_0, Ri_0) investigated. The profiles plotted correspond to $n = 120, 300$ and 480 , and black and red profiles correspond to the cases where $\theta_0 = 0^\circ$ and $\theta_0 = 20^\circ$ respectively.

A striking feature of the compared density-profile pairs is the level of similarity in each pair, particularly for the cases in Figures 6.9 (a) and (b), corresponding to higher values of Ri_0 . In the large Ri_0 regime, we have $\delta H_2 \sim Ri_0^{-1}$, and so the vortex rings are only able to weakly scour the interface, upwardly entraining only a small amount of lower-layer fluid across the interface. This implies that, for large Ri_0 , the dynamical differences of the ring scouring the interface between the cases $\theta_0 = 0^\circ$ and 20° will only have a small effect on the differences in the composition of mixed fluid after each interaction.

If this is the case, for large Ri_0 , we can expect that the majority of the mixing induced will be associated with the previously-described convective instabilities. This would imply that the mixing efficiency of sustained periodic forcing is controlled predominantly by the $Ri_g^{(n)}(z)$ profile (introduced in Equation (6.19)) in the upper layer. For $4 \leq Ri_0 \leq 12$, Olsthoorn and Dalziel (2015) observe the stratification maintaining a self-similar form, with a weakly stratified upper layer of constant density gradient. For the cases $\theta_0 = 0^\circ$, Figure 6.10 (a) plots

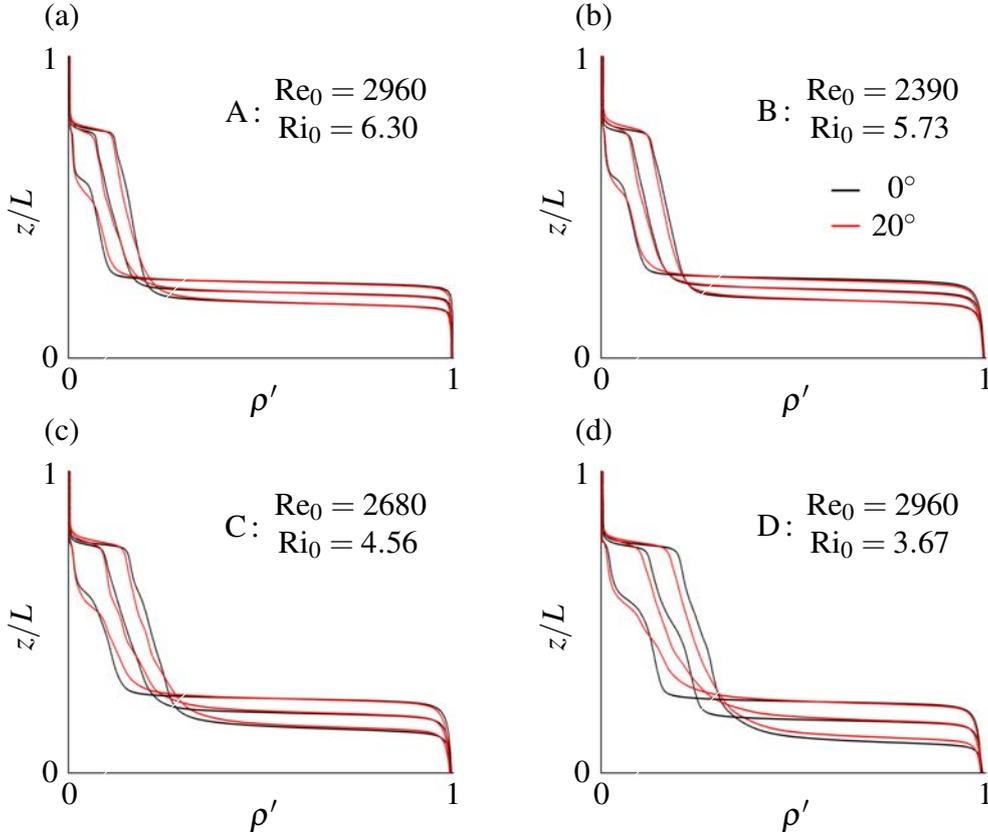


Fig. 6.9 Non-dimensionalised density profiles after 120, 300 and 480 ring generations for the cases (a) (A0, A20), (b) (B0, B20), (c) (C0, C20), and (d) (D0, D20) respectively. Black and red profiles correspond to the cases where $\theta_0 = 0^\circ$ and $\theta_0 = 20^\circ$ respectively.

the mean upper-layer gradient Richardson number, $\overline{Ri_{g,U}^{(n)}}$, which we define as

$$\overline{Ri_{g,U}^{(n)}} = \frac{1}{H_B - H_2^{(n)} - a} \int_{H_2^{(n)} + a/2}^{H_B - a/2} Ri_g^{(n)}(z) dz. \quad (6.22)$$

Additionally, Figure 6.10 (b) plots the standard deviation of $Ri_{g,U}^{(n)}$, given by

$$\text{std}(Ri_{g,U}^{(n)}) = \frac{1}{\sqrt{H_B - H_2^{(n)} - a}} \sqrt{\int_{H_2^{(n)} + a/2}^{H_B - a/2} \left(Ri_g^{(n)}(z) - \overline{Ri_{g,U}^{(n)}} \right)^2 dz}, \quad (6.23)$$

to demonstrate how closely $Ri_{g,U}^{(n)}$ approximates to exhibiting a constant gradient. Note that we expect

$$\overline{Ri_{g,U}^{(n)}} \sim (H_B - H_2)^{-1},$$

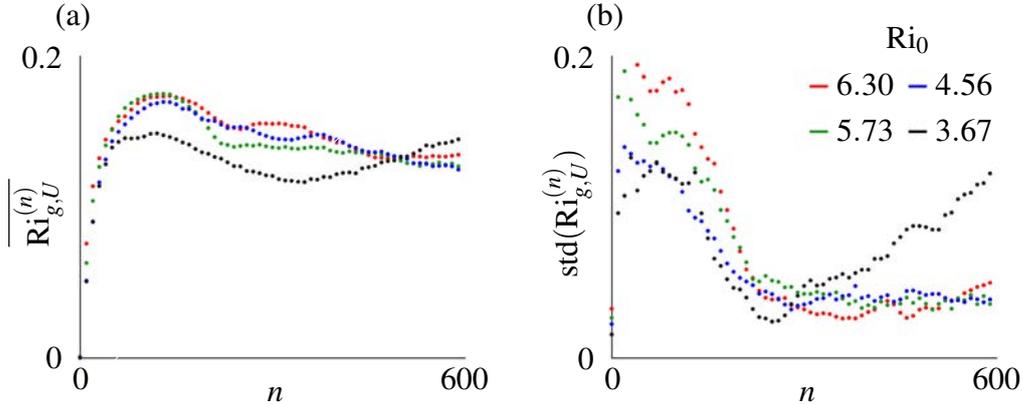


Fig. 6.10 Plots of (a) $\overline{\text{Ri}_{g,U}^{(n)}}$, and (b) $\text{std}(\text{Ri}_{g,U}^{(n)})$, against ring generation n for $\theta_0 = 0^\circ$. Cases A-D correspond to the labels given in Table 6.2 and the black profiles in Figures 6.9 (a)-(d).

as the mixing equilibrium ensures that vortex rings are always able to upwardly entrain some fluid from the lower layer across the interface.

The high level of overlap between cases A-C of $\overline{\text{Ri}_{g,U}^{(n)}}$ is suggestive that, after the IAP for $\text{Ri}_0 \gtrsim 4$, the mixing equilibrium discussed in Section 6.3.3 fixes $\overline{\text{Ri}_{g,U}^{(n)}}$ to evolve in a way that is Ri_0 -independent to a high approximation. Assuming in this regime that convective instabilities are responsible for the majority of the mixing, this may explain the high, Ri_0 -independent mixing efficiency reported by Olsthoorn and Dalziel (2015). For case D (corresponding to $\text{Ri}_0 = 3.67$), an increase is observed in $\text{std}(\text{Ri}_{g,U}^{(n)})$, which appears to be associated with the strengthening of the stratification towards the bottom of the upper layer, exhibited in Figure 6.9 (d). This is suggestive of a transition in the non-penetrative regime where the mixing induced by the ring interacting with the interface becomes a more significant fraction of the total mixing as Ri_0 decreases.

Another feature observed in the presented density-profile pairs in Figure 6.9 is that, in the $\theta_0 = 0^\circ$ case, the periodically forced system is more effective at transporting buoyancy towards the top of the upper layer. Conversely, in the $\theta_0 = 20^\circ$ case, vortex rings appear to be more effective at inducing mixing towards the bottom of the lower layer, mixing with the fluid upwardly entrained across the interface. Due to the non-zero bulk horizontal momentum of the vortex ring for the $\theta_0 = 20^\circ$ case, it seems reasonable to expect that obliquely propagating rings will be more effective at generating overturning events towards the bottom of the upper layer, relative to the $\theta_0 = 0^\circ$ case. In turn, the mixing produced by these overturning events would lead to an accumulation of more evenly mixed fluid at the bottom of the upper layer, and less convective mixing as the overturning-mixed fluid would already be in close

proximity to its neutral buoyancy level immediately after its generation. This suggested mechanism would explain the observed decrease in δRi_T for the $\theta_0 = 20^\circ$ case relative to the $\theta_0 = 0^\circ$ case, shown in Figure 6.8 (a), and would be consistent with the larger discrepancies between compared profile pairs being exhibited for smaller values of Ri_0 .

Unfortunately, the range of Re_0 investigated here is not large enough to be able to attribute features of the evolution of the stratification directly to the Reynolds number. However, as Re_0 increases, we do expect that the vortex ring would become more susceptible to developing the Widnall instability, or other instabilities, as it propagates through the weakly-stratified upper layer. Such instabilities may lead to some mixing occurring before the ring reaches the interface, possibly accounting for the variations in $d\rho/dz$ in the upper-layer exhibited in Figure 6.7 (b) and Figures 6.9 (c) and (d). To confirm this, further experiments would need to be conducted that probe the evolution of the vortex ring in this stage of the flow.

6.4 Summary

In this chapter, we have conducted experiments in which vortex rings were generated periodically, mixing an initially two-layered density stratification. Experiments were classified into ‘penetrative’ and ‘non-penetrative’ regimes, corresponding to whether downward entrainment of upper-layer fluid into the lower layer occurred in addition to upward entrainment of lower-layer fluid into the upper layer. It was demonstrated that, in both regimes, the measured mixing efficiencies associated with the sustained forcing of vortex rings in the QSS was significantly higher than mixing efficiencies associated with the mixing induced by a relatively small number of vortex-ring generations. For the non-penetrative regime, this called into question the association of the mixing efficiency $\eta_c = 0.42$ determined by Olsthoorn and Dalziel (2015) with the mixing efficiency of a ring–interface interaction between otherwise homogeneous layers.

In the non-penetrative regime, the stratification evolves such that, after the initial-adjustment phase (IAP), an energetic balance is maintained between the ring-induced transport of fluid from the top of the upper-layer to the interface, ring-induced mixing at the interface, and convective instabilities that develop as the core and ring-mixed fluid rise to find their new neutral buoyancy levels. Figure 6.11 sketches the evolution of the horizontally-averaged density profile $\bar{\rho}(z)$ during a vortex ring mixing event after the IAP. Over the range of parameters considered, the energetic balance locks the evolution of the stratification into a state corresponding to a constant rate of upward entrainment across the interface (proportional to Ri_0^{-1}), and constant mixing efficiency.

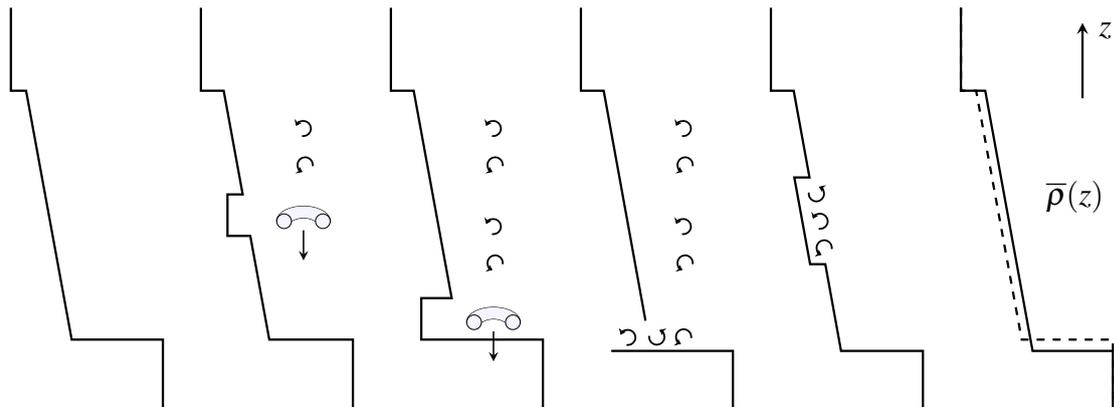


Fig. 6.11 A simplified sketch of the horizontally-averaged density profile $\bar{\rho}(z)$ during the lifespan of a vortex ring mixing event after the IAP. Each vortex ring entrains fluid from the lower layer, after which convective mixing occurs as the ring-transported and ring-mixed fluid rise to find their new neutral buoyancy levels.

We propose that the high mixing efficiency observed in the QSS should be attributed to the combination of the mixing induced by the ring interacting with the interface and the subsequent convective instabilities that develop after each ring interaction with the interface. This proposition, in combination with our measurements of η_0^{40} in Figure 6.5 (b) (which fell in the range $0.15 < \eta_0^{40} < 0.33$ over the parameters investigated), suggests that the mixing efficiency η of the ring–interface interaction between two homogeneous layers is significantly lower than the reported value of $\eta_c = 0.42$. Moreover, the precise functional dependency of η on Ri_0 remains unclear, let alone its dependencies on other likely-important control parameters, including Re_0 , θ_0 and the interfacial thickness, δz . Considerations relating to future research in characterising η are discussed in Section 7.3.

Chapter 7

Conclusions

This chapter begins by presenting a summary to this thesis, given in Section 7.1. Then, in Section 7.2, we relate our findings to Turner’s mixing box and the ring–eddy analogy. Potential avenues for future research are suggested in Section 7.3, after which a final reflection on this thesis is given in Section 7.4.

7.1 Thesis summary

The primary motivation of the research presented in this thesis was to understand how the dynamics and mixing induced by a vortex ring, interacting with a sharply-stratified density interface, change with respect to the initial propagation angle θ_0 of the vortex ring towards the interface. Olsthoorn and Dalziel (2017) identified a Crow-like instability for the ring–interface interaction at $\theta_0 = 0^\circ$, which was seen to be effective at transferring the energy of the ring to smaller scales for irreversible mixing. They associated the instability with the unexpectedly high mixing efficiency of $\eta_c = 0.42$, reported earlier by Olsthoorn and Dalziel (2015) for a periodic sequence of vortex rings mixing an initially two-layered stratification. This called into question the usefulness of relating the ring–interface interaction at $\theta_0 = 0^\circ$ to the ring–eddy analogy introduced by Maxworthy (1972) and Linden (1973), as it remained unclear whether the associated dynamics and entrainment mechanisms were highly unique in the broader family of eddy interactions with density interfaces. As such, in this thesis we chose to restrict our attention to cases with $\theta_0 \leq 25^\circ$, to investigate how breaking azimuthal symmetry by introducing a small to moderate ring-propagation angle would alter the dynamics and mixing properties of the ring–interface interaction.

This thesis has presented an experimental investigation into the dynamics and mixing properties of obliquely propagating vortex rings interacting with a density interface. Before presenting our results, an extensive review of background theory and literature relevant to

our investigation was carried out in Chapter 2, which enabled us to infer three-dimensional features of the ring–interface interaction from our two-dimensional measurements. Then, in Chapter 3, we discussed the experimental apparatus, setup and procedures we made use of to conduct the experiments presented in this thesis.

The dynamics of ring–interface interactions was investigated using two-dimensional PIV and LIF measurements to probe the velocity and density fields in the central plane of the flow. Observations from these experiments were discussed in Chapter 5. To reduce the relative measurement error (RME) across all our velocity measurements, we developed a ‘multi-frame’ PIV algorithm, that locally optimises the time separation ΔT between camera frames to achieve the RME reduction. This algorithm, in addition to our particular implementation of the algorithm for our experiments, was introduced in Chapter 4. Although multi-frame-algorithm variants have been previously developed by Pereira et al. (2004), Hain and Kähler (2007) and Persoons and O’Donovan (2011), we believe our algorithm is the first that takes into account having only a single suitable image-pair time step available to construct low RME velocity measurements in the fastest moving regions of the flow. As a result, our algorithm is uniquely suited to PIV setups where the ratio between the two smallest available image-pair time steps is greater than two, such as the setup for the PIV/LIF experiments presented in this thesis.

Our PIV/LIF experiments presented in Chapter 5 represent the first systematic study on the dynamics of vortex rings obliquely impacting a density interface. For both non-penetrative and penetrative ring–interface interactions, a wide range of dynamical behaviour was observed across the range of propagation angles ($\theta_0 \leq 20^\circ$) investigated. Comparing ring–interface interactions between the $\theta_0 = 0^\circ$ and 5° cases, the break in azimuthal symmetry associated with non-zero θ_0 leads to the triggering of different instability mechanisms that may have significant consequences for the entrainment and small-scale mixing that occurs. Indeed, our measurements of η_0^{40} and η_0^{80} in Figure 6.5 (b) are suggestive that the mixing efficiency η of the ring–interface interaction may have a sensitive dependence on its control parameters (Re_0, Ri_0, θ_0) for the range of parameters investigated, possibly in addition to other hidden dependencies. As such, it appears unlikely that η can be characterised by having a parametric dependence on its control parameters, though it has been commonplace to attempt to do so for other flows (Caulfield, 2021).

With the initial aim of probing the mixing efficiency of the ring–interface interaction, experiments were conducted in which a periodic sequence of vortex rings were generated to mix an initially two-layered stratification. Mixing-efficiency measurements made after a small number of vortex-ring generations were shown to be consistently lower than the measurements made after the sustained forcing of hundreds of vortex rings, indicating that

additional mixing mechanisms may be at play after a finite number of ring generations. This was found to be the case.

By analysing in detail the evolution of the stratification for non-penetrative mixing, a balancing mechanism was identified between two mixing mechanisms. Firstly, the ring-induced upward entrainment of lower-layer fluid across the density interface (inhibiting the ability of subsequent vortex rings to entrain fluid upwardly across the interface). Secondly, the ring-induced downward transport of low-density fluid from the top of the upper layer to the interface and the subsequent convective mixing that occurs as the ring-transported and ring mixed fluid rise to find their new neutral buoyancy levels in the upper layer (leading to increases in the ring density of subsequent vortex rings, thus enhancing their ability to entrain fluid upwardly across the interface). This balancing mechanism was found to regulate the observed constant upward-entrainment volume (proportional to Ri_0^{-1} for $Ri_0 \gtrsim 3.5$) across the interface per ring generation, in addition to the mean gradient Richardson number $\overline{Ri}_{g,U}$ in the upper layer, which was seen to evolve seemingly independently of Ri_0 for $Ri_0 \gtrsim 4$ for $\theta_0 = 0^\circ$. We proposed in Chapter 6 that, for $Ri_0 \gtrsim 4$, the occurrence of convective mixing, in addition to the Ri_0 -independent evolution of $\overline{Ri}_{g,U}$, were responsible for both the constant mixing-efficiency regime identified by Olsthoorn and Dalziel (2015) and for their reported mixing efficiency of $\eta_c = 0.42$ being so high.

7.2 Turner's mixing box and the ring–eddy analogy

The mixing-box experiments of Turner (1968) conclude that the dominant mechanism for entrainment across the density interface is the largest eddies intermittently interacting with the interface in localised events, thus motivating the ring–eddy analogy. Figure 7.1 (a) sketches the mixing-box experiment during an entrainment event. Relating our ring-mixing experiments in the non-penetrative regime to the mixing-box experiment when grid-induced stirring occurs only in the upper layer, it is reasonable to assume that Turner's largest eddies will entrain fluid across the interface and produce partially-mixed fluid with a distribution of densities with $0 < \rho' \lesssim 0.2$. This partially-mixed fluid is then swept into the upper layer by the less energetic eddies in the flow and mixed perfectly by the grid-generated turbulence.

Figure 7.1 (b) sketches the ring-mixing experiment during an entrainment event in the non-penetrative mixing regime, after the initial-adjustment phase (IAP). As we have seen in Chapter 6, ring-mixing events in the IAP lead to an accumulation of partially-mixed fluid that ultimately occupies the entirety of the upper-layer region, after which the upper layer exhibits a weak density gradient. Subsequently, vortex rings efficiently transport fluid from the top of the upper layer to the interface, with a large fraction of the kinetic energy (KE) lost

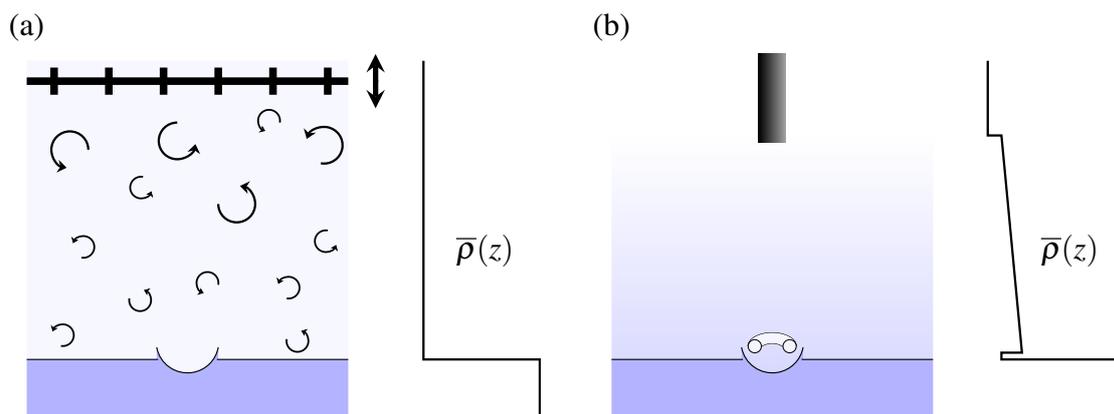


Fig. 7.1 Sketches of the entrainment events in (a) grid-mixing experiments, and (b) ring-mixing experiments after the initial-adjustment phase. Horizontally-averaged density profiles $\bar{\rho}(z)$ are provided for both sketches.

transferred to available potential energy (APE). After the entrainment event, the convectively unstable ring-transported and ring-mixed fluid rise to find their new neutral buoyancy levels, leading to highly efficient mixing as much of the APE in the system is directly converted to background potential energy (BPE).

For both types of mixing experiment, the entrainment event at the interface is not necessarily an efficient mixing event, and much of the total mixing that occurs happens in the upper layer away from the interface. The fundamental difference between the two experiments is the upper-layer mixing process. Grid mixing supplies a lot of KE to the upper layer, a small fraction of which is able to entrain fluid across the interface. The rest of the KE is highly effective at homogenising the upper layer, but a lot of this energy is dissipated without increasing the BPE of the system, and so the mixing efficiency is very low. Conversely, in our ring-mixing experiments, almost all of the mixing that occurs away from the interface is convective mixing of the ring-transported and ring-mixed fluid. This leads to a high overall mixing efficiency of the system, made possible by the inability to perfectly mix the upper layer.

For non-penetrative mixing, our measurements of η_0^{40} suggest that the mixing efficiency η of the ring–interface interaction takes a broad range of values over the range of parameters investigated, with $0.15 < \eta < 0.33$. While the error associated with these measurements is non-negligible due to the precision attainable with the conductivity probe, the wide range of measured values suggest that η is sensitive to the particular instability mechanisms of the ring–interface interaction that transfer energy to the smaller scales for irreversible mixing. If this sensitivity exists and is significant, then careful interpretation should be given in future when attempting to relate the mixing properties of vortex rings to general eddies.

7.3 Avenues for future research

There are a wide range of potential avenues to explore for future research that follow naturally from the research presented in this thesis. Those that appear to be the most natural next steps forward from our perspective are discussed in this section.

7.3.1 Multi-frame PIV algorithm

The multi-frame PIV algorithm we introduced in Section 4.3 only attempts to locally optimise the image-pair time separation ΔT across all regions of the flow (with the aim of obtaining a low relative measurement error for all velocity scales in the flow). To reduce the finest spatial resolutions over which velocity gradients can be accurately resolved, an additional functionality could be incorporated into our multi-frame algorithm that would aim to locally optimise the interrogation window size n_I and spacing δ_I . In a similar way to our optimisation strategy for ΔT , the optimisation of n_I and δ_I could be made by first interrogating image pairs using the largest values of n_I and δ_I selected for use. Subsequently, criteria based on the obtained pixel displacement and displacement gradient measurements could be applied iteratively to determine which regions of the flow required further interrogation at successively smaller values of n_I and δ_I .

The particular implementation of our multi-frame PIV algorithm for our PIV/LIF experiments proved to be sufficient for our purposes, though there are a number of ways it could be further improved. For example, a measurement-validation criterion based on the in-plane divergence $\nabla \cdot \mathbf{u}_{2D}$ could be incorporated to ensure that measurements made in regions with high $\nabla \cdot \mathbf{u}_{2D}$ only made use of the smallest interrogation time step available. This would help to reduce measurement errors associated with the pattern-matching algorithms used having to account for a large change in the number of particle images in a given region of the flow.

A natural extension of the implementation of our PIV algorithm would be to allow the interrogation points selected for further interrogation at smaller image-pair time steps to be disconnected on the sensor array. In this way, disconnected regions exhibiting high velocities in the flow could be interrogated utilising smaller time steps without interrogating the slower moving regions between them. This extension could significantly reduce the total computational time required to process a sequence of PIV images using our algorithm, particularly for turbulent flows where the turbulence is not localised.

As discussed in Section 4.5, it would be useful to use synthetic PIV images to test our algorithm against the pre-existing PIV functionality in DigiFlow. Such tests could also

be used to understand how to optimise the parameters associated with our measurement-validation and selection criteria for different types of flow and imaging conditions.

7.3.2 Ring–interface interaction dynamics

In a similar way to the experiments of Walker et al. (1987) for vortex rings normally impacting a solid wall (discussed in Section 2.3.1) and complementing the experimental observations of Dahm et al. (1989), it would be interesting to conduct dye experiments to investigate the dependence of the ring–interface interaction at $\theta_0 = 0^\circ$ on Re_0 for a large set of Reynolds number in the range $100 < Re_0 < 2000$, for a small set of Ri_0 (corresponding to non-penetrative and penetrative cases). The easiest way to vary Re_0 would be to vary the ring diameter by using ring tubes with a range of different diameters. After calibrating each ring tube so that the controllable piston parameters can be related to the ring propagation velocity, Ri_0 could be maintained after varying Re_0 by adjusting the buoyancy difference between the two layers accordingly. While such dye experiments may seem primitive compared to current technological capabilities, they have the advantage of providing three-dimensional visualisations of the flow with relative ease, making it possible to investigate a large number of cases without requiring extensive computational resources to process and analyse data. By conducting separate experiments either dyeing one fluid layer or dyeing the vortex ring, such experiments would help characterise the Re_0 dependence on the instability mechanisms that develop. We would expect over the suggested range of Re_0 that the ring–interface interaction ranges from being almost entirely laminar in nature, to exhibiting the Crow-like instability identified by Olsthoorn and Dalziel (2017).

Additionally, dye experiments such as those described above could be used to investigate the influence of θ_0 on the evolution of the ring–interface interaction for given (Re_0, Ri_0) . By conducting experiments for a similar range of Re_0 and Ri_0 suggested above, it would be possible to assess (to a qualitative degree) the sensitivity of the instability mechanisms that develop on θ_0 for a range of Re_0 and Ri_0 . Such an investigation would provide some clarity over how much variation should be expected in quantities such as the entrainment rate and the mixing efficiency to apparently small changes in the control parameters of the ring–interface interaction. As a start, we suggest first considering small angles of θ_0 such as those investigated in this thesis, representing a small deviation in θ_0 from the vertical case.

For high-resolution velocity and density measurements of the ring–interface interaction, the PIV approach developed by Olsthoorn and Dalziel (2017) could be used to obtain three-dimensional velocity fields for a wider range of Ri_0 and Re_0 at $\theta_0 = 0^\circ$, provided the phase-locking of instabilities between consecutive experiments is a feature that can be repeated. For $\theta_0 > 0^\circ$, the same PIV approach could be attempted, though it would rely

entirely on the instabilities in the interaction being highly repeatable between consecutive experiments. Otherwise, it may be worth waiting until in-plane scanning technology, such as that developed by Partridge et al. (2019), is capable of scanning the flow at a fast enough rate to resolve the three-dimensional evolution of the flow before revisiting the ring–interface interaction. It is clear from our two-dimensional observations presented in Chapter 5 that complex, three-dimensional dynamics are at play in the evolution of the flow that we have not been able to fully observe. For now, making use of other experimental techniques such as shadowgraphy is likely to provide further useful insights into how the flow evolves, and so are encouraged.

In this thesis we have chosen to restrict our attention to cases where $\theta_0 \leq 25^\circ$. Experiments we have chosen not to include in this thesis looked at higher propagation angles, up to $\theta_0 = 75^\circ$. Intermediate values of θ_0 ($30^\circ \lesssim \theta_0 \lesssim 60^\circ$) tended to produce overturning mixing events, with lower-layer fluid entrained into the centre of the vortex ring on the side of C_1 (defined in Section 5.1) leading to much of the mixing that was observed. For $\theta_0 = 75^\circ$ and $Ri_0 \gtrsim 1$, the vortex ring sweeps interfacial fluid along the interface, causing it to accumulate closely behind the ring as it continues to propagate, leading to the formation of an interfacial wave. This interfacial wave accelerates the ring on the side of C_1 , causing the trajectory of the vortex ring to deviate away from the lower-layer region. Clearly much more work is needed to understand the dynamics of the ring–interface interaction at the higher values of θ_0 not considered in this thesis, and the implications these dynamics have on the mixing produced.

With continuous advancements being made in computational capabilities, high resolution numerical simulations could be carried out to investigate the evolution of the ring–interface interaction. Previous simulations for the $\theta_0 = 0^\circ$ case have been carried out using the vortex-sheet method introduced by Stock et al. (2008), which could be adapted to investigate cases with $\theta_0 > 0^\circ$. Alternatively, DNS could be used to model the ring–interface interaction, with the flow initialised by prescribing an initial vorticity distribution for the vortex ring.

7.3.3 Conductivity probe design

Our analysis of the conductivity probe we made use of, which is presented in Appendix C, reveals a design flaw associated with the probe that appears to be capable of introducing important systematic measurement errors when the probe is traversed through and samples a region with a high electrical conductivity gradient, such as a sharply stratified density interface. The design flaw concerns the gap inside the probe head, between the exit of the probe-tip region and the inner electrode. If the fluid in this gap is not flushed through effectively by the siphoning of the probe, it is possible for the conductivity of the gap fluid to

be substantially lower than the conductivity of the fluid passing through the probe tip. This can lead to the probe-measured conductivity being significantly lower than the probe-tip-fluid conductivity.

If a new design of conductivity probe could ensure a direct electrical connection is made between the probe-tip fluid and the inner electrode, the design flaw described above would be overcome. The testing of a redesigned probe could be carried out in the Rayleigh-Taylor tank (designed by Dalziel (1993) and housed in DAMTP), which has the advantage of easily being able to produce (almost) perfectly sharp density interfaces. Then, the ability of the redesigned probe in resolving density interfaces of varying interfacial thicknesses (produced by leaving the initial interface to diffuse) could be assessed in a precise way. If a redesign of the probe is able to improve the probe performance sufficiently, the newly designed probe may be capable of directly measuring the mixing efficiency η of the ring–interface interaction to a high level of accuracy.

7.3.4 Eddy-induced mixing

The observed sharpening of the density interface at the start of the IAP in our ring-mixing experiments is suggestive that the mixing induced during a ring–interface interaction depends significantly on the (non-dimensionalised) interfacial thickness $\delta z/a$. Indeed, if $\delta z/a = 0$, it seems fully reasonable to expect that there would be a critical Richardson number Ri_c such that the mixing efficiency $\eta = 0$ for $Ri_0 > Ri_c$. However, real density interfaces diffuse, and so even for $Ri_0 \gg 1$, the ring will be able to entrain fluid from the diffuse interface and generate some mixing. It is important to clarify how $\delta z/a$ will influence the ring-induced mixing, and so we suggest that experiments should be carried out to assess this dependency. Note however that obtaining a high level of accuracy in such experiments may require the design of a new generation of conductivity probe, such as that described in the previous subsection.

To obtain accurate measurements of the entrainment volume $A\delta H_2$ of the ring–interface interaction as a function of the various control parameters, our periodic-mixing experiments could be adapted to have a period τ_G of weak grid mixing after every N ring generations, to sweep the ring-mixed fluid at the bottom of the upper layer and mix it perfectly into the upper-layer. Using the periods τ_R and τ_S defined in Section 6.1.1, the time t_n of the n th vortex-ring generation would then be given by

$$t_n = n\tau_R + \left\lfloor \frac{n}{N} \right\rfloor (\tau_S + \tau_G - \tau_R). \quad (7.1)$$

Such an experiment would require a calibration of the mixing grid to determine suitable values for the grid-stirring frequency n and the period τ_G required to homogenise the upper layer, as well as the fraction of the total entrainment across the interface attributable to the mixing grid. Additionally, changes in Ri_0 across a given experiment associated with increases in the upper-layer density should be taken into account when relating measurements of the entrainment volume to Ri_0 . It is worth clarifying that by introducing grid mixing, one loses the ability to measure the mixing efficiency of the ring–interface interaction, as this requires measuring the density distribution of the ring-mixed fluid before it is mixed further.

7.4 Final words

This thesis represents a modest but significant step forward in understanding the dynamics of vortex rings obliquely impacting a density interface, and the associated mixing produced. We have demonstrated a sensitive dependence of the interaction dynamics on θ_0 , supporting the notion that the azimuthal symmetry of the $\theta_0 = 0^\circ$ case produces an eddy interaction with unique instability mechanisms that may lead to their mixing properties being importantly different in the broader family of eddy–interface interactions. We have also demonstrated that the mixing efficiency η of the ring–interface interaction in isolation is much lower than the reported value of $\eta_c = 0.42$, with our measurements indicating $0.15 < \eta < 0.33$ over the range of parameters investigated. Consequently, it appears that the mixing properties of vortex rings at density interfaces are sensitive to their control parameters over the range of Re_0 considered, and so caution should be taken when relating the properties of eddy-induced mixing to stratified turbulent flows.

References

- S Advait, K. V. Manu, A. Tinaikar, U. K. Chetia, and S. Basu. Interaction of vortex ring with a stratified finite thickness interface. *Physics of Fluids*, **29**(9):093602, 2017.
- D. G. Akhmetov. *Vortex Rings*. Springer-Verlag Berlin Heidelberg, 2009.
- P. J. Archer, T. G. Thomas, and G. N. Coleman. The instability of a vortex ring impinging on a free surface. *Journal of Fluid Mechanics*, **642**:79–94, 2010.
- L. P. Bernal and J. T. Kwon. Vortex ring dynamics at a free surface. *Physics of Fluids A: Fluid Dynamics*, **1**:449–451, 1989.
- N. Bethke and S. B. Dalziel. Resuspension onset and crater erosion by a vortex ring interacting with a particle layer. *Physics of Fluids*, **24**:063301, 2012.
- R. J. Bristol, P. S. Ortega, P. S. Marcus, and Ö. Savaş. On cooperative instabilities of parallel vortex pairs. *Journal of Fluid Mechanics*, **517**:331–358, 2004.
- C. P. Caulfield. Layering, instabilities, and mixing in turbulent stratified flows. *Annual review of Fluid Mechanics*, **53**:113–145, 2021.
- M. Cheng, J. Lou, and L.-S. Luo. Numerical study of a vortex ring impacting a flat wall. *Journal of Fluid Mechanics*, **660**:430–455, 2010.
- E. C. W. Clarke and D. N. Glew. Evaluation of the thermodynamic functions for aqueous sodium chloride from equilibrium and calorimetric measurements below 154°C. *Journal of Physics and Chemistry Reference Data*, **14**:489–610, 1985.
- L. D. Couch and P. S. Kreuger. Experimental investigation of vortex rings impinging on inclined surfaces. *Experiments in Fluids*, **51**:1123–1138, 2011.
- S. C. Crow. Stability theory for a pair of trailing vortices. *AIAA Journal*, **8**:2172–79, 1970.
- W. J. A. Dahm, C. M. Scheil, and G. Tryggvason. Dynamics of vortex interaction with a density interface. *Journal of Fluid Mechanics*, **205**:1–43, 1989.
- S. B. Dalziel. Rayleigh-taylor instability: experiments with image analysis. *Dynamics of Atmospheres and Oceans*, **20**:127–153, 1993.
- S.B. Dalziel. Digiflow user guide. *Dalziel Research Partners*, 2017.
- P. A. Davidson. *Turbulence in rotating, stratified and electrically conducting fluids*. Cambridge University Press, 2013.

- M. S. Davies Wykes and S. B. Dalziel. Efficient mixing in stratified flows: experimental study of a rayleigh-taylor unstable interface with an otherwise stable stratification. *Journal of Fluid Mechanics*, **756**:1027 – 1057, 2014.
- M. S. Davies Wykes, G. O. Hughes, and S. B. Dalziel. On the meaning of mixing efficiency for buoyancy-driven mixing in stably stratified flows. *Journal of Fluid Mechanics*, **781**: 261 – 275, 2015.
- N. Didden. On the formation of vortex rings: Rolling-up and production of circulation. *J. Appl. Math. and Phys.*, **30**:101–116, 1979.
- S. Ersoy and J. D. A. Walker. The boundary layer due to a three-dimensional vortex loop. *Journal of Fluid Mechanics*, **185**:569–598, 1987.
- H. J. S. Fernando and R. R. Long. On the nature of the entrainment interface of a two-layered fluid subjected to zero-mean-shear turbulence. *Journal of Fluid Mechanics*, **151**:21–53, 1985.
- R. Ferrari and C. Wunsch. Ocean circulation kinetic energy: Reservoirs, sources and sinks. *Annual review of Fluid Mechanics*, **41**:253 – 282, 2009.
- L. E. Fraenkel. Examples of steady vortex rings of small cross-section in an ideal fluid. *Journal of Fluid Mechanics*, **51**:119–135, 1972.
- P. O. Gendron, F. Avaltroni, and K. J. Wilkinson. Diffusion coefficients of several rhodamine derivatives as determined by pulsed field gradient–nuclear magnetic resonance and fluorescence correlation spectroscopy. *Journal of fluorescence*, 18(6):1093–1101, 2008.
- W. R. Geyer, M. E. Scully, and D. K. Ralston. Quantifying vertical mixing in estuaries. *Environmental Fluid Mechanics*, 8.5:495–509, 2008.
- M. Gharib and A. Weigand. Experimental studies of vortex disconnection and connection at a free surface. *Journal of Fluid Mechanics*, **321**:59–86, 1996.
- M. Gharib, E. Rambod, and K. Shariff. A universal time scale for vortex ring formation. *Journal of Fluid Mechanics*, **360**:121–140, 1998.
- S. I. Green. *Fluid Vortices*. Kluwer Academic Publishers, 1995.
- M. C. Gregg. Mixing efficiency in the ocean. *Journal of Fluid Mechanics*, **10**:443–473, 2018.
- R. Hain and C. J. Kähler. Fundamentals of multiframe particle image velocimetry (piv). *Experiments in Fluids*, **42**:575–587, 2007.
- H. von Helmholtz. Über integrale der hydrodynamischen gleichungen, welche der wirbelbewegung entsprechen. *J. für die reine und angewandte Mathematik*, **55**:25–55, 1858.
- M. J. M. Hill. On a spherical vortex. *Phil. Trans. R. Soc. Lond.*, pages 213–245, 1894.
- J. M. Holford and P. F. Linden. Turbulent mixing in a stratified fluid. *Dynamics of Atmospheres and Oceans*, **30**:173–198, 1999.

- J. C. R. Hunt and J. M. R. Graham. Free stream turbulence near plane boundaries. *Journal of Fluid Mechanics*, **84**:209–235, 1978.
- J. C. R. Hunt, A. A. Wray, and P. Moin. Eddies, streams and convergence zones in turbulent flows. *Centre for turbulence research report CTR-S88*, pages 193–208, 1988.
- G. N. Ivey, K. B. Winter, and J. R. Koseff. Density stratification, turbulence, but how much mixing? *Annual Review of Fluid Mechanics*, **40**:169–184, 2008.
- R. Kerswell. Elliptical instability. *Annual Review of Fluid Mechanics*, **34**:83–113, 2002.
- S. Kida and M. Takaoka. Vortex reconnection. *Annual Review of Fluid Mechanics*, **26**:169–189, 1994.
- C.-H. Kruttsch. Über eine experimentell beobachtete erscheinung an wirbelringen bei ihrer translatorischen bewegung in wirklichen flüssigkeiten. *Annalen der Physik*, **427**:497–523, 1939.
- K. Kuehn, M. Moeller, M. Schulz, and D. Safelippo. Vortex ring refraction at large froude numbers. *Physical Review E - Statistical, Nonlinear and Soft Matter Physics*, **82**:016312, 2010.
- H. Lamb. *Hydrodynamics*. Cambridge University Press, 1932.
- A.S.P. Lefauve. *Waves and turbulence in sustained stratified shear flows*. PhD Thesis, University of Cambridge, 2018.
- T. Leweke, S. Le Dizès, and C.H.K. Williamson. Dynamics and instabilities of vortex pairs. *Annual Review of Fluid Mechanics*, **48**:507–541, 2016.
- C. Liess and N. Didden. Experimente zum einfluß der anfangsbedingungen auf die instabilität von ringwirbeln. *ZAMM - Journal of Applied Mathematics and Mechanics*, **56**:206–208, 1976.
- T. T. Lim. An experimental study of a vortex ring interacting with an inclined wall. *Experiments in Fluids*, **7**:453–463, 1989.
- T. T. Lim. A note on the cause of rebound in the head on collision of a vortex ring with a wall. *Experiments in Fluids*, **12**:41–48, 1991.
- P. F. Linden. The interaction of a vortex ring with a sharp density interface: a model for turbulent entrainment. *Journal of Fluid Mechanics*, **60**:467–480, 1973.
- P. F. Linden and J. S. Turner. The formation of ‘optimal’ vortex rings, and the efficiency of propulsion devices. *Journal of Fluid Mechanics*, **427**:61–72, 2001.
- C. H. Liu. Vortex simulation of unsteady shear flow induced by a vortex ring. *Computers & Fluids*, **31**:183–207, 2002.
- H. J. Lugt and S. Ohring. The oblique rise of a viscous vortex ring toward a deformable free surface. *Meccanica*, **29**:313–329, 1994.

- D. L. Marcus and J. B. Bell. Numerical simulation of a viscous vortex ring interaction with a density interface. *Physics of Fluids*, **6**:1505–1514, 1992.
- T. Maxworthy. The structure and stability of vortex rings. *Journal of Fluid Mechanics*, **51**: 15–32, 1972.
- T. Maxworthy. Some experimental studies of vortex rings. *Journal of Fluid Mechanics*, **81**: 465–495, 1977.
- T.C. McDougall. On the elimination of refractive-index variations in turbulent density-stratified liquid flows. *Journal of Fluid Mechanics*, **93**:83–96, 1979.
- H.K. Moffatt. The degree of knottedness of tangled vortex lines. *Journal of Fluid Mechanics*, **35**:117–129, 1969.
- K. Mohseni and M. Gharib. A model for universal time scale of vortex ring formation. *Physics of Fluids*, **10**:2436–2438, 1998.
- D. W. Moore and P. G. Saffman. The motion of a vortex filament with axial flow. *Phil. Trans. Roy. Soc. Lon.*, **272**:403–429, 1972.
- D. W. Moore and P. G. Saffman. Axial flow in laminar trailing vortices. *Proc. R. Soc. Lond. A.*, **333**:491–508, 1973.
- D. W. Moore and P. G. Saffman. The instability of a straight vortex filament in a strain field. *Proc. R. Soc. Lond. A.*, **346**:413–425, 1975.
- T. H. New, S. Shi, and B. Zang. Some observations on vortex-ring collisions upon inclined surfaces. *Experiments in Fluids*, **57**, 2016.
- J. Norbury. A family of steady vortex rings. *Journal of Fluid Mechanics*, **57**:417–431, 1973.
- R. L. F. Oglethorpe, C. P. Caulfield, and A. W. Woods. Spontaneous layering in stratified turbulent taylor-couette flow. *Journal of Fluid Mechanics*, **721**:R3, 2013.
- S. Ohring and H. J. Lugt. Interaction of an obliquely rising vortex ring with a free surface in a viscous fluid. *Meccanica*, **31**:623–655, 1996.
- J. Olsthoorn and S. B. Dalziel. Vortex-ring-induced stratified mixing. *Journal of Fluid Mechanics*, **781**:113–126, 2015.
- J. Olsthoorn and S. B. Dalziel. Three-dimensional visualisation of the interaction of a vortex ring with a stratified interface. *Journal of Fluid Mechanics*, **820**:549–579, 2017.
- J. Olsthoorn and S. B. Dalziel. Vortex-ring-induced stratified mixing: mixing model. *Journal of Fluid Mechanics*, **837**:129–146, 2018.
- P. Orlandi, P. Egermann, and E. J. Hopfinger. Vortex rings descending in a stratified fluid. *Physics of Fluids*, **10**:2819–2827, 1998.
- Y.-G. Park, J. A. Whitehead, and A. Gnanadeskian. Turbulent mixing in stratified fluids : layer formation and energetics. *Journal of Fluid Mechanics*, **279**:279–311, 1994.

- J. L. Partridge, A. Lefauve, and S. B. Dalziel. A versatile scanning method for volumetric measurements of velocity and density fields. *Measurement Science and Technology*, **30**(5), 2019.
- W. R. Peltier and C. P. Caulfield. Mixing efficiency in stratified shear flows. *Annual review of Fluid Mechanics*, **35**:135 – 167, 2003.
- F. Pereira, A. Ciarravano, G. P. Romano, and F. Di Felice. Adaptive multi-frame piv. *In Proceedings of the 12th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal*, 2004.
- T. Persoons and T. S. O’Donovan. High dynamics velocity range particle image velocimetry using multiple pulse separation imaging. *Sensors*, **11**:1–18, 2011.
- S. B. Pope. *Turbulent Flows*. Cambridge University Press, 2000.
- A.K. Prasad, R.J. Adrian, C.C. Landreth, and P.W. Offutt. Effect of resolution on the speed and accuracy of particle image velocimetry interrogation. *Experiments in Fluids*, **13**: 105–116, 1992.
- D. I. Pullin. Vortex ring formation at tube and orifice openings. *The Physics of Fluids*, **22**: 401–403, 1979.
- Y. Quayan and C. K. Chu. The nonlinear interaction of vortex rings with a free surface. *Acta Mechanica Sinica*, **13**:120–129, 1997.
- George William Rayfield and F Reif. Quantized vortex rings in superfluid helium. *Physical Review*, 136(5A):A1194, 1964.
- A. Revuelta, A. Sánchez, and A. Liñán. Confined axisymmetric laminar jets with large expansion ratios. *Journal of Fluid Mechanics*, **456**:319–352, 2002.
- M. Rosenfeld, E. Rambod, and M. Gharib. Circulation and formation number of laminar vortex rings. *Journal of Fluid Mechanics*, **376**:297–318, 1998.
- P. G. Saffman. The number of waves on unstable vortex rings. *Journal of Fluid Mechanics*, **84**:625 – 639, 1978.
- P.G. Saffman. The velocity of viscous vortex rings. *Studies. Appl. Maths*, **49**:371–380, 1970.
- P.G. Saffman. *Vortex Dynamics*. Cambridge Monographs on Mechanics and Applied Mathematics, 1995.
- M. M. Scase and S. B. Dalziel. An experimental study of the bulk properties of vortex rings translating through a stratified fluid. *European Journal of Mechanics - B/Fluids*, **25** (3): 302 – 320, 2006.
- J. W. Shan, D. B. Lang, and P. E. Dimotakis. Scalar concentration measurements in liquid-phase flows with pulsed lasers. *Experiments in Fluids*, **36**:268–273, 2004.
- K. Shariff and A. Leonard. Vortex rings. *Annual Review of Fluid Mechanics*, **24**:235–279, 1992.

- A. Shrivat, C. Cenedese, and C. P. Caulfield. Entrainment and mixing dynamics of surface-stress-driven stratified flow in a cylinder. *Journal of Fluid Mechanics*, **691**:498–517, 2012.
- M. Song, L. P. Bernal, and G. Tryggvason. Head-on collision of a large vortex ring with a free surface. *Physics of Fluids A: Fluid Dynamics*, **4**:1457–1466, 1992.
- M. J. Stock. *A Regularised Inviscid Vortex Sheet Method for Three Dimensional Flows with Density Interfaces*. PhD Thesis, University of Michigan, 2006.
- M. J. Stock, W. J. A. Dahm, and G. Tryggvason. Impact of a vortex ring on a density interface using a regularized inviscid vortex sheet method. *Journal of Computational Physics*, **227**(21):9021–9043, 2008.
- G. G. Stokes. *On the effect of the internal friction of fluids on the motion of pendulums*, volume **9**. Pitt Press Cambridge, 1851.
- J. D. Swearingen, J. D. Crouch, and R. A. Handler. Dynamics and stability of a vortex ring impacting a solid boundary. *Journal of Fluid Mechanics*, **297**:1–28, 1995.
- R. Tailleux. On the energetics of stratified turbulent mixing, irreversible thermodynamics, boussinesq models and the ocean heat engine controversy. *Journal of Fluid Mechanics*, **638**:339–382, 2009.
- W. Thomson. Vi - on vortex motions. *Trans. Roy. Soc. Edin.*, **225**:217–260, 1868.
- W. Thomson. Xxiv. vibrations of a columnar vortex. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, **10.61**:155–168, 1880.
- Johannes Töger, Mikael Kanski, Marcus Carlsson, Sándor J Kovács, Gustaf Söderlind, Håkan Arheden, and Einar Heiberg. Vortex ring formation in the left ventricle of the heart: analysis by 4d flow mri and lagrangian coherent structures. *Annals of biomedical engineering*, **40**(12):2652–2662, 2012.
- C.-Y. Tsai and S. E. Widnall. The stability of short waves on a straight vortex filament in a weak externally imposed strain field. *Journal of Fluid Mechanics*, **73**:721–733, 1976.
- J. S. Turner. The influence of molecular diffusivity on turbulent entrainment across a density interface. *Journal of Fluid Mechanics*, **33**:639–656, 1968.
- P. A. Tyvand and T. Miloh. Axisymmetric interaction between a vortex ring and a free surface. *Physics of Fluids*, **6**:224–238, 1994.
- C. W. van Atta. Vortex ring instability and collapse in a stably stratified fluid. *Experiments in Fluids*, **7**:197–200, 1989.
- R. Verzicco and P. Orlandi. Normal and oblique collisions of a vortex ring with a wall. *Meccanica*, **29**:383–391, 1994.
- J. D. A. Walker, C. R. Smith, A. W. Cerra, and T. L. Doligalski. The impact of a vortex ring on a wall. *Journal of Fluid Mechanics*, **181**:99–140, 1987.

- R.C. Weast, M.J. Astle, and W. H. Beyer. *CRC Handbook of Chemistry and Physics, 65th edition*. CRC Press, 1984.
- J. Westerweel. Fundamentals of digital particle image velocimetry. *Measurement Science and Technology*, **8**:1379–1392, 1997.
- S. E. Widnall and J. P. Sullivan. On the stability of vortex rings. *Proc. R. Soc. Lon.*, **332**:335–353, 1973.
- S. E. Widnall and C.-Y. Tsai. The instability of the thin vortex ring of constant vorticity. *Phil. Trans. R. Soc. Lond.*, **287**:273–305, 1977.
- S. E. Widnall, D. B. Bliss, and C.-Y. Tsai. The instability of short waves on a vortex ring. *Journal of Fluid Mechanics*, **66**:35–47, 1974.
- B. Wilson and B. Smith. Uncertainty on piv mean and fluctuating velocity due to bias and random errors. *Measurement Science and Technology*, **24**, 2013.
- K. B. Winters, P. N. Lombard, J. J. Riley, and E. A. D’Asaro. Available potential energy and mixing in density-stratified fluids. *Journal of Fluid Mechanics*, **289**:115–128, 1995.
- A. W. Woods, C. P. Caulfield, J. R. Landel, and A. Kuesters. Non-invasive turbulent mixing across a density interface in a turbulent taylor–couette flow. *Journal of fluid mechanics*, **663**:347–357, 2010.
- M. G. Worster and H. E. Huppert. Time-dependent density profiles in a filling box. *Journal of Fluid Mechanics*, **132**:457–466, 1983.
- C. Wunsch and R. Ferrari. Vertical mixing, energy, and the general circulation of the oceans. *Annual Review of Fluid Mechanics*, **36**:281–314, 2004.
- H. Xu and E. Bodenschatz. Motion of inertial particles with size larger than kolmogorov scale in turbulent flows. *Physica D, Nonlinear Phenomena*, **237**:2095–2100, 2008.
- KWB Yeo, JY Koh, J Long, and TH New. Flow transitions in collisions between vortex-rings and density interfaces. *Journal of Visualization*, pages 1–9, 2020.
- C. Zhang, L. Shen, and D. K. P. Yue. The mechanism of vortex connection at a free surface. *Journal of Fluid Mechanics*, **384**:207–241, 1999.

Appendix A

Kelvin Wave properties for different azimuthal wavenumbers

To restrict our attention to the modes that are of most interest to us, it is instructive to consider the properties of waves with different azimuthal wavenumbers. This can be done effectively by considering the deformation of the vortex column. From (2.36), the column surface r_s is given by

$$r_s = a - \frac{u_r}{\alpha} \cos(kz) \cos(\alpha t - m\phi). \quad (\text{A.1})$$

Consider a slice of the column along a plane $z = z_0$. From (A.1) we find that the rate of change of the slice area is given by:

$$\int_0^{2\pi} \dot{r} d\phi = \begin{cases} 2\pi u_r \cos(kz_0) \sin(\alpha t) & m = 0 \\ 0 & m \neq 0 \end{cases}. \quad (\text{A.2})$$

Equation (A.2) highlights $m = 0$ as an exceptional case, corresponding to “varicose” modes. For such modes, the disturbance to the column is axisymmetric, with the column broadening and narrowing periodically in the z -direction. Additional insight can be made by considering the centroid velocity of our z_0 -column slice. The x -component is given by

$$\frac{d}{dt} \int_0^{2\pi} x_s d\phi = \frac{d}{dt} \int_0^{2\pi} r_s \cos \phi d\phi = \begin{cases} \frac{\pi u_r}{2} \cos(kz_0) \sin(\alpha t) & |m| = 1 \\ 0 & |m| \neq 1 \end{cases}. \quad (\text{A.3})$$

Equation (A.3) demonstrates that for $|m| = 1$, the core centreline deviates from the z -axis in an oscillatory fashion. The family of waves associated with $|m| = 1$ are “bending” modes, their name being lent to the nature of deformation of the vortex column. Modes with $|m| \geq 2$

are called “fluted” modes. For such modes, the vortex boundary deforms into a petal-like shape, with m peaks.

Appendix B

Tracer Particle Calculations

B.1 Estimation of the Kolmogorov Lengthscale

The turbulent energy cascade describes the process of larger eddies becoming unstable and breaking up, transferring their energy to successively smaller eddies. Eventually the motion for small eddies becomes stable as molecular viscosity becomes effective in dissipating the kinetic energy (Pope (2000)). The size of the largest eddies that are stable due to viscous dissipation is governed by the kinematic viscosity ν and the dissipation rate ε . Using these parameters, a unique lengthscale can be created to characterise the size of these eddies. This is the Kolmogorov lengthscale, η , which is given by

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

Making the assumption that the length and velocity scales associated with ε are the vortex tube diameter $2R$ and the vortex ring propagation velocity U , we obtain the scaling

$$\varepsilon = \nu \mathbf{u} \cdot \nabla^2 \mathbf{u} \sim \nu \frac{U^2}{4R^2}.$$

The kinematic viscosity $\nu = 10^{-6} \text{ cm}^2/\text{s}$, the vortex tube diameter $2R = 4 \text{ cm}$ and from the bike pump calibrations in Section 3.2, the range of propagation velocities in our experiments is $5 \text{ cm s}^{-1} \leq U \leq 10 \text{ cm s}^{-1}$. Taking the maximal value of U , we obtain an estimate for the dissipation rate in our experiments:

$$\varepsilon_{\max} \sim 10^{-6} \times \frac{10^{-2}}{(4 \times 10^{-2})^2} = \frac{25}{4} \times 10^{-6} \text{ m}^2 \text{ s}^{-3}. \quad (\text{B.1})$$

Using B.1, we estimate the minimum value of the Kolmogorov lengthscale in our experiments to be

$$\eta_{\min} \sim \left(\frac{4 \times 10^{-12}}{25} \right)^{\frac{1}{4}} \approx 6.3 \times 10^{-4} \text{ m}. \quad (\text{B.2})$$

B.2 Order of magnitude calculation of the Stokes Number

From Section 3.3.4, as the tracer particles we use have diameter $d_p = 5 \times 10^{-5} \text{ m} < \eta_{\min}$, the expression given in (3.6) for the Stokes number St_K is appropriate for characterising the ability of the particles to closely follow fluid streamlines.

As ρ_p is close to the density of freshwater, we take $\rho_p/\rho_f \approx 1$. Substituting this approximation, the particle diameter d_p and the expression for η_{\min} given in equation (B.2), we obtain

$$St_K = \frac{1}{18} \left(\frac{\rho_p}{\rho_f} \right) \left(\frac{d_p}{\eta} \right)^2 \approx \frac{1}{18} \times 1 \times \frac{5 \times 10^{-5}}{6.3 \times 10^{-4}} \sim O(10^{-4}) \ll 1. \quad (\text{B.3})$$

As $St_K \ll 1$, we confirm that the particles we use act as passive tracers in the flow, validating their use for PIV.

B.3 Calculation of the settling velocity

From equation (3.7), the settling velocity \mathbf{U}_g is given by

$$\mathbf{U}_g = d_p^2 \frac{(\rho_p - \rho_f)}{18\mu} \mathbf{g}.$$

As discussed in Section 3.3.4, the density distribution for the particles we use was found to be bimodal, with strong peaks at $\rho_{p-} = 1.023 \text{ g/cm}^3$ and $\rho_{p+} = 1.025 \text{ g/cm}^3$. As the density distribution is narrow, for convenience in calculating \mathbf{U}_g we take $\rho_p = 1.024 \text{ g/cm}^3$. Taking the units of ρ_p and ρ_f to be g/cm^3 , substituting ρ_p and d_p into (B.3), we obtain

$$\begin{aligned} U_g &= \left(5 \times 10^{-5} \right)^2 \times \frac{(1.024 - \rho_f) \times 10^3}{18 \times 10^{-3}} \times -9.8 \text{ m/s} \\ &= \frac{25 \times 9.8}{18} \times (\rho_f - 1.024) \times 10^{-4} \times 10^3 \times 3.6 \times 10^3 \text{ cm/h} \\ \Rightarrow U_g &= 4.9 (\rho_f - 1.024) \times 10^3 \text{ cm/h}. \end{aligned}$$

Appendix C

Analysis of the conductivity probe

As discussed in Section 3.4, using the conductivity probe, electrical conductivity measurements are made by sampling the conductivity of the circuit (hereafter referred to as the ‘probe circuit’) that is established between the two electrodes that form the probe when fluid is drawn through the probe tip. Additional steps were incorporated into our experimental procedure for our periodic mixing experiments presented in Chapter 6, so that the response from the probe more closely reflects its desired performance, which is to sample instantaneously the electrical conductivity of the fluid in the stratification at the height of the entrance to the probe-tip opening. To clarify the relationship between the vertical profiles of the probe-measured electrical conductivity and the electrical conductivity profile of the fluid in the stratification, we give a detailed analysis of the conductivity probe in this appendix.

C.1 Probe circuit

First, we clarify the relationship between the probe measurements and the electrical resistivity of the fluid that closes the probe circuit, after which we discuss the relationship between the circuit-closing fluid and the fluid stratification in the tank. Figures C.1 (a) and (b) illustrates the probe tip, secured to the rest of the probe. There are three distinct regions of fluid that require separate consideration as the electrical current flows from the outer electrode to the inner electrode. Firstly, the volume of fluid carrying current from the outer electrode, up to the probe-tip entrance. Secondly, the volume of fluid occupying the probe-tip opening, and thirdly, the volume of fluid carrying current from the exit of the probe-tip opening to the inner electrode. These regions are illustrated in Figure C.1 (a), and we denote parameters associated with each of these regions with subscripts O (outer), T (tip) and I (inner) respectively.

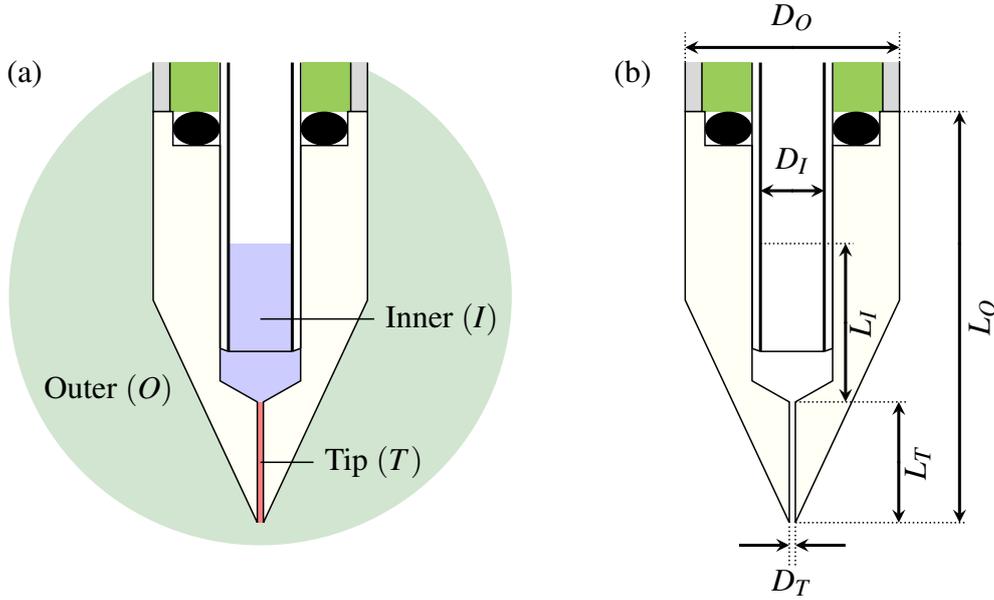


Fig. C.1 Illustrations of the conductivity probe tip, secured to the inner and outer electrodes that form the conductivity probe. (a) The shaded regions illustrates the different regions that form the circuit-closing fluid, where the conductivity of the fluid plays a role in setting the electrical conductivity of the circuit. (b) Illustration of the lengthscales L_k , D_k that set the electrical resistivity R_k for each of the circuit-closing regions of fluid. For a full schematic of the conductivity probe, see Figure 3.13.

Each region has associated lengthscales L_k , D_k ($k = O, T, I$) corresponding to the distance the electrical current travels and the cross-sectional area through which it travels respectively. These are illustrated in Figure C.1 (b). Due to the much greater electrical conductivity of the electrode materials relative to the NaCl solutions we use in our experiments, the current will preferentially travel up the probe through the inner-electrode material rather than the probe-internal fluid. It follows therefore that $L_I \sim D_I$, as D_I is the confining lengthscales in the probe interior. For the purposes of this discussion we take $L_I = 2D_I$, which is appropriate for the specifications of the probe we use.

For a homogeneous fluid with electrical conductivity γ , the electrical resistivity R_k of each region is given approximately by

$$R_k(\gamma) \sim \frac{L_k}{\gamma D_k^2}. \quad (\text{C.1})$$

As the circuit-closing fluid is in series within the rest of the circuit, the electrical conductivity C of the circuit is given by

$$C = \frac{1}{R_c + R_O + R_T + R_I}, \quad (\text{C.2})$$

where R_c is the total resistance associated with the circuit excluding the circuit-closing fluid. As discussed in Section 3.4, the value for R_c can be set using a choice of resistors contained within the bridge circuit, and is chosen based on the range of salinities used for a given set of experiments to optimise the signal to noise ratio of the probe measurements. Variation in R_c over a day to day timescale is responsible for the small drift between measured and calibration-calculated values of the fluid density, the effect of which we mitigate using the methodology outlined in Section 6.1.3.

From Equation (C.2), variation in C will occur primarily due to variation in resistivity of the circuit-closing fluid. From Equation (C.1), for the specifications of the conductivity probe, we have $R_O/R_T \sim 3.4 \times 10^{-3}$ and $R_I/R_T \sim 2.0 \times 10^{-2}$. Therefore, provided the conductivity of the fluid in the tip region is not significantly higher than the conductivity of the fluid in either the inner or outer region, variation in the conductivity of the circuit will be dominated by the resistivity of the fluid passing through the tip region.

Figure C.2 plots the electrical conductivity γ of NaCl solution for the range of salinities used in our experiments. To reduce the ratio γ_2/γ_1 between the conductivities of the lower and upper-layer solutions across all our experiments, we avoid using fresh water to form our upper layer, instead opting for a weak salinity solution with density $\rho_1 = 1.002 \pm 0.0003$, corresponding to $\gamma_1 = 8.7 \pm 0.6$. In doing so, we reduced the maximum for γ_2/γ_1 across our experiments from 100 to 12, thus reducing the contributions of R_O and R_I relative to R_T by almost a factor of 10.

As $\gamma_2/\gamma_1 \leq 12$ for all our periodic mixing experiments and $R_O/R_T \sim 3.4 \times 10^{-3}$ for a homogeneous fluid, it follows that $(R_O/R_T)_{\max} \sim 4.2 \times 10^{-2}$ provides an upper bound for R_O/R_T realisable in our experiments. As the sharpest density interfaces exhibited in our experiments are approximately 1cm thick (which is equal to the length of the probe tip), we expect the values realised for R_O/R_T throughout our experiments are sufficiently small that we can neglect the effect of fluid in the outer-probe region impacting our conductivity measurements.

Though it would have been possible to reduce the maximum γ_2/γ_1 further by prescribing larger values for ρ_1 , ρ_2 , there was little room for improvement before the probe voltage C_V would begin to clip as it attains its maximum measurable value (we discuss this limiting factor in Section 3.4). Choosing a stronger resistor in the bridge circuit would have resolved the clipping issue, but at the expense of the signal to noise ratio of the probe voltage, which would have reduced the precision of our measurements for $\delta PE_R^{(n)}$ by at least a factor of two given the choice of resistors available. We deemed this loss in precision as unacceptable, and therefore opted to prescribe the values for ρ_1 given above.

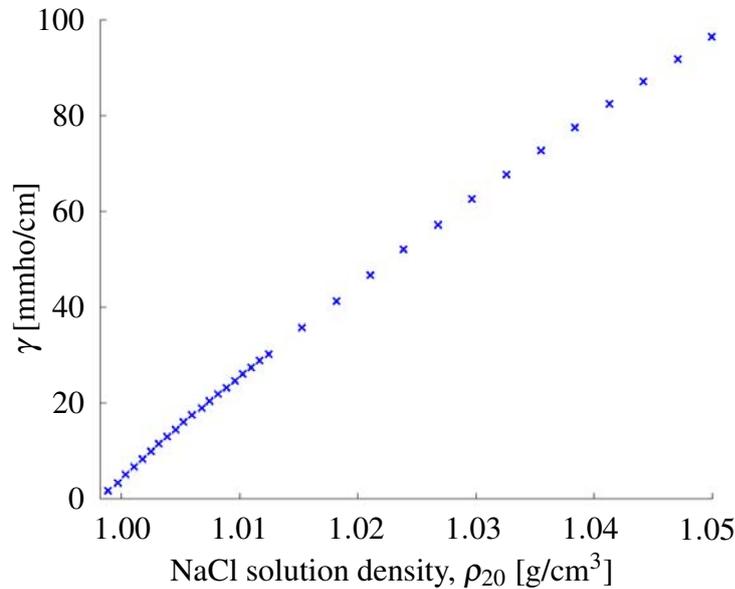


Fig. C.2 Plot of the electrical conductivity γ against the density ρ_{20} of NaCl solution at 20°C, covering the range of ρ_{20} used in our experiments. Data plotted from values tabulated by Weast et al. (1984).

C.2 Sampling the density stratification

We now discuss how the fluid occupying the tip and inner regions will vary as the probe traverses through and samples the stratification, and the effect this has on the conductivity measurements. When the probe samples the stratification, fluid is siphoned through the probe tip by a hydrostatic pressure head, controlled using a pinch valve (as discussed in Section 3.4). The siphoning volume flux $Q_S \sim 0.1 - 0.2 \text{ cm}^3/\text{s}$, and as the volume of the probe-tip region $V_T = \pi D_T^2 L_T / 4 \approx 2.83 \times 10^{-4} \text{ cm}^3$, fluid in the probe-tip region is entirely replaced by siphoning approximately every $t_S = 2 \text{ ms}$. As the probe is traversed at speed $U_S^\downarrow = 10 \text{ mm/s}$ when sampling and $U_S t_S \ll 1 \text{ mm}$, we can treat the replacement of fluid occupying the probe-tip region as instantaneous.

C.2.1 Probe-tip region

The siphoning of fluid by the probe induces a velocity field in the vicinity of the probe-tip entrance as fluid is drawn into the probe interior. The composition of fluid occupying the probe-tip region at a given instant in time will predominantly depend on the density stratification local to the probe-tip entrance, but will also carry some dependence on Q_S and U_S^\downarrow . In particular, the maximum vertical extent beneath the probe-tip entrance at which fluid

is drawn into the probe will set a precision limit δz for our conductivity measurements. It is important to verify that this precision limit is smaller than the vertical lengthscale of strongly density-stratified regions in our experiments.

To obtain an order-of-magnitude estimate for δz , we model the traversing and sampling of the probe as a point sink of strength Q_S traversed at speed U_S^\downarrow vertically downwards in a homogeneous fluid that is otherwise at rest. In the reference frame of the point sink, the sink-induced speed at a distance R is given by $U(R) = Q_S/(4\pi R^2)$.

Consider the Lagrangian trajectory of a parcel of fluid at height $Z(t)$ that is initially a distance H directly beneath the probe tip. The evolution equation for $Z(t)$ is given by

$$\frac{dZ}{dt} = \frac{Q_S}{4\pi(H - U_S^\downarrow t - Z)^2}. \quad (\text{C.3})$$

Equation (C.3) can be non-dimensionalised by taking $\zeta = Z/H$, $\tau = U_S^\downarrow t/H$ and $Q = Q_S/(4\pi H^2 U_S)$, giving

$$\frac{d\zeta}{d\tau} = \frac{Q}{(1 - \tau - \zeta)^2}, \quad (\text{C.4})$$

where $0 \leq \zeta, \tau \leq 1$. It is possible to obtain an analytic solution for Equation (C.4), giving an implicit equation for ζ ,

$$\tau = f(Q, \zeta) = \frac{\tan(\zeta/\sqrt{Q})}{\sqrt{Q} + \tan(\zeta/\sqrt{Q})} - \zeta. \quad (\text{C.5})$$

Figure C.3 (a) plots the solution for ζ as a function of τ . Unsurprisingly, the substantial majority of the elevation gained by the fluid parcel occurs in the near vicinity of the point sink. It is easy to show that the fluid parcel coincides with the point sink for $(\tau, \zeta) = (\tau_0, \zeta_0)$, where $\tau_0 + \zeta_0 = 1$. Using Equation (C.5), it follows that

$$(\tau_0, \zeta_0) = \left(1 - \frac{\pi}{2}\sqrt{Q}, \frac{\pi}{2}\sqrt{Q}\right).$$

For comparison to our experiments, we take $Q_S = 0.2\text{cm}^3/\text{s}$, $H = 500\text{mm}$ and $U_S^\downarrow = 10\text{mm/s}$, giving $Q \approx 6.4 \times 10^{-6}$. The dimensional total elevation height of the fluid parcel is then given by

$$Z_0 = H\zeta_0 = \frac{\sqrt{\pi}}{4} \sqrt{\frac{Q_S}{U_S^\downarrow}} \approx 2.0\text{mm}. \quad (\text{C.6})$$

The value obtained for Z_0 provides an estimate for the total elevation height of fluid in the tank due to the siphoning and traversing of the conductivity probe. From Equation (C.5), over 98% of the total fluid elevation is estimated to occur when the fluid is within 1mm of

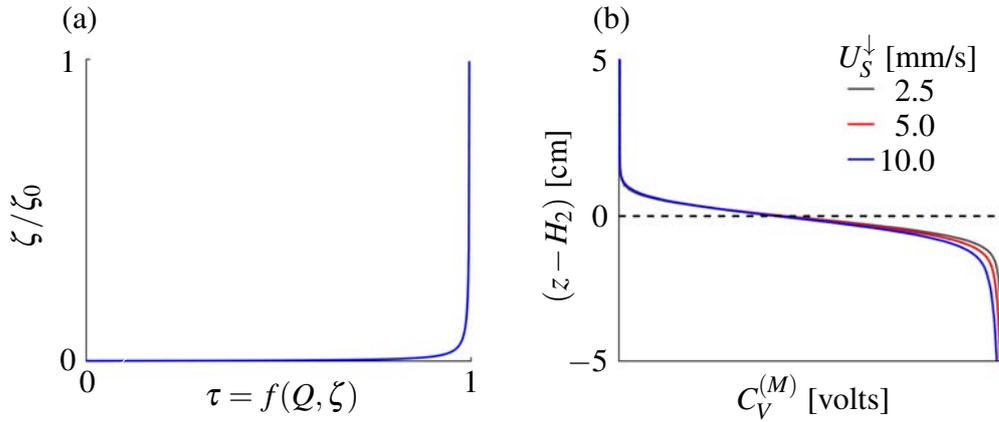


Fig. C.3 (a) From our point-sink model, plot of the dimensionless elevation height ζ normalised by the height ζ_0 obtained when the fluid parcel makes contact with the point sink. The plot is given against the non-dimensionalised time $\tau = f(Q, \zeta)$, given by Equation (C.5). (b) Plots of the probe-measured voltage profiles $C_V^{(M)}$, taken before the start of an experiment for three values of U_S^\downarrow . Plots are given against the height $(z - H_2)$ [cm] relative to the prescribed interface.

the probe-tip entrance. Based on this estimate, we can assume that for our experiments, the flow induced by the probe siphoning as it traverses has a negligible impact on the evolution of the stratification, and the composition of fluid occupying the probe-tip region at a given time is made up of fluid from an $O(1)$ mm-thick layer of fluid centred at the height of the probe-tip entrance.

The functional form of Equation (C.6) indicates a dependence of Z_0 on the siphoning volume flux and the probe traverse speed. For fixed $Q_S \approx 0.1 \text{ cm}^3/\text{s}$ and different values of U_S^\downarrow , Figure C.3 (b) plots conductivity profiles obtained when sampling a diffuse density interface of ~ 2 cm thickness. The differences between the profiles for $-1 < z - H_2 < 0$ can be attributed in part to the increase in Z_0 for decreasing U_S^\downarrow . This attribution is consistent with the profiles exhibiting higher values for the conductivity at a given height for lower values of U_S^\downarrow , due to the ability of the probe to withdraw fluid from further beneath the probe-tip entrance for the lower traverse speeds.

However, Figure C.3 (b) also exhibits discrepancies that cannot be attributed to selective withdrawal. These discrepancies are observed throughout the lower layer, both in the probe-measured lower-layer conductivity gradients dC_2/dz and the average conductivities $\overline{C_2}$. The conductivity gradients converge to 0 for decreasing U_S^\downarrow , suggesting that the lower layer is uniform rather than weakly stratified. This was important to verify to validate our tank-filling procedure. We also see that $\overline{C_2}$ increases for decreasing U_S^\downarrow . Our analysis above suggests that

the differences in $\overline{C_2}$ are large enough that they cannot be attributable to the fluid occupying the outer or tip regions of the circuit-closing fluid, and so they must be attributable to differences in the composition of fluid occupying the inner-probe region for different U_S^\downarrow .

C.2.2 Inner-probe region

Two time parameters play an important role in the composition of fluid occupying the inner-probe region as the probe traverses through and samples the stratification. Firstly, the time t_I taken for an inner-probe-region-equivalent volume of fluid passes through the inner-probe region. This is given by

$$t_I = \frac{\pi D_I^2 L_I}{4Q_S} \sim 0.1 - 0.2\text{s}, \quad (\text{C.7})$$

where the lengthscales D_I , L_I are illustrated in Figure C.1 (b), and Q_S varied between $0.1\text{cm}^3/\text{s}$ and $0.2\text{cm}^3/\text{s}$ across our experiments. Over time t_I , the probe is traversed $1 - 2\text{mm}$, which is similar to the measurement precision of evaluating the conductivity at a given height, set by the elevation height Z_0 estimated by Equation (C.6).

Secondly, the time t_F taken for the volume of fluid occupying the inner-probe region at a given time to be completely flushed through by the fluid sampled by the probe. While it is clear that the ratio $\tau_F = t_F/t_I \geq 1$ by volume conservation, the range of values τ_F might take in typical experimental setups making use of the conductivity probe is less clear. If $\tau_F \gg 1$, it is conceivable that a non-negligible disparity could arise between the probe-measured conductivity and the conductivity of the fluid at the probe-tip entrance. Such a disparity seems most likely to arise as the probe traverses through and samples a sharp density interface separated by two uniform layers, with conductivities γ_1 , γ_2 such that $\gamma_2/\gamma_1 \gg 1$. In this case, after the probe tip passes through the interface, if the typical inner-probe fluid conductivity $\gamma_I \sim \gamma_1$, then the inner-probe fluid could play a non-negligible role in setting the electrical conductivity of the probe circuit, given in Equation (C.2).

The plausibility of this scenario motivates considering the dynamics of the probe-siphoned fluid as it enters the inner-probe region so that an estimate on the range of values τ_F is able to take can be obtained. For the case of the probe traversing through and sampling a density interface, once the pass tip passes through the density interface, the inner-probe dynamics can be characterised by a dense fountain discharging vertically upwards into a radially-confined duct. Unfortunately, we were unable to find any previously published studies on such flows, and an investigation into the evolution of such a flow is not presented in this thesis.

For the case of the probe traversing through and sampling a uniform stratification, the inner-probe dynamics is similar to that investigated by Revuelta et al. (2002), who studied

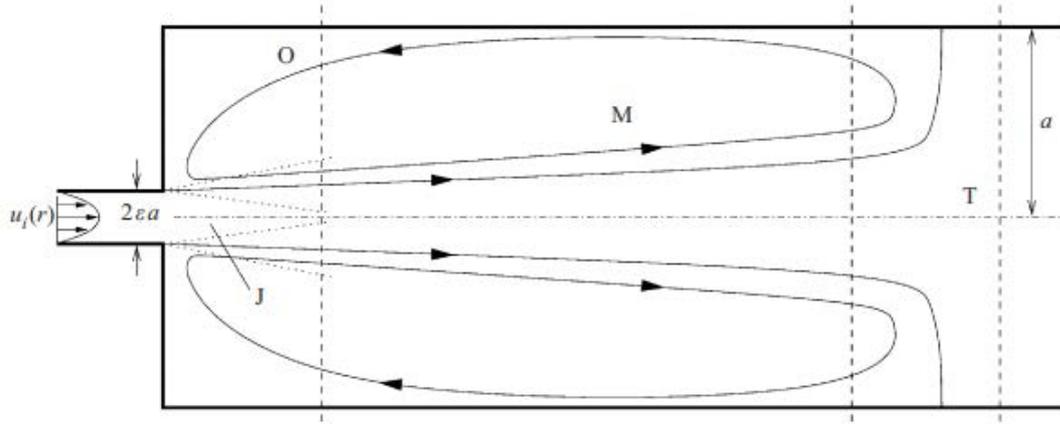


Fig. C.4 Schematic of a radially confined axisymmetric jet flow. Figure duplicated from Revuelta et al. (2002).

round laminar jets entering a radially confined duct. This flow depends mainly on two parameters: the ‘expansion ratio’ $\varepsilon = D_T/D_I$, and the ‘jet Reynolds number’, Re_I , which for our experiments are given by $\varepsilon \approx 0.13$ and

$$Re_I = \frac{U_T D_T}{\nu} \approx \frac{[1.41 - 2.83] \times (3 \times 10^{-4})}{10^{-6}} \sim 420 - 850. \quad (C.8)$$

Figure C.4 shows a schematic of the axisymmetric flow investigated by Revuelta et al. (2002). The schematic illustrates a ‘main recirculating region’, indicated by the closed streamlines on either side of the incoming jet. The recirculating region extends from the jet opening into the confined region a characteristic lengthscale l_R , which is given by

$$l_R = \frac{1}{2} D_T \varepsilon Re_I = 8.2 - 16.6 \text{ mm} \quad (C.9)$$

for the conductivity probe.

Note that as $l_R/D_I > 2$, the recirculating region in the probe (which we assume exists if uniform-density fluid is being siphoned) makes contact with the entirety of the portion of the inner electrode in contact with the inner-probe region. This indicates the importance of flushing through the recirculating fluid quickly as the probe traverses through a density interface, as the electrical current established between the two electrodes will pass through the recirculating region. Assuming a fixed expansion ratio, the jet Reynolds number will control the rate at which fluid from the recirculating region is entrained into the jet. For the conductivity probe, this will in turn control the flushing ratio τ_F .

We now return to the discrepancy in conductivity profiles presented in Figure C.3 (b) for $z - H_2 < -1$. Note that, for these cases, the volume flux $Q_S \approx 0.1 \text{ cm}^3/\text{s}$, giving $\text{Re}_I \approx 420$. As the probe traverses through and samples a density interface, if the probe fountain is unable to quickly flush through the recirculating upper-layer fluid occupying the inner-probe region with lower layer fluid, then there may be a period of time where the average conductivity of fluid occupying the inner-probe region is significantly lower than the conductivity of fluid occupying the probe-tip region. If the ratio $R_I/R_T \sim O(10^{-1})$ over this period of time, then the probe-measured conductivity will incur a non-negligible decrease due to the increased value of R_I . This is consistent with the conductivity profiles presented in Figure C.3 (b). Indeed, for smaller values of U_S^\downarrow , as the probe-tip reaches a given height in the lower layer, more time would have elapsed since the probe-tip passed through the density interface, giving more time to flush the upper-layer fluid occupying the inner-probe region with lower-layer fluid.

As $U_T = 4Q_S/(\pi D_T^2)$, prescription of Re_I for inner-probe-fluid entrainment is entirely controllable through Q_S for given D_T . Control over Q_S was made by the vertical positioning of the pinch valve relative to the free surface of the stratification. Most of our experiments were conducted with $Q_S \approx 0.2 \text{ cm}^3/\text{s}$, giving $\text{Re}_I \approx 850$. For these experiments, the probe-measured conductivity was almost uniform throughout the entirety of the lower layer, suggesting that for $\text{Re}_I \sim 850$ the probe jet is able to rapidly flush fluid occupying the inner-probe region.

The rest of our experiments were conducted with $Q_S \approx 0.1 \text{ cm}^3/\text{s}$, despite making no re-positioning of the pinch valve or changing the prescribed height of the free surface. The reduction in Q_S we attributed to a small blockage in the probe tip, though this was not confirmed. To validate our data from these experiments, some of these experiments were conducted using alternating traverse speeds for consecutive probe/thermistor traverses, to assess whether the reduced ability of the probe jet at flushing fluid in the inner-probe region had an impact on derived quantities, such as the potential energy of the system. Figures C.5 (a) and (b) respectively plot the salinity excess $S^{(n)}$ normalised by the mean value of $S^{(n)}$, and the increase in potential energy $\text{PE}^{(n)} - \text{PE}^{(0)}$, against the traverse number n . Both plots are derived using the probe-voltage-adjustment methodology outlined in Section 6.1.3. The plots indicate that, despite the increased probe-measured salinity excess for $U_S^\downarrow = 5 \text{ mm/s}$, the choice of U_S^\downarrow had a negligible impact on the measured rate of increase in potential energy.

C.3 Summary

We have carried out an analysis on the performance of the conductivity probe as it traverses through and samples a density stratification. The fluid occupying the probe-tip

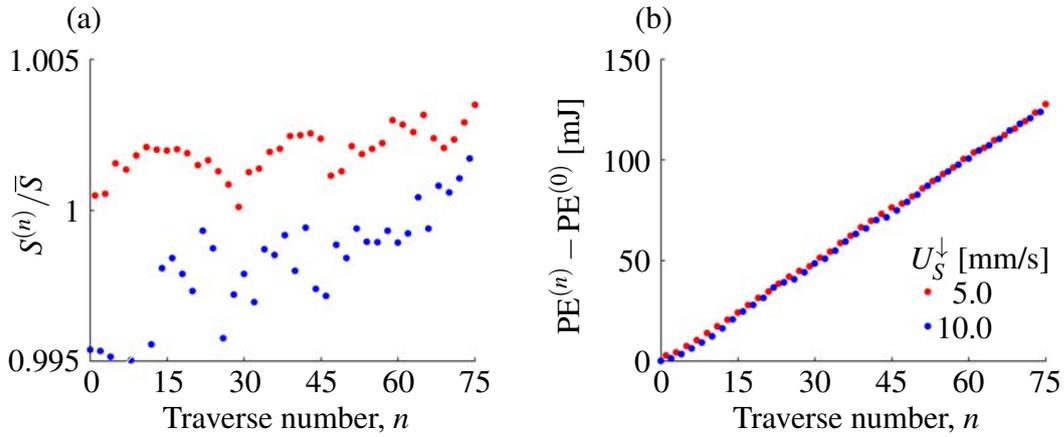


Fig. C.5 (a) Plot of the derived salinity excess $S^{(n)}$ normalised by the mean derived salinity excess \bar{S} against the traverse number n . (b) Plot of the potential energy gain $PE^{(n)} - PE^{(0)}$ in mJ against the traverse number n . Both plots are taken from a single experiment, where alternating probe/thermistor profiles are taken at traverse speeds $U_S^\downarrow = 10\text{mm/s}$ (blue) and $U_S^\downarrow = 5\text{mm/s}$ (red).

region was shown to consist of fluid occupying an $O(1)$ mm-thick layer of fluid centred around the probe-tip, demonstrating that the conductivity of fluid occupying the probe tip is similar to the horizontally-averaged conductivity at the height of the probe-tip opening, assuming the fluid is at rest. We argue that a non-negligible discrepancy between the probe-measured conductivity and the conductivity of fluid at the entrance to the probe-tip opening can arise due to the dynamics of the inner-probe region. This discrepancy can be reduced significantly by appropriate prescription of the jet Reynolds number, which can be controlled through the probe-siphoning volume flux.

Appendix D

Ring-induced heat fluxes

In this appendix we demonstrate that the ring-induced heat fluxes to the temperature stratification in our periodic mixing experiments could be treated as negligible. This appendix complements our discussion in Section 6.2 on the relative importance of different mechanisms that impacted the change in potential energy of the fluid stratification in our periodic mixing experiments.

Mechanical energy input: Using a Norbury model, the kinetic energy of each vortex ring is given by

$$\delta\text{KE} = C_{\text{KE}} \left(\frac{1}{2} \rho_0 U_p^2 \right) \left(\frac{4\pi}{3} \left(\frac{a}{2} \right)^3 \right) \quad (\text{D.1})$$

For $\rho_0 \approx 1002\text{kg/m}^3$, $U_p \approx 0.06\text{m/s}$, $a \approx 0.051\text{m}$ and $C_{\text{KE}} \approx 7$, we have $\delta\text{KE} \lesssim 10^{-3}\text{J}$. Therefore, firing 750 vortex rings, the total mechanical energy input from the vortex rings ME is in the range $0.5\text{J} < \text{ME} < 1\text{J}$.

Heat of dissipation: Suppose (as a worst case scenario) that all the kinetic energy was being converted to heat via viscous dissipation, and this heat contribution was being made only to a 1cm thick horizontal layer of the stratification. The volume of this layer is $V_{\text{layer}} = 4.8\text{l}$, which has a mass of approximately 4800g. Therefore the energy contribution per gram is

$$\frac{\text{ME}}{V_{\text{layer}}} = (0.5 - 1)/4800\text{J/g} = (1 - 2) \times 10^{-4}\text{J/g}$$

This would contribute to a temperature change of $\sim 10^{-4}/4.2\text{K} \sim 10^{-5} - 10^{-4}\text{K}$, which would make a negligible contribution to the density stratification.

Heat of dilution: For sodium chloride, the integral heat of dilution ΔH_d is the total enthalpy change per mole of NaCl in a solution which is then infinitely diluted. For a solution with initial state molality $m[\text{mol/kg}]$, this is related to the ‘apparent molal relative enthalpy’

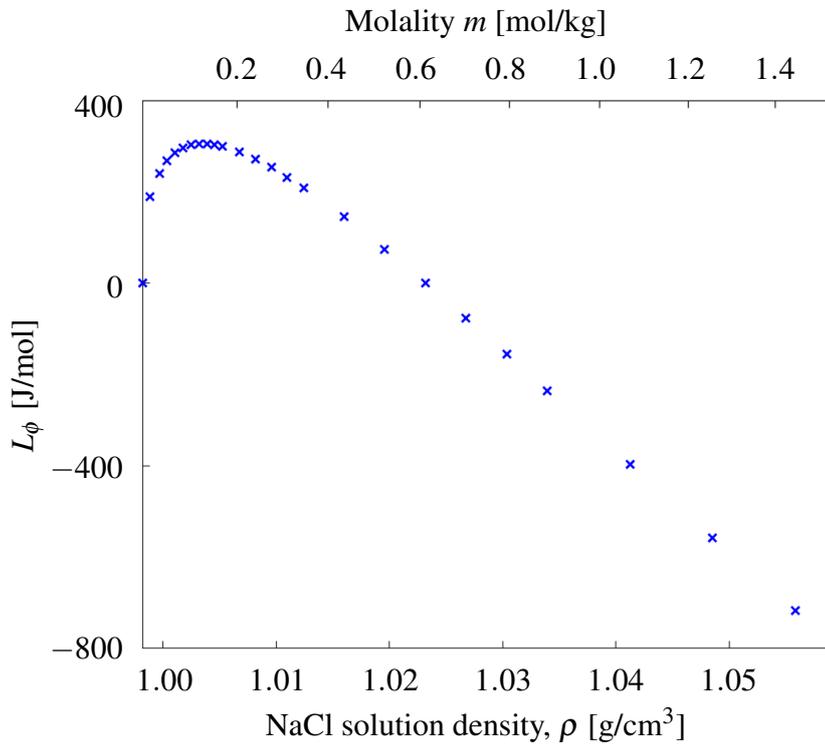


Fig. D.1 Plot of L_ϕ for NaCl solution at 20°C against fluid density and the molality. Data plotted from Tables 22A and 22B of Clarke and Glew (1985).

$L_\phi(m)$ [J/mol] by

$$\Delta H_d(m \rightarrow 0) = -L_\phi(m).$$

When a solution of initial molality m_1 is diluted to a final molality m_2 , the integral heat of dilution is given by

$$\Delta H_d(m_1 \rightarrow m_2) = L_\phi(m_2) - L_\phi(m_1). \quad (\text{D.2})$$

Values for L_ϕ at 20°C are given in Table 22 of Clarke and Glew (1985). The tabulated data is given in Figure D.1, plotting L_ϕ against both the fluid density and the ‘molality’, which is the number of moles of solute present in 1kg of solvent.

For low Richardson number mixing experiments, m ranges from 0.09 to 0.31; (weight % from 0.5 to 1.8), over which L_ϕ ranges from 230 to 300, and the gradient of L_ϕ is steepest at $m = 0.31$. Assuming the volume of fluid V_R entrained by the vortex ring can be approximated as a sphere of diameter a , then the volume is given by $V_R \approx 70\text{cm}^3$. Therefore the mass of the entrained fluid volume is about 70g.

Suppose 100g of NaCl solution is diluted from initial molality m_1 to final molality m_2 , where $\Delta H_d(m_1 \rightarrow m_2) = -70\text{J/mol}$. An upper bound on this heat change is given by

$$70\text{J/mol} \times 0.31\text{mol/kg} \times 0.1\text{kg} = 2.17\text{J},$$

which is a heat change of 0.0217J/g for a 100g solution, corresponding to a temperature change of $5 \times 10^{-3}\text{K}$. If this heat change is spread over a 1cm thick layer of our stratification, the temperature change is 10^{-4}K , which has a negligible effect on the density stratification.

